

Advanced Quantitative Research Methodology, Lecture Notes: **Introduction**¹

Gary King
GKing.Harvard.edu

January 25, 2016

¹©Copyright 2016 Gary King, All Rights Reserved.

Who takes this course?

Who takes this course?

- Most Gov Dept grad students doing empirical work, the 2nd course in their methods sequence (Gov2001)

Who takes this course?

- Most Gov Dept grad students doing empirical work, the 2nd course in their methods sequence (Gov2001)
- Undergrads (Gov1002)

Who takes this course?

- Most Gov Dept grad students doing empirical work, the 2nd course in their methods sequence (Gov2001)
- Undergrads (Gov1002)
- Grad students from other departments (Gov2001)

Who takes this course?

- Most Gov Dept grad students doing empirical work, the 2nd course in their methods sequence (Gov2001)
- Undergrads (Gov1002)
- Grad students from other departments (Gov2001)
- Non-Harvard students, visitors, faculty, & others (online through the Harvard Extension school StatE-200)

Who takes this course?

- Most Gov Dept grad students doing empirical work, the 2nd course in their methods sequence (Gov2001)
- Undergrads (Gov1002)
- Grad students from other departments (Gov2001)
- Non-Harvard students, visitors, faculty, & others (online through the Harvard Extension school StatE-200)
- $\rightsquigarrow$ Some of the best experiences here: getting to know people in other fields

How much math will you scare us with?

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth
 - Use rigorous statistical theory — when needed

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth
 - Use rigorous statistical theory — when needed
 - Insure we understand the intuition — always

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth
 - Use rigorous statistical theory — when needed
 - Insure we understand the intuition — always
 - Always traverse from theoretical foundations to practical applications

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth
 - Use rigorous statistical theory — when needed
 - Insure we understand the intuition — always
 - Always traverse from theoretical foundations to practical applications
 - Includes “how to” computation

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth
 - Use rigorous statistical theory — when needed
 - Insure we understand the intuition — always
 - Always traverse from theoretical foundations to practical applications
 - Includes “how to” computation
 - $\rightsquigarrow$ Fewer proofs, more concepts, better practical knowledge

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth
 - Use rigorous statistical theory — when needed
 - Insure we understand the intuition — always
 - Always traverse from theoretical foundations to practical applications
 - Includes “how to” computation
 - $\rightsquigarrow$ Fewer proofs, more concepts, better practical knowledge
- Do you have the background for this class?

How much math will you scare us with?

- All math requires two parts: **proof** and **concepts & intuition**
- Different classes emphasize:
 - **Baby Stats**: dumbed down proofs, vague intuition
 - **Math Stats**: rigorous mathematical proofs
 - **Practical Stats**: deep concepts and intuition, proofs when needed
 - Goal: how to do empirical research, in depth
 - Use rigorous statistical theory — when needed
 - Insure we understand the intuition — always
 - Always traverse from theoretical foundations to practical applications
 - Includes “how to” computation
 - $\rightsquigarrow$ Fewer proofs, more concepts, better practical knowledge
- Do you have the background for this class? **A Test: What's this?**

$$b = (X'X)^{-1}X'y$$

What's this course about?

What's this course about?

- Specific statistical methods for many research problems

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know
- How to write a publishable scholarly paper

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know
- How to write a publishable scholarly paper
- All the practical tools of research — theory, applications, simulation, programming, word processing, plumbing, whatever is useful

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know
- How to write a publishable scholarly paper
- All the practical tools of research — theory, applications, simulation, programming, word processing, plumbing, whatever is useful
- $\rightsquigarrow$ Outline and class materials:

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know
- How to write a publishable scholarly paper
- All the practical tools of research — theory, applications, simulation, programming, word processing, plumbing, whatever is useful
- $\rightsquigarrow$ Outline and class materials:

j.mp/G2001



What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know
- How to write a publishable scholarly paper
- All the practical tools of research — theory, applications, simulation, programming, word processing, plumbing, whatever is useful
- $\rightsquigarrow$ Outline and class materials:

j.mp/G2001



- The syllabus gives topics, not a weekly plan.

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know
- How to write a publishable scholarly paper
- All the practical tools of research — theory, applications, simulation, programming, word processing, plumbing, whatever is useful
- $\rightsquigarrow$ Outline and class materials:

j.mp/G2001



- The syllabus gives topics, not a weekly plan.
- We will go as fast as possible subject to everyone following along

What's this course about?

- Specific statistical methods for many research problems
 - How to learn (or create) new methods
 - Inference: Using facts you know to learn about facts you don't know
- How to write a publishable scholarly paper
- All the practical tools of research — theory, applications, simulation, programming, word processing, plumbing, whatever is useful
- $\rightsquigarrow$ Outline and class materials:

j.mp/G2001



- The syllabus gives topics, not a weekly plan.
- We will go as fast as possible subject to everyone following along
- We cover different amounts of material each week

Requirements

Requirements

① Weekly assignments

Requirements

- 1 Weekly assignments
 - Readings (via Perusall), assignments

Requirements

① Weekly assignments

- Readings (via Perusall), assignments
- Take notes, read carefully, don't skip equations

Requirements

- 1 Weekly assignments
 - Readings (via Perusall), assignments
 - Take notes, read carefully, don't skip equations
- 2 One “publishable” coauthored paper. (Easier than you think!)

Requirements

- ① Weekly assignments
 - Readings (via Perusall), assignments
 - Take notes, read carefully, don't skip equations
- ② One “publishable” coauthored paper. (Easier than you think!)
 - Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards

Requirements

① Weekly assignments

- Readings (via Perusall), assignments
- Take notes, read carefully, don't skip equations

② One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications

Requirements

① Weekly assignments

- Readings (via Perusal), assignments
- Take notes, read carefully, don't skip equations

② One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.

Requirements

① Weekly assignments

- Readings (via Perusal^l), assignments
- Take notes, read carefully, don't skip equations

② One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”

Requirements

① Weekly assignments

- Readings (via Perusal^l), assignments
- Take notes, read carefully, don't skip equations

② One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”
- You won't be alone: you'll work with each other and us

Requirements

① Weekly assignments

- Readings (via Perusall), assignments
- Take notes, read carefully, don't skip equations

② One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”
- You won't be alone: you'll work with each other and us

③ Participation and collaboration:

Requirements

① Weekly assignments

- Readings (via Perusal), assignments
- Take notes, read carefully, don't skip equations

② One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”
- You won't be alone: you'll work with each other and us

③ Participation and collaboration:

- Do assignments in groups: “Cheating” is encouraged, so long as you write up your work on your own.

Requirements

1 Weekly assignments

- Readings (via Perusal), assignments
- Take notes, read carefully, don't skip equations

2 One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”
- You won't be alone: you'll work with each other and us

3 Participation and collaboration:

- Do assignments in groups: “Cheating” is encouraged, so long as you write up your work on your own.
- Participating in a conversation >> Evesdropping

Requirements

1 Weekly assignments

- Readings (via Perusal), assignments
- Take notes, read carefully, don't skip equations

2 One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”
- You won't be alone: you'll work with each other and us

3 Participation and collaboration:

- Do assignments in groups: “Cheating” is encouraged, so long as you write up your work on your own.
- Participating in a conversation >> Evesdropping
- Use collaborative learning tools (we'll introduce)

Requirements

① Weekly assignments

- Readings (via Perusal), assignments
- Take notes, read carefully, don't skip equations

② One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”
- You won't be alone: you'll work with each other and us

③ Participation and collaboration:

- Do assignments in groups: “Cheating” is encouraged, so long as you write up your work on your own.
- Participating in a conversation >> Evesdropping
- Use collaborative learning tools (we'll introduce)
- Build class camaraderie: prepare, participate, help others

Requirements

1 Weekly assignments

- Readings (via Perusall), assignments
- Take notes, read carefully, don't skip equations

2 One “publishable” coauthored paper. (Easier than you think!)

- Many class papers have been published, presented at conferences, become dissertations or senior theses, and won many awards
- Undergrads have often had professional journal publications
- Draft submission and replication exercise helps a lot.
- See “Publication, Publication”
- You won't be alone: you'll work with each other and us

3 Participation and collaboration:

- Do assignments in groups: “Cheating” is encouraged, so long as you write up your work on your own.
- Participating in a conversation >> Evesdropping
- Use collaborative learning tools (we'll introduce)
- Build class camaraderie: prepare, participate, help others

4 Focus, like I will, on learning (not grades): Especially when we work on papers, I will treat you like a colleague, not a student

Got Questions?

Got Questions?

- Participate in Collaborative Annotation, Chat, Email

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-ter-

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-ter-rupt

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-ter-rupt me as often as necessary

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-ter-rupt me as often as necessary
- (Got a dumb question? Assume you are the smartest person in class and you eventually will be!)

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-ter-rupt me as often as necessary
- (Got a dumb question? Assume you are the smartest person in class and you eventually will be!)
- When are Gary's office hours?

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-ter-rupt me as often as necessary
- (Got a dumb question? Assume you are the smartest person in class and you eventually will be!)
- When are Gary's office hours?
 - (Big secret: The point of office hours is to prevent students from visiting at other times)

Got Questions?

- Participate in Collaborative Annotation, Chat, Email
- Browse archive of previous years' emails (Note which now-famous scholar is asking the question!)
- In-ter-rupt me as often as necessary
- (Got a dumb question? Assume you are the smartest person in class and you eventually will be!)
- When are Gary's office hours?
 - (Big secret: The point of office hours is to prevent students from visiting at other times)
 - Come whenever you like; if you can't find me or I'm in a meeting, come back, talk to my assistant in the office next to me, or email any time

What is the field of statistics?

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state*-istics” (circa 1662)

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state*-istics” (circa 1662)
- **A new field:**

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state*-istics” (circa 1662)
- **A new field:**
 - **Mid-1930s:** Experiments and random assignment

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state*-istics” (circa 1662)
- **A new field:**
 - **Mid-1930s:** Experiments and random assignment
 - **1950s:** The modern theory of inference

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state*-istics” (circa 1662)
- **A new field:**
 - **Mid-1930s:** Experiments and random assignment
 - **1950s:** The modern theory of inference
 - **In your lifetime:** Modern causal inference

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state*-istics” (circa 1662)
- **A new field:**
 - **Mid-1930s:** Experiments and random assignment
 - **1950s:** The modern theory of inference
 - **In your lifetime:** Modern causal inference
 - **Since you started school:** Part of a monumental societal change, “big data”, and the march of quantification through academic, professional, commercial, government and policy fields.

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state*-istics” (circa 1662)
- **A new field:**
 - **Mid-1930s:** Experiments and random assignment
 - **1950s:** The modern theory of inference
 - **In your lifetime:** Modern causal inference
 - **Since you started school:** Part of a monumental societal change, “big data”, and the march of quantification through academic, professional, commercial, government and policy fields.
- **The number of new methods is increasing fast**

What is the field of statistics?

- **An old field:** statistics originates in the study of politics and government: “*state-istics*” (circa 1662)
- **A new field:**
 - **Mid-1930s:** Experiments and random assignment
 - **1950s:** The modern theory of inference
 - **In your lifetime:** Modern causal inference
 - **Since you started school:** Part of a monumental societal change, “big data”, and the march of quantification through academic, professional, commercial, government and policy fields.
- **The number of new methods is increasing fast**
- **Most important methods originate *outside* the discipline of statistics** (random assignment, experimental design, survey research, machine learning, MCMC methods, ...). Statistics: abstracts, proves formal properties, generalizes, and distributes results back out.

What is the subfield of political methodology?

What is the subfield of political methodology?

- A **relative of other social science methods subfields** — econometrics, psychological statistics, biostatistics, chemometrics, sociological methodology, cliometrics, etc. — with many cross-field connections

What is the subfield of political methodology?

- A **relative of other social science methods subfields** — econometrics, psychological statistics, biostatistics, chemometrics, sociological methodology, cliometrics, etc. — with many cross-field connections
- **Heavily interdisciplinary**, like the discipline of political science

What is the subfield of political methodology?

- A **relative of other social science methods subfields** — econometrics, psychological statistics, biostatistics, chemometrics, sociological methodology, cliometrics, etc. — with many cross-field connections
- **Heavily interdisciplinary**, like the discipline of political science
- The **crossroads for other disciplines**, and a great place to learn about methods broadly.

What is the subfield of political methodology?

- A **relative of other social science methods subfields** — econometrics, psychological statistics, biostatistics, chemometrics, sociological methodology, cliometrics, etc. — with many cross-field connections
- **Heavily interdisciplinary**, like the discipline of political science
- The **crossroads for other disciplines**, and a great place to learn about methods broadly.
- **Second largest APSA section**, essentially **no unemployment**

What is the subfield of political methodology?

- A **relative of other social science methods subfields** — econometrics, psychological statistics, biostatistics, chemometrics, sociological methodology, cliometrics, etc. — with many cross-field connections
- **Heavily interdisciplinary**, like the discipline of political science
- The **crossroads for other disciplines**, and a great place to learn about methods broadly.
- **Second largest APSA section**, essentially **no unemployment**
- Part of a massive **change in the evidence base of the social sciences**:
(a) surveys, (b) end of period government stats, and (c) one-off studies of people, places, or events ~> numerous new types and huge quantities of (big) data

What is the subfield of political methodology?

- A **relative of other social science methods subfields** — econometrics, psychological statistics, biostatistics, chemometrics, sociological methodology, cliometrics, etc. — with many cross-field connections
- **Heavily interdisciplinary**, like the discipline of political science
- The **crossroads for other disciplines**, and a great place to learn about methods broadly.
- **Second largest APSA section**, essentially **no unemployment**
- Part of a massive **change in the evidence base of the social sciences**: (a) surveys, (b) end of period government stats, and (c) one-off studies of people, places, or events $\rightsquigarrow$ numerous new types and huge quantities of (big) data
- Helped give rise to **data science**, encompassing statistics, CS, and a substantive field, and **computational social science**

Course strategy

Course strategy

- We could teach you the latest and greatest methods,

Course strategy

- We could teach you the latest and greatest methods, but when you graduate **they will be old**

Course strategy

- We could teach you the latest and greatest methods, but when you graduate **they will be old**
- We could teach you all the methods that might prove useful during your career,

Course strategy

- We could teach you the latest and greatest methods, but when you graduate **they will be old**
- We could teach you all the methods that might prove useful during your career, but when you graduate **you will be old**

Course strategy

- We could teach you the latest and greatest methods, but when you graduate **they will be old**
- We could teach you all the methods that might prove useful during your career, but when you graduate **you will be old**
- Instead, we teach you the **fundamentals**, the underlying **theory of inference**, from which statistical models are developed:

Course strategy

- We could teach you the latest and greatest methods, but when you graduate **they will be old**
- We could teach you all the methods that might prove useful during your career, but when you graduate **you will be old**
- Instead, we teach you the **fundamentals**, the underlying **theory of inference**, from which statistical models are developed:
 - We will **reinvent** existing methods by creating them from scratch.

Course strategy

- We could teach you the latest and greatest methods, but when you graduate **they will be old**
- We could teach you all the methods that might prove useful during your career, but when you graduate **you will be old**
- Instead, we teach you the **fundamentals**, the underlying **theory of inference**, from which statistical models are developed:
 - We will **reinvent** existing methods by creating them from scratch.
 - We will learn: its easy to **invent** new methods too, when needed.

Course strategy


- We could teach you the latest and greatest methods, but when you graduate **they will be old**
- We could teach you all the methods that might prove useful during your career, but when you graduate **you will be old**
- Instead, we teach you the **fundamentals**, the underlying **theory of inference**, from which statistical models are developed:
 - We will **reinvent** existing methods by creating them from scratch.
 - We will learn: its easy to **invent** new methods too, when needed.
 - The fundamentals help us pick up new methods created by others.

Course strategy


- We could teach you the latest and greatest methods, but when you graduate **they will be old**
- We could teach you all the methods that might prove useful during your career, but when you graduate **you will be old**
- Instead, we teach you the **fundamentals**, the underlying **theory of inference**, from which statistical models are developed:
 - We will **reinvent** existing methods by creating them from scratch.
 - We will learn: its easy to **invent** new methods too, when needed.
 - The fundamentals help us pick up new methods created by others.
- This helps us separate the conventions from underlying statistical theory. (How to get an F in Econometrics: follow advice from Psychometrics. Works in reverse too, even when the foundations are identical.)

e.g.,: How to fit a line to a scatterplot?

e.g.,: How to fit a line to a scatterplot?




e.g.,: How to fit a line to a scatterplot?




- visually (tends to be principle components)

e.g.,: How to fit a line to a scatterplot?




- visually (tends to be principle components)
- a rule (least squares, least absolute deviations, etc.)

e.g.,: How to fit a line to a scatterplot?



- visually (tends to be principle components)
- a rule (least squares, least absolute deviations, etc.)
- criteria (unbiasedness, efficiency, sufficiency, admissibility, etc.)


e.g.,: How to fit a line to a scatterplot?




- visually (tends to be principle components)
- a rule (least squares, least absolute deviations, etc.)
- criteria (unbiasedness, efficiency, sufficiency, admissibility, etc.)
- **from a theory of inference, and for a substantive purpose** (like causal estimation, prediction, etc.)

Software options

Software options




Software options



- We'll use **R** — a free open source program, a commons, a movement

Software options



- We'll use **R** — a free open source program, a commons, a movement
- **Zelig**, which simplifies R and helps you up the steep slope fast (see ZeligProject.org)


Goal: Quantities of Interest

Goal: Quantities of Interest

Summaries vs. Inference (using facts you know to learn facts you don't know)

Goal: Quantities of Interest

Summaries vs. Inference (using facts you know to learn facts you don't know)



What is this?

What is this?



What is this?



- Now you know what a **model** is. (Its an abstraction.)

What is this?



- Now you know what a **model** is. (Its an abstraction.)
- Is this model true?

What is this?



- Now you know what a **model** is. (Its an abstraction.)
- Is this model true?
- Are models ever true or false?

What is this?



- Now you know what a **model** is. (Its an abstraction.)
- Is this model true?
- Are models ever true or false?
- Are models ever realistic or not?

What is this?



- Now you know what a **model** is. (Its an abstraction.)
- Is this model true?
- Are models ever true or false?
- Are models ever realistic or not?
- Are models ever useful or not?

What is this?



- Now you know what a **model** is. (Its an abstraction.)
- Is this model true?
- Are models ever true or false?
- Are models ever realistic or not?
- Are models ever useful or not?
- Models of dirt on airplanes, vs models of aerodynamics

Statistical Models: Variable Definitions

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable
 - Y is $n \times 1$.

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable
 - Y is $n \times 1$.
 - y_i , a number (after we know it)

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable
 - Y is $n \times 1$.
 - y_i , a number (after we know it)
 - Y_i , a random variable (before we know it)

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable
 - Y is $n \times 1$.
 - y_i , a number (after we know it)
 - Y_i , a random variable (before we know it)
 - Commonly misunderstood: a “dependent variable” can be

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable
 - Y is $n \times 1$.
 - y_i , a number (after we know it)
 - Y_i , a random variable (before we know it)
 - Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable
 - Y is $n \times 1$.
 - y_i , a number (after we know it)
 - Y_i , a random variable (before we know it)
 - Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set
 - the random variable *for each unit i* .

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable

- Y is $n \times 1$.
- y_i , a number (after we know it)
- Y_i , a random variable (before we know it)
- Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set
 - the random variable *for each unit i* .

- Explanatory variables

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable

- Y is $n \times 1$.
- y_i , a number (after we know it)
- Y_i , a random variable (before we know it)
- Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set
 - the random variable *for each unit i* .

- Explanatory variables

- aka “covariates,” “independent,” or “exogenous” variables

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable

- Y is $n \times 1$.
- y_i , a number (after we know it)
- Y_i , a random variable (before we know it)
- Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set
 - the random variable *for each unit i* .

- Explanatory variables

- aka “covariates,” “independent,” or “exogenous” variables
- $X = \{x_{ij}\}$ is $n \times k$ (observations by variables)

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable

- Y is $n \times 1$.
- y_i , a number (after we know it)
- Y_i , a random variable (before we know it)
- Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set
 - the random variable *for each unit i* .

- Explanatory variables

- aka “covariates,” “independent,” or “exogenous” variables
- $X = \{x_{ij}\}$ is $n \times k$ (observations by variables)
- A set of columns (variables): $X = \{x_1 \dots, x_k\}$

Statistical Models: Variable Definitions

- **Dependent (or “outcome”) variable**

- Y is $n \times 1$.
- y_i , a number (after we know it)
- Y_i , a random variable (before we know it)
- Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set
 - the random variable *for each unit i* .

- **Explanatory variables**

- aka “covariates,” “independent,” or “exogenous” variables
- $X = \{x_{ij}\}$ is $n \times k$ (observations by variables)
- A set of columns (variables): $X = \{x_1, \dots, x_k\}$
- Row (observation) i : $x_i = \{x_{i1}, \dots, x_{ik}\}$

Statistical Models: Variable Definitions

- Dependent (or “outcome”) variable

- Y is $n \times 1$.
- y_i , a number (after we know it)
- Y_i , a random variable (before we know it)
- Commonly misunderstood: a “dependent variable” can be
 - a column of numbers in your data set
 - the random variable *for each unit i* .

- Explanatory variables

- aka “covariates,” “independent,” or “exogenous” variables
- $X = \{x_{ij}\}$ is $n \times k$ (observations by variables)
- A set of columns (variables): $X = \{x_1, \dots, x_k\}$
- Row (observation) i : $x_i = \{x_{i1}, \dots, x_{ik}\}$
- X is fixed (not random).

Equivalent Linear Regression Notation

Equivalent Linear Regression Notation

- Standard version

Equivalent Linear Regression Notation

- Standard version

$$Y_i = x_i\beta + \epsilon_i \quad = \text{systematic} + \text{stochastic}$$

Equivalent Linear Regression Notation

- Standard version

$$Y_i = x_i\beta + \epsilon_i \quad = \text{systematic} + \text{stochastic}$$
$$\epsilon_i \sim f_N(0, \sigma^2)$$

Equivalent Linear Regression Notation

- Standard version

$$Y_i = x_i\beta + \epsilon_i \quad = \text{systematic} + \text{stochastic}$$
$$\epsilon_i \sim f_N(0, \sigma^2)$$

- Alternative version

Equivalent Linear Regression Notation

- Standard version

$$Y_i = x_i\beta + \epsilon_i \quad = \text{systematic} + \text{stochastic}$$
$$\epsilon_i \sim f_N(0, \sigma^2)$$

- Alternative version

$$Y_i \sim f_N(\mu_i, \sigma^2) \quad \text{stochastic}$$

Equivalent Linear Regression Notation

- Standard version


$$Y_i = x_i\beta + \epsilon_i \quad = \text{systematic} + \text{stochastic}$$
$$\epsilon_i \sim f_N(0, \sigma^2)$$

- Alternative version

$$Y_i \sim f_N(\mu_i, \sigma^2) \quad \text{stochastic}$$
$$\mu_i = x_i\beta \quad \text{systematic}$$


Understanding the Alternative Regression Notation

Understanding the Alternative Regression Notation



where $\mu_i = X_i\beta$.

Understanding the Alternative Regression Notation



where $\mu_i = X_i\beta$.

A Test: Is a histogram of y a test of normality?

Generalized Alternative Notation for Most Models

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

where

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

where

Y_i random outcome variable

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

where

Y_i random outcome variable

$f(\cdot)$ probability density

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

where

Y_i random outcome variable

$f(\cdot)$ probability density

θ_i a systematic feature of the density that varies over i

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha) \quad \text{stochastic}$$
$$\theta_i = g(X_i, \beta) \quad \text{systematic}$$

where

Y_i random outcome variable

$f(\cdot)$ probability density

θ_i a systematic feature of the density that varies over i

α ancillary parameter (feature of the density constant over i)

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

where

Y_i random outcome variable

$f(\cdot)$ probability density

θ_i a systematic feature of the density that varies over i

α ancillary parameter (feature of the density constant over i)

$g(\cdot)$ functional form

Generalized Alternative Notation for Most Models

$$\begin{array}{ll} Y_i \sim f(\theta_i, \alpha) & \text{stochastic} \\ \theta_i = g(\textcolor{red}{X}_i, \beta) & \text{systematic} \end{array}$$

where

Y_i random outcome variable

$f(\cdot)$ probability density

θ_i a systematic feature of the density that varies over i

α ancillary parameter (feature of the density constant over i)

$g(\cdot)$ functional form

$\textcolor{red}{X}_i$ explanatory variables

Generalized Alternative Notation for Most Models

$$Y_i \sim f(\theta_i, \alpha) \quad \text{stochastic}$$
$$\theta_i = g(X_i, \beta) \quad \text{systematic}$$

where

Y_i random outcome variable

$f(\cdot)$ probability density

θ_i a systematic feature of the density that varies over i

α ancillary parameter (feature of the density constant over i)

$g(\cdot)$ functional form

X_i explanatory variables

β effect parameters

Forms of Uncertainty

Forms of Uncertainty

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

Forms of Uncertainty

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

Forms of Uncertainty

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

- **Estimation uncertainty:** Lack of knowledge of β and α . Vanishes as n gets larger.

Forms of Uncertainty

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic

$$\theta_i = g(X_i, \beta)$$

systematic

- **Estimation uncertainty**: Lack of knowledge of β and α . Vanishes as n gets larger.
- **Fundamental uncertainty**: Represented by the stochastic component. Exists no matter what the researcher does; no matter how large n is.

Forms of Uncertainty

$$Y_i \sim f(\theta_i, \alpha)$$

stochastic


$$\theta_i = g(X_i, \beta)$$

systematic


- **Estimation uncertainty**: Lack of knowledge of β and α . Vanishes as n gets larger.
- **Fundamental uncertainty**: Represented by the stochastic component. Exists no matter what the researcher does; no matter how large n is.
- (A Test: If you know the model, is $R^2 = 1$?)

Systematic Components: Examples

Systematic Components: Examples




Systematic Components: Examples




- $E(Y_i) \equiv \mu_i = X_i\beta = \beta_0 + \beta_1X_{1i} + \cdots + \beta_kX_{ki}$

Systematic Components: Examples




- $E(Y_i) \equiv \mu_i = X_i\beta = \beta_0 + \beta_1 X_{1i} + \dots + \beta_k X_{ki}$
- $\Pr(Y_i = 1) \equiv \pi_i = \frac{1}{1+e^{-x_i\beta}}$

Systematic Components: Examples




- $E(Y_i) \equiv \mu_i = X_i\beta = \beta_0 + \beta_1 X_{1i} + \dots + \beta_k X_{ki}$
- $\Pr(Y_i = 1) \equiv \pi_i = \frac{1}{1 + e^{-x_i\beta}}$
- $V(Y_i) \equiv \sigma_i^2 = e^{x_i\beta}$

Systematic Components: Examples




- $E(Y_i) \equiv \mu_i = X_i\beta = \beta_0 + \beta_1 X_{1i} + \cdots + \beta_k X_{ki}$
- $\Pr(Y_i = 1) \equiv \pi_i = \frac{1}{1+e^{-x_i\beta}}$
- $V(Y_i) \equiv \sigma_i^2 = e^{x_i\beta}$
- Interpretation:

Systematic Components: Examples




- $E(Y_i) \equiv \mu_i = X_i\beta = \beta_0 + \beta_1 X_{1i} + \cdots + \beta_k X_{ki}$
- $\Pr(Y_i = 1) \equiv \pi_i = \frac{1}{1+e^{-x_i\beta}}$
- $V(Y_i) \equiv \sigma_i^2 = e^{x_i\beta}$
- Interpretation:
 - Each is a **class of functional forms**

Systematic Components: Examples



- $E(Y_i) \equiv \mu_i = X_i\beta = \beta_0 + \beta_1X_{1i} + \cdots + \beta_kX_{ki}$
- $\Pr(Y_i = 1) \equiv \pi_i = \frac{1}{1+e^{-x_i\beta}}$
- $V(Y_i) \equiv \sigma_i^2 = e^{x_i\beta}$
- Interpretation:
 - Each is a **class of functional forms**
 - Set β and it picks out one **member of the class**

Systematic Components: Examples



- $E(Y_i) \equiv \mu_i = X_i\beta = \beta_0 + \beta_1 X_{1i} + \cdots + \beta_k X_{ki}$
- $\Pr(Y_i = 1) \equiv \pi_i = \frac{1}{1 + e^{-x_i\beta}}$
- $V(Y_i) \equiv \sigma_i^2 = e^{x_i\beta}$
- Interpretation:
 - Each is a **class of functional forms**
 - Set β and it picks out one **member of the class**
 - β in each is an “effect parameter” vector, with different meaning

Systematic Components: Examples

Systematic Components: Examples

- We will:

Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)

Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class

Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class
 - Be uncertain (Because of sampling, measurement, & fundamental uncertainties) about:

Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class
 - Be uncertain (Because of sampling, measurement, & fundamental uncertainties) about:
 - the family $\rightsquigarrow$ model dependence

Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class
 - Be uncertain (Because of sampling, measurement, & fundamental uncertainties) about:
 - the family $\rightsquigarrow$ model dependence
 - the member of the chosen family $\rightsquigarrow$ sampling error

Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class
 - Be uncertain (Because of sampling, measurement, & fundamental uncertainties) about:
 - the family $\rightsquigarrow$ model dependence
 - the member of the chosen family $\rightsquigarrow$ sampling error
- If we choose the family of functional forms wrong, we:

Systematic Components: Examples



- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class
 - Be uncertain (Because of sampling, measurement, & fundamental uncertainties) about:
 - the family $\rightsquigarrow$ model dependence
 - the member of the chosen family $\rightsquigarrow$ sampling error
- If we choose the family of functional forms wrong, we:
 - Have specification error, and potentially bias

Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class
 - Be uncertain (Because of sampling, measurement, & fundamental uncertainties) about:
 - the family $\rightsquigarrow$ model dependence
 - the member of the chosen family $\rightsquigarrow$ sampling error
- If we choose the family of functional forms wrong, we:
 - Have specification error, and potentially bias
 - Still get the best [linear,logit,etc] approximation to the correct functional form.


Systematic Components: Examples

- We will:
 - Assume a class of functional forms from theory (each form is flexible and maps out many potential relationships)
 - Estimate β from data to choose a member of the class
 - Be uncertain (Because of sampling, measurement, & fundamental uncertainties) about:
 - the family $\rightsquigarrow$ model dependence
 - the member of the chosen family $\rightsquigarrow$ sampling error
- If we choose the family of functional forms wrong, we:
 - Have specification error, and potentially bias
 - Still get the best [linear,logit,etc] approximation to the correct functional form.
 - May be close or far from the truth:




Overview of Stochastic Components

Overview of Stochastic Components




- Normal — continuous, unimodal, symmetric, unbounded

Overview of Stochastic Components




- Normal — continuous, unimodal, symmetric, unbounded
- Log-normal — continuous, unimodal, skewed, bounded from below by zero

Overview of Stochastic Components



- Normal — continuous, unimodal, symmetric, unbounded
- Log-normal — continuous, unimodal, skewed, bounded from below by zero
- Bernoulli — discrete, binary outcomes

Overview of Stochastic Components



- Normal — continuous, unimodal, symmetric, unbounded
- Log-normal — continuous, unimodal, skewed, bounded from below by zero
- Bernoulli — discrete, binary outcomes
- Poisson — discrete, countably infinite on the nonnegative integers

Choosing systematic and stochastic components

Choosing systematic and stochastic components

- If one is bounded, so is the other

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set
- **What if we don't know the DGP (& we usually don't)?**

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set
- **What if we don't know the DGP (& we usually don't)?**
 - **The problem:** model dependence

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set
- **What if we don't know the DGP (& we usually don't)?**
 - **The problem:** model dependence
 - **Our first approach:** make “reasonable” assumptions and check fit (& other observable implications of the assumptions)

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set
- **What if we don't know the DGP (& we usually don't)?**
 - **The problem:** model dependence
 - **Our first approach:** make “reasonable” assumptions and check fit (& other observable implications of the assumptions)
 - **Later:**

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set
- **What if we don't know the DGP (& we usually don't)?**
 - **The problem:** model dependence
 - **Our first approach:** make “reasonable” assumptions and check fit (& other observable implications of the assumptions)
 - **Later:**
 - Avoid it: relax assumptions (functional form, distribution, etc)

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set
- **What if we don't know the DGP (& we usually don't)?**
 - **The problem:** model dependence
 - **Our first approach:** make “reasonable” assumptions and check fit (& other observable implications of the assumptions)
 - **Later:**
 - Avoid it: relax assumptions (functional form, distribution, etc)
 - Detect remaining model dependence

Choosing systematic and stochastic components

- If one is bounded, so is the other
- If the stochastic component is bounded, the systematic component must be globally nonlinear (tho possibly locally linear)
- All modeling decisions are about the **data generation process** — how the information made its way from the world (including how the world produced the data) to your data set
- **What if we don't know the DGP (& we usually don't)?**
 - **The problem:** model dependence
 - **Our first approach:** make “reasonable” assumptions and check fit (& other observable implications of the assumptions)
 - **Later:**
 - Avoid it: relax assumptions (functional form, distribution, etc)
 - Detect remaining model dependence
 - Remove model dependence: preprocess data (via matching, etc.)

Probability as a Model of Uncertainty

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$
 - 3 If $z_1, \dots, z_k$ are mutually exclusive events,

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$
 - 3 If $z_1, \dots, z_k$ are mutually exclusive events,

$$\Pr(z_1 \cup \dots \cup z_k) = \Pr(z_1) + \dots + \Pr(z_k),$$

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$
 - 3 If $z_1, \dots, z_k$ are mutually exclusive events,

$$\Pr(z_1 \cup \dots \cup z_k) = \Pr(z_1) + \dots + \Pr(z_k),$$

- The first two imply $0 \leq \Pr(z) \leq 1$

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$
 - 3 If $z_1, \dots, z_k$ are mutually exclusive events,

$$\Pr(z_1 \cup \dots \cup z_k) = \Pr(z_1) + \dots + \Pr(z_k),$$

- The first two imply $0 \leq \Pr(z) \leq 1$
- Axioms are not assumptions; they can't be wrong.

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$
 - 3 If $z_1, \dots, z_k$ are mutually exclusive events,

$$\Pr(z_1 \cup \dots \cup z_k) = \Pr(z_1) + \dots + \Pr(z_k),$$

- The first two imply $0 \leq \Pr(z) \leq 1$
- Axioms are not assumptions; they can't be wrong.
- From the axioms come *all* rules of probability theory.

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$
 - 3 If $z_1, \dots, z_k$ are mutually exclusive events,

$$\Pr(z_1 \cup \dots \cup z_k) = \Pr(z_1) + \dots + \Pr(z_k),$$

- The first two imply $0 \leq \Pr(z) \leq 1$
- Axioms are not assumptions; they can't be wrong.
- From the axioms come *all* rules of probability theory.
- **A Test: what happens if we said $\Pr(\text{sample space}) = 1$**

Probability as a Model of Uncertainty

- $\Pr(y|M) = \Pr(\text{data}|\text{Model})$, where $M = (f, g, X, \beta, \alpha)$.
- 3 *axioms* define the function $\Pr(\cdot|\cdot)$:
 - 1 $\Pr(z) \geq 0$ for some event z
 - 2 $\Pr(\text{sample space}) = 1$
 - 3 If $z_1, \dots, z_k$ are mutually exclusive events,

$$\Pr(z_1 \cup \dots \cup z_k) = \Pr(z_1) + \dots + \Pr(z_k),$$

- The first two imply $0 \leq \Pr(z) \leq 1$
- Axioms are not assumptions; they can't be wrong.
- From the axioms come *all* rules of probability theory.
- **A Test: what happens if we said $\Pr(\text{sample space}) = 1$**
- Rules can be applied *analytically* or via *simulation*.

Simulation is used to:

Simulation is used to:

- 1 Solve probability problems

Simulation is used to:

- 1 Solve probability problems
- 2 Evaluate estimators

Simulation is used to:

- ① Solve probability problems
- ② Evaluate estimators
- ③ Calculate features of probability densities

Simulation is used to:

- ① Solve probability problems
- ② Evaluate estimators
- ③ Calculate features of probability densities
- ④ Transform statistical results into quantities of interest

Simulation is used to:

- 1 Solve probability problems
- 2 Evaluate estimators
- 3 Calculate features of probability densities
- 4 Transform statistical results into quantities of interest
- 5 Get the right answer: students get the right answer far more frequently by using simulation than math

What is simulation?

What is simulation?

Survey Sampling

Simulation

What is simulation?

Survey Sampling

1. Learn about a population by taking a random sample from it

Simulation

1. Learn about a distribution by taking random draws from it

What is simulation?

Survey Sampling

1. Learn about a population by taking a random sample from it
2. Use the random sample to estimate a feature of the population

Simulation

1. Learn about a distribution by taking random draws from it
2. Use the random draws to approximate a feature of the distribution

What is simulation?

Survey Sampling

1. Learn about a population by taking a random sample from it
2. Use the random sample to estimate a feature of the population
3. The estimate is arbitrarily precise for large n

Simulation

1. Learn about a distribution by taking random draws from it
2. Use the random draws to approximate a feature of the distribution
3. The approximation is arbitrarily precise for large M

What is simulation?

Survey Sampling

1. Learn about a population by taking a random sample from it
2. Use the random sample to estimate a feature of the population
3. The estimate is arbitrarily precise for large n
4. Example: estimate the mean of the population

Simulation

1. Learn about a distribution by taking random draws from it
2. Use the random draws to approximate a feature of the distribution
3. The approximation is arbitrarily precise for large M
4. Example: Approximate the mean of the distribution

What's the probability that these goats are real?

What's the probability that these goats are real?



What's the probability that these goats are real?



HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



- They're real goats! (I took the photos myself)

HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



- They're real goats! (I took the photos myself)
- Key points:

HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



- They're real goats! (I took the photos myself)
- Key points:
 - You had uncertainty in the DGP (data generation process)

HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



- They're real goats! (I took the photos myself)
- Key points:
 - You had uncertainty in the DGP (data generation process)
 - To convince you, you needed more info about the DGP

HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



- They're real goats! (I took the photos myself)
- Key points:
 - You had uncertainty in the DGP (data generation process)
 - To convince you, you needed more info about the DGP
 - Probability: Assumes DGP (from which we can simulate)

HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



- They're real goats! (I took the photos myself)
- Key points:
 - You had uncertainty in the DGP (data generation process)
 - To convince you, you needed more info about the DGP
 - Probability: Assumes DGP (from which we can simulate)
 - Statistical Inference: The process of learning about the DGP

HEY, WHO DO YOU THINK YOU'RE LOOKING AT!?



- They're real goats! (I took the photos myself)
- Key points:
 - You had uncertainty in the DGP (data generation process)
 - To convince you, you needed more info about the DGP
 - Probability: Assumes DGP (from which we can simulate)
 - Statistical Inference: The process of learning about the DGP
 - $\rightsquigarrow$ Make others justify their choice of a DGP!

Monte Hall's Let's Make a Deal

Monte Hall's Let's Make a Deal

You have a choice among 3 doors. Behind a random door is a **car**; behind the other two are **goats**. You choose one at random. Monte peeks behind the other two doors and opens the one (or one of the two) with the **goat** and asks if you'd like to switch your door for the other door that hasn't been opened yet. Should you switch?

Monte Hall's Let's Make a Deal

You have a choice among 3 doors. Behind a random door is a **car**; behind the other two are **goats**. You choose one at random. Monte peeks behind the other two doors and opens the one (or one of the two) with the **goat** and asks if you'd like to switch your door for the other door that hasn't been opened yet. Should you switch?

```
sims <- 1000
WinNoSwitch <- 0
WinSwitch <- 0
doors <- c(1, 2, 3)
```

Monte Hall's Let's Make a Deal

You have a choice among 3 doors. Behind a random door is a **car**; behind the other two are **goats**. You choose one at random. Monte peeks behind the other two doors and opens the one (or one of the two) with the **goat** and asks if you'd like to switch your door for the other door that hasn't been opened yet. Should you switch?

```
sims <- 1000
WinNoSwitch <- 0
WinSwitch <- 0
doors <- c(1, 2, 3)
for (i in 1:sims) {
  WinDoor <- sample(doors, 1)
  choice <- sample(doors, 1)
  if (WinDoor == choice)                                # no switch
    WinNoSwitch <- WinNoSwitch + 1
  doorsLeft <- doors[doors != choice]                   # switch
  if (any(doorLeft == WinDoor))
    WinSwitch <- WinSwitch + 1
}
```

Monte Hall's Let's Make a Deal

You have a choice among 3 doors. Behind a random door is a **car**; behind the other two are **goats**. You choose one at random. Monte peeks behind the other two doors and opens the one (or one of the two) with the **goat** and asks if you'd like to switch your door for the other door that hasn't been opened yet. Should you switch?

```
sims <- 1000
WinNoSwitch <- 0
WinSwitch <- 0
doors <- c(1, 2, 3)
for (i in 1:sims) {
  WinDoor <- sample(doors, 1)
  choice <- sample(doors, 1)
  if (WinDoor == choice)                                # no switch
    WinNoSwitch <- WinNoSwitch + 1
  doorsLeft <- doors[doors != choice]                   # switch
  if (any(doorsLeft == WinDoor))
    WinSwitch <- WinSwitch + 1
}
cat("Prob(Car | no switch)=", WinNoSwitch/sims, "\n")
cat("Prob(Car | switch)=", WinSwitch/sims, "\n")
```

Let's Make a Deal

$\Pr(\text{car} \text{No Switch})$	$\Pr(\text{car} \text{Switch})$
.324	.676
.345	.655
.320	.680
.327	.673

The Birthday Problem

The Birthday Problem

Given a room with 24 randomly selected people, what is the probability that at least two have the same birthday?

The Birthday Problem

Given a room with 24 randomly selected people, what is the probability that at least two have the same birthday?

```
sims <- 1000
people <- 24
alldays <- seq(1, 365, 1)
sameday <- 0
```


The Birthday Problem

Given a room with 24 randomly selected people, what is the probability that at least two have the same birthday?

```
sims <- 1000
people <- 24
alldays <- seq(1, 365, 1)
sameday <- 0
for (i in 1:sims) {
  room <- sample(alldays, people, replace = TRUE)
  if (length(unique(room)) < people) # same birthday
    sameday <- sameday+1
}
```

The Birthday Problem

Given a room with 24 randomly selected people, what is the probability that at least two have the same birthday?

```
sims <- 1000
people <- 24
alldays <- seq(1, 365, 1)
sameday <- 0
for (i in 1:sims) {
  room <- sample(alldays, people, replace = TRUE)
  if (length(unique(room)) < people) # same birthday
    sameday <- sameday+1
}

cat("Probability of >=2 people having the same birthday:", sameday/sims, "\n")
```

The Birthday Problem


Given a room with 24 randomly selected people, what is the probability that at least two have the same birthday?

```
sims <- 1000
people <- 24
alldays <- seq(1, 365, 1)
sameday <- 0
for (i in 1:sims) {
  room <- sample(alldays, people, replace = TRUE)
  if (length(unique(room)) < people) # same birthday
    sameday <- sameday+1
}
```





```
cat("Probability of >=2 people having the same birthday:", sameday/sims, "\n")
```

Four runs: .538, .550, .547, .524

Computing Probabilities from PDFs







Computing Probabilities from PDFs



Required for a PDF:





Computing Probabilities from PDFs



Required for a PDF:

1 $\sum_y P(y) = 1$


Computing Probabilities from PDFs



Required for a PDF:

① $\sum_y P(y) = 1$ or $\int_{-\infty}^{\infty} P(y) dy = 1$





Computing Probabilities from PDFs



Required for a PDF:

- ① $\sum_y P(y) = 1$ or $\int_{-\infty}^{\infty} P(y) dy = 1$
- ② $P(y) \geq 0$ for every y

Computing Probabilities from PDFs







Required for a PDF:

- 1 $\sum_y P(y) = 1$ or $\int_{-\infty}^{\infty} P(y) dy = 1$
- 2 $P(y) \geq 0$ for every y

PDF Calculations:

Computing Probabilities from PDFs




Required for a PDF:

- 1 $\sum_y P(y) = 1$ or $\int_{-\infty}^{\infty} P(y) dy = 1$
- 2 $P(y) \geq 0$ for every y

PDF Calculations:

$$\bullet \Pr(a \leq Y \leq b) = \int_a^b P(y) dy$$

Computing Probabilities from PDFs




Required for a PDF:

- ① $\sum_y P(y) = 1$ or $\int_{-\infty}^{\infty} P(y) dy = 1$
- ② $P(y) \geq 0$ for every y

PDF Calculations:

- $\Pr(a \leq Y \leq b) = \int_a^b P(y) dy$
- For discrete: $\Pr(y) = P(y)$

Computing Probabilities from PDFs



Required for a PDF:

- ① $\sum_y P(y) = 1$ or $\int_{-\infty}^{\infty} P(y) dy = 1$
- ② $P(y) \geq 0$ for every y

PDF Calculations:

- $\Pr(a \leq Y \leq b) = \int_a^b P(y) dy$
- For discrete: $\Pr(y) = P(y)$
- For continuous: $\Pr(y) = 0$ (why?)

What you should know about every pdf

What you should know about every pdf

- The **assignment** of a probability or probability density to every conceivable value of Y_i

What you should know about every pdf

- The **assignment** of a probability or probability density to every conceivable value of Y_i
- The **first principles**

What you should know about every pdf

- The **assignment** of a probability or probability density to every conceivable value of Y_i
- The **first principles**
- How to **use** the final expression (but not necessarily the full derivation)

What you should know about every pdf

- The **assignment** of a probability or probability density to every conceivable value of Y_i
- The **first principles**
- How to **use** the final expression (but not necessarily the full derivation)
- How to **simulate** from the density

What you should know about every pdf


- The **assignment** of a probability or probability density to every conceivable value of Y_i
- The **first principles**
- How to **use** the final expression (but not necessarily the full derivation)
- How to **simulate** from the density
- How to **compute** features of the density such as its “moments”

What you should know about every pdf


- The **assignment** of a probability or probability density to every conceivable value of Y_i
- The **first principles**
- How to **use** the final expression (but not necessarily the full derivation)
- How to **simulate** from the density
- How to **compute** features of the density such as its “moments”
- How to **verify** that the final expression is indeed a proper density

Uniform Density on the interval $[0, 1]$

Uniform Density on the interval $[0, 1]$




Uniform Density on the interval $[0, 1]$



First Principles about the process that generates Y_i is such that


Uniform Density on the interval $[0, 1]$



First Principles about the process that generates Y_i is such that

- Y_i falls in the interval $[0, 1]$ with probability 1: $\int_0^1 P(y)dy = 1$


Uniform Density on the interval $[0, 1]$



First Principles about the process that generates Y_i is such that

- Y_i falls in the interval $[0, 1]$ with probability 1: $\int_0^1 P(y)dy = 1$
- $\Pr(Y \in (a, b)) = \Pr(Y \in (c, d))$ if $a < b$, $c < d$, and $b - a = d - c$.


Uniform Density on the interval $[0, 1]$



First Principles about the process that generates Y_i is such that

- Y_i falls in the interval $[0, 1]$ with probability 1: $\int_0^1 P(y)dy = 1$
- $\Pr(Y \in (a, b)) = \Pr(Y \in (c, d))$ if $a < b$, $c < d$, and $b - a = d - c$.
- Is it a pdf? How do you know?


Uniform Density on the interval $[0, 1]$



First Principles about the process that generates Y_i is such that

- Y_i falls in the interval $[0, 1]$ with probability 1: $\int_0^1 P(y)dy = 1$
- $\Pr(Y \in (a, b)) = \Pr(Y \in (c, d))$ if $a < b$, $c < d$, and $b - a = d - c$.
- Is it a pdf? How do you know?
- How to simulate?

Uniform Density on the interval $[0, 1]$




First Principles about the process that generates Y_i is such that


- Y_i falls in the interval $[0, 1]$ with probability 1: $\int_0^1 P(y)dy = 1$
- $\Pr(Y \in (a, b)) = \Pr(Y \in (c, d))$ if $a < b$, $c < d$, and $b - a = d - c$.
- Is it a pdf? How do you know?
- How to simulate? `runif(1000)`

Bernoulli pmf

Bernoulli pmf




Bernoulli pmf




- **First principles** about the process that generates Y_i :

Bernoulli pmf




- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and

Bernoulli pmf




- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and
 - The 2 outcomes are **exhaustive**

Bernoulli pmf




- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and
 - The 2 outcomes are **exhaustive**
- In this simple case, we will compute features analytically and by simulation.

Bernoulli pmf




- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and
 - The 2 outcomes are **exhaustive**
- In this simple case, we will compute features analytically and by simulation.
- Mathematical expression for the pmf

Bernoulli pmf




- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and
 - The 2 outcomes are **exhaustive**
- In this simple case, we will compute features analytically and by simulation.
- Mathematical expression for the pmf
 - $\Pr(Y_i = 1|\pi_i) = \pi_i$, $\Pr(Y_i = 0|\pi_i) = 1 - \pi_i$

Bernoulli pmf




- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and
 - The 2 outcomes are **exhaustive**
- In this simple case, we will compute features analytically and by simulation.
- Mathematical expression for the pmf
 - $\Pr(Y_i = 1|\pi_i) = \pi_i$, $\Pr(Y_i = 0|\pi_i) = 1 - \pi_i$
 - The parameter π happens to be interpretable as a probability

Bernoulli pmf



- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and
 - The 2 outcomes are **exhaustive**
- In this simple case, we will compute features analytically and by simulation.
- Mathematical expression for the pmf
 - $\Pr(Y_i = 1|\pi_i) = \pi_i, \Pr(Y_i = 0|\pi_i) = 1 - \pi_i$
 - The parameter π happens to be interpretable as a probability
 - $\implies \Pr(Y_i = y|\pi_i) = \pi_i^y (1 - \pi_i)^{1-y}$

Bernoulli pmf



- **First principles** about the process that generates Y_i :
 - Y_i has 2 **mutually exclusive** outcomes; and
 - The 2 outcomes are **exhaustive**
- In this simple case, we will compute features analytically and by simulation.
- Mathematical expression for the pmf
 - $\Pr(Y_i = 1|\pi_i) = \pi_i, \Pr(Y_i = 0|\pi_i) = 1 - \pi_i$
 - The parameter π happens to be interpretable as a probability
 - $\implies \Pr(Y_i = y|\pi_i) = \pi_i^y (1 - \pi_i)^{1-y}$
 - Alternative notation: $\Pr(Y_i = y|\pi_i) = \text{Bernoulli}(y|\pi_i) = f_b(y|\pi_i)$

Features of the Bernoulli: analytically

Features of the Bernoulli: analytically

- Expected value:

Features of the Bernoulli: analytically

- Expected value:

$$E(Y) = \sum_{\text{all } y} yP(y)$$

Features of the Bernoulli: analytically

- Expected value:

$$\begin{aligned} E(Y) &= \sum_{\text{all } y} yP(y) \\ &= 0 \Pr(0) + 1 \Pr(1) \end{aligned}$$

Features of the Bernoulli: analytically

- Expected value:

$$\begin{aligned} E(Y) &= \sum_{\text{all } y} yP(y) \\ &= 0 \Pr(0) + 1 \Pr(1) \\ &= \pi \end{aligned}$$

Features of the Bernoulli: analytically

- Expected value:

$$\begin{aligned} E(Y) &= \sum_{\text{all } y} yP(y) \\ &= 0 \Pr(0) + 1 \Pr(1) \\ &= \pi \end{aligned}$$

- Variance:

Features of the Bernoulli: analytically

- Expected value:

$$\begin{aligned} E(Y) &= \sum_{\text{all } y} yP(y) \\ &= 0 \Pr(0) + 1 \Pr(1) \\ &= \pi \end{aligned}$$

- Variance:

$$V(Y) = E[(Y - E(Y))^2] \quad (\text{The definition})$$

Features of the Bernoulli: analytically

- Expected value:

$$\begin{aligned} E(Y) &= \sum_{\text{all } y} yP(y) \\ &= 0 \Pr(0) + 1 \Pr(1) \\ &= \pi \end{aligned}$$

- Variance:

$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \end{aligned}$$

Features of the Bernoulli: analytically

- Expected value:

$$\begin{aligned} E(Y) &= \sum_{\text{all } y} yP(y) \\ &= 0 \Pr(0) + 1 \Pr(1) \\ &= \pi \end{aligned}$$

- Variance:

$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \\ &= E(Y^2) - \pi^2 && \text{(An easier version)} \end{aligned}$$

Features of the Bernoulli: analytically

- Expected value:

$$\begin{aligned} E(Y) &= \sum_{\text{all } y} yP(y) \\ &= 0 \Pr(0) + 1 \Pr(1) \\ &= \pi \end{aligned}$$

- Variance:

$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \\ &= E(Y^2) - \pi^2 && \text{(An easier version)} \end{aligned}$$

- How do we compute $E(Y^2)$?

Expected values of functions of random variables

Expected values of functions of random variables

$$E[g(Y)] = \sum_{\text{all } y} g(y)P(y)$$

Expected values of functions of random variables

$$E[g(Y)] = \sum_{\text{all } y} g(y)P(y)$$

or

Expected values of functions of random variables

$$E[g(Y)] = \sum_{\text{all } y} g(y)P(y)$$

or

$$E[g(Y)] = \int_{-\infty}^{\infty} g(y)P(y)$$

Expected values of functions of random variables

$$E[g(Y)] = \sum_{\text{all } y} g(y)P(y)$$

or

$$E[g(Y)] = \int_{-\infty}^{\infty} g(y)P(y)$$

For example,

Expected values of functions of random variables

$$E[g(Y)] = \sum_{\text{all } y} g(y)P(y)$$

or

$$E[g(Y)] = \int_{-\infty}^{\infty} g(y)P(y)$$

For example,

$$E(Y^2) = \sum_{\text{all } y} y^2 P(y)$$

Expected values of functions of random variables

$$E[g(Y)] = \sum_{\text{all } y} g(y)P(y)$$

or

$$E[g(Y)] = \int_{-\infty}^{\infty} g(y)P(y)$$

For example,

$$\begin{aligned} E(Y^2) &= \sum_{\text{all } y} y^2 P(y) \\ &= 0^2 \Pr(0) + 1^2 \Pr(1) \end{aligned}$$

Expected values of functions of random variables

$$E[g(Y)] = \sum_{\text{all } y} g(y)P(y)$$

or

$$E[g(Y)] = \int_{-\infty}^{\infty} g(y)P(y)$$

For example,

$$\begin{aligned} E(Y^2) &= \sum_{\text{all } y} y^2 P(y) \\ &= 0^2 \Pr(0) + 1^2 \Pr(1) \\ &= \pi \end{aligned}$$

Variance of the Bernoulli (uses above results)

Variance of the Bernoulli (uses above results)

$$V(Y) = E[(Y - E(Y))^2] \quad (\text{The definition})$$

Variance of the Bernoulli (uses above results)

$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \end{aligned}$$

Variance of the Bernoulli (uses above results)

$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \\ &= \pi - \pi^2 \end{aligned}$$

Variance of the Bernoulli (uses above results)

$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \\ &= \pi - \pi^2 \\ &= \pi(1 - \pi) \end{aligned}$$

Variance of the Bernoulli (uses above results)


$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \\ &= \pi - \pi^2 \\ &= \pi(1 - \pi) \end{aligned}$$

This makes sense:

Variance of the Bernoulli (uses above results)

$$\begin{aligned} V(Y) &= E[(Y - E(Y))^2] && \text{(The definition)} \\ &= E(Y^2) - E(Y)^2 && \text{(An easier version)} \\ &= \pi - \pi^2 \\ &= \pi(1 - \pi) \end{aligned}$$

This makes sense:



How to Simulate from the Bernoulli with parameter π

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$
- Set π to a particular value

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$
- Set π to a particular value
- Set $y = 1$ if $u < \pi$ and $y = 0$ otherwise

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$
- Set π to a particular value
- Set $y = 1$ if $u < \pi$ and $y = 0$ otherwise
- In R:

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$
- Set π to a particular value
- Set $y = 1$ if $u < \pi$ and $y = 0$ otherwise
- In R:

```
sims <- 1000      # set parameters
bernpi <- 0.2
u <- runif(sims)   # uniform sims
y <- as.integer(u < bernpi)
y                  # print results
```

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$
- Set π to a particular value
- Set $y = 1$ if $u < \pi$ and $y = 0$ otherwise
- In R:

```
sims <- 1000      # set parameters
bernpi <- 0.2
u <- runif(sims)  # uniform sims
y <- as.integer(u < bernpi)
y                # print results
```

Running the program gives:

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$
- Set π to a particular value
- Set $y = 1$ if $u < \pi$ and $y = 0$ otherwise
- In R:

```
sims <- 1000      # set parameters
bernpi <- 0.2
u <- runif(sims)   # uniform sims
y <- as.integer(u < bernpi)
y                  # print results
```

Running the program gives:

```
0 0 0 1 0 0 1 1 0 0 1 1 1 0 ...
```

How to Simulate from the Bernoulli with parameter π

- take one draw u from a uniform density on the interval $[0,1]$
- Set π to a particular value
- Set $y = 1$ if $u < \pi$ and $y = 0$ otherwise
- In R:

```
sims <- 1000      # set parameters
bernpi <- 0.2
u <- runif(sims)  # uniform sims
y <- as.integer(u < bernpi)
y                # print results
```

Running the program gives:

0 0 0 1 0 0 1 1 0 0 1 1 1 0 ...

- What can we do with the simulations?

Binomial Distribution

Binomial Distribution

First principles:

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

$$P(Y = y|\pi) = \binom{N}{y} \pi^y (1 - \pi)^{N-y}$$

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

$$P(Y = y|\pi) = \binom{N}{y} \pi^y (1 - \pi)^{N-y}$$

Explanation:

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

$$P(Y = y|\pi) = \binom{N}{y} \pi^y (1 - \pi)^{N-y}$$

Explanation:

- $\binom{N}{y}$ because (1 0 1) and (1 1 0) are both $y = 2$.

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

$$P(Y = y|\pi) = \binom{N}{y} \pi^y (1 - \pi)^{N-y}$$

Explanation:

- $\binom{N}{y}$ because (1 0 1) and (1 1 0) are both $y = 2$.
- π^y because y successes with π probability each (product taken due to independence)

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

$$P(Y = y|\pi) = \binom{N}{y} \pi^y (1 - \pi)^{N-y}$$

Explanation:

- $\binom{N}{y}$ because (1 0 1) and (1 1 0) are both $y = 2$.
- π^y because y successes with π probability each (product taken due to independence)
- $(1 - \pi)^{N-y}$ because $N - y$ failures with $1 - \pi$ probability each

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

$$P(Y = y|\pi) = \binom{N}{y} \pi^y (1 - \pi)^{N-y}$$

Explanation:

- $\binom{N}{y}$ because (1 0 1) and (1 1 0) are both $y = 2$.
- π^y because y successes with π probability each (product taken due to independence)
- $(1 - \pi)^{N-y}$ because $N - y$ failures with $1 - \pi$ probability each
- Mean $E(Y) = N\pi$

Binomial Distribution

First principles:

- N iid Bernoulli trials, $y_1, \dots, y_N$
- The trials are independent
- The trials are identically distributed
- We observe $Y = \sum_{i=1}^N y_i$

Density:

$$P(Y = y|\pi) = \binom{N}{y} \pi^y (1 - \pi)^{N-y}$$

Explanation:

- $\binom{N}{y}$ because (1 0 1) and (1 1 0) are both $y = 2$.
- π^y because y successes with π probability each (product taken due to independence)
- $(1 - \pi)^{N-y}$ because $N - y$ failures with $1 - \pi$ probability each
- Mean $E(Y) = N\pi$
- Variance $V(Y) = \pi(1 - \pi)/N$.

How to simulate from the Binomial distribution

How to simulate from the Binomial distribution

- To simulate from the Binomial(π ; N):

How to simulate from the Binomial distribution

- To simulate from the Binomial(π ; N):
 - Simulate N independent Bernoulli variables, $Y_1, \dots, Y_N$, each with parameter π

How to simulate from the Binomial distribution

- To simulate from the Binomial(π ; N):
 - Simulate N independent Bernoulli variables, $Y_1, \dots, Y_N$, each with parameter π
 - Add them up: $\sum_{i=1}^N Y_i$

How to simulate from the Binomial distribution

- To simulate from the Binomial(π ; N):
 - Simulate N independent Bernoulli variables, $Y_1, \dots, Y_N$, each with parameter π
 - Add them up: $\sum_{i=1}^N Y_i$
- What can you do with the simulations?

Where to get uniform random numbers

Where to get uniform random numbers

- Random is not haphazard (e.g., Benford's law)

Where to get uniform random numbers

- Random is not haphazard (e.g., Benford's law)
- Random number generators are perfectly predictable (what?)

Where to get uniform random numbers

- Random is not haphazard (e.g., Benford's law)
- Random number generators are perfectly predictable (what?)
- We use pseudo-random numbers which have (a) digits that occur with 1/10th probability, (b) no time series patterns, etc.

Where to get uniform random numbers


- Random is not haphazard (e.g., Benford's law)
- Random number generators are perfectly predictable (what?)
- We use pseudo-random numbers which have (a) digits that occur with 1/10th probability, (b) no time series patterns, etc.
- How to create real random numbers?

Where to get uniform random numbers


- Random is not haphazard (e.g., Benford's law)
- Random number generators are perfectly predictable (what?)
- We use pseudo-random numbers which have (a) digits that occur with 1/10th probability, (b) no time series patterns, etc.
- How to create real random numbers?
- Some chips now use quantum effects

Discretization for random draws from discrete pmfs

Discretization for random draws from discrete pmfs




Discretization for random draws from discrete pmfs




- Divide up PDF into a grid

Discretization for random draws from discrete pmfs




- Divide up PDF into a grid
- Approximate probabilities by trapezoids

Discretization for random draws from discrete pmfs




- Divide up PDF into a grid
- Approximate probabilities by trapezoids
- Map $[0,1]$ uniform draws to the proportion area in each trapezoid

Discretization for random draws from discrete pmfs




- Divide up PDF into a grid
- Approximate probabilities by trapezoids
- Map $[0,1]$ uniform draws to the proportion area in each trapezoid
- Return midpoint of each trapezoid

Discretization for random draws from discrete pmfs





- Divide up PDF into a grid
- Approximate probabilities by trapezoids
- Map $[0,1]$ uniform draws to the proportion area in each trapezoid
- Return midpoint of each trapezoid
- More trapezoids $\rightsquigarrow$ better approximation

Discretization for random draws from discrete pmfs




- Divide up PDF into a grid
- Approximate probabilities by trapezoids
- Map $[0,1]$ uniform draws to the proportion area in each trapezoid
- Return midpoint of each trapezoid
- More trapezoids $\rightsquigarrow$ better approximation
- (Works for a few dimensions, but infeasible for many)

Inverse CDF: drawing from arbitrary continuous pdfs




Inverse CDF: drawing from arbitrary continuous pdfs




- From the pdf $f(Y)$, compute the cdf:
 $\Pr(Y \leq y) \equiv F(y) = \int_{-\infty}^y f(z) dz$

Inverse CDF: drawing from arbitrary continuous pdfs




- From the pdf $f(Y)$, compute the cdf:
$$\Pr(Y \leq y) \equiv F(y) = \int_{-\infty}^y f(z) dz$$
- Define the inverse cdf $F^{-1}(y)$, such that $F^{-1}[F(y)] = y$

Inverse CDF: drawing from arbitrary continuous pdfs




- From the pdf $f(Y)$, compute the cdf:
$$\Pr(Y \leq y) \equiv F(y) = \int_{-\infty}^y f(z) dz$$
- Define the inverse cdf $F^{-1}(y)$, such that $F^{-1}[F(y)] = y$
- Draw random uniform number, U

Inverse CDF: drawing from arbitrary continuous pdfs




- From the pdf $f(Y)$, compute the cdf:
$$\Pr(Y \leq y) \equiv F(y) = \int_{-\infty}^y f(z) dz$$
- Define the inverse cdf $F^{-1}(y)$, such that $F^{-1}[F(y)] = y$
- Draw random uniform number, U
- Then $F^{-1}(U)$ gives a random draw from $f(Y)$.

Using Inverse CDF to Improve Discretization Method




Using Inverse CDF to Improve Discretization Method




- Refined Discretization Method:

Using Inverse CDF to Improve Discretization Method



- **Refined Discretization Method:**
 - Choose interval randomly as above (based on area in trapezoids)


Using Inverse CDF to Improve Discretization Method



- **Refined Discretization Method:**

- Choose interval randomly as above (based on area in trapezoids)
- Draw a number within each trapezoid by the inverse CDF method applied to the trapezoidal approximation.

Using Inverse CDF to Improve Discretization Method



- **Refined Discretization Method:**
 - Choose interval randomly as above (based on area in trapezoids)
 - Draw a number within each trapezoid by the inverse CDF method applied to the trapezoidal approximation.
- **Drawing random numbers from arbitrary multivariate densities: now an enormous literature**

Normal Distribution

Normal Distribution

- Many different first principles

Normal Distribution

- Many different first principles
- A common one is the central limit theorem

Normal Distribution

- Many different first principles
- A common one is the central limit theorem
- The univariate normal density (with mean μ_i , variance σ^2)

Normal Distribution

- Many different first principles
- A common one is the central limit theorem
- The univariate normal density (with mean μ_i , variance σ^2)

$$N(y_i|\mu_i, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2\sigma^2}\right)$$

Normal Distribution

- Many different first principles
- A common one is the central limit theorem
- The univariate normal density (with mean μ_i , variance σ^2)

$$N(y_i|\mu_i, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2\sigma^2}\right)$$

- The stylized normal: $f_{stn}(y_i|\mu_i) = N(y|\mu_i, 1)$

Normal Distribution

- Many different first principles
- A common one is the central limit theorem
- The univariate normal density (with mean μ_i , variance σ^2)

$$N(y_i|\mu_i, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2\sigma^2}\right)$$

- The stylized normal: $f_{stn}(y_i|\mu_i) = N(y|\mu_i, 1)$

$$f_{stn}(y|\mu_i) = (2\pi)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2}\right)$$

Normal Distribution

- Many different first principles
- A common one is the central limit theorem
- The univariate normal density (with mean μ_i , variance σ^2)

$$N(y_i|\mu_i, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2\sigma^2}\right)$$

- The stylized normal: $f_{stn}(y_i|\mu_i) = N(y|\mu_i, 1)$

$$f_{stn}(y|\mu_i) = (2\pi)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2}\right)$$

- The standardized normal: $f_{sn}(y_i) = N(y_i|0, 1) = \phi(y_i)$

Normal Distribution

- Many different first principles
- A common one is the central limit theorem
- The univariate normal density (with mean μ_i , variance σ^2)

$$N(y_i|\mu_i, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2\sigma^2}\right)$$

- The stylized normal: $f_{stn}(y_i|\mu_i) = N(y_i|\mu_i, 1)$

$$f_{stn}(y_i|\mu_i) = (2\pi)^{-1/2} \exp\left(\frac{-(y_i - \mu_i)^2}{2}\right)$$

- The standardized normal: $f_{sn}(y_i) = N(y_i|0, 1) = \phi(y_i)$

$$f_{sn}(y_i) = (2\pi)^{-1/2} \exp\left(\frac{-y_i^2}{2}\right)$$

Multivariate Normal Distribution

Multivariate Normal Distribution

- Let $Y_i \equiv \{Y_{1i}, \dots, Y_{ki}\}$ be a $k \times 1$ vector, jointly random:

Multivariate Normal Distribution

- Let $Y_i \equiv \{Y_{1i}, \dots, Y_{ki}\}$ be a $k \times 1$ vector, jointly random:

$$Y_i \sim N(y_i | \mu_i, \Sigma)$$

Multivariate Normal Distribution

- Let $Y_i \equiv \{Y_{1i}, \dots, Y_{ki}\}$ be a $k \times 1$ vector, jointly random:

$$Y_i \sim N(y_i | \mu_i, \Sigma)$$

where μ_i is $k \times 1$ and Σ is $k \times k$. For $k = 2$,

Multivariate Normal Distribution

- Let $Y_i \equiv \{Y_{1i}, \dots, Y_{ki}\}$ be a $k \times 1$ vector, jointly random:

$$Y_i \sim N(y_i | \mu_i, \Sigma)$$

where μ_i is $k \times 1$ and Σ is $k \times k$. For $k = 2$,

$$\mu_i = \begin{pmatrix} \mu_{1i} \\ \mu_{2i} \end{pmatrix} \quad \Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$$

Multivariate Normal Distribution

- Let $Y_i \equiv \{Y_{1i}, \dots, Y_{ki}\}$ be a $k \times 1$ vector, jointly random:

$$Y_i \sim N(y_i | \mu_i, \Sigma)$$

where μ_i is $k \times 1$ and Σ is $k \times k$. For $k = 2$,

$$\mu_i = \begin{pmatrix} \mu_{1i} \\ \mu_{2i} \end{pmatrix} \quad \Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$$

- Mathematical form:

Multivariate Normal Distribution

- Let $Y_i \equiv \{Y_{1i}, \dots, Y_{ki}\}$ be a $k \times 1$ vector, jointly random:

$$Y_i \sim N(y_i | \mu_i, \Sigma)$$

where μ_i is $k \times 1$ and Σ is $k \times k$. For $k = 2$,

$$\mu_i = \begin{pmatrix} \mu_{1i} \\ \mu_{2i} \end{pmatrix} \quad \Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$$

- Mathematical form:

$$N(y_i | \mu_i, \Sigma) = (2\pi)^{-k/2} |\Sigma|^{-1/2} \exp \left[-\frac{1}{2} (y_i - \mu_i)' \Sigma^{-1} (y_i - \mu_i) \right]$$

Multivariate Normal Distribution

- Let $Y_i \equiv \{Y_{1i}, \dots, Y_{ki}\}$ be a $k \times 1$ vector, jointly random:

$$Y_i \sim N(y_i | \mu_i, \Sigma)$$

where μ_i is $k \times 1$ and Σ is $k \times k$. For $k = 2$,

$$\mu_i = \begin{pmatrix} \mu_{1i} \\ \mu_{2i} \end{pmatrix} \quad \Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}$$

- Mathematical form:

$$N(y_i | \mu_i, \Sigma) = (2\pi)^{-k/2} |\Sigma|^{-1/2} \exp \left[-\frac{1}{2} (y_i - \mu_i)' \Sigma^{-1} (y_i - \mu_i) \right]$$

- Simulating *once* from this density produces k numbers. Special algorithms are used to generate normal random variates (in R, `mvrnorm()`, from the MASS library).

Multivariate Normal Distribution

- Moments: $E(Y) = \mu_i$, $V(Y) = \Sigma$, $\text{Cov}(Y_1, Y_2) = \sigma_{12} = \sigma_{21}$.

Multivariate Normal Distribution


- Moments: $E(Y) = \mu_i$, $V(Y) = \Sigma$, $\text{Cov}(Y_1, Y_2) = \sigma_{12} = \sigma_{21}$.
- $\text{Corr}(Y_1, Y_2) = \frac{\sigma_{12}}{\sigma_1 \sigma_2}$

Multivariate Normal Distribution


- Moments: $E(Y) = \mu_i$, $V(Y) = \Sigma$, $\text{Cov}(Y_1, Y_2) = \sigma_{12} = \sigma_{21}$.
- $\text{Corr}(Y_1, Y_2) = \frac{\sigma_{12}}{\sigma_1 \sigma_2}$
- Marginals:

$$N(Y_1 | \mu_1, \sigma_1^2) = \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} N(y_i | \mu_i, \Sigma) dy_2 dy_3 \cdots dy_k$$


Truncated bivariate normal examples (for β^b and β^w)



(a) 0.5 0.5 0.15 0.15 0



(b) 0.1 0.9 0.15 0.15 0



(c) 0.8 0.8 0.6 0.6 0.5

Parameters are μ_1 , μ_2 , σ_1 , σ_2 , and ρ .

Stop here

We will stop here this year and skip to the next set of slides.
Please refer to the slides below for further information on probability densities and random number generation; they offer more sophisticated .

Beta (continuous) density

Beta (continuous) density

- Used to model proportions.

Beta (continuous) density

- Used to model proportions.
- We'll use it first to generalize the Binomial distribution

Beta (continuous) density


- Used to model proportions.
- We'll use it first to generalize the Binomial distribution
- y falls in the interval $[0,1]$

Beta (continuous) density

- Used to model proportions.
- We'll use it first to generalize the Binomial distribution
- y falls in the interval $[0,1]$
- Takes on a variety of flexible forms, depending on the parameter values:

Beta (continuous) density

- Used to model proportions.
- We'll use it first to generalize the Binomial distribution
- y falls in the interval $[0,1]$
- Takes on a variety of flexible forms, depending on the parameter values:



Standard Parameterization

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1 - y)^{\beta-1}$$

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1 - y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1 - y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

$$\Gamma(x) = \int_0^{\infty} z^{x-1} e^{-z} dz$$

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1 - y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

$$\Gamma(x) = \int_0^{\infty} z^{x-1} e^{-z} dz$$

For integer values of x , $\Gamma(x + 1) = x! = x(x - 1)(x - 2) \cdots 1$.

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1 - y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

$$\Gamma(x) = \int_0^{\infty} z^{x-1} e^{-z} dz$$

For integer values of x , $\Gamma(x + 1) = x! = x(x - 1)(x - 2) \cdots 1$.

Non-integer values of x produce a continuous interpolation. In R or gauss:
`gamma(x)` ;

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1-y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

$$\Gamma(x) = \int_0^{\infty} z^{x-1} e^{-z} dz$$

For integer values of x , $\Gamma(x+1) = x! = x(x-1)(x-2)\cdots 1$.

Non-integer values of x produce a continuous interpolation. In R or gauss:
`gamma(x)` ;

Intuitive?

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1 - y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

$$\Gamma(x) = \int_0^{\infty} z^{x-1} e^{-z} dz$$

For integer values of x , $\Gamma(x + 1) = x! = x(x - 1)(x - 2) \cdots 1$.

Non-integer values of x produce a continuous interpolation. In R or gauss:
`gamma(x)` ;

Intuitive? The moments help some:

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1 - y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

$$\Gamma(x) = \int_0^{\infty} z^{x-1} e^{-z} dz$$

For integer values of x , $\Gamma(x + 1) = x! = x(x - 1)(x - 2) \cdots 1$.

Non-integer values of x produce a continuous interpolation. In R or gauss:
`gamma(x)` ;

Intuitive? The moments help some:

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

Standard Parameterization

$$\text{Beta}(y|\alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} y^{\alpha-1} (1-y)^{\beta-1}$$

where, $\Gamma(x)$ is the gamma function:

$$\Gamma(x) = \int_0^{\infty} z^{x-1} e^{-z} dz$$

For integer values of x , $\Gamma(x+1) = x! = x(x-1)(x-2)\cdots 1$.

Non-integer values of x produce a continuous interpolation. In R or gauss:
`gamma(x)` ;

Intuitive? The moments help some:

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

$$V(Y) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$$

Alternative parameterization

Alternative parameterization

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

Alternative parameterization

Set $\mu = E(Y) = \frac{\alpha}{(\alpha+\beta)}$ and $\frac{\mu(1-\mu)\gamma}{(1+\gamma)} = V(Y) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$,

Alternative parameterization

Set $\mu = E(Y) = \frac{\alpha}{(\alpha+\beta)}$ and $\frac{\mu(1-\mu)\gamma}{(1+\gamma)} = V(Y) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$, solve for α and β and substitute in.

Alternative parameterization

Set $\mu = E(Y) = \frac{\alpha}{(\alpha+\beta)}$ and $\frac{\mu(1-\mu)\gamma}{(1+\gamma)} = V(Y) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$, solve for α and β and substitute in.

Result:

Alternative parameterization

Set $\mu = E(Y) = \frac{\alpha}{(\alpha+\beta)}$ and $\frac{\mu(1-\mu)\gamma}{(1+\gamma)} = V(Y) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$, solve for α and β and substitute in.

Result:

$$\text{beta}(y|\mu, \gamma) = \frac{\Gamma(\mu\gamma^{-1} + (1-\mu)\gamma^{-1})}{\Gamma(\mu\gamma^{-1})\Gamma[(1-\mu)\gamma^{-1}]} y^{\mu\gamma^{-1}-1} (1-y)^{(1-\mu)\gamma^{-1}-1}$$

Alternative parameterization

Set $\mu = E(Y) = \frac{\alpha}{(\alpha+\beta)}$ and $\frac{\mu(1-\mu)\gamma}{(1+\gamma)} = V(Y) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$, solve for α and β and substitute in.

Result:

$$\text{beta}(y|\mu, \gamma) = \frac{\Gamma(\mu\gamma^{-1} + (1-\mu)\gamma^{-1})}{\Gamma(\mu\gamma^{-1})\Gamma[(1-\mu)\gamma^{-1}]} y^{\mu\gamma^{-1}-1} (1-y)^{(1-\mu)\gamma^{-1}-1}$$

where now $E(Y) = \mu$ and γ is an index of variation that varies with μ .

Alternative parameterization

Set $\mu = E(Y) = \frac{\alpha}{(\alpha+\beta)}$ and $\frac{\mu(1-\mu)\gamma}{(1+\gamma)} = V(Y) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$, solve for α and β and substitute in.

Result:

$$\text{beta}(y|\mu, \gamma) = \frac{\Gamma(\mu\gamma^{-1} + (1-\mu)\gamma^{-1})}{\Gamma(\mu\gamma^{-1})\Gamma[(1-\mu)\gamma^{-1}]} y^{\mu\gamma^{-1}-1} (1-y)^{(1-\mu)\gamma^{-1}-1}$$

where now $E(Y) = \mu$ and γ is an index of variation that varies with μ .

Reparameterization like this will be key throughout the course.

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

How to simulate

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

How to simulate

(First principles are easy to see from this too.)

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

How to simulate

(First principles are easy to see from this too.)

- Begin with N Bernoulli trials with parameter π_j , $j = 1, \dots, N$ (not necessarily independent or identically distributed)

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

How to simulate

(First principles are easy to see from this too.)

- Begin with N Bernoulli trials with parameter π_j , $j = 1, \dots, N$ (not necessarily independent or identically distributed)
- Choose $\mu = E(\pi_j)$ and γ

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

How to simulate

(First principles are easy to see from this too.)

- Begin with N Bernoulli trials with parameter π_j , $j = 1, \dots, N$ (not necessarily independent or identically distributed)
- Choose $\mu = E(\pi_j)$ and γ
- Draw $\tilde{\pi}$ from $\text{Beta}(\pi|\mu, \gamma)$ (without this step we get Binomial draws)

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

How to simulate

(First principles are easy to see from this too.)

- Begin with N Bernoulli trials with parameter π_j , $j = 1, \dots, N$ (not necessarily independent or identically distributed)
- Choose $\mu = E(\pi_j)$ and γ
- Draw $\tilde{\pi}$ from $\text{Beta}(\pi|\mu, \gamma)$ (without this step we get Binomial draws)
- Draw N Bernoulli variables $\tilde{z}_j$ ($j = 1, \dots, N$) from $\text{Bernoulli}(z_j|\tilde{\pi})$

Useful if the binomial variance is not approximately $\pi(1 - \pi)/N$.

How to simulate

(First principles are easy to see from this too.)

- Begin with N Bernoulli trials with parameter π_j , $j = 1, \dots, N$ (not necessarily independent or identically distributed)
- Choose $\mu = E(\pi_j)$ and γ
- Draw $\tilde{\pi}$ from $\text{Beta}(\pi|\mu, \gamma)$ (without this step we get Binomial draws)
- Draw N Bernoulli variables $\tilde{z}_j$ ($j = 1, \dots, N$) from $\text{Bernoulli}(z_j|\tilde{\pi})$
- Add up the $\tilde{z}$'s to get $y = \sum_j^N \tilde{z}_j$, which is a draw from the beta-binomial.

Beta-Binomial Analytics

Beta-Binomial Analytics

Recall:

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

- Derive the joint density of y and π . Then

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

- Derive the joint density of y and π . Then
- Average over the unknown π dimension

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

- Derive the joint density of y and π . Then
- Average over the unknown π dimension

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

- Derive the joint density of y and π . Then
- Average over the unknown π dimension

Hence, the beta-binomial (or extended beta-binomial):

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

- Derive the joint density of y and π . Then
- Average over the unknown π dimension

Hence, the beta-binomial (or extended beta-binomial):

$$\text{BB}(y_i|\mu, \gamma) = \int_0^1 \text{Binomial}(y_i|\pi) \times \text{Beta}(\pi|\mu, \gamma) d\pi$$

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

- Derive the joint density of y and π . Then
- Average over the unknown π dimension

Hence, the beta-binomial (or extended beta-binomial):

$$\begin{aligned} \text{BB}(y_i|\mu, \gamma) &= \int_0^1 \text{Binomial}(y_i|\pi) \times \text{Beta}(\pi|\mu, \gamma) d\pi \\ &= \int_0^1 \Pr(y_i, \pi|\mu, \gamma) d\pi \end{aligned}$$

Beta-Binomial Analytics

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B) \Pr(B)$$

Plan:

- Derive the joint density of y and π . Then
- Average over the unknown π dimension

Hence, the beta-binomial (or extended beta-binomial):

$$\begin{aligned}\text{BB}(y_i|\mu, \gamma) &= \int_0^1 \text{Binomial}(y_i|\pi) \times \text{Beta}(\pi|\mu, \gamma) d\pi \\ &= \int_0^1 P(y_i, \pi|\mu, \gamma) d\pi \\ &= \frac{N!}{y_i!(N-y_i)!} \prod_{j=0}^{y_i-1} (\mu + \gamma j) \prod_{j=0}^{N-y_i-1} (1 - \mu + \gamma j) \prod_{j=0}^{N-1} (1 + \gamma j)\end{aligned}$$


Poisson Distribution

Poisson Distribution

- Begin with an observation period:


Poisson Distribution

- Begin with an observation period:



Poisson Distribution


- Begin with an observation period:



- All assumptions are about the events that occur between the start and when we observe the count. The process of event generation is assumed not observed.

Poisson Distribution


- Begin with an observation period:



- All assumptions are about the events that occur between the start and when we observe the count. The process of event generation is assumed not observed.
- 0 events occur at the start of the period

Poisson Distribution


- Begin with an observation period:



- All assumptions are about the events that occur between the start and when we observe the count. The process of event generation is assumed not observed.
- 0 events occur at the start of the period
- Only observe number of events at the end of the period

Poisson Distribution


- Begin with an observation period:



- All assumptions are about the events that occur between the start and when we observe the count. The process of event generation is assumed not observed.
- 0 events occur at the start of the period
- Only observe number of events at the end of the period
- No 2 events can occur at the same time

Poisson Distribution

- Begin with an observation period:



- All assumptions are about the events that occur between the start and when we observe the count. The process of event generation is assumed not observed.
- 0 events occur at the start of the period
- Only observe number of events at the end of the period
- No 2 events can occur at the same time
- $\Pr(\text{event at time } t \mid \text{all events up to time } t - 1)$ is constant for all t .

Poisson Distribution

- First principles imply:



$$\text{Poisson}(y|\lambda) = \begin{cases} \frac{e^{-\lambda} \lambda^{y_i}}{y_i!} & \text{for } y_i = 0, 1, \dots \\ 0 & \text{otherwise} \end{cases}$$

Poisson Distribution

- First principles imply:



$$\text{Poisson}(y|\lambda) = \begin{cases} \frac{e^{-\lambda} \lambda^{y_i}}{y_i!} & \text{for } y_i = 0, 1, \dots \\ 0 & \text{otherwise} \end{cases}$$

- $E(Y) = \lambda$

Poisson Distribution

- First principles imply:

-

$$\text{Poisson}(y|\lambda) = \begin{cases} \frac{e^{-\lambda} \lambda^{y_i}}{y_i!} & \text{for } y_i = 0, 1, \dots \\ 0 & \text{otherwise} \end{cases}$$

- $E(Y) = \lambda$
- $V(Y) = \lambda$

Poisson Distribution

- First principles imply:

-

$$\text{Poisson}(y|\lambda) = \begin{cases} \frac{e^{-\lambda} \lambda^{y_i}}{y_i!} & \text{for } y_i = 0, 1, \dots \\ 0 & \text{otherwise} \end{cases}$$

- $E(Y) = \lambda$
- $V(Y) = \lambda$
- That the variance goes up with the mean makes sense, but should they be equal?


Poisson Distribution

- First principles imply:

-

$$\text{Poisson}(y|\lambda) = \begin{cases} \frac{e^{-\lambda} \lambda^{y_i}}{y_i!} & \text{for } y_i = 0, 1, \dots \\ 0 & \text{otherwise} \end{cases}$$

- $E(Y) = \lambda$
- $V(Y) = \lambda$
- That the variance goes up with the mean makes sense, but should they be equal?



Poisson Distribution

Poisson Distribution

- If we assume Poisson dispersion, but $Y|X$ is *over-dispersed*, standard errors are too small.

Poisson Distribution

- If we assume Poisson dispersion, but $Y|X$ is *over-dispersed*, standard errors are too small.
- If we assume Poisson dispersion, but $Y|X$ is *under-dispersed*, standard errors are too large.

Poisson Distribution

- If we assume Poisson dispersion, but $Y|X$ is *over-dispersed*, standard errors are too small.
- If we assume Poisson dispersion, but $Y|X$ is *under-dispersed*, standard errors are too large.
- How to simulate? We'll use canned random number generators.

Gamma Density

- Used to model durations and other nonnegative variables

Gamma Density

- Used to model durations and other nonnegative variables
- We'll use first to generalize the Poisson

Gamma Density

- Used to model durations and other nonnegative variables
- We'll use first to generalize the Poisson
- Parameters: $\phi > 0$ is the mean and $\sigma^2 > 1$ is an index of variability.

Gamma Density

- Used to model durations and other nonnegative variables
- We'll use first to generalize the Poisson
- Parameters: $\phi > 0$ is the mean and $\sigma^2 > 1$ is an index of variability.
- Moments: mean $E(Y) = \phi > 0$ and

Gamma Density

- Used to model durations and other nonnegative variables
- We'll use first to generalize the Poisson
- Parameters: $\phi > 0$ is the mean and $\sigma^2 > 1$ is an index of variability.
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \phi(\sigma^2 - 1)$

Gamma Density

- Used to model durations and other nonnegative variables
- We'll use first to generalize the Poisson
- Parameters: $\phi > 0$ is the mean and $\sigma^2 > 1$ is an index of variability.
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \phi(\sigma^2 - 1)$

$$\text{gamma}(y|\phi, \sigma^2) = \frac{y^{\phi(\sigma^2-1)^{-1}-1} e^{-y(\sigma^2-1)^{-1}}}{\Gamma[\phi(\sigma^2-1)^{-1}](\sigma^2-1)^{\phi(\sigma^2-1)^{-1}}}$$

Negative Binomial

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \sigma^2 \phi$

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \sigma^2 \phi$
- Allows *over-dispersion*: $V(Y) > E(Y)$.

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \sigma^2 \phi$
- Allows *over-dispersion*: $V(Y) > E(Y)$.
- As $\sigma^2 \rightarrow 1$, $\text{NegBin}(y|\phi, \sigma^2) \rightarrow \text{Poisson}(y|\phi)$ (i.e., small σ^2 makes the variation from the gamma vanish)

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \sigma^2 \phi$
- Allows *over-dispersion*: $V(Y) > E(Y)$.
- As $\sigma^2 \rightarrow 1$, $\text{NegBin}(y|\phi, \sigma^2) \rightarrow \text{Poisson}(y|\phi)$ (i.e., small σ^2 makes the variation from the gamma vanish)

How to simulate (and first principles)

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \sigma^2 \phi$
- Allows *over-dispersion*: $V(Y) > E(Y)$.
- As $\sigma^2 \rightarrow 1$, $\text{NegBin}(y|\phi, \sigma^2) \rightarrow \text{Poisson}(y|\phi)$ (i.e., small σ^2 makes the variation from the gamma vanish)

How to simulate (and first principles)

- Choose $E(Y) = \phi$ and σ^2

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \sigma^2 \phi$
- Allows *over-dispersion*: $V(Y) > E(Y)$.
- As $\sigma^2 \rightarrow 1$, $\text{NegBin}(y|\phi, \sigma^2) \rightarrow \text{Poisson}(y|\phi)$ (i.e., small σ^2 makes the variation from the gamma vanish)

How to simulate (and first principles)

- Choose $E(Y) = \phi$ and σ^2
- Draw $\tilde{\lambda}$ from $\text{gamma}(\lambda|\phi, \sigma^2)$.

Negative Binomial

- Same logic as the beta-binomial generalization of the binomial
- Parameters $\phi > 0$ and dispersion parameter $\sigma^2 > 1$
- Moments: mean $E(Y) = \phi > 0$ and variance $V(Y) = \sigma^2 \phi$
- Allows *over-dispersion*: $V(Y) > E(Y)$.
- As $\sigma^2 \rightarrow 1$, $\text{NegBin}(y|\phi, \sigma^2) \rightarrow \text{Poisson}(y|\phi)$ (i.e., small σ^2 makes the variation from the gamma vanish)

How to simulate (and first principles)

- Choose $E(Y) = \phi$ and σ^2
- Draw $\tilde{\lambda}$ from $\text{gamma}(\lambda|\phi, \sigma^2)$.
- Draw Y from $\text{Poisson}(y|\tilde{\lambda})$, which gives one draw from the negative binomial.

Negative Binomial Derivation

Negative Binomial Derivation

Recall:

Negative Binomial Derivation

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B)\Pr(B)$$

Negative Binomial Derivation

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B)\Pr(B)$$

$$\text{NegBin}(y|\phi, \sigma^2) = \int_0^\infty \text{Poisson}(y|\lambda) \times \text{gamma}(\lambda|\phi, \sigma^2) d\lambda$$

Negative Binomial Derivation

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B)\Pr(B)$$

$$\begin{aligned}\text{NegBin}(y|\phi, \sigma^2) &= \int_0^\infty \text{Poisson}(y|\lambda) \times \text{gamma}(\lambda|\phi, \sigma^2) d\lambda \\ &= \int_0^\infty P(y, \lambda|\phi, \sigma^2) d\lambda\end{aligned}$$

Negative Binomial Derivation

Recall:

$$\Pr(A|B) = \frac{\Pr(AB)}{\Pr(B)} \implies \Pr(AB) = \Pr(A|B)\Pr(B)$$

$$\begin{aligned}\text{NegBin}(y|\phi, \sigma^2) &= \int_0^\infty \text{Poisson}(y|\lambda) \times \text{gamma}(\lambda|\phi, \sigma^2) d\lambda \\ &= \int_0^\infty P(y, \lambda|\phi, \sigma^2) d\lambda \\ &= \frac{\Gamma\left(\frac{\phi}{\sigma^2-1} + y_i\right)}{y_i! \Gamma\left(\frac{\phi}{\sigma^2-1}\right)} \left(\frac{\sigma^2-1}{\sigma^2}\right)^{y_i} (\sigma^2)^{\frac{-\phi}{\sigma^2-1}}\end{aligned}$$