


Probability



John N. Guidi
 Lecture Notes - Spring 1998
 MIT Course 18.313
 Professor Gian-Carlo Rota
 Volume 1 of 2
 (Lectures 1–21)

This publication of *Probability: Volumes 1 & 2* (combined) is dedicated to John N. Guidi (1954-2012) whose almost "verbatim" notes in Prof. Gian-Carlo Rota's courses in 1998 at MIT faithfully reproduce both the content and erudition of Rota's famous lectures just before Rota's premature death in 1999.

John N. Guidi
Visiting Scholar in Applied Mathematics
Department of Mathematics
Room 2-392
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
Email: guidi@mit.edu

Copyright ©1998 John N. Guidi. All rights reserved.

These *Lecture Notes* originated from the lectures presented by Gian-Carlo Rota, Professor of Applied Mathematics, for course 18.313 - Probability, which he taught at MIT, during the Spring 1998 semester. The *Lecture Notes* were produced from notes I made during class, audio recordings I made of lectures, as well as clarifications and expansions I made of the material presented, after the fact. Although no attempt was made to provide a literal transcription of the lectures, an effort was made to present the material here in the same spirit and I hope to have introduced only a limited number of errors. The *Lecture Notes* are comprised of a two volume set. Volume 1 contains lectures 1–22, volume 2 contains lectures 22–35 and super class lectures 1–4.

Sara Billey, Assistant Professor of Applied Mathematics, who was the course director and whose recitations I attended, presented lectures 20, 21, and 27. The other recitation instructors were Alex Perlin and Ioana Popescu. Their efforts clearly contributed to the success of this course.

I was absent from lecture 25 and my thanks go to Jeff Lieberman, who taped the lecture for me and provided his notes, as well as to Ira Gerhardt and Rick Monte, who also provided their notes of this class.

During the course, I found that preparation of these *Lecture Notes* was a particularly useful way to really understand the material (or, as Professor Rota is fond of saying, “*to really rub it in*”). My gratitude goes to him for imparting his enthusiasm and sharing his profound knowledge of mathematics, as well as for presenting these superb lectures on probability.

John N. Guidi
October 15, 1998

Within the text, pagination is of the form *lecture_date.page*. Note that a different convention is used in both the table of contents and the index, where pagination is of the form *lecture_number.page*. The table of contents provides the mapping between *lecture_number* and *lecture_date*. For example, page 1.1 in the table of contents and index corresponds to page 2/4/98.1 in the text, where the topic anomaly detection is discussed.

Contents

Research Problems

Volume 1

Lecture 1 - [2/4/98]

1. Unsolved Problems of Probability	1.1
1.1 Cluster Analysis	1.1
1.2 Anomaly Detection	1.1
1.3 Hard Spheres	1.1
1.4 Self Avoiding Random Walk	1.2
2. Sample Space	1.3
2.1 Events	1.3
3. Probability	1.3
4. Consequences of the Axioms of Probability	1.4
5. Independence	1.5
6. Example - Coin Tossing (Bernoulli Process)	1.7

Lecture 2 - [2/6/98]

1. Sample Spaces - The Bernoulli Process (cont'd)	2.1
2. Review	2.1
3. Continuity Property of Probability	2.3
4. Independence - The Fundamental Notion of Probability	2.4
5. Bernoulli Process	2.6
5.1 Sample Points	2.6
5.2 Events	2.6
5.3 Probability	2.7
6. Example - Probability that a Run of n Heads ever occurs	2.9

Lecture 3 - [2/9/98]

1. The Bernoulli Process (cont'd)	3.1
2. Borel-Cantelli Lemma	3.2
3. Tail Event	3.8
4. Kolmogorov Zero-One Law	3.8
4.1 Example	3.9
5. Philosophical Interpretation of Probability	3.10

Lecture 4 - [2/11/98]

1. The Theory of Distribution and Occupancy (beg'g)	4.1
2. Maxwell-Boltzmann Sample Space	4.1
3. Quantities	4.2
3.1 Number of Permutations of a Set	4.2
3.2 Binomial Coefficient	4.2
3.3 Lower Factorial	4.3
3.4 Difference Operator	4.3
3.5 Rising Factorial	4.4
3.6 Backwards Difference Operator	4.4
3.7 Multiset Coefficients	4.4
3.8 Multinomial Coefficient	4.4
4. Binomial Theorem	4.4
5. Examples - Maxwell-Boltzmann Statistics	4.5
6. Distribution Interpretation	4.6
7. Occupancy Interpretation	4.6
8. Occupation Numbers	4.7
9. Proof of Binomial Theorem	4.7
10. Probability of Finding all Occupation Numbers	4.9
10.1 Multinomial Coefficient	4.9
11. Balls into Boxes	4.10
12. Dispositions	4.11

Lecture 5 - [2/13/98]

1. The Theory of Distribution and Occupancy (cont'd)	5.1
2. Maxwell-Boltzmann Statistics	5.1
2.1 Occupancy Interpretation	5.2
3. Occupation Number - Maxwell-Boltzmann Statistics	5.2
4. Dispositions	5.3
4.1 Occupation numbers!dispositions	5.4

4.2	Binomial identity	5.4
4.3	Probability Independent of Given Occupation Numbers . . .	5.5
5.	Bose-Einstein Statistics	5.7
5.1	Number of Sample Points - Proof 1	5.7
5.2	Number of Sample Points - Proof 2	5.8
5.3	Number of Sample Points - Proof 3	5.8
5.4	Interpretation	5.10
5.5	Sampling with Replacement	5.10
5.6	Multiset	5.10

Lecture 6 - [2/17/98]

1.	The Theory of Distribution and Occupancy (cont'd)	6.1
1.1	Maxwell-Boltzmann Statistics	6.1
1.2	Dispositions	6.1
1.3	Bose-Einstein Sample Space	6.2
2.	Fermi-Dirac Sample Space	6.4
3.	Probabilistic Identities	6.5
3.1	Binomial Coefficients	6.5
3.2	Multiset Coefficients	6.5
4.	Inclusion-Exclusion Principle	6.7
5.	Rényi's Principle	6.9
5.1	An Application	6.9
6.	Matching - Application of Rényi's Principle	6.12

Lecture 7 - [2/18/98]

1.	Inclusion/Exclusion Principle (cont'd)	7.1
2.	Example	7.1
2.1	Matching	7.1
2.2	All Boxes Occupied in Maxwell-Boltzmann Statistics	7.4
2.3	At Least One Box Occupied in Bose-Einstein Statistics . . .	7.6
3.	Probability that Exactly One Event Happens	7.8

Lecture 8 - [2/20/98]

1.	Random Variables (Integer)	8.1
2.	Probability Distribution	8.1
3.	Identically Distributed Random Variables	8.2
4.	Independent Random Variables	8.4

5.	Maxwell-Boltzmann Statistics	8.5
5.1	Occupancy Numbers Not Independent	8.5
5.2	Position Random Variables	8.6
6.	Abstract Definition of Maxwell-Boltzmann Statistics	8.7
7.	Famous Random Variables in the Bernoulli Process	8.7
8.	Outcome of n^{th} Toss	8.7
9.	Number of Heads in First n Tosses	8.8
9.1	Proof of the Binomial Theorem	8.8
10.	Waiting Time for First Head	8.9
10.1	Geometric Distribution	8.12
11.	Waiting Time for k^{th} Head	8.12
11.1	Negative Binomial Distribution	8.13

Lecture 9 - [2/23/98]

1.	Random Variables (Integer) (cont'd)	9.1
2.	Event of Null Probability	9.1
3.	Review	9.2
4.	Gaps between Successive Heads	9.3
5.	Sampling	9.6
5.1	Sampling with Replacement	9.6
5.2	Sampling without Replacement	9.7
6.	Sampling with Replacement	9.8
6.1	Binomial Theorem	9.8
7.	Sampling without Replacement	9.9
7.1	The Wrong Way - with Fermi-Dirac Statistics	9.10
7.2	Hypergeometric Distribution	9.11
8.	Infinite Sampling with Replacement	9.12
9.	Joint Distribution	9.13

Lecture 10 - [2/25/98]

1.	Random Variables: Joint Distribution and Expectation	10.1
2.	Stochastic Process	10.2
3.	Joint Distribution	10.3
3.1	Occupation Numbers in Maxwell-Boltzmann	10.3
3.2	Combinatorial Identity	10.5
3.3	Multivariate Hypergeometric Distribution	10.5
3.4	Vandermonde's Identity	10.7
4.	Expectation - a Fundamental Concept	10.8

5. Additivity of Expectations	10.8
6. Expectation	10.10
6.1 Maxwell-Boltzmann	10.10
6.2 Multivariate Hypergeometric Distribution	10.12
6.3 Bernoulli	10.14

Lecture 11 - [2/27/98]

1. Conditional Probability	11.1
2. Conditional Probability is a Probability	11.2
3. Example	11.4
3.1 Maxwell-Boltzmann Statistics	11.4
3.2 Bernoulli Process	11.6
4. Basic Techniques for Working with Conditional Probabilities	11.7
4.1 Rule of Successive Probabilities	11.7
4.2 Law of Alternatives	11.8
4.3 Law of Successive Conditioning	11.11

Lecture 12 - [3/2/98]

1. Conditional Probability (cont'd)	12.1
2. Review	12.1
3. Sample Spaces on Probability Trees	12.3
3.1 Vertices for Events	12.4
3.2 Edges for Conditional Probabilities	12.4
4. Bayesian Theory	12.5
5. Classical Example - The Test Problem	12.6
6. Probability Tree - Maxwell-Boltzmann with Replacement	12.9

Lecture 13 - [3/6/98]

1. Conditional Probability (cont'd)	13.1
2. Motivating Examples	13.1
3. Probability Tree - Bernoulli Process	13.4
3.1 Examples	13.4

Lecture 14 - [3/9/98]

1. Law of Alternatives for Conditional Probability	14.1
1.1 Examples - Bernoulli Process	14.2

2. Problem of Scientific Induction	14.8
3. Bayes' Law	14.8

Lecture 15 - [3/11/98]

1. Digressions	15.1
1.1 Expectation with Finite and Infinite Sample Spaces	15.1
1.2 Occupancy -vs- Distribution Perspectives	15.4
1.3 Bayes' Law	15.9
2. Bayes' Law	15.10
3. Uniform Prior	15.11
3.1 Fermi-Dirac Statistics	15.11
3.2 Binomial Identity	15.13

Lecture 16 - [3/13/98]

1. Bayesian Theory (cont'd)	16.1
2. Binomial Identity	16.1
2.1 Fermi-Dirac Position of First Check	16.1
2.2 Fermi-Dirac Position of Second Check	16.3
3. Bayes' Theorem for Sampling without Replacement	16.5
3.1 Uniform Prior	16.6
3.2 Conjugate Prior	16.6
4. Bayes' theorem!sampling with replacement	16.9
4.1 Uniform Prior	16.10

After Notes Lecture AN16 - [3/16/98]

1. Conjugate Priors	AN16.1
2. Sampling without Replacement	AN16.1
2.1 Uniform Prior	AN16.2
2.2 Conjugate Prior	AN16.2
3. Posterior's for Given Conjugate Prior's	AN16.4
4. Sampling with Replacement	AN16.6
4.1 Uniform Prior	AN16.7
4.2 Conjugate Prior	AN16.8
4.3 Discrete Beta Function Identity Series	AN16.8

Lecture 17 - [3/16/98]

1. Continuous Random Variables	17.1
1.1 General Definition of a Random Variable	17.1
1.2 Independence of Random Variables	17.2
1.3 Cumulative Distribution Function	17.2
1.4 Density of Continuous Random Variable	17.3
2. Example - The Uniform (Dirichlet) Process	17.5
3. Given Occupation Numbers in Uniform Process Intervals	17.7
4. Order Statistics	17.9
4.1 Cumulative Distribution	17.10
4.2 Density	17.11
5. Continuous Random Variable Taking a Single Value	17.12

Lecture 18 - [3/18/98]

1. Continuous Random Variables (cont'd)	18.1
2. Characterization of Density	18.1
3. Cumulative Distribution	18.2
4. Median of a Random Variable	18.3
5. Expectation of a Continuous Random Variable	18.5
6. Example - Uniform (Dirichlet) Distribution	18.5
6.1 Density - the Hard Way	18.8
6.2 Density - the Easy Way	18.9
6.3 Beta Function Integral	18.11
7. Gaps between Points Dropped on a Line	18.12
7.1 A Purely Probabilistic Argument	18.12
8. Expectation of the Order Statistics	18.14
8.1 Binomial Identity	18.14

Lecture 19 - [3/20/98]

1. Order Statistics	19.1
2. The k^{th} Order Statistic of the Dirichlet Process	19.3
2.1 Events	19.3
2.2 Probability	19.3
2.3 Density Function	19.4
3. Density Curves	19.4
4. Joint Distribution of 2 Continuous Random Variables	19.5
5. Marginal Density	19.6
6. Joint Density as Derivative of Joint Cumulative Distribution	19.7

Lecture 20 (Sara Billey) - [3/30/98]

1. Join Distribution	20.1
2. Join Density	20.1
3. Example - Dropping Points Uniformly on a Triangle	20.1
4. Joint Density and Joint Distribution	20.6
4.1 Order Statistics of the Uniform Process	20.6
5. Expectation	20.9
6. Expectation of Order Statistics	20.11

Lecture 21 (Sara Billey) - [4/1/98]

1. In Class Review (no notes)	21.1

Volume 2

Lecture 22 - [4/6/98]

1. Continuous Conditional Probability	22.1
2. Dirichlet Process	22.2
2.1 Cumulative Distribution	22.3
2.2 Density	22.3
3. Example - Probability of Events with Order Statistics	22.4
3.1 Integrating the Joint Density	22.5
4. Conditional Probability	22.7
4.1 By the Book	22.7
4.2 By Intuition	22.7
5. Example - Using Inclusion-Exclusion Principle	22.9
6. Conditional Density	22.11

Lecture 23 - [4/8/98]

1. Dirichlet Process - Dropping Two Points	23.1
1.1 Intuition	23.1
1.2 Rigorous Computation	23.2
2. Dirichlet Process - Dropping Arbitrary Number of Points	23.3
2.1 Intuition	23.3
2.2 Rigorous Computation	23.4
3. Continuous Version of the Law of Alternatives	23.5
3.1 Proof	23.5

4. Gaps in the Uniform Process	23.8

Lecture 24 - [4/10/98]

1. Continuous Analogue of the Law of Alternatives	24.1
2. Example - Needles on a Stick	24.1
3. Example - Bernoulli Process: Is the Coin Fair?	24.7
3.1 Continuous Analogue of Bayes' Law	24.8
3.2 Uniform Prior gives Beta Function Integral	24.11
3.3 Density Plots	24.13
4. Confidence Intervals	24.14

Lecture 25 - [4/13/98]

1. Beta Function	25.1
2. Continuous Law of Alternatives	25.1
3. Continuous Analogue of Bayes' Law	25.1
4. Use of Uniform Density	25.2
5. Confidence Intervals	25.3
6. Bayes' Estimate	25.4
7. Application - the Laplace Law of Succession	25.5
8. Conjugate Prior	25.6
9. Bayes' Estimate using Conjugate Prior	25.7
10. The Algebra of Probability Distributions	25.9
11. Density of a Function in Terms of Inverse Function	25.13
12. Complicated Densities	25.14

Lecture 26 - [4/15/98]

1. The Algebra of Probability Densities (cont'd)	26.1
2. Densities	26.1
3. Joint Densities	26.3
3.1 Identity Analogous to Conditional Probability	26.3
4. Joint Density with Complicated Random Variables	26.5
5. Joint Density of Sum of Random Variables	26.6
6. Convolution - Joint Density of Sum of Independent Random Variables	26.7
7. Joint Density of Independent, Uniformly Distributed Random Variables	26.9
7.1 Indicator Random Variables	26.9

8. Cauchy's Functional Equation	26.14

Lecture 27 (Sara Billey) - [4/17/98]

1. Examples	27.1
1.1 Record Values	27.1
1.2 Expectation of Random Variable Plus a Constant	27.4
1.3 Runners on a Track	27.8
1.4 Inspector's Paradox	27.11

Lecture 28 - [4/22/98]

1. The Poisson Process	28.1
2. Back to the Bernoulli Process	28.2
3. Continuous Waiting Time	28.4
3.1 Memoryless Property	28.4
3.2 Cauchy's Functional Formula	28.5
3.3 Density of a Memoryless Waiting Time	28.6
3.4 Expectation of a Memoryless Waiting Time	28.7
4. Exponential Distribution	28.8
5. Sample Space of the Poisson Process	28.9
5.1 Physicists' Definition of the Poisson Process	28.10
5.2 Rare Events	28.10
5.3 Poisson Event	28.11
6. Probability of Poisson Process	28.12
6.1 Main Point	28.13
6.2 Axiom	28.14
7. Justification of Derivation of Exponential Distribution	28.15

Lecture 29 - [4/24/98]

1. Poisson Process - Motivation	29.1
2. The Four Fundamental Stochastic Processes	29.2
2.1 Bernoulli Process	29.2
2.2 Uniform Process	29.2
2.3 Poisson Process	29.2
2.4 Processes Pertaining to the Normal Distribution	29.2
3. Poisson Events	29.2
3.1 Poisson Probability	29.2
4. Random Function	29.3

5. Poisson Probability Distribution with Intensity α	29.4
6. Number of Blips in Disjoint Intervals is Independent	29.4
7. Expectation of a Poisson Process Random Function	29.5
8. The 7 Fundamental Properties of the Poisson Process	29.6
9. Property 1 - Law of Rare Events	29.6
10. Property 2 - Limit of Uniform Process	29.9
11. Property 3 - Memorylessness	29.10
11.1 Gamma Distribution	29.11

Lecture 30 - [4/29/98]

1. The 7 Fundamental Properties of the Poisson Process (cont'd) . . .	30.1
2. Property 4 - Uniform Process Obtained by Conditioning Poisson Process	30.4
2.1 Example - Needles on a Stick	30.6
3. Property 5 - Schrödinger Randomization	30.9
3.1 Maxwell-Boltzmann Statistics from Conditioning Poisson Pro- cess	30.10
3.2 Any Problem of Occupation Numbers	30.13
4. Property 6 - Poisson Blips in 2 Colors	30.17

Lecture 31 - [5/1/98]

1. Property 6 - Poisson Blips in 2 Colors (cont'd)	31.1
2. Property 7 - Laplace Transform	31.4
3. Random Walk	31.5
4. Time of First Return to Origin	31.7
4.1 z Transform	31.10
5. Continuous Analogue of Random Walk	31.13

Lecture 32 - [5/4/98]

1. Random Walk	32.1
2. Standard Normal Distribution	32.2
3. Standard Normal Distributed Random Variable	32.5
4. Variance of a Random Variable	32.6
5. Views of a Random Variable	32.7
5.1 A Random Phenomenon	32.7
5.2 Result of a Search	32.7
5.3 Measurement of an Imperfect Quantity	32.7

6. Variance Addition Formula	32.8

Lecture 33 - [5/6/98]

1. Variance (cont'd)	33.1
2. Density Plots	33.1
3. Properties of Variance	33.2
4. Examples	33.2
4.1 Bernoulli Process	33.2
4.2 Poisson Process	33.5
4.3 Exponential Random Variable	33.6
5. Expectation of a Linear Function of a Random Variable	33.8
6. Standard Deviation of a Linear Function of a Random Variable	33.8
7. Standardized (Normalized) Random Variable	33.8
8. Conformation of Properties of Normalized Random Variables	33.10
8.1 Expectation	33.10
8.2 Variance	33.10
9. Density of Linear Function of Normalized Random Variables	33.12
10. n Measurements of the Same Quantity	33.14
11. Statistics in One Easy Lesson	33.15
12. Likelihood	33.17

Lecture 34 - [5/8/98]

1. Normal Distribution (cont'd)	34.1
2. Basic Rule of Statistics	34.2
3. Bayes' Law for Densities	34.3
4. Priors	34.5
4.1 Honest Choice	34.5
4.2 Dishonest Choice	34.5
5. Measure	34.6
6. Prediction of Average of n Measurements	34.7
7. Probability Tables of Standard Normal Random Variable	34.8
8. Example - Fair Coin?	34.9
9. The Central Limit Theorem	34.11
9.1 Justification 1 - Psychological	34.11
9.2 Example - Coin Tosses to Choose Airline Seats	34.13

Lecture 35 - [5/11/98]

1. Gaussian (Normal) Distribution	35.1
2. Normalizing Random Variables	35.2
3. Gaussian (Normal) Distribution Justifications	35.3
3.1 The Central Limit Theorem (cont'd)	35.3
3.2 Maxwell-Einstein Derivation	35.4
3.3 Wiener's Characterization	35.7
4. The Law of Large Numbers	35.12
5. The Strong Law of Large Numbers	35.13

Super Class Lecture SC1 - [2/27/98]

1. Proof of Inclusion-Exclusion Principle	SC1.1
1.1 Indicator Random Variable	SC1.1
2. More on Rényi's Principle	SC1.5
3. Circuit Theoretic Interpretation of Boolean Operations	SC1.7
4. Algebra of Partitions	SC1.10

Super Class Lecture SC2 - [3/13/98]

1. Probability Trees	SC2.1
2. Urn Models	SC2.2
2.1 Pólya Urn Model	SC2.3
3. Balls Distinguishable, Boxes Indistinguishable	SC2.4
4. Balls Indistinguishable, Boxes Indistinguishable	SC2.5
4.1 Partition of an Integer into Summands	SC2.5
5. Reluctant Functions	SC2.7
6. Balls in the Box Form a Rooted Tree	SC2.8
6.1 Disposition	SC2.8
6.2 Abel Polynomials	SC2.8
7. Partitions	SC2.9
7.1 Bell Numbers	SC2.9
7.2 Stirling Number's of the Second Kind	SC2.10

Super Class Lecture SC3 - [4/10/98]

1. Needles on a Stick (again)	SC3.1
2. Buffon Needle Problem	SC3.2
2.1 From Needles to Curves	SC3.4
2.2 Computing π	SC3.5

3. Convex Set	SC3.7
4. Sylvester's Theorem	SC3.7
5. Measure	SC3.9
6. Measure of Lines in the Plane	SC3.11
7. Computing Perimeters	SC3.12

Super Class Lecture SC4 - [4/24/98]

1. Entropy and Information	SC4.1
2. Partition	SC4.1
3. Fundamental Properties of Entropy	SC4.3
3.1 Independent Partitions	SC4.3
3.2 Finer the Partition, Bigger the Entropy	SC4.5
3.3 Entropy of Meet	SC4.5
4. Maximum Entropy Principle	SC4.6
5. Information Search	SC4.7
6. Famous Problem of Scales	SC4.9
7. Shannon Coding Theorem	SC4.12

Index

Within the text, pagination is of the form *lecture_date.page*. Note that a different convention is used in both the table of contents and the index, where pagination is of the form *lecture_number.page*. The table of contents provides the mapping between *lecture_number* and *lecture_date*. For example, page 1.1 in the table of contents and index corresponds to page 2/4/98.1 in the text, where the topic anomaly detection is discussed.

Research Problems

1. Given a family of intersections of events, where the probability of each intersection is the product, determine if the events are independent 2.5
2. Show that the sample space and probability defined for the Bernoulli Process define a unique probability on all elements 3.1
3. Let A_n be event that all even number of tosses have equal numbers of 0's and 1's. Determine the probability that infinitely many of the A_n occur 3.7
4. Prove the Kolmogorov Zero-one Law 3.8
5. Provide an elegant simplification for the expression that gives the first occupation number in Bose-Einstein statistics 5.10
6. Provide a satisfactory explanation of the identity relating binomial and multiset coefficients 6.6
7. Given a number of random variables, of which certain subsets of intersections are independent, determine whether the random variables are independent 8.5
8. Relate Bose-Einstein statistics to dispositions and explain, with respect to sampling 9.11
9. Given a uniform process on the interval from zero to an unknown endpoint, compute the Bayes' estimate and determine which priors make sense 25.8
10. Find a nice formula for the density of n independent, uniformly distributed random variables 26.14

The idea of probability is one of the great ideas developed in this century + you can find it in all the sciences.


Unsolved Problems of Probability

- Cluster Analysis
 - Anomaly Detection
- } There are literally millions, perhaps billions, of dollars to be made if someone were to make an important scientific advance in any of these problems.

These problems are not just hard to solve - they are hard to state. We know roughly what we want, but we don't know precisely what we are going to look at.

Cluster Analysis

You are given sets of multidimensional data in the plane.



Somehow, these data cluster. You want a taxonomical criteria for separating data into clusters.

How do they do it today?

Like this.

They take a pencil and go like this.

That's all we know to this day.

There is no categorization theory.

Anomaly Detection

A very closely related problem.

You have data (for example, income tax data) and you want some sort of criteria for spotting anomalous data - someone who is cheating.

The IRS has allocated \$8 Billion



These are problems that are very much talked about in our time.

Some older problems of probability that come closer to the exact sciences:

- Hard Spheres - the name has been given to this problem by physicists.

The hard spheres problem is best stated for the plane.

You have a supply of identical pennies and you drop them on the carpet. What is the probability that no two will overlap?



Q: How big is the carpet?
A: A big carpet

2/4/98.2



So, the probability that the two of them overlap is unknown can be very difficult. It has been solved for $n=3$ + $n=4$ pennies.

This is a typical example of an extremely simple problem in probability for which very little is known.

In 3-dimensions, this would be spheres. People working in statistical mechanics would love to have the solution to the analogous problem.

• Self avoiding random walk

Another famous problem that arises in chemistry.



You have a grid of points (in the plane, or in space) w/ integer coordinates.

You start at the origin.

You have a fair die w/ 4 sides.

You flip it and, according to which side it shows, you take a step in the corresponding direction. Etc.

This is called a random walk.

It is one of the most studied stochastic processes

Self avoiding random walk problem is, again, an extremely simple problem to state.

It asks for the probability that after n steps, the random walk never visits a site that it has already occupied.

The basic definition of probability is the notion of a sample space.

Sample Space

plus

• Ω = a set } which, by abuse of speech, is sometimes called a sample space

• a family of subsets of Ω , called events, say $\mathcal{E}$

Events satisfy the following properties:

1) ϕ is an event
↳ empty set

2) if A is an event, then so is $A^c = \Omega - A$

3) if A and B are events (i.e., if $A, B \in \mathcal{E}$)
then $A \cap B$ is also an event

↳ $A+B$ are members of the family of events.

4) if $A_1, A_2, A_3, \dots$ is a finite or infinite sequence of disjoint events (i.e., $A_i \cap A_j = \phi$, $i \neq j$) then
 $A_1 \cup A_2 \cup A_3 \cup \dots$ is also an event

This is called the Countable Additivity Axioms

A sample space is a set together w/ a family of events

To describe a sample space means not just to describe a set Ω , but a set Ω together w/ a family of subsets (called events), which satisfy these properties.

I'm not giving you an example, because this whole course will be examples of this.

A probability P on a sample space Ω (with events $\mathcal{E}$) is a function that assigns to every event a number between 0 and 1 (i.e., $0 \leq P(A) \leq 1$ for $A \in \mathcal{E}$) subject to:

axioms

↳ whole sample space
(1) $P(\Omega) = 1$

(2) $P(A^c) = 1 - P(A)$

(3) if $A_1, A_2, A_3, \dots$ is a finite or infinite sequence of disjoint events, then:

$$P(A_1 \cup A_2 \cup A_3 \cup \dots) = P(A_1) + P(A_2) + P(A_3) + \dots$$

This completes the definition of the notion of probability.

A sample space with a probability is sometimes called a probability space.

Q: $P(\emptyset)$?

A: By combining axioms (1) & (2) we have:

$$\left. \begin{array}{l} A = \Omega \\ A^c = \emptyset \end{array} \right\} \Rightarrow \begin{array}{l} P(A^c) = 1 - P(A) \\ P(\emptyset) = 1 - 1 \\ = 0 \end{array}$$


Consequences of the axioms of probability

Example $P(A \cup B) = P(A) + P(B) - P(A \cap B)$

it is understood that
 $A+B$ are events

This is intuitively clear.


If you visualize events as sets and if you visualize probability as some sort of measure, as in area, then you have:



$A \cup B$ $\equiv$
 $A \cap B$ $\equiv$

If you add $P(A) + P(B)$, you have overcounted the intersection.

Proof



$$A \cup B = \underbrace{(A \cap B^c) \cup (B \cap A^c) \cup (A \cap B)}_{\text{these 3 events are all disjoint}}$$

Thus, By axiom 3:

$$P(A \cup B) = P(A \cap B^c) + P(B \cap A^c) + P(A \cap B)$$

$$A = (A \cap B^c) \cup (A \cap B)$$

$$B = (B \cap A^c) \cup (A \cap B)$$

$$P(A) = P(A \cap B^c) + P(A \cap B)$$

$$P(B) = P(B \cap A^c) + P(A \cap B)$$

again, by
axiom 3, as
sets are disjoint

Therefore:

term appears twice

$$P(A) + P(B) = \underbrace{P(A \cap B^c) + P(B \cap A^c) + P(A \cap B) + P(A \cap B)}_{P(A \cup B)}$$

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

This is a typical manipulation of events.

The next point is the most important concept of our book.

So far, everything can be read as sets.

But we come to the notion that set theory alone is not adequate.


We come to the notion of independent events.Basic notion of probabilityIndependenceEvents A & B (in a sample space) are said to be independent when

$$P(A \cap B) = P(A)P(B)$$

We will see many interpretations of this notion.

Observation:

- If A & B are independent events, then A & B^c are also independent events

Proof

$$A \cap B^c = A - (A \cap B)$$

the subtraction of 2 sets is well defined only when the second set is a subset of the first.

$$P(A \cap B^c) = P(A - A \cap B) \quad \leftarrow \text{we'll show this next}$$

~~$$= P(A) - P(A \cap B)$$~~

$$= P(A) - P(A \cap B)$$

$$= P(A) - P(A)P(B)$$


$$= P(A)[1 - P(B)]$$

$$P(A \cap B^c) = P(A)P(B^c) \quad (\text{by axiom 2})$$

A & B^c are independent events

(A & B are independent events, thus: $P(A \cap B) = P(A)P(B)$)

- If A & C are events and if $C \subseteq A$
then $P(A-C) = P(A) - P(C)$



$$A-C = A \cap C^c$$

$$A = (A \cap C^c) \cup C$$

union of disjoint events

$$P(A) = P(A \cap C^c) + P(C) \quad (\text{by axiom 2})$$

$$= P(A-C) + P(C) \quad (\text{from above})$$

$$\therefore P(A-C) = P(A) - P(C)$$

Q/A: if you wish, you could condense this to:

$$A = (A-C) \cup C$$

That's it my friends.
From now on we see examples of this.
That's it for the grammar

{we'll discuss later the use of one word}
versus the other.

2/4/98.7

Example 1 Coin Tossing (= Bernoulli {sample space process})

Biased coin that gives Heads w/ probability p ($0 \leq p \leq 1$)
Tails " " $q = 1 - p$

If we toss the coin infinitely many times, what will happen?

Now you say - we can't toss the coin infinitely many times.

By imagining tossing the coin infinitely many times, you include all possible finite tosses of the coin. Therefore, all questions about tossing the coin 1 thousand versus 1 million times are included in our study of tossing the coin infinitely many times.

Here we come to the crucial step. We set up a sample space that reflects all properties of what can happen when we toss a coin infinitely many times. We can work this to answer any question about tossing the coin. These questions can be extremely complicated.

Let me give you an example of one question we will answer shortly.
What is the probability that a run of 3 Heads will occur before a run of 7 Tails?

A run is a continuous, uninterrupted sequence of Heads or of Tails.

The event that a run of 3 Heads will occur before a run of 7 Tails has a definite probability and our mathematical setup should enable us to compute this probability easily.

What is our mathematical setup?

Here we come to a beginning of understanding of those mysterious properties we have been writing so far.

- Sample space Ω $\leftarrow$ consists of all possible outcomes of infinite sequences of tosses of the coin.

$\omega \in \Omega$

$\leftarrow$ a point in the sample space is called a sample point

The sample space for coin tossing for the Bernoulli process consists of all possible sequential tosses of the coin at integer times.

We represent this by ω with an infinite sequence that tells us whether you get Heads or Tails. We could use H/T, but it has become much more customary to use 1/0.

1 = Head

0 = Tail

$$\omega = (1011100 \dots)$$

2/4/98.8

So a sample point is just an infinite sequence of 0 + 1.
All possible sequences of 0 + 1 are sample points.

We will sometimes write a sample point as:

$$\omega = (\omega_1, \omega_2, \omega_3, \dots)$$

$$\text{where } \omega_n = \begin{cases} 0 & \text{tail} \\ 1 & \text{head} \end{cases}$$

Observe that a sample point does not look like a point.
The points of probability do not look like a point of geometry.
In the case of coin tossing, the sample point is an infinite sequence of 0 + 1. This is a long way from being a point.

Q/A: If you consider 0/1 as binary expansion, something very interesting happens. You get the Cantor Set.

Have we defined a sample space? No!

We have to specify which subsets of the sample space are the events.

Intuitively, that corresponds to specifying all possible events whose probability we may ever wish to compute.

How shall we do this?

We first single out the most obvious event

• Event

H_n = event that n^{th} trial results in Heads

Translated into set theoretic language:

H_n = the set of all sample points $\omega \in \Omega$ with property that $\omega_n = 1$

The idea of probability is not to view this as a set.

It is to view it as something that either happens or does not happen.

In probability, you have a novel way of visualizing sets.

You visualize this as the event that the n^{th} toss resulted in Heads is psychologically very different from thinking the set of all sequences of 0 + 1 with the $n^{\text{th}} = 1$. Quite different.

Learning to rethink all of set theory is exactly what you are supposed to do in this course. It's a new way of thinking. That's why this course is hard.

Sample Spaces - The Bernoulli Process (cont'd)

A sample space is the basic structure that models probabilistic phenomena. To set up sample space for various probabilistic phenomena has the property, because it is, in fact, sample spaces, that you will be able to compute the answers to obvious, real life questions about these phenomena.

sample space : Ω = a set plus a family $\mathcal{E}$ of subsets of Ω , called events

You specify, inside the sample space, a family of subsets called events. These are the subsets whose probabilities might be computed.

Why doesn't it take every subset to be an event?
We'll discuss this later.

Events satisfy the following properties:

- (1) $\emptyset \in \mathcal{E}$
- (2) if A is an event, then $A^c \in \mathcal{E}$
- (3) if $A, B \in \mathcal{E}$, then $A \cap B \in \mathcal{E}$
- (4) if $A_1, A_2, \dots$ is a finite or infinite sequence of pairwise disjoint events (i.e., if $A_n \in \mathcal{E}$ for all n) then $\bigcup_n A_n \in \mathcal{E}$

Events specifying a sample space means:

- 1) specifying the set Ω
 - and 2) somehow specifying which subsets are the events
- This specification must be made in such a way that the events satisfy the above properties. It will take us all term to see how this is done. This is an entirely new concept. Don't expect to see the entire meaning of this concept when you see one or two examples. You have to see 60.

Q: For finite sets, is there any reason why you wouldn't take all of the subsets for your events?

A: If Ω is finite set, then it is almost always the case that the events are also.

2/6/98.2

Probability: $P: \mathcal{E} \rightarrow [0, 1]$

A function P that associates to every event a number bt 0 and 1 with the properties:

$$(1) P(\emptyset) = 0 \quad \leftarrow \text{probability that nothing happens is 0}$$

$$(2) P(A^c) = 1 - P(A)$$

Probability that something does not happen is $(1 - \text{probability that it does happen})$

Last time we stated things in terms of sets.

Now we rephrase things in probabilistic language.

$\hookrightarrow$ another way of interpreting set theory.

{Countable
Additivity
Axiom}

$$(3) \text{ if } A_1, A_2, \dots \text{ is a finite or infinite sequence of pairwise disjoint events, then } P\left(\bigcup_n A_n\right) = \sum_n P(A_n)$$


pairwise disjoint = any 2 events are disjoint

These axioms are intuitive, what is not so quickly clear, at first, is why is it that we must admit the possibility of an infinite sequence here. That we have to learn from examples, which are coming.

We have seen last time that from these axioms, one derives all the intuitive and obvious properties of probability.

$$\text{Ex: if } B \supseteq A \text{ then } P(B) \geq P(A)$$

if A happens, then B happens



$$B = (B \cap A^c) \cup A$$

$$P(B) = \underbrace{P(B \cap A^c)}_{\geq 0} + P(A) \quad (\text{by axiom 3})$$

≥ 0 , since $P: \mathcal{E} \rightarrow [0, 1]$

$$\therefore P(B) \geq P(A)$$

This fact is a consequence of the axioms.

2/6/98.3

Another fact that can be obtained as a consequence of the axioms, which may be on the quiz:


Ex: If we have a sequence of events such that:

$$B_1 \subseteq B_2 \subseteq B_3 \subseteq \dots$$

$$\text{then } P(\bigcup_n B_n) = \lim_{n \rightarrow \infty} P(B_n)$$

This follows from countable additivity

This is sometimes called the Continuity Property of Probability



Written as the union of pairwise disjoint events:

$$\bigcup_n B_n = (B_n - B_{n-1}) \cup (B_{n-1} - B_{n-2}) \cup \dots \cup (B_3 - B_2) \cup (B_2 - B_1) \cup B_1$$

The countable additivity axiom gives:

$$P(\bigcup_n B_n) = P(B_n - B_{n-1}) + P(B_{n-1} - B_{n-2}) + \dots + P(B_3 - B_2) + P(B_2 - B_1) + P(B_1)$$

On [2/4/98.6] we proved that if $A + C$ are events and if $C \subseteq A$,
then $P(A - C) = P(A) - P(C)$

Thus we have:

$$P(\bigcup_n B_n) = \underbrace{P(B_n) - P(B_{n-1})} + \underbrace{P(B_{n-1}) - P(B_{n-2})} + \dots + \underbrace{P(B_3) - P(B_2)} + \underbrace{P(B_2) - P(B_1)} + P(B_1)$$

$$P(\bigcup_n B_n) = P(B_n)$$

Thus, as $n \rightarrow \infty$,

$$P(\bigcup_n B_n) = \lim_{n \rightarrow \infty} P(B_n)$$

2/6/98.4

It will be interesting later to see examples where these axioms are not satisfied.
Perfectly innocent probabilities that do not satisfy these axioms.

This is the abstract grammatical scheme of probability, on which all of probability theory is based.

To this, we add:

⊗ The Fundamental Notion of Probability - ^{Notion of} Independence

Events $A + B \leftarrow \left\{ \begin{array}{l} \text{I pass over, in silence, in a given sample space and with a} \\ \text{are said to be} \\ \text{independent when} \end{array} \right. \left. \begin{array}{l} \text{given probability} \end{array} \right\}$

$$P(A \cap B) = P(A)P(B)$$

$\uparrow$ probability of
their joint happening

This is a great discovery of mathematics. That the intuitive idea of 2 events being independent is rendered mathematically with this extremely simple equation.
It's an incredible discovery.

That this corresponds, as we will convince ourselves gradually in the next few weeks, exactly to our intuitive notion of 2 things happening independent of each other. It's a great thing. We tend to dismiss things because we are jaded. But every once in a while you should reflect on how extraordinary those discoveries are. You should recover your sense of wonder from time to time.

One remark that I omitted last time:

This is the definition for independence of two events. What about 3 events?

More generally, events A, B, C are independent when any two of them are independent:

$$\left\{ \begin{array}{l} P(A \cap B) = P(A)P(B), \quad P(A \cap C) = P(A)P(C), \quad P(B \cap C) = P(B)P(C) \\ \text{AND} \\ P(A \cap B \cap C) = P(A)P(B)P(C) \end{array} \right. \leftarrow \begin{array}{l} \text{(probability of all 3 events happening} \\ \text{is the product of the probabilities)} \end{array}$$

For 3 events to be independent, it's not enough for any 2 to be independent.
You have to add this last condition.

Q: Couldn't we just use the last condition: $P(A \cap B \cap C) = P(A)P(B)P(C)$?

A: No. The last condition does not imply the others.

You can construct artificial sample space and artificial probabilities to show that all three conditions are needed. Can show that only some, not all, conditions can be satisfied by artificially assigning probabilities to sample sizes.

2/6/98.5

For more than 3 events, similarly, for all possible intersections of subsets, when you take the probability, it is the product.

Similarly - for any finite family of events [It's pedantic to write this down]

• $\in$ all possible $\cap \Rightarrow P(\cdot) = \text{product}$

Q/A: For 4 events to be independent:

all pairs
all triples
all quadruples } $P(\cdot) = \text{product}$

Otherwise, no independence.

Open question:

It is an open question as to which of these conditions imply which.
In other words, to build an implication table, if I take a certain family of all possible intersections, does it imply the others?
This has not been settled to this day.

That's it.

The intuitive meaning of independence will gradually be found.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



- Bernoulli Process - mathematical model for all questions pertaining to coin tossing. Any question pertaining to coin tossing is answered by computation with the Bernoulli Process. Some of them can be extremely difficult.

Sample points $\omega = (\omega_1, \omega_2, \dots)$ $\leftarrow$ infinite sequences of 0 and 1

$$\omega_n = \begin{cases} 1 \\ 0 \end{cases}$$

$\leftarrow$ (all possible histories of tossing a coin
(List all possible histories of an infinite sequence of tosses.)

Sample space Ω = set of all sample points
plus
Secondly, we have to define events

Events

Define events in 2 steps.

- 1) Pick out events whose probability is obvious

H_n = event that n th toss is Head (1)

The set of all sample points whose n th entry is 1, but we don't say it that way.

As a set, it's an infinite set. $\leftarrow$ {there are ∞ many sample points whose n th entry is 1}

- 2) Take these elementary events + take all events obtained by taking $\cup, \cap, \complement$ and unions of infinite sequences of disjoint events. In all possible ways.

Any subset of Ω obtained from $H_1, H_2, \dots$ by taking $\cup, \cap, \complement$, and unions of infinite sequences of disjoint events is an event. $\leftarrow$ By definition

You may ask - Doesn't this include all subsets of a sample space?

This is an interesting question.

To prove that this does not include all subsets, you need the axiom of choice.

The subsets in (2) above are the only subsets whose probabilities we can hope to compute - because they are obtained from elementary events by operations where the probability can be computed.

Probability for the Bernoulli Process

$$1) P(H_n) = p \quad \Rightarrow P(T_n) = 1 - p = q$$

↑ the bias of the coin

2) Every finite set of H_n is independent

e.g., $P(H_2 \cap H_3) = P(H_2)P(H_3)$

$$P(H_1 \cap H_2^c \cap \dots \cap H_{10}) = P(H_1)P(H_2^c) \dots P(H_{10})$$

} by definition of independence

↑
{ extension of [2/4/98.5]. If $A + B$ are independent events, then $A + B^c$ are independent events. }

Intuitively, we tell this system that successive tosses of the coin are independent events.

We have the following result, which we are not going to prove:

Theorem (1) and (2) uniquely specify P on all events

Once you assume (1) + (2) and the axioms of probability then the one and only one way of computing the probability of each event is set.

We will not prove this. It's hopelessly dull. Proof as research problem.

In this way, we have defined the Bernoulli Process.

What do we do with it?

Example: Let's compute the probability of some event related to coin tossing, using this model.

X $A =$ event that a run of $\overset{||}{\text{Two Heads}}$ occurs before a run of $\overset{oo}{\text{Two Tails}}$

this is not well stated
instead:



✓ $A =$ event that the first run of Z is a run of $\overset{||}{Z \text{ Heads}}$

What is $P(A)$?

2/6/98.8

We have to express, some how, the event A in terms of the events H_n and their intersections. From [2/6/98.7], all we know is:

$$(1) P(H_n) = p$$

(2) Every finite set of H_n is independent

$$\text{e.g., } P(H_1 \cap H_2) = P(H_1)P(H_2)$$

So we need to reduce A to H_n and its intersections.

$$A = (H_1 \cap H_2) \cup (H_1 \cap T_2 \cap H_3 \cap H_4) \cup (H_1 \cap T_2 \cap H_3 \cap T_4 \cap H_5 \cap H_6) \cup \dots \leftarrow \text{exhaustively enumerate all that start w/ Head}$$

$$\cup (T_1 \cap H_2 \cap H_3) \cup (T_1 \cap H_2 \cap T_3 \cap H_4 \cap H_5) \cup (T_1 \cap H_2 \cap T_3 \cap H_4 \cap T_5 \cap H_6 \cap H_7) \cup \dots \leftarrow \text{w/ Tail}$$

This accounts for all possibilities.

We have expressed event A as a union of intersections of elementary events H_n .

Any of these intersections, we know the probability of (from (2) above) since they are independent.

Furthermore, if you take any two of the union terms, they are disjoint.

Why are they disjoint? They differ by the number of H_n , so they can't happen together.

Therefore the Axiom of Countable Additivity allows us to compute the probability by adding these infinite events.

Now you see where the axiom of countable additivity comes in.

If we didn't have this axiom, we couldn't even compute the probability of this event A .

$$P(A) = p^2 + (pg)p^2 + (pg)^2 p^2 + \dots + (pg)^n p^2 + \dots$$

$$+ gp^2 + g(pg)p^2 + g(pg)^2 p^2 + \dots + g(pg)^n p^2 + \dots$$

$$= p^2 (1 + pg + (pg)^2 + \dots + (pg)^n + \dots)$$

$$+ gp^2 (1 + pg + (pg)^2 + \dots + (pg)^n + \dots)$$

$$= p^2 \left(\frac{1}{1-pg} \right) + gp^2 \left(\frac{1}{1-pg} \right)$$

geometric series

$$1 + pg + (pg)^2 + \dots = \frac{1 - (pg)^{\infty}}{1 - pg} \rightarrow 0$$

$$P(A) = \frac{p^2 + gp^2}{1 - pg}$$

Probability that the first run of length 2 is a run of 2 Heads.

2/6/98.9

How do you double check this?

Suppose the coin is fair. Then $P(A)$ better be $\frac{1}{2}$.
If you don't see this, I can not explain it.

The Principle of Ignorance

$$\begin{aligned} p=q=\frac{1}{2} \Rightarrow P(A) &= \frac{p^2 + p^3}{1-p^2} \\ &= \frac{p^2 (1+p)}{(1+p)(1-p)} \\ &= \frac{1}{2} \checkmark \end{aligned}$$

Even in computing one of the simplest questions about coin tossing, we have used the full power of the axioms.

Let's do another one.
Jack it up.

Example: B = event that a run of 9 Heads ever occurs

$$P(B) = 1 - P(B^c)$$

$\uparrow$ B^c = event that 9 successive Heads never occurs.

If the event B never occurs, then it doesn't occur in the first 9 tosses.

It doesn't occur in tosses 1-9
10-18
19-27
etc.

B_1 = event that $w_1 = w_2 = \dots = w_9 = 1$

B_2 = event that $w_{10} = w_{11} = \dots = w_{18} = 1$

B_3 = event that $w_{19} = w_{20} = \dots = w_{27} = 1$
etc.

I spared them disjointly.
It looks a little weird at first, because it can happen that by straddling you can get 9 successive Heads

If B doesn't occur, then none of $B_1, B_2, \dots$ can occur.

2/6/98.10

$$\left. \begin{aligned} B &\supseteq B_1 \cup B_2 \cup B_3 \cup \dots \\ B^c &\subseteq B_1^c \cap B_2^c \cap B_3^c \cap \dots \end{aligned} \right\}$$



$$A \supseteq C \Rightarrow A^c \subseteq C^c$$

Recall [2/6/98.2]

If $C \supseteq A$ then $P(C) \geq P(A)$

Applied to above gives:

$$P(B^c) \leq P(B_1^c \cap B_2^c \cap B_3^c \cap \dots)$$

all disjoint

So they are independent

$$= P(B_1^c) P(B_2^c) P(B_3^c) \dots$$

Note that:

$$P(B_n) = p^9$$

$$P(B_n^c) = 1 - P(B_n)$$

$$= 1 - p^9$$

$$P(B^c) \leq (1 - p^9)(1 - p^9)(1 - p^9) \dots$$

$$p \in [0, 1] \quad \text{If } p \neq 0, (1 - p^9) < 1$$

$$\lim_{n \rightarrow \infty} P(B^c) = 0$$

Therefore, with probability $P(B) = 1$ (in the limit) a run of 9 Heads occurs. Thus we conclude:

$$P(B) = 1$$



2/6/98.11

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



We can even jazz it up.

I say that this is not a strong enough statement.

Suppose 9 Heads occurs only once.

Then you start the next one. It doesn't know what happened before it. It occurs again.

It can not happen just once. Because then you repeat the argument from this point on. So it occurs twice. But then you can start again. Etc.

So it occurs infinitely often.

C = event that a run of 9 Heads occurs infinitely often

Claim $P(C) = 1$

Proof

C^c = event that a run of 9 Heads occurs finitely many times

If the run of 9 Heads occurs only finitely many times, then it doesn't occur from a certain toss onward.

(i.e., no run ||||||| occurs after n^{th} toss, for some n)

But if we start tossing the coin after the n^{th} toss, it's as if we started from scratch.

And we've already proved (previous page) that with probability 1 there has to be one occurrence of the run of |||||||.

(Contradicts previous argument unless $P(C^c) = 0$)

Therefore $P(C) = 1$

The event (that a run of 9 Heads occurs infinitely often) occurs with certainty.

- Note - instead of 9 Heads, we could have chosen any finite pattern whatsoever. Exactly the same argument works. The same pattern will work for any pattern whatsoever.

Take any finite pattern (e.g., 00100111).

The probability that this pattern ever occurs is 1

What's more, this pattern will occur infinitely often

This is the famous assertion that if a monkey starts typing randomly on a computer, sooner or later the monkey will type the bible. And, it will type it infinitely often.

Q/A: Usually $\cap$ takes precedence over $\cup$.
 $\uparrow$ not always the case

The Bernoulli Process (cont'd)

{After this we begin the theory of
distribution of occupancy.}

The Bernoulli Process, w/ the sample space Ω ,
consists of all sample points which are infinite

- $\Omega = \{ \omega = (\omega_1, \omega_2, \omega_3, \dots), \omega_n = \{ \overset{1 = \text{Heads}}{\underset{0 = \text{Tails}}{0, 1}} \} \}$
 $\uparrow$ sample point = infinite sequence of tosses

Philosophically - sample points are all the possibilities
events are all the properties

We try to use sample points as little as possible.

We try to reason about everything in terms of properties $\leftarrow$ i.e., in terms of events

Events

- $\mathcal{E}$: take all subsets of $\overset{\text{sample space}}{\Omega}$ "generated" by $H_n, n=1,2,\dots$

By generated, we mean we can take
 $\cup, \cap, \complement$, and infinite unions of sequences that
are disjoint.

Anything that can be obtained that way is an event.
There is no other description of an event - unless you
want to go into excruciating detail for mathematical
specialists.

Probability

$\bullet P(H_n) = p$

furthermore

any finite set of H_n 's is independent

[Research Problem]

These data determine a unique probability on all the elements.

which satisfies the 3 axioms
of probability.

This is the model of anything you want whatsoever from coin tossing.
That's the Bernoulli Process.

Set
50 SHEETS
100 SHEETS
200 SHEETS

22-141
22-142
22-144



2/9/98.2

We then proceeded to establish the following fact:

The probability that a finite run $\underbrace{111 \dots 11}_k$ occurs infinitely often equals one.

This we established by a very simple argument, which we will shortly summarize again.

A highly non-trivial result, as we remarked at the end of our last meeting. It means that if a monkey starts typing on the computer, the monkey will type the bible infinitely often, or any message whatsoever.

We also remarked that exactly the same argument that we use to establish that the probability that a run occurs infinitely often equals one can be used to show that: the probability that any pattern whatsoever of 0's + 1's, provided it is finite, occurs infinitely often is one.

Exactly the same reasoning.

By the same reasoning, we establish the Borel-Cantelli Lemma.
The probability that any finite pattern, say 010011...01, occurs infinitely often equals one.

The proof is exactly like we did it for a run of 1's [2/6/98.9-11]
It is intuitively obvious that it is the same proof.

Let's restate this in terms of sets.

You realize now that if you switch from probabilistic language to set theoretic language, this is harder to state.

This is the probability of what event?

The event that a given pattern occurs infinitely often.

→ the set of all sample points that contain infinitely many occurrences of that pattern

Now you begin to see that constructing the probability from set theoretic language is more complicated than for probabilistic language. Even though the two languages are completely similar.

You say:

"The event that the pattern 01110 occurs infinitely often"
You have an intuitive feeling for that.

But if you say:

"The set of all sample points that contain infinitely many occurrences of this pattern"

It's not as intuitive. That's the way it is.

So we are gradually shifting our language.

Let's do a little grammar.

This is pure mathematics at it's best (or worst) depending on what your view is.
The pure mathematician at this point will say:

Just a minute. Is this really the most general way of stating the Borel - Cantelli theorem? Can we jazz it up and make it comprehensible? Even though it's the same statement?
Sure.

General Formulation of Borel-Cantelli Lemma

This is where you find it in some probability books that never explain to you what is going on. You have to figure it out.

Given Ω and P { as usual, we pass over in silence the description of events. Of course there are events. }
Sample space with probability P

Given a sequence $A_1, A_2, A_3, \dots$ of independent events

Assume that $\sum_n P(A_n) = \infty$

Then, with probability one, infinitely many events A_n occur.

Where did we get this? The secret crib.

The secret crib goes like this.

Imagine the Bernoulli process. Imagine a pattern of length k .

$\overbrace{\dots\dots\dots}^k$ $\dots\dots\dots$ $\dots\dots\dots$
 A_1 A_2 A_3 } non overlapping
 A_1 is event that pattern occurs in first k tosses
 A_2 " " " " " " second k tosses
etc.

This is uneconomical because you do not consider the possibility of the pattern occurring in between.

But, as we saw last time, it is enough to establish the result we want.

The events $A_1, A_2, A_3, \dots$ are secretly copied from the list of ? of events of non overlapping successive patterns.

In this case, of course, since the pattern of 1's is probability p^k then the sum of the probabilities is infinite.

$$P(A_n) = p^k$$

$$\sum_n P(A_n) = \infty$$

This is the general formulation of the Borel-Cantelli Lemma.

Any ^{independent} event whatsoever, which has any pattern, and the sum of the probabilities is divergent, then infinitely many of them occur.

Since we are on a pure math kick, let's really rub it in. Now you say "Excuse me. What does it mean for infinitely many of the A_n to occur?" We know intuitively. How do you this mathematically? This is how you say it:

First, we give it a name

Say $A_{i.o.}$ is the event that infinitely many of the A_n occur
 $\uparrow$ infinitely often

{ How do we know $A_{i.o.}$ is an event?
 We have to express it in terms of elementary events }

Then

$$A_{i.o.} = \bigcap_{k=1}^{\infty} \left(\bigcup_{n \geq k} A_n \right)$$

We will convince ourselves in a minute that that expression is exactly the event that infinitely many of the events A_n occur. There are many ways of convincing ourselves of what I just said.

One of them is the extremely pure way. You just take this - you don't care what it means.

Lets take the complement and see what this means.

De Morgan Laws

← "There is nothing wrong with not knowing something. What's wrong is not admitting that you don't know. That is wrong" *

$$(A \cup B)^c = A^c \cap B^c$$

$$(A \cap B)^c = A^c \cup B^c$$

You convince yourself of De Morgan Laws by drawing little pictures

The complement of $A_{i.o.}$ gives:

$$A_{i.o.}^c = \left[\bigcap_{k=1}^{\infty} \left(\bigcup_{n \geq k} A_n \right) \right]^c$$

$$= \bigcup_{k=1}^{\infty} \left(\bigcap_{n \geq k} A_n^c \right)$$

Let's write this out:

2/9/98.5

$$= \left(\bigcap_{n \geq 1} A_n^c \right) \cup \left(\bigcap_{n \geq 2} A_n^c \right) \cup \left(\bigcap_{n \geq 3} A_n^c \right) \cup \dots$$

Let's read the first 3 terms:

none of the events A_n occur

none of the events occur after the first k tosses

none of the events occur after the second k tosses

How can we read the infinite union?

It means that none of the events occurs after some point or other, which means that only finitely many of the events occur.

Therefore, the complement of our set is the event that only finitely many of the A_n occur.

Therefore $A_{i\infty}$ is the event that infinitely many occur. There you are.

$$P(A_{i\infty}^c) = P\left(\left(\bigcap_{n \geq 1} A_n^c \right) \cup \left(\bigcap_{n \geq 2} A_n^c \right) \cup \left(\bigcap_{n \geq 3} A_n^c \right) \cup \dots \right)$$

Well, we don't know whether these will be disjoint.

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

$$P(A \cup B) \leq P(A) + P(B)$$

$$P(A_{i\infty}^c) \leq P\left(\bigcap_{n \geq 1} A_n^c \right) + P\left(\bigcap_{n \geq 2} A_n^c \right) + P\left(\bigcap_{n \geq 3} A_n^c \right) + \dots$$

At this point, the mathematician asks: we haven't used the assumption that these A_n are independent.


It has to come in somewhere.

$$P\left(\bigcap_{n \geq 1} A_n^c \right) = \prod_{n \geq 1} P(A_n^c)$$

$$= \prod_{n \geq 1} (1 - P(A_n))$$

$$= \prod_{n \geq 1} (1 - P(A_n)) + \prod_{n \geq 2} (1 - P(A_n)) + \prod_{n \geq 3} (1 - P(A_n)) + \dots$$

For $0 \leq x \leq 1$,
 $1 - x \leq e^{-x}$



Maclaurin series for e^{-x} →

$$1 - x \leq 1 - x + \underbrace{\frac{x^2}{2!} - \frac{x^3}{3!}}_{\geq 0} + \underbrace{\frac{x^4}{4!} - \frac{x^5}{5!}}_{\geq 0} + \dots$$

every pair ≥ 0

$$\underbrace{\hspace{10em}}_{\geq 0}$$

$$\therefore 1 - x \leq e^{-x}$$

$$\begin{aligned} P(A_{i_0}^c) &\leq \prod_{n \geq 1} (1 - P(A_n)) + \prod_{n \geq 2} (1 - P(A_n)) + \prod_{n \geq 3} (1 - P(A_n)) + \dots \\ &\leq \prod_{n \geq 1} e^{-P(A_n)} + \prod_{n \geq 2} e^{-P(A_n)} + \prod_{n \geq 3} e^{-P(A_n)} + \dots \\ &= e^{-\sum_{n \geq 1} P(A_n)} + e^{-\sum_{n \geq 2} P(A_n)} + e^{-\sum_{n \geq 3} P(A_n)} + \dots \end{aligned}$$

But the sum of the probabilities of A_n is assumed to be infinite:

$$\sum_n P(A_n) = \infty$$

$$= 0 + 0 + 0 + \dots$$

a countable number of times

Therefore,

$$P(A_{i_0}^c) = 0$$

Therefore,

$$P(A_{i_0}) = 1, \text{ as desired}$$

That is what mathematicians make out of the Borel-Cantelli Lemma.

If you think about it concretely, you must always think about the patterns. But with the lemma couched this way, you can apply it to a number of situations.

Q: What are the A_n above?

A: A_n are events with the properties:

1) A_n are independent

$$2) \sum_n P(A_n) = \infty$$

sum of the probabilities is infinite.

Recap:

• First, we showed Borel-Cantelli with a crib (concrete example)

• Then, we proved Borel-Cantelli abstractly. The advantage of doing it abstractly is that you can find unexpected examples like the following

[Research Problem] this is pretty cute

Example: ^{Bernoulli Process} Let A_n = event that at $2n^{\text{th}}$ ^{even} toss, the number of 0's equals the number of 1's.

Here is a sequence of events.

What is the probability that infinitely many of the A_n occur?

Note that the A_n are not independent.

Q/A: why?

Deferred. If they are independent, so much the better. But it's pointless to spend time showing something is not true.

If they were independent - better. But we don't want to worry about it.

$$A_{\infty} = \bigcap_{k=1}^{\infty} \left(\bigcup_{n \geq k} A_n \right)$$

This is the event that infinitely many times, you have equalization of Heads + Tails.

Now this is a sensible and important question to ask.

Q: Will this happen infinitely often or not?

A: The answer is very interesting

$$P(A_{\infty}) = \begin{cases} 0 & \text{if } p \neq \frac{1}{2} \\ 1 & \text{if } p = \frac{1}{2} \end{cases}$$

↑ for an absolutely fair coin, you get infinitely often.

If the coin is just a little bit unfair - then no.

2/9/98.8

Before we leave this preliminary study of the Bernoulli Process, let's examine one of the most extraordinary properties of coin tossing.
A property that goes against intuition:

Kolmogorov zero-one law

Let's philosophize about A_{∞} . ^{infinitely often}

If we leave out any finite number of the A_n , then the event A_{∞} has the same solution.

Also, this event A_{∞} doesn't depend on what happened at those coin tosses

→ The probability is not affected.

An event B in the Bernoulli Process is said to be a tail event when B is independent of $H_1, H_2, \dots, H_n$ for every n .

What are examples of tail events?

The occurrence of a pattern ever.

The event that you eventually get a run of 15 Heads is a tail event.

It's probability doesn't change if the probability is dependent on the first coin toss.

Anything to do w/ eventually is likely to be a tail event

[Research Problem]

• Kolmogorov Zero-One Law

Every Tail Event has probability 0 or 1 } not particularly hard, but hopelessly dull

50 SHEETS
100 SHEETS
200 SHEETS



Example

$$1 + \frac{1}{2} + \frac{1}{3} + \frac{1}{4} + \dots = \infty$$

Harmonic Series is divergent

Keep the terms of the Harmonic Series, but instead of making them all +, I make them + or -, at random.

You all know this.
If you didn't, you do now.

The probability of this series will converge

Let's make this more precise:

$$1 \pm \frac{1}{2} \pm \frac{1}{3} \pm \frac{1}{4} \pm \dots$$

$$\omega_1 = \begin{cases} 1 & \text{if } + \\ 0 & \text{if } - \end{cases} \quad \omega_2 = \begin{cases} 1 \\ 0 \end{cases} \quad \omega_3 = \begin{cases} 1 \\ 0 \end{cases}$$

You look at each entry of the sample point and put + or - according to whether entry is Heads or Tails.

Sample point:

$$\omega = (\omega_1, \omega_2, \omega_3, \dots)$$

Rule for assigning the signs.
Toss a fair coin.

$$\text{So } P(+)=P(-)=\frac{1}{2}$$

Assign the n^{th} sign depending on what the n^{th} toss is.
You get a random series

C = event that series is convergent

What is $P(C)$?

We apply the Kolmogorov Zero-One Law.

The convergence of the series is not affected if you cross out the first n terms, for any n .

A divergent series - if you cross out first 15,000 terms is still divergent.

Therefore, C is a Tail Event $\leftarrow$ in the sense of the Kolmogorov Zero-One Law

$$\therefore P(C) = \begin{cases} 1 \\ \text{or} \\ 0 \end{cases} \leftarrow \text{I leave it to you to figure out which one it is. Actually it is:}$$

$$P(C) = 1 \quad \text{Rademacher's Theorem}$$

2/9/98, 10

Now you say - But that's impossible. What if I make up some weirder series. Ex:

$$\sum \frac{(\sin n\pi / \log n)}{n^{1/2}}$$

If you assign +/- at random, the series converges by Kolmogorov Zero-One Law with either probability 0 or 1.
No matter how weird the series.

If you establish that the probability is stated at .00001 then you are through, probability is 1.

Q/A: "Anything that makes sense is an event"

Lastly, before we completely leave the Bernoulli Process:

Observe that we have so far, passed over in silence, the phenomena of events of probability 0

To say that $P(A) = 0$ does not mean that the event is impossible.
For example

event $W = (1111 \dots)$ has $P(W) = 0$ if $0 < p < 1$ } But it may occur
probability zero

Similarly, if we assert that $P(\text{event}) = 1$, this does not mean the event is certain.

The point we are making now is what is the philosophical interpretation of probability?

You must be aware that there are events with:

$P(\text{event}) = 0$ - not impossible } but pretty close
 $P(\text{event}) = 1$ - not certain

This concludes our preliminary study of the Bernoulli Process.

50 SHEETS
100 SHEETS
200 SHEETS

22-141
22-142
22-144



Theory of Distribution and Occupancy

This is a big chapter of probability.
And it's very broad and buttery - this stuff has tremendous applications.
What's it all about?

Probabilistic analysis of the following situation:

We have a set of k balls (set of balls = B)
 n boxes (set of boxes = U (urns))



then we place the balls into the boxes, subject to various restrictions.

And we ask for the probability of various events happening.

You think this is easy?

You'll see in a minute, this can be awfully tough.

Some fundamental results of physics depend crucially on estimating the probability of certain events relating to placements of balls into boxes.

In fact, many people whisper behind my back that all I did in my life is place balls into boxes.

There are really 3 cases:

- Maxwell-Boltzmann statistics statistics here is another word for sample space
 k ordinary balls, n boxes and you place balls into boxes.
- Bose-Einstein statistics
Strikingly different from Maxwell-Boltzmann, as you will see,
The balls are indistinguishable. Identical balls. You will see by example that this does happen in real life.
- Fermi-Dirac statistics
Identical balls, but
no more than one ball per box Identical balls.

All these occur in real life,

They also occur in the wonderful theory land of Quantum Mechanics.
Hence the names. They were first identified in Q.M., but there are examples from concrete, everyday life, as you will see.


For the next couple of lectures, we will study these things to death.
And you have to learn it cold.



I will guarantee you one thing.
Every word in this course will be used later in your life.
You are not wasting any time.
On the other hand, you can not afford to fall behind. If you fall one week behind, you will be unable to follow the lecture.
You should take careful notes. Part of your training is to take careful notes of this course. The material in the lectures is strongly at variance with the material in the so-called text, which I am trying to rewrite.
A lot of this material is not covered in any textbook. It is in research papers. This is not the thing you are going to find in any old book.
If you miss these classes, fine. You won't find this material anywhere else. For the kind of tuition you pay, that's what you deserve.
Otherwise, you take any old probability book + go to Oshkosh College. It will cost you half as much.
Take down everything I say. And treat it for later references. You may get a job on one remark in this course. I'm not joking.

The theory of distribution and occupancy (beg'g)

This is a central chapter of probability and one that is very easy to visualize. Because this has to do w/ placing balls into boxes.
The setup is:



You place the balls into boxes in an arbitrary way.

Every ball goes into one box.

The same box may receive several balls.

Ω_{MB} = Maxwell-Boltzmann sample space

This is a finite sample space

sample space: a sample point is a function ω from the set B to the set U .

{ Placing balls into boxes is taking a function from the set of balls }
to the set of boxes.

How many such functions are there?

k balls $\left\{ \begin{array}{l} \text{The first ball can be placed in any of } n \text{ boxes,} \\ \text{The second ball can be placed in any of } n \text{ boxes,} \\ \text{etc} \end{array} \right.$

Therefore, there are n^k sample points

In the absence of any further information, we will assume all sample points have the same probability.

$$P(\omega) = \frac{1}{n^k}$$

event consisting of one sample point.

Events: all subsets of Ω_{MB}

as happens for all finite sample spaces, where the events are all subsets

$$\begin{aligned} \text{If } A \text{ is an event then } P(A) &= \sum_{\omega \in A} P(\omega) \leftarrow \left\{ \begin{array}{l} \text{computed in the almost} \\ \text{obvious customary fashion} \end{array} \right\} \\ &= \frac{|A|}{n^k} \leftarrow \begin{array}{l} \text{number of elements} \\ \text{of } A \end{array} \end{aligned}$$

All this is completely obvious.

In order to perform some interesting computations with the Maxwell-Boltzmann sample space, also known as Maxwell-Boltzmann statistics, as we said last time, we need to review some classical quantities. We assume, as known, the following facts:

$n!$ = number of permutations of a set with n elements

Read: "binomial" n^k $\binom{n}{k}$ = binomial coefficient
number of $\underbrace{k \text{ subsets}}_{\text{subset w/ } k \text{ elements}}$ of an n set

There are formulas for the binomial coefficient.
Let's introduce one formula for the binomial coefficient:

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}$$

Note that:

$$0! = 1$$

The number of subsets of the empty set:

$$\binom{0}{0} = 1$$

↑ the empty set $\emptyset$ has one subset.
Itself.

We also use the following notation:

Lower factorial

$$(n)_k = \overbrace{n(n-1)(n-2)\dots(n-k+1)}^{k \text{ terms}}$$

Read:
"n lower factorial k"

We have the binomial coefficient:

$$\begin{aligned} \binom{n}{k} &= \frac{n!}{(n-k)!} \left(\frac{1}{k!} \right) \\ &= \frac{(n)_k}{k!} \end{aligned}$$

Observe that the derivative $D = \frac{d}{dn}$ is:

$$D(n^k) = k n^{k-1}$$

Similarly, we can define the difference operator $\Delta f(n) = f(n+1) - f(n)$.
There is a very close analogy to the derivative with the difference operator.

$$\begin{aligned} \Delta(n)_k &= (n+1)_k - (n)_k \\ &= (n+1) \underbrace{n \dots (n-k+2)}_{(n)_{k-1}} - \underbrace{n(n-1) \dots (n-k+2)(n-k+1)}_{(n)_{k-1}} \\ &= ((n+1) - (n-k+1)) (n)_{k-1} \\ \Delta(n)_k &= k (n)_{k-1} \end{aligned}$$

Rising factorial

$$\langle n \rangle_k = \overbrace{n(n+1)(n+2)\cdots(n+k-1)}^{k \text{ terms}}$$

Read: "Bra n ket k"

Backwards Difference Operator

$$\nabla f(n) = f(n) - f(n-1)$$

$$\begin{aligned}\nabla \langle n \rangle_k &= \langle n \rangle_k - \langle n-1 \rangle_k \\ &= \underbrace{n(n+1)\cdots(n+k)(n+k-1)}_{\langle n \rangle_{k+1}} - (n-1) \underbrace{n \cdots (n+k)}_{\langle n \rangle_k} \\ &= [(n+k) - (n-1)] \langle n \rangle_k \\ &= k \langle n \rangle_k\end{aligned}$$

$$\nabla \langle n \rangle_k = k \langle n \rangle_k$$

Multiset Coefficients - We define the multiset coefficients. Soon to be interpreted.

$$\left\langle \begin{matrix} n \\ k \end{matrix} \right\rangle = \frac{\langle n \rangle_k}{k!}$$

"bra n over k ket"

These are quantities that will be interpreted extensively over the next couple of lectures. I thought it better for me to write them out so you know what is coming.

Multinomial coefficient

$$\binom{k}{i_1, i_2, \dots, i_n} = \begin{cases} 0 & \text{if } i_1 + i_2 + \dots + i_n \neq k \\ \frac{k!}{i_1! i_2! \cdots i_n!} & \text{otherwise} \end{cases}$$

Binomial Theorem

$$(a+b)^n = \sum_{i=0}^n \binom{n}{i} a^i b^{n-i}$$

In the course of these lectures, we will see 17 different proofs of the binomial theorem.

Let's go back to Maxwell-Boltzmann statistics and do some computation.


Example with Ω_{MB}

month, day

A = event that no two persons in room 4-370 have the same birthday

$P(A)$?

The answer is obtained by setting up the Maxwell-Boltzmann statistics, where the balls are the persons and the boxes are the birthdays.



Two people could have the same birthday.
Each person has only one birthday.

The way it is set up is a function from persons to birthdays.
The event that we want is the event that a sample point will be a function that has the property that no two balls will be in the same box.

How many sample points are there in Ω_{MB} with no two balls in the same box?



Let's count.

You take the 1st ball and it can go into any one of the boxes.
There are n ways of doing things.

Now we take the second ball. It can go into any box except the box occupied by the 1st ball. Which means there are $(n-1)$ ways of doing things.

$$\underbrace{n(n-1)(n-2) \cdots (n-k+1)}_{k \text{ terms}} = (n)_k$$

Therefore, the probability of our event is:

$$P(A) = \frac{(n)_k}{n^k}$$

Let's go back to the title of this lecture.

The title of this lecture is The Theory of Distribution and Occupancy

Why distribution and occupancy?


Distribution is the interpretation of a function from the set B to the set U

$$B \rightarrow U$$

which consists of placing the balls of B into the boxes of U .

There is, however, another interpretation of this same sample space.

Occupancy interpretation



Now, the function from the balls to the boxes means that you take the letter b and place it in position 1
Then letter a in position 2
etc.

b a c a b b d

You visualize a function as a word made of an alphabet.
This is mathematically the same, but a completely different visualization

That's the occupancy interpretation of Maxwell-Boltzmann statistics.

We can ask: what's the probability that a word starts with ab ?

B = event that a word starts with ab

$$1 \rightarrow a$$

$$2 \rightarrow b$$

remaining $k-2$ letters can go anywhere

Therefore, there are n^{k-2} ways of making such a word.

And each one has probability $\frac{1}{n^k}$.

Thus:


$$P(B) = \frac{n^{k-2}}{n^k} = \frac{1}{n^2}$$

2/11/98.7

Now let's look at a further problem.

For every placement of balls into boxes, we have occupation numbers.
Namely:

Occupation Numbers



we label the boxes $1, 2, \dots, n$.
(we could label these $a, b, \dots, c$, if you wish.)

θ_1 = number of balls in box 1

for every sample point (for every way of placing the balls into the boxes) there are a certain number of balls in box 1.
That's the occupation number, by definition, of box 1.

θ_2 = number of balls in box 2
etc.

What is $P(\theta_1 = i)$?

Placing the balls into the boxes at random, what's the probability that i balls fall into box 1.

First, out of the set of balls, we choose i balls.

This is done in $\binom{k}{i}$ ways. Read: "binomial k i "

Then the remaining balls must go into boxes labelled 2 to n .

Therefore, you have:

$(n-1)^{k-i}$ ways of placing the remaining balls into the remaining boxes.

Since each of these has probability $\frac{1}{n^k}$, we have:

$$P(\theta_1 = i) = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k}$$

Let's fool around with this a bit.

We notice that the events:

$$(\theta_1 = 0) \cup (\theta_1 = 1) \cup (\theta_1 = 2) \cup \dots \cup (\theta_1 = k) = \Omega_{MB}$$

The union of these events is the whole sample space.

These events are disjoint.

If you have 3 balls in the 1st box, you can't have 2 balls in the 1st box.



Therefore: the probability of the union of these events is the sum of the probabilities.

$$P(\theta_1=0) + P(\theta_1=1) + P(\theta_1=2) + \dots + P(\theta_1=k) = P(\Omega_{MB})$$

Q/A: Why do we not have $\dots P(\theta_1=k+1) + \dots + P(\theta_1=n)$?

Because there are only k balls. n is the number of boxes.

We all fall into this trap. We exchange the balls for the boxes. Sometimes it's hard to set up the problem.

We already have the answer to $P(\theta_1=i)$.
So this is:

$$\sum_{i=0}^k \frac{\binom{k}{i} (n-1)^{k-i}}{n^k} = 1$$

Let's interpret this.

$$\sum_{i=0}^k \binom{k}{i} (n-1)^{k-i} = n^k$$

We just proved the binomial theorem.
The 1st of 17 proofs.

Why is this the binomial theorem?

Because $n^k = (\underbrace{(n-1)}_{\text{this is the } a} + \underbrace{1}_{\text{this is the } b})^k$

$$\sum_{i=0}^k \binom{k}{i} (n-1)^{k-i} = ((n-1) + 1)^k$$

Q: Isn't this the special case of the binomial theorem?

A: It's special, but you easily see that it implies the general form.
It's only seemingly special.
You split the set of boxes into 2.



Next, let's do an extension of this problem.
What's the probability of having the 1st box have i balls and the 2nd box have j balls?

From the set of k balls, we take i balls and place them in box 1.
You are left w/ $k-i$ balls; from these, you take j balls and place them in box 2.
You are left w/ $k-i-j$ balls, which have to go into boxes labelled 3 to n .
That means that you have only $n-2$ boxes into which $k-i-j$ balls can be placed.

$$P((\theta_1=i) \cap (\theta_2=j)) = \frac{\binom{k}{i} \binom{k-i}{j} (n-2)^{k-i-j}}{n^k}$$

Next - let's now compute the following probability.
Let us find all the occupation numbers.

$$P((\theta_1=i_1) \cap (\theta_2=i_2) \cap \dots \cap (\theta_n=i_n)) = \binom{k}{i_1} \binom{k-i_1}{i_2} \binom{k-i_1-i_2}{i_3} \dots \binom{k-i_1-i_2-\dots-i_{n-1}}{i_n} \frac{1}{n^k}$$

probability of each sample point

First, from the set of k balls, take i_1 balls and place them in the 1st box.
We are left w/ $k-i_1$ balls. From these, we pick a subset of i_2 and place them in the 2nd box.
We are left w/ $k-i_1-i_2$ balls. From these, we pick a subset of i_3 and place them in the 3rd box.

if $i_1 + i_2 + \dots + i_n \neq k$

then you got a contradiction.

The probability of the event is 0 because no ball can be thrown away.

We can simplify the above expression in an extraordinary manner.
Let's write the expressions for the binomial coefficients:

$$= \frac{k!}{i_1! (k-i_1)!} \cdot \frac{(k-i_1)!}{i_2! (k-i_1-i_2)!} \cdot \frac{(k-i_1-i_2)!}{i_3! (k-i_1-i_2-i_3)!} \dots \frac{(k-i_1-i_2-\dots-i_{n-1})!}{i_n! (k-i_1-i_2-\dots-i_{n-1}-i_n)!} \cdot \frac{1}{n^k}$$

Now, miracles of cancellation occur.

0! = 1

$$= \frac{k!}{i_1! i_2! \dots i_n!} \cdot \frac{1}{n^k} \quad \left(\text{this is the multinomial coefficient} \right)$$

[2/11/98, 4]

$$= \binom{k}{i_1, i_2, \dots, i_n} \frac{1}{n^k}$$

This is the probability of us finding all of the occupation numbers.

This is really the interpretation of the multinomial coefficient.

The number of ways of placing balls into boxes when the occupation number of each box is assigned.

$$\binom{K}{i_1, i_2, \dots, i_n} = \text{number of ways of placing } K \text{ balls into } n \text{ boxes} \\ \text{with } i_1 \text{ balls in box 1, } i_2 \text{ balls in box 2, etc.}$$

This is a very important result.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



In real life, you get Maxwell-Boltzmann statistics thrown at you under various disguises. Being an applied mathematician means being apply to recognize them.

Someone comes in w/ some weirdo statement - but it's nothing but balls into boxes. Nobody can teach you that. You have to learn it by yourself the hard way.

For example, a physicist comes in and says: "Oh, we have particles and states. We want to know how many ways these particles can be in these states".
The particles are the balls.
The states are the boxes.

The only mistakes that I make and that you will make, probably, is to sometimes confuse the balls w/ the boxes.

(*) { The way you test yourself (Always remember this)
A ball can only go into one box.
A box may receive several balls. }

This is the test for particles and states:

You can have several particles in the same state.

But to each particle there is only one state.

Which means:

particles are the balls
states are the boxes

Or else, someone comes in w/ questions of linguistics.

You have an alphabet.

What's the probability that the letter "a" does not occur in a word of length n ?

It's just balls into boxes.

It is the occupancy interpretation.

In this case, the letters are the boxes
the positions are the balls

1 2 3 ... ?

$\boxed{a}$ $\boxed{b}$... $\boxed{z}$

We can go on and on w/ all sorts of interpretations.
Chemistry, management, jobs and people, etc.

2/11/98.11

We have time to go to our next sample space in our theory of distribution and occupancy. This is the sample space of dispositions.
This is another finite sample space.

Dispositions $\mathcal{R}_D$

We have, again, K balls and n boxes. {Before you are through w/ this course, you will remember this case.}

$0 \quad 0 \quad \dots \quad 0 \quad K$

$\square \quad \square \quad \dots \quad \square \quad n$

A disposition is a funny thing.

It is the placement of the balls into the boxes - but, after placing the balls into the boxes, then you linearly order the balls in each box.

So a disposition is a placement of the balls into the boxes plus a linear order of the balls in the boxes.

Placement of B to U , but the set of balls in each box is linearly ordered.

What would be a model of disposition?
My favorite model of disposition is:

balls $\rightarrow$ flags


boxes $\rightarrow$ poles.

You have to raise the flags on different poles.
And the flags on a pole are linearly ordered - some are higher than others.

This is a disposition.

The number of ways of putting k flags on n poles is the number of dispositions of k flags. Because you have to arrange linearly the flags on each pole.

You see the idea. It's not just balls into boxes. It's placing balls into boxes with some extra structure.



k flags
 n poles



2/11/98, 12

A Mickey Mouse example with 2 balls and 3 boxes.
Example: $K=2$, $n=3$

Let's divide the blackboard into 2 and compare Maxwell-Boltzmann with Dispositions.

MB : b

D : a b

1	2	3
a	b	.
a	.	b
b	a	.
b	.	a
.	a	b
.	b	a
ab	.	.
.	ab	.
.	.	ab

1	2	3
a	b	.
a	.	b
b	a	.
b	.	a
.	a	b
.	b	a
ab	.	.
ba	.	.
.	ab	.
.	ba	.
.	.	ab
.	.	ba

So far, it looks just like ordinary placement of balls into boxes,

But now we have this.
ab and ba are different now.

These are different dispositions because of the poles.

Maxwell-Boltzmann has
 $n^k = 3^2 = 9$
sample points

Each member of this list is a sample point of the space of dispositions.

As I told you already, sample points of probability do not look like points.

How many have we got? What will be the formula?

Let's conclude today's meeting with the formula for the number of dispositions.


Read: "n, rising factorial k"

There are $\langle n \rangle_k$ dispositions of k balls into n boxes

Proof


This is pretty cute. Watch.

I draw the boxes this way



I take the 1st ball and put it into a box.


Now, this act of placing the first ball into the box is the same as if the box were to be split into 2 boxes. Because the successive ball will give a different disposition, according to whether it goes before or after.



Therefore, it's as if we have one more box for the next ball, to get the disposition.

1st ball — n ways


2nd ball:



The second ball can go any of the $n-1$ boxes plus any of these two, which is $n+1$.

Wherever the 2nd ball goes, it can be assumed to split the portion of the box where it goes into two. Because the remaining balls have a choice to go between.

So even if it goes in the same box as the 1st ball, it cuts the box into two:



So every ball cuts a box into two.

$$\begin{aligned} \# \text{ of dispositions of } k \text{ balls into } n \text{ boxes} &= \overbrace{n(n+1)(n+2) \dots (n+k-1)}^{k \text{ terms}} \\ &= \langle n \rangle_k \end{aligned}$$

So there are $\langle n \rangle_k$ ways of disposing k balls into n boxes.

Therefore, we define in Ω_D (the sample space of disposition) a probability by setting:

$$P_D(\omega) = \frac{1}{\langle n \rangle_k}$$

probability of
disposition of a
single sample point

And, we have again that the probability of any event A is:

$$P(A) = \sum_{\omega \in A} P_D(\omega) = \frac{|A|}{\langle n \rangle_k}$$

Next time we will see other models for dispositions.

I would appreciate if you think of some interesting other models for disposition. Because I only have two models in my book. I'd love to have 4 or 5.



You have a new problem set today. It's extremely difficult. 5 problems were taken from research papers. Don't wait until the last night, the night before the problem set is due. Work on it. And try to do them jointly as much as possible. If you can get a sucker to work out these problems for you, you deserve an A. There are very few people on this campus willing to spend his or her time working out these problems. Remember, the problems in this course are completely different from the quizzes. The quizzes test strictly the material from the lectures. The problems are, in general, expansions of material in the lectures. When you study for the quiz, when the time comes, you will have to study your notes. The quiz will cover strictly the material we cover in the lectures.

As I say, this is an intensive course. With the kind of tuition you pay, I feel morally obliged to give you the hardest problems I can find. What else can I do? Otherwise it's not worth your tuition. So I spent 4 hours yesterday going through research papers preparing your problems. I'm not kidding.

I've given out a number of research problems. The research problems you've received so far are not hard and well worth working on. You can consult any published material for the research problems. You can crib them, if you can find them written up somewhere.

Q/A: I'd love to have a really elegant proof of the Kolmogorov Zero-One Law. Something cute. So I can put it in my book.

The Theory of Distribution and Occupancy (cont'd)

We've seen so far 2 finite sample spaces that pertain to the theory of distribution and occupancy.

1) Maxwell-Boltzmann ^(sample space) statistics Ω_{MB}

In this sample space, you have n^k sample points.

$$P_{MB}(\omega) = \frac{1}{n^k}$$

$\uparrow$
sample point

The sample points correspond to the functions from a set of k balls to a set of n urns:



A sample point is a function from the set of balls to the set of urns.

Such functions can be given a dual interpretation.

The distribution interpretation and the occupancy interpretation.

The distribution interpretation is the usual interpretation as particles going to states.

The occupancy interpretation is the following:

the balls are numbered $1, \dots, k$

the boxes are viewed as the letters of an alphabet

Last time we considered the sample points corresponding to the situation where:

$$|B| = 2$$

$$|U| = 3$$

Last time, we considered this from a distribution point of view.

Let's now do the same from an occupancy point of view

1 2

a b c

list of sample points from an occupancy point of view:

aa
ab
ac
ba
bb
bc
ca
cb
cc

words of the sample space

Mathematically the same, but psychologically different

eventually translates into a mathematical statistic

Let's write down some of the results we obtained for Maxwell-Boltzmann statistics

Occupation number of the i^{th} box:

$$P(\theta_i = i) = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k}$$

← numerator is the number of sample points.

← $\frac{1}{n^k}$ is probability of a single sample point.

recall that probability of event A is:

$$P(A) = \frac{|A|}{n^k} \leftarrow \text{number of elements in } A$$

This is the ^{$P(\theta_i = i)$} occupation number distribution, which is an important datum about Maxwell-Boltzmann, as well as balls into boxes. We also computed:

$$P(\theta_1 = i_1) \wedge (\theta_2 = i_2) \wedge \dots \wedge (\theta_n = i_n) = \frac{k!}{i_1! i_2! \dots i_n!} \cdot \frac{1}{n^k}$$

if $i_1 + i_2 + \dots + i_n = k$

This term is also known as the multinomial coefficient:

$$\binom{k}{i_1, i_2, \dots, i_n}$$

As a matter of fact, the computation is the combinatorial definition of the multinomial coefficient.

2) Now let's do the same w/ the sample space of dispositions

$\Omega_D = \text{Dispositions}$ k balls into n boxes

We have seen that the number of elements in the sample space is:

$$|\Omega_D| = \langle n \rangle_k = n(n+1)(n+2) \dots (n+k-1)$$

This gives a combinatorial meaning to the rising factorial.

Therefore, we define the disposition probability of a sample point as:

$$P_D(\omega) = \frac{1}{\langle n \rangle_k}$$

Let's now perform computations of occupation numbers for dispositions. Occupation numbers are Very Important Data. Because that's what you usually observe.

2/13/98.4

Confusing balls w/ boxes is a very common mistake. There are situations where you don't know which are the balls + which are the boxes. That is called applied mathematics. Sometimes this is extremely difficult. Sometimes, the balls become boxes.

$$P_D(\theta_1 = i) = \binom{k}{i} \square \langle n-1 \rangle_{k-i}$$

binomial k i :
we place i
balls into box 1

$\langle n \rangle_k$

we are left w/ $k-i$ balls,
which can be disposed in
any of the other $n-1$ boxes.

$\langle n-1 \rangle_{k-i}$

I haven't ordered the balls in box 1 yet.
Remember, a disposition doesn't just mean
put the balls into the boxes. In each
box, you order the balls. If you order
them differently, it is different.
Therefore, I have to count the number of
ways of ordering the balls in box 1.
This is just $i!$

$$P_D(\theta_1 = i) = \frac{\binom{k}{i} i! \langle n-1 \rangle_{k-i}}{\langle n \rangle_k}$$

It's an interesting exercise, which I did and I leave for you to work out that:

$$P(\theta_1 = 0) + P(\theta_1 = 1) + \dots + P(\theta_1 = k) = 1$$

and you get a fancy combinatorial identity proved cheapo.
You don't have any computations.

$$\sum_{i=0}^k \frac{\binom{k}{i} i! \langle n-1 \rangle_{k-i}}{\langle n \rangle_k} = 1$$

$$\boxed{\sum_{i=0}^k \binom{k}{i} i! \langle n-1 \rangle_{k-i} = \langle n \rangle_k}$$

a combinatorial alternative
for the rising factorial
 $\langle n \rangle_k = n(n+1) \dots (n+k-1)$

2/13/98.5

Now, we could do the occupation numbers of box 1 + box 2 simultaneously. As we have done it already for Maxwell-Boltzmann statistics. But that's kind of dull + you have the picture anyway.

Rather, let's jump ahead to the most striking and unexpected phenomenon pertaining to dispositions.

Namely, let's consider this event - that each of the boxes is given an occupation number.

$$P((\theta_1 = i_1) \cap (\theta_2 = i_2) \cap \dots \cap (\theta_n = i_n)) = ?$$

$$\text{where } i_1 + i_2 + \dots + i_n = k$$

$$(\text{otherwise } P = 0)$$

Let's compute it brutally and then pretend we have a simple proof. Once we get a result, then we immediately see there ought to have been a simpler way of getting it.

Take i_1 balls and dispose in box 1.
I'm left w/ $k - i_1$ balls and dispose i_2 in box 2.
I'm left w/ $k - i_1 - i_2$ balls and dispose i_3 in box 3.
etc

$$= \binom{k}{i_1} i_1! \binom{k-i_1}{i_2} i_2! \binom{k-i_1-i_2}{i_3} i_3! \dots \binom{k-i_1-i_2-\dots-i_{n-1}}{i_n} i_n! \cdot \frac{1}{\langle n \rangle_k}$$

Let's simplify this. A wonder of wonders happens.

$$= \frac{k!}{i_1! (k-i_1)!} i_1! \frac{(k-i_1)!}{i_2! (k-i_1-i_2)!} i_2! \frac{(k-i_1-i_2)!}{i_3! (k-i_1-i_2-i_3)!} i_3! \dots \frac{(k-i_1-i_2-\dots-i_{n-1})!}{i_n! (k-i_1-i_2-\dots-i_n)!} i_n! \cdot \frac{1}{\langle n \rangle_k}$$

$0! = 1$

$$= \frac{k!}{\langle n \rangle_k} \left\{ \leftarrow \text{Q/A: Note that } \frac{k!}{\langle n \rangle_k} = \frac{1}{\langle \frac{n}{k} \rangle} \leftarrow \text{multiset coefficient} \right.$$

You're spiling my act.

What is remarkable about this result is that it is independent of the occupation numbers.

Balls into boxes with any prescription of occupation numbers, the probability is not dependent on $i_1, i_2, \dots, i_n$.



This is a very fundamental fact

2/13/98.6

The probability

$$P((\theta_1 = i_1) \cap (\theta_2 = i_2) \cap \dots \cap (\theta_n = i_n)) = \frac{k!}{\langle n \rangle_k}$$

does not depend on $i_1, i_2, \dots, i_n$

in striking contrast to Maxwell-Boltzmann, where probability of this event sure does depend on $i_1, i_2, \dots, i_n$

Easy Proof

Now that we have this result, you say "There must be a reasonable way of seeing that"

Now let's do it easy.

The easy way of seeing that, most people don't get the first time.

Write down all the permutations:

$\{1, 2, \dots, k\}$

$k!$ {
1 2 3 4 5 ... k
2 1 3 4 5 ... k
3 1 2 4 5 ... k
:
k 1 2 3 4 5 ... k

You are given $i_1, i_2, \dots, i_n$
Count the first i_1 and then I put in a stopper (1).
Then I count the next i_2 and then I put in a stopper, etc.
I do this in exactly the same position for all of these.

$k!$ {

i_1	i_2	i_n
1 2	3 4 5	... k
2 1	3 4 5	... k
2 3	1 4 5	... k
:	:	:

When I've done this, I've done all possible dispositions of balls into boxes using the given occupation numbers. Because they are lined up inside the stoppers, as pieces of all possible permutations.

This says that the number of possible ways of disposing the balls into the boxes with given occupation numbers is $k!$.

Because you know where to put the stoppers.

Easy proof: Write down all permutations then put in the stoppers.

This gives all possible dispositions.

Each of $k!$ dispositions has probability $\frac{1}{\langle n \rangle_k}$
so $P((\theta_1 = i_1) \cap \dots \cap (\theta_n = i_n)) = \frac{k!}{\langle n \rangle_k}$

50 SHEETS
22-141
100 SHEETS
22-142
200 SHEETS
22-144



Let's now go into the 3rd sample space in the theory of distribution and occupancy.

Bose - Einstein statistics

• • • • • k balls B
— — — — — n boxes U

In this case, the balls are indistinguishable

That means that you can visualize a sample point by simply putting checks (✓) corresponding to the number of balls in each box.
All that matters is how many balls are in each box.

In other words, the assignment of the occupation numbers to each box completely determines the placement of indistinguishable balls into the boxes.

All you know is how many balls you have.
So there is a check (✓)



provided that the number of checks (✓) add up to k

Before we interpret this, let's ask the question:

How many such placements are there of indistinguishable balls into boxes?

We've already done that. Look, that's why we did dispositions.
Look at this result

$$P(\theta_1 = i_1) \cap (\theta_2 = i_2) \cap \dots \cap (\theta_n = i_n) = \frac{k!}{\langle n \rangle_k}$$

for dispositions

That means that all possible distributions of occupation numbers are equally likely in the space of dispositions.

Therefore, how many possible distributions of occupation numbers are there?

One over that, i.e.,

$$\frac{1}{\frac{k!}{\langle n \rangle_k}} = \frac{\langle n \rangle_k}{k!}$$

Because they are all equally likely

From this result, we get that the number of possible $i_1, i_2, \dots, i_n$ such that $i_1 + i_2 + \dots + i_n = k$ is:

$$\frac{\langle n \rangle_k}{k!}$$

But this is exactly the number of ways of putting indistinguishable balls into boxes.

Because putting indistinguishable balls into boxes is the same as assigning numbers $i_1, i_2, \dots, i_n$ such that $i_1 + i_2 + \dots + i_n = k$.

And we just saw this.

$$\Omega_{BE} \text{ has } \frac{\langle n \rangle_k}{k!} \text{ sample points} \quad |\Omega_{BE}| = \frac{\langle n \rangle_k}{k!}$$

↑
Bose-Einstein

$$\frac{\langle n \rangle_k}{k!} = \left\langle \begin{matrix} n \\ k \end{matrix} \right\rangle$$


is called the multiset coefficient

$$P_{BE}(\omega) = \frac{1}{\left\langle \begin{matrix} n \\ k \end{matrix} \right\rangle}$$

Let's prove this result by an entirely different method.
Re-prove it.

Another method to determine $|\Omega_{BE}|$

I will write down $n+k-1$ boxes




Now I do the following. Out of these $n+k-1$ boxes, I choose $n-1$ boxes.
How many ways can I do that?

There are $\binom{n+k-1}{n-1}$ subsets of $n-1$ boxes

So I pick from this set of $n+k-1$ boxes, a subset of $n-1$ boxes.
For each of the $n-1$ boxes I pick, what do I do to the boxes?

I do this. If I pick this box, I turn it 90° .
And I stop calling it a box - I call it a stopper.



So I pick the $n-1$ boxes and turn them into stoppers.
That way, I am left with: $n+k-1$
 $- (n-1)$
 $\hline k$ boxes


Each set of boxes is between stoppers.

So that the process of determining how many boxes are between stoppers is the same as the process of checking ($\checkmark$) boxes.

Therefore, this is the same as placing k indistinguishable balls into n boxes.

But, now you say "The stoppers are completely different."
Well let's see.

$$\binom{n+k-1}{n-1} = \binom{n+k-1}{k}$$



Identity (choosing the complementary set)

$$= \frac{(n+k-1)_k}{k!}$$

$$= \frac{(n+k-1)(n+k-2) \dots (n)}{k!}$$

(the numerator is just $\langle n \rangle_k$ w/ terms in the opposite order)

$$|\Omega_{BE}| = \frac{\langle n \rangle_k}{k!} = \langle n \rangle_k$$

So we got the same results.

Let's see a third proof of $|\Omega_{BE}| = \frac{\langle n \rangle_k}{k!}$

This is very close to the argument we've been using. But there is a variant and this is an important result.

I have the indistinguishable balls and I put them into boxes.

I do this in 2 steps:

1) I dispose the balls into the boxes $\langle n \rangle_k$

2) then I divide by $k!$, as the balls are indistinguishable and there are $k!$ permutations of each of these positions.

$$|\Omega_{BE}| = \frac{\langle n \rangle_k}{k!}$$

2/13/98, 10

Step 2 identifies any dispositions which have the same occupation numbers.

But 2 dispositions have the same occupation numbers if they differ by a permutation.

It is an interesting exercise you should do when you go home to see why this doesn't work w/ Maxwell-Boltzmann statistics. In other words, why can't you just:

- 1) place the balls into the boxes by Maxwell-Boltzmann and then
- 2) divide by $k!$

to get k indistinguishable balls in n boxes.

I won't tell you why this doesn't work, except you don't get integers. But, convince yourself of the reason why it doesn't work. Do it as an example w/ 2 balls + 3 boxes, and you'll see.

What about an interpretation?

I don't want to be very sophisticated.

Let's compute occupation numbers.

$$P_{BE}(\theta_i = i) = \frac{\binom{n-1}{k-i}}{\binom{n}{k}}$$

$P_{BE}(w)$ - probability of sample point in BE

I pick a subset of i balls.
Since the balls are indistinguishable,
that can be done in only one way.

The remaining $k-i$ balls (all indistinguishable) are disposed in $n-1$ boxes.

Research Problem

Find evident simplifications of this expression that I can put into my book.
This expression can be simplified in many ways.
Find something elegant.

This result is very important.


When particle physicists first started their game, they thought that all particles were Maxwell-Boltzmann. It's absurd to think that two particles were indistinguishable. But, then they computed the occupation numbers of states and they agreed in the formula.

Not w/ Maxwell-Boltzmann formula, but w/ Bose-Einstein formula.
So they were reluctantly lead to the conclusion that the particles are indistinguishable from this statistical fact.



Concrete Interpretation of Bose-Einstein statistics

An occupancy interpretation:




The interpretation of the sample points is this.

You have a big urn in which you have n balls labelled $a, b, \dots, c$
 You extract a ball,
 register the letter,
 put the ball back.

Sampling with replacement

You extract k balls with replacement.
 Then you add up how many a 's you got, how many b 's, etc.
 Irrespective of the order.



This is called a multiset.

↑ Like a set, except elements may appear w/ a multiplicity.

It is a situation that occurs very often in applications.

(In the 19th Century, multisets were called combinations.)

A multiset is an element of the Bose-Einstein sample space.
 Simply - a set of elements, where some element may occur more than once. The idea of a multiset carries the connotation that the same element may appear twice.

Just like there is an algebra of sets, there is also an algebra of multisets. If you ask me, it's one of those evolutionary choices, that the foundations of mathematics have developed in terms of sets rather than in terms of multisets.

2/13/98.12

You could very easily conceive of another evolutionary tree where the idea of multisets occurs before the idea of sets.
Perfectly conceivable.

Then logic would have developed differently, etc.

The idea of multisets is just as natural as the idea of sets.

Sets - obey the laws of Boolean Algebra

Multisets - obey the laws of another algebra, which I don't have time to go into.

Last question:

$$P_{BE} \left((\theta_1 = i_1) \wedge (\theta_2 = i_2) \wedge \dots \wedge (\theta_n = i_n) \right)$$

How many sample points in this event? 1

One. Because the assignment of the occupation numbers completely determines the dispositions.
(assuming that $i_1 + i_2 + \dots + i_n = k$, of course)

$$\therefore P_{BE} \left((\theta_1 = i_1) \wedge (\theta_2 = i_2) \wedge \dots \wedge (\theta_n = i_n) \right) = \frac{1}{\binom{n}{k}} \quad \leftarrow \text{Read: "One over multiset } n \text{ } k \text{"}$$

And so, we are almost finished.

We are only left w/ Fermi-Dirac statistics, which we will deal with next time.

Distribution and Occupancy (cont'd)

There are 4 main sample spaces in distribution and occupancy.
So far, we have covered 3.
Let us cover the last one.

To review:

I) Ω_{ms} : k balls into n boxes

$$P_{ms}(\omega) = \frac{1}{n^k}$$

probability of a sample point - namely, the placement of the balls into the boxes.

$$P((\theta_1 = i_1) \cap (\theta_2 = i_2) \cap \dots \cap (\theta_n = i_n)) = \underbrace{\binom{k}{i_1, i_2, \dots, i_n}}_{\text{multinomial coefficient}} \frac{1}{n^k}$$

probability of a given occupation number

And we computed the probability of several other events in this sample space.

II) Ω_D = dispositions

k balls into n boxes
and balls in each box are linearly ordered

Two models for this are:

1) balls are flags
boxes are poles } you place the flags on the poles.
Clearly, the flags on each pole are linearly ordered.

Two different linear orders give you two different dispositions.

2) balls are cars
boxes are lanes at the entrance of a turnpike } It's not enough for cars to choose a lane, but the cars line up linearly in front of the entrance.
This is the case for dispositions.

3) Further models?

I will give you a prize for further examples of dispositions.
I'd like some more examples.

$$P_D(\omega) = \frac{1}{\langle n \rangle_k}$$

In the case of dispositions, we have seen the remarkable fact that the distribution of occupation numbers is the same, no matter what the occupation numbers are.

$$P_D((\theta_1=i_1) \cap (\theta_2=i_2) \cap \dots \cap (\theta_n=i_n)) = \frac{k!}{\langle n \rangle_k}$$

under the assumption that
 $i_1 + i_2 + \dots + i_n = k$

$$= \binom{n}{k}^{-1}$$

(otherwise, it is 0)

We have reasoned this out in a couple of ways [2/13/98, 4-7].

III) Ω_{BE} - Bose-Einstein sample space

the balls are indistinguishable

We have seen, by reasoning in terms of dispositions [2/13/98, 7-10], that:

$$P_{BE}(\omega) = \frac{1}{\binom{n}{k}} \leftarrow \text{multiset coefficient}$$

The interpretations of Bose-Einstein sample space are several. Let's review a couple.

distribution interpretation - views the balls as being put into the boxes

occupancy interpretation - views boxes as letters of the alphabet and the balls as places where we put the letters. Sample points are viewed as words,

{ I don't know a good occupancy interpretation for dispositions.
 If you can think of one, I'd appreciate it.
 I don't know a good way to visualize this }

Distribution interpretation of Bose - Einstein

You have the boxes, with checks in the boxes.
The number of checks adding up to k .

✓ ✓✓ ✓ ... ✓✓

The balls being indistinguishable, all the data is how many balls are in each box.

Therefore, the distribution interpretation of BE can be interpreted as follows:

A sample point ω is simply the possible distribution of the i 's, i.e., the occupation numbers.

From a distribution point of view, the elements of BE are simply possible occupation numbers.

That's how we obtained $P_{BE}(\omega) = \frac{1}{\langle n \rangle_k}$

Occupancy interpretation of Bose-Einstein

The occupancy interpretation is a little more visualizable.
View the boxes as letters:

$\overbrace{a \quad b \quad \dots \quad c}^n$

n labelled balls are in an urn.

A sample point in Bose-Einstein sample space means:

do this k times $\left\{ \begin{array}{l} \text{pick out one ball from the urn,} \\ \text{register which ball it is,} \\ \text{put the ball back in the urn.} \end{array} \right.$

At the end, you have the data of how many times you have extracted ball "a", how many times ball "b", etc.
Altogether k extractions.

This is called, as we said last time, a multiset.

You can view the boxes as a set.

You are taking subsets of the the set, but a subset where some elements may be repeated.

This is a natural concept.

For example, suppose you use as letters $a, b, \dots, c$ and you consider them as variables.

Then you ask: How many monomials of degree k can be made w/ these variables?

And finally:
IV) Ω_{FD} - Fermi-Dirac sample space

again k balls and n boxes

balls are indistinguishable

all occupation numbers $\Theta_i = \begin{cases} 0 \\ 1 \end{cases}$

↑ this is called the Pauli exclusion principle.

Look at the illustrious names that are associated w/ these sample spaces, or statistics, as they are called.

How many sample points are there?

— ✓ ✓ . . . — ✓

Each box either has a check (✓) or not.

The balls are indistinguishable.

You check k of the n boxes.

Which means you are taking a subset of k boxes out of the n boxes.

So a sample point ω in Fermi-Dirac statistics is just a subset.

And physicists make a big deal of this.

A sample point ω is a k -subset of U .
Therefore:

$$P_{FD}(\omega) = \frac{1}{\binom{n}{k}}$$

Let's do, as an exercise, the following:

$$P_{FD}(\theta_1 = 0) = \frac{\binom{n-1}{k}}{\binom{n}{k}}$$

you take a subset of all the boxes and you leave out the first box (since $\theta_1 = 0$), which means you pick a k -subset from $n-1$ boxes.

$$P_{FD}(\theta_1 = 1) = \frac{\binom{n-1}{k-1}}{\binom{n}{k}}$$

you pick one ball and put it in the first box. Since the balls are indistinguishable, it doesn't matter which, you just check the box. The rest $(k-1)$ go into the rest of the boxes.

$$\underbrace{(\theta_1 = 0) \cup (\theta_1 = 1)}_{\text{independent events}} = \underbrace{\Omega_{FD}}_{\text{whole sample space}}$$

Then we have:

$$P(\theta_1 = 0) + P(\theta_1 = 1) = 1$$

$$\binom{n-1}{k} + \binom{n-1}{k-1} = \binom{n}{k}$$

This is a probabilistic proof of this identity.

Let's do similarly w/ Bose-Einstein:

$$P_{BE}(\theta_1 = i) = \frac{\langle n-1 \rangle_{k-i}}{\langle n \rangle_k}$$

[2/12/98.10]

There's only one way of placing i indistinguishable balls in the i th box. That means the remaining $k-i$ balls are placed, in a Bose-Einstein way, in the remaining $n-1$ boxes.

You people don't ask enough questions. I'm not giving out chocolate bars. I never had a class full of geniuses before.

$$\underbrace{(\theta_1 = 0) \cup (\theta_1 = 1) \cup \dots \cup (\theta_1 = k)}_{\text{all disjoint}} = \Omega_{BE}$$

↑ since box 1 has some number, or none, of balls.

Taking probabilities:

$$P_{BE}(\theta_1 = 0) + P_{BE}(\theta_1 = 1) + \dots + P_{BE}(\theta_1 = k) = 1$$

which gives:

$$\langle n-1 \rangle_k + \langle n-1 \rangle_{k-1} + \langle n-1 \rangle_{k-2} + \dots + \langle n-1 \rangle_0 = \langle n \rangle_k$$

There is a strange analogy between the multiset coefficient $\langle n \rangle_k$

counts the number of multisets of size k from a set of size n .

and the binomial coefficient $\binom{n}{k}$.

counts the number of subsets of size k from a set of size n .

There is a very strange analogy between the two, which is one of the deepest problems of mathematics.

Negative sets are multisets.

Watch.

$$\begin{aligned} \binom{-n}{k} &= \frac{\overbrace{(-n)(-n-1) \cdots (-n-k+1)}^{k \text{ terms}}}{k!} \\ &= (-1)^k \frac{(n)(n+1) \cdots (n+k-1)}{k!} \end{aligned}$$

$$= (-1)^k \frac{\langle n \rangle_k}{k!}$$

$$\boxed{\binom{-n}{k} = (-1)^k \langle n \rangle_k}$$

There is no satisfactory explanation of this identity.

If you get one, you get an A and you are excused from any further work in this course.

It's a very deep problem that no one has been able to figure out. What is the interpretation of this identity?

Now - we want to do some deeper computations in these sample spaces.

And to this end, we must develop some more powerful techniques for computing probability.

At this point, I was going to just give you the Principle of Inclusion/Exclusion, then I thought - "Gee, they pay such high tuition. I can't give the something straight out of the book."

So, I will give something not in the book (not in any book), which is the general theory for establishing identities of probability.

Mr. Guidi, take this down carefully. 😊

Rényi's Principle

Let's see some examples of identities holding in probability spaces.
The first example we have is:

$$P(A_1 \cup A_2) = P(A_1) + P(A_2) - P(A_1 \cap A_2) \quad [2/14/98.4-5]$$

This is an identity we established in the first day of class.

Another identity is: Let's try to expand it and get a simple identity

$$\begin{aligned} P(A_1 \cup A_2 \cup A_3) &= P(A_1 \cup (A_2 \cup A_3)) \quad \text{associativity of } \cup \text{ of sets} \\ &\quad \text{applying identity for union of 2 sets} \\ &= P(A_1) + P(A_2 \cup A_3) - P(A_1 \cap (A_2 \cup A_3)) \end{aligned}$$

$$\begin{aligned} &\quad \downarrow \text{distributive law} \\ &= P(A_1) + P(A_2) + P(A_3) - P(A_2 \cap A_3) - P((A_1 \cap A_2) \cup (A_1 \cap A_3)) \end{aligned}$$

$$= P(A_1) + P(A_2) + P(A_3) - P(A_2 \cap A_3) - P(A_1 \cap A_2) - P(A_1 \cap A_3) + P(A_1 \cap A_2 \cap A_3)$$

$$\begin{aligned} P(A_1 \cup A_2 \cup A_3) &= P(A_1) + P(A_2) + P(A_3) \\ &\quad - P(A_1 \cap A_2) - P(A_1 \cap A_3) - P(A_2 \cap A_3) \\ &\quad + P(A_1 \cap A_2 \cap A_3) \end{aligned}$$

if you intersect twice w/ A_1 , you might as well skip one of the A_i .

For the event $A_1 \cup A_2 \cup A_3$ we have the above inclusion/exclusion formula

Now the inclusion/exclusion principle is the generalization of this formula to n events.

It can be proved brutally by plugging in. We could do that and there is a proof in the book.

But then, there are so many other formulas that come up and each time you have to prove them from scratch.

We are led to ask: Is there a super method that leads to all possible formulas pertaining to events and probabilities?

And it is an amazing fact that there is a super method which will check any formula.

By way of guidance, let's write the Inclusion-Exclusion Principle in general:

Inclusion-Exclusion Principle

↑ why they call it a "principle" beats me.
It's just a formula.

$A_1, A_2, \dots, A_n$ = events in a sample space Ω ,
w/ probability P

Then the probability that at least one of the events occurs is given by:
↑ which means the union

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



(*)

$$P(A_1 \cup A_2 \cup \dots \cup A_n)$$

$$= \underbrace{\sum_i P(A_i)}_{\text{sum of probability of each individual event}} - \underbrace{\sum_{i < j} P(A_i \cap A_j)}_{\text{subtract the probabilities of all the pairs. Except you have to be very careful to count each pair only once. E.g., don't want to count both } A_1 \cap A_2 \text{ and } A_2 \cap A_1} + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots + \dots$$

$$\pm P(A_1 \cap A_2 \cap \dots \cap A_n)$$

↑ period

subtract the probabilities of all the pairs.
Except you have to be very careful to count each pair only once.
E.g., don't want to count both $A_1 \cap A_2$ and $A_2 \cap A_1$
All the pairs are counted only once.
You could take:

$$- \frac{1}{2} \sum_{\text{all pairs}} P(A_i \cup A_j)$$

but dividing by 2 gets you into trouble.
Instead, we limit the sum w/ $\sum_{i < j}$

It's a very powerful formula, as we are going to see.

Q: Can you generalize for infinite events?

A: Yes, but it's never been satisfactorily generalized.
I'll come back to that

The inclusion/exclusion formula is also an example of an identity holding in a probability space. And we'll soon see that there are other, more complicated, identities that have been developed to solve more complicated problems. There are oozes of identities.

There are also inequalities.
For example:

$$P(A_1 \cup A_2) \leq P(A_1) + P(A_2)$$

which is immediate from the formula: $P(A \cup B) = P(A) + P(B) - P(A \cap B)$

So there are also complicated inequalities.

Now, we will see that there is one method for checking every identity and every inequality in a probability space.

Let us first state Rényi's Principle in a sort of philosophical way and later we will prove it.

Rényi's Principle

I don't like the word involving, but I see no way out here.

An identity or inequality involving n events $A_1, A_2, \dots, A_n$ in a probability space holds for all $A_1, A_2, \dots, A_n$

if and only if

it is true in the special case where each $A_i = \begin{cases} \phi \\ \Omega \end{cases}$

each A_i is either the null set or the whole sample space

This is extraordinary

In order to accept a formula, you only have to check it when the A_i 's are the whole sample space and null.

If it's true under these cases, then it's true forever.

So you can write down the most complicated inequality or identity and it's a cinch to check them.

That's Rényi's Principle.

It's a decision procedure to check all possible inequalities and identities.

We will prove it later. Now, let's apply it.

Application - A Mickey Mouse application

Let's prove:

$$P(A_1 \cup A_2) = P(A_1) + P(A_2) - P(A_1 \cap A_2)$$

Let's pretend we don't know this. Let's apply Rényi's Principle to check it.

We only to check the following cases:

(1) $A_1 = \phi$ $A_2 = \phi$	(2) $A_1 = \phi$ $A_2 = \Omega$	(3) $A_1 = \Omega$ $A_2 = \phi$	(4) $A_1 = \Omega$ $A_2 = \Omega$
$0 = 0 + 0 - 0$ ✓	$1 = 0 + 1 - 0$ ✓	$1 = 1 + 0 - 0$ ✓ (symmetrical w/ (2))	$1 = 1 + 1 - 1$ ✓

Therefore, identity is true for every event, says Rényi.
You got away w/ murder

- Let's check next the inclusion-exclusion formula by this Principle
Verification of (by Rényi's Principle):

$$(*) P(A_1 \cup A_2 \cup \dots \cup A_n) = \sum_i P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots + \dots \pm P(A_1 \cap A_2 \cap \dots \cap A_n).$$

First, let me remind you that:

$$(1-1)^k = 0$$

In our 1st of 17 proofs of the binomial theorem [2/11/98, 7-8], we showed that:

$$\sum_{i=0}^k \binom{k}{i} (n-1)^{k-i} = n^k$$

Let $n = 1-1$. This gives:

$$\sum_{i=0}^k \binom{k}{i} (-1)^{k-i} = 0$$

$$\pm \binom{k}{0} \mp \binom{k}{1} \pm \binom{k}{2} \mp \dots + \binom{k}{k} = 0$$

sign dependent
on whether
 k is even (+)
or odd (-)

$$\binom{k}{0} = 1 \leftarrow \text{the empty set } \phi$$

Move 1st term to RHS.

Multiply both sides by -1 , but only if k is even.

This gives:

$$\boxed{\binom{k}{1} - \binom{k}{2} + \dots \pm \binom{k}{k} = 1}$$

We are going to use this.

By Rényi's Principle, in order to verify the identity $(*)$, all you have to do is choose some of the A_i 's to be the whole sample space Ω and some to be ϕ .

By symmetry, we choose the 1st k A_i 's to be Ω and the remaining $n-k$ A_i 's to be ϕ . This is symmetric.

We need only check $(*)$ when:

$$A_1 = A_2 = \dots = A_k = \Omega$$

$$A_{k+1} = \dots = A_n = \phi$$

} if we check it and it is true for all k , then it has to be true for all events.

$k=0$

$$A_1 = A_2 = \dots = A_n = \phi$$

$$0 = 0 \quad \leftarrow \text{LHS} = \text{RHS of } (**)$$

 $k > 0$

$$\text{LHS} = 1 \quad \leftarrow \text{since at least one of the } A_i = \Omega \text{ (the whole space)}$$

$$1 \stackrel{?}{=} \underbrace{\sum_i P(A_i)}_{k \text{ of } A_i \text{ equal } \Omega. \text{ So } k \text{ of } P(A_i)=1. \text{ The other } n-k \text{ have } P(A_i)=0. \text{ So this is } k.} - \underbrace{\sum_{i < j} P(A_i \cap A_j)}_{\text{Take all pairs that are the whole space. Because if one of them is } \phi, \text{ then } \Omega = \phi. \text{ So you take any pair of } A_1, A_2, \dots, A_k, \text{ since } A_1 = A_2 = \dots = A_k = \Omega. \text{ Any such pair has } P(A_i \cap A_j) = P(\Omega) = 1. \text{ How many pairs are there out of a set of } k \text{ elements? } \binom{k}{2}} + \underbrace{\sum_{i < j < k} P(A_i \cap A_j \cap A_k)}_{\text{Here, you take all the triples of the } k \text{ } A_i\text{'s which have the whole space } \Omega. \text{ Otherwise } \Omega = \phi \text{ and } P(\phi) = 0. \text{ There are } \binom{k}{3} \text{ ways to choose such triples}} - \dots$$

$$= k - \binom{k}{2} + \binom{k}{3} - \dots \pm \binom{k}{k}$$

On the previous page, we've shown that

$$\binom{k}{1} - \binom{k}{2} + \binom{k}{3} - \dots \pm \binom{k}{k} = 1$$

$\stackrel{=k}{=}$

Therefore, RHS above = 1

$$1 = 1 \quad \checkmark$$

Therefore, using Rényi's Principle, we have proved this fantastic inclusion/exclusion formulae with no work!



2/17/98, 12

Actually, an application of In-Ex Principle.

Before going on to fancier formulas, let's have some applications of Rényi's Principle.

Beginning w/ the most famous one, which is the problem of the matching.

• Application of Rényi's Principle - Matching Problem

I have 2 identical decks of cards.

I shuffle one.

Then I place them, one deck beside the other.

Then I turn over the ~~cards~~ top card in each deck simultaneously.

What is the probability that I get at least one match?

Or, equivalently, what's the probability I get no matches.

200 years ago, the Chevalier de Montmort, who is one of the earliest workers in probability, thought that this problem, together w/ the inclusion/exclusion principle, was "at the extreme limit of human reasoning." Now, we have it in 18.313 after 2 weeks.

What are we really talking about?

We have a set of n elements.

$1 \ 2 \ \dots \ n$

And we take an arbitrary permutation

$i_1 \ i_2 \ i_3 \ i_4 \ \dots \ i_n$

We are asking for the probability that the random permutation will have a fixed point.

A = event that a random permutation has at least one fixed point.

There are many other ways of stating this.
For example:

The secretary collects grades from the students, but he mixes up the grades. What is the probability that a student will come out with his or her correct grade?

Or you have a party and a gentleman gives you a ticket for your top hat. (It's 50 years ago). And he mixes up the tags. What is the probability that the gentleman will give you back your own top hat?

Next time, we will apply the inclusion/exclusion principle to see how easy it is to solve these problems.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



The inclusion-exclusion principle (cont'd)

Last time, we stated w/o proof Rényi's Principle.

We are going to see it again.

You can verify, from Rényi's Principle, one of the most important formulas in probability theory. Namely, the inclusion-exclusion principle.

The inclusion-exclusion principle says the following:

$A_1, A_2, \dots, A_n$ are any events (in a probability space - that goes w/o saying)

The inclusion-exclusion principle allows you to express the probability of one of the events happening, in terms of the probability of certain subsets of the events.

$$(*) \quad P(A_1 \cup A_2 \cup \dots \cup A_n) = \sum_{i=1}^n P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) \\ - \dots + \dots + (-1)^{n+1} P(A_1 \cap A_2 \cap \dots \cap A_n)$$

probability that one of the events happens is the union of the events

As I said, 200 years ago it was considered the utmost achievement - the limit of human reasoning.

It has extraordinarily vast applications for all sorts of situations.

So now let's see a few of them.

First, the problem of the matching.

Example: Matching

At commencement, every student gets a degree.

But the secretary mixed up the degrees, so that the president hands out the degrees in the wrong order.

What's the probability that at least one student will get his/her own degree?

$$\Omega = \{\omega = \text{permutations of the set } \{1, 2, \dots, n\}\}$$

Ω = space of all sample points, which are permutations of the set $\{1, 2, \dots, n\}$

A sample pt ω is a permutation

$$P(\omega) = \frac{1}{n!}$$

probability of a sample point

← Because in the absence of any further information, we give all permutations the same probability.

A = event that a random permutation has at least one fixpoint ($\omega(i) = i$)

ω is a permutation

$\omega(i)$ is the i th element of the permutation ω

$$P(A) = ?$$

We compute this probability by the inclusion-exclusion principle.

A_1 = event that $\omega(1)=1$ (i.e., that element 1 is a fixpoint)

A_2 = event that $\omega(2)=2$ (" " " 2 " " ")
etc.

Then, of course:

$$A = A_1 \cup A_2 \cup \dots \cup A_n$$

$P(A)$ is the probability that a permutation has at least one fixpoint.

↑ maybe more

So we compute $P(A)$ by the inclusion-exclusion principle, which I am about to cover.

$$P(A_1) = \frac{(n-1)!}{n!}$$

probability that a permutation sends 1 into 1

With 1 fixed, for 2 through n there are $(n-1)!$ permutations

each permutation has probability $\frac{1}{n!}$

$$= \frac{1}{n}$$

Similarly:

$$P(A_2) = \frac{(n-1)!}{n!} = \frac{1}{n}$$

Now let's compute probability of A_1 and A_2 :

$$P(A_1 \cap A_2) = \frac{(n-2)!}{n!}$$

you are producing the numbers 3 through n in an arbitrary way.

want 1 to remain fixed and 2 to remain fixed in the permutations

$$= \frac{1}{n(n-1)}$$

Similarly, of course:

$$P(A_i \cap A_j) = \frac{1}{n(n-1)}, \quad i < j$$

By the same reasoning.

It is clear that this kind of reasoning can be generalized:

$$P(A_i \cap A_j \cap A_k) = \frac{(n-3)!}{n!} \quad \leftarrow \text{you are permuting } n-3 \text{ points arbitrarily.}$$

$i < j < k$

i^{th} point is fixed
and j^{th} point is fixed
and k^{th} point is fixed

$$= \frac{1}{n(n-1)(n-2)}, \quad i < j < k$$

With these data, we can now apply the inclusion-exclusion principle.

By ^[2/18/98.1] (*), we have:

$$\begin{aligned} P(A) &= \sum_{i=1}^n P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots + \dots \\ &= n \cdot \frac{1}{n} - \binom{n}{2} \frac{1}{n(n-1)} + \binom{n}{3} \frac{1}{n(n-1)(n-2)} - \dots + \dots + (-1)^{n+1} \frac{1}{n!} \end{aligned}$$

n terms, each one equal to $\frac{1}{n}$. all unordered pairs possible $\binom{n}{2}$, each one equal to $\frac{1}{n(n-1)}$. all unordered triples possible $\binom{n}{3}$. etc.

$$\binom{n}{k} = \frac{n!}{k!(n-k)!} = \frac{(n)_k}{k!} \quad \left. \vphantom{\binom{n}{k}} \right\} \text{different ways of simplifying the binomial coefficient}$$

$$= 1 - \frac{(n)_2}{2!} \frac{1}{n(n-1)} + \frac{(n)_3}{3!} \frac{1}{n(n-1)(n-2)} - \dots + \dots \pm \frac{1}{n!}$$

$$P(A) = 1 - \frac{1}{2!} + \frac{1}{3!} - \dots + \dots \pm \frac{1}{n!}$$

That's our probability.

Is there a simple way of computing this probability?

Sure.

2/18/98.4

You recall from 18.01:

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \dots$$

set $x = -1$

$$e^{-1} = 1 - 1 + \frac{1}{2!} - \frac{1}{3!} + \dots$$

It just so happens that the Taylor formula for e^x converges extremely fast. So that after 4 terms, you have a 4 digit approximation.

So our probability $P(A)$ is approximately equal to:

$$P(A) \approx 1 - e^{-1}$$

$$P(A) = 1 - \frac{1}{e} = .6321$$

For all practical purposes, this is the answer

We have a remarkable conclusion:

The probability of at least one match, in this matching problem, is independent of the number degrees/cards/etc.

This is paradoxical at first, but facts are facts.

Whether $n = 52$ or $n = 320$, the probability is practically the same.

Let's do another one:

Example Ω_{MB} I want the probability that all the boxes are occupied.

$$\text{event } ((\theta_1 > 0) \cap (\theta_2 > 0) \cap \dots \cap (\theta_n > 0))$$

The complement of this event, by DeMorgan's Law:

$$(\theta_1 > 0)^c \cup (\theta_2 > 0)^c \cup \dots \cup (\theta_n > 0)^c$$

where $(\theta_i > 0)^c = (\theta_i = 0) \leftarrow$ if you don't see that, I quit.

$$\text{call this event } A = \underbrace{(\theta_1 = 0)}_{A_1} \cup \underbrace{(\theta_2 = 0)}_{A_2} \cup \dots \cup \underbrace{(\theta_n = 0)}_{A_n}$$

Let's compute $P(A)$ by the inclusion-exclusion principle.

2/18/98.5

In order to apply the inclusion-exclusion principle, we compute the probabilities of various intersections of these events.

Let's do that.

$$P(A_i) = \frac{(n-1)^k}{n^k}$$

← box i empty means that all the k balls go into $n-1$ boxes
 ← each sample point has probability $\frac{1}{n^k}$
 probability that box i is zero

Next:

$$P(A_i \cap A_j) = \frac{(n-2)^k}{n^k}$$

← same idea. All the balls go into the other boxes.
 $i < j$

etc.

Now we are ready to apply the inclusion-exclusion principle:

$$\begin{aligned}
 P(A) &= \sum_{i=1}^n P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots + \dots \\
 &= n \frac{(n-1)^k}{n^k} - \binom{n}{2} \frac{(n-2)^k}{n^k} + \binom{n}{3} \frac{(n-3)^k}{n^k} - \dots + \dots
 \end{aligned}$$

Further simplification is possible, but not much.

In fact, the simplification of this expression is a cottage industry.

There are many ways of simplifying it.

It doesn't simplify like the previous ones.

I'm tempted to simplify it, but we'd spend 15 minutes simplifying it and it's not worth it at this stage.

Finally, we want probability of the event that all the boxes are occupied.

This is the event A^c .

And we obtain this as:

$$P(A^c) = 1 - P(A)$$

Now let me pull a fast one on you.

Example Ω_{BE}

Let A be the same event, namely that at least one box is empty.

$$A = \underbrace{(\theta_1=0)}_{A_1} \cup \underbrace{(\theta_2=0)}_{A_2} \cup \dots \cup \underbrace{(\theta_n=0)}_{A_n}$$

What is $P_{BE}(A)$?

Do it:

$$P(A_i) = \frac{\langle n-1 \rangle_k}{\langle n \rangle_k} \leftarrow \begin{array}{l} \text{all } k \text{ indistinguishable balls go into} \\ n-1 \text{ boxes. The number of ways} \\ \text{this can be done is the multiset} \\ \text{coefficient } \langle n-1 \rangle_k. \end{array}$$

probability that box i is empty

$$P(A_i \cap A_j) = \frac{\langle n-2 \rangle_k}{\langle n \rangle_k} \leftarrow \begin{array}{l} \text{number of ways of placing the } k \text{ indistinguishable} \\ \text{balls into } n-2 \text{ boxes.} \end{array}$$

$i < j$

etc.

Therefore, by the mutual exclusion principle:

$$\begin{aligned} P(A) &= \sum_{i=1}^n P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots + \dots \\ &= n \frac{\langle n-1 \rangle_k}{\langle n \rangle_k} - \binom{n}{2} \frac{\langle n-2 \rangle_k}{\langle n \rangle_k} + \binom{n}{3} \frac{\langle n-3 \rangle_k}{\langle n \rangle_k} - \dots + \dots \end{aligned}$$

event that at least one box is empty.

That seems to be the answer.

Now - where did I pull a fast one?

This is correct.

This is stupid. Because I could have done it this way.

The balls are indistinguishable, I pick n balls and put each one in one box. They don't know.

That's the fundamental principle. Balls do not think.

So this is very deep actually.

We often think that they think.

We take any n of the balls and put each in a separate box.

How many ways can that be done? One.

Because the balls are indistinguishable. You put a check (✓) on it.

We are left with $k-n$ balls, which we put in any way in the boxes.

A = event at least one box is empty

A^c = event all the boxes are occupied

$$P(A^c) = \frac{\binom{n}{k-n}}{\binom{n}{k}}$$

That means I take the remaining $k-n$ balls and put them into n boxes.
(You have to have more balls than boxes)
 $k \geq n$

Therefore:

$$\begin{aligned} P(A) &= 1 - P(A^c) \\ &= 1 - \frac{\binom{n}{k-n}}{\binom{n}{k}} \end{aligned}$$

Taking the series we computed for $P(A)$ on the previous page and combining it with the above—we have proved a very interesting identity, with which you can tease your friends.

$$\binom{n}{1} = n \quad \binom{n-1}{k} - \binom{n}{2} \frac{\binom{n-2}{k}}{\binom{n}{k}} + \binom{n}{3} \frac{\binom{n-3}{k}}{\binom{n}{k}} - \dots + \dots = 1 - \frac{\binom{n}{k-n}}{\binom{n}{k}}$$

$$\binom{n}{k} = \binom{n}{k-n} + \sum_{i=1}^n (-1)^{n+i} \binom{n}{i} \binom{n-i}{k}$$

How do you prove that? If you don't know the secret, it's a little tough.

There is an apocryphal story about von Neumann being given this problem (or a variant of it) and immediately he came up with this. But then one of his friends gave him this and he says "How did you write that that fast?"

Now let's go back to Rényi's Principle and get some deeper identities for probabilities.

Let's jazz this up.

Again, given events $A_1, A_2, \dots, A_n$, we can ask other questions about these events.

The inclusion-exclusion principle gives you the probability that at least one of the events happens.


What if we need to compute the probability that exactly one of the events happens?

When you have a problem like that, what is step 1?
Give it a name.

Let w_1 ← the letter "w" = event that exactly one of the A_i 's happens
 w_2 = event that exactly two of the A_i 's happens
 etc.

Let's derive a formula for the probabilities of these events.
 Let's start w/ w_1 , which will give us a good orientation.

w_1 , in set theoretic terms, is the set of all sample points that belongs to exactly one of the A_i . We have the following picture



/// = w_1 - exactly 1 event
 ... = w_2 - exactly 2 events
 x x x = w_3 - exactly 3 events

There are many situations where we need to compute this. Very practical situations. I'd be interested if you come up with some interesting practical examples. Life and death examples, for me to put in my book, so people buy the book.

$$P(w_1) = ?$$

I want to first define some simplifying summations. This is standard notation. Let me motivate it this way. Go back to the inclusion-exclusion principle. You give a name to the first sum, you give a name to the second sum, etc. It's useful to do that because the same sums occur in the expressions for the probabilities of the w_i .

$$S_0 = 1$$

$$S_1 = \sum_{i=1}^n P(A_i)$$

$$S_2 = \sum_{i < j} P(A_i \cap A_j)$$

S_2 is the sum of probabilities of all pairs of intersections.

etc.

Observe that $S_1, S_2, \dots$ are not the probabilities of any events. They are just numbers.

In terms of S_i , the inclusion-exclusion principle can be written as follows:

$$P(A_1 \cup A_2 \cup \dots \cup A_n) = S_1 - S_2 + S_3 - \dots + \dots$$

Another way of writing the inclusion-exclusion principle.

And this gives us a lead on how to compute the probability of W_1 .

$$(*) \quad P(W_1) \stackrel{?}{=} S_1 - 2S_2 + 3S_3 - \dots + \dots$$

probability that
exactly one of
the events happens

Let's check whether this is true or not.

Let's check this really quickly.

How do we check this?

We check it by Rényi's Principle

In other words, this formula is true if it is true when the events A_i are allowed to be only either the whole sample space (Ω) or the empty set (ϕ).

If it checks for these cases, we win.

Keep your fingers crossed.

By Rényi's Principle:

expression being symmetric,

set $A_1 = \Omega$ $A_2 = \dots = A_n = \phi$

LHS of (*) = 1 $\leftarrow P(W_1)$ = probability that exactly one of the A_i 's occurs.

Every sample pt ω belongs to exactly one of the A_i 's.

W_1 is the event that only one of the A_i 's happens.

If A_1 is Ω and the rest of the A_i are ϕ , then for sure exactly one of the A_i happens.

So $P(W_1) = 1$

2/18/98.10

$S_1 = 1$, since $P(A_1) = 1$ and for $i = 2, \dots, n$,
 $P(A_i) = 0$.

What about S_2 ? S_2 involves the intersection of 2 events. But one of them will be the empty set.
 Therefore:

$$P(A_i \cap A_j) = 0, \quad \forall i, j$$

Same idea holds for $S_3, S_4, \dots$

$$P(w_1) \stackrel{?}{=} S_1 - 2S_2 + 3S_3 - \dots + \dots$$

$$1 \stackrel{\checkmark}{=} 1 - 0 + 0 - \dots + \dots$$

Good. It checks, at least for the first case.

Next:

$$\text{set } A_1 = A_2 = \Omega \quad A_3 = \dots = A_n = \phi$$

Let's see what happens.

What is the probability that exactly one of the events A_i happens? Zero

$$\uparrow A_1 = A_2 = \Omega$$

$$S_1 = 2$$

$$S_2 = \sum_{i < j} P(A_i \cap A_j) = 1$$

$$S_3 = S_4 = \dots = \phi \leftarrow \text{since at least one of the events in each intersection is } \phi.$$

$$P(w_1) \stackrel{?}{=} S_1 - 2S_2 + 3S_3 - \dots + \dots$$

$$0 \stackrel{\checkmark}{=} 2 - 2 + 0 - \dots + \dots$$

It checks.

So now we jump to the general case:

$$\text{set } A_1 = A_2 = \dots = A_k = \Omega$$

$$A_{k+1} = \dots = A_n = \phi$$

} if it checks for this, for all $k \geq 3$,
 then it's true by Rényi's Principle.

$$P(w_1) = 0 \leftarrow w_1 \text{ is an impossible event, as more than one of the } A_i \text{'s happens.}$$



$$P(W_i) \stackrel{?}{=} S_1 - 2S_2 + 3S_3 - \dots + \dots$$

$$(**) \quad 0 \stackrel{?}{=} K - 2\binom{K}{2} + 3\binom{K}{3} - \dots + \dots$$

If you have K events w/ $A_i = \Omega$ then there are

$\binom{K}{2}$ terms $A_i \cap A_j = \Omega$, the rest will be empty

$\binom{K}{3}$ terms $A_i \cap A_j \cap A_k = \Omega$, the rest will be empty

etc

Is the RHS $= 0$?

It is 0 for the following reason.

I'm sorry, but I have to resort to a non probabilistic argument.

I could use a probabilistic argument, but it's too long.

$$(1+x)^K = 1 + \binom{K}{1}x + \binom{K}{2}x^2 + \binom{K}{3}x^3 + \binom{K}{4}x^4 + \dots$$

differentiate w.r.t x :

$$K(1+x)^{K-1} = 0 + \binom{K}{1} + 2\binom{K}{2}x + 3\binom{K}{3}x^2 + 4\binom{K}{4}x^3 + \dots$$

Set $x = -1$

$$0 = K - 2\binom{K}{2} + 3\binom{K}{3} - 4\binom{K}{4} + \dots$$

So the RHS of $(**)$ is 0.

And we have:

$$0 \stackrel{?}{=} 0$$

By Rényi's Principle, it verifies the formula:

$$P(W_i) = S_1 - 2S_2 + 3S_3 + \dots - \dots$$

Done.

I leave it to you to prove the general formula:

$$P(W_k) = \sum_{i=0}^{n-k} (-1)^i \binom{k+i}{i} S_{i+k}$$

Pretty cute.

Verify by Rényi's Principle and you get a binomial coefficient identity, which reduces to the one I just covered, when you simplify.

You get these very fancy formula cheapo.

Next time, you ask - How does Rényi's Principle work? Why does it hold?

Because for your next problem set, I'm going to pull all sorts of inequalities out of the literature, which were proved by weirder methods, and you re-prove them by using Rényi's Principle.

So now you can prove any theorem. You can go to IEEE Transactions on Reliability Theory, which is about inclusion-exclusion formula, exclusively, and you can prove everything. You can trivialize all the papers there.

Q: When did Rényi come up w/ this?

A: He told it to me in 1969.



Last time we used Rényi's Principle to verify the inclusion-exclusion principle. You are wondering, now, how to prove Rényi's Principle. Well, I decided to write up the proof and give it out as a bulletin, so you can read it, understand it, and present it in class. Why? Because it's boring. You can understand it better by watching the various steps in the proof. The point in writing is required material, but there's no point in my spending a half an hour writing the proof.

Let's go on to another extremely important chapter.

Random Variables (integer)

for the moment, we will deal w/ one class of random variables, which we call integer random variables.

I will not keep saying integer. It will be understood that the random variables we deal w/ are integer random variables — until further notice.

First, the definition. And in this case, the definition conceals the meaning of the topic. Definitions conceal more than they reveal. You can not understand anything by reading these definitions. Similarly, it's the only hope with which you can understand the meaning.

Official Definition:

A random variable X is a function from a probability space to the integers, with the property that, for every integer n ,

$$\{\omega : X(\omega) = n\} = (X = n) \text{ is an event}$$

The set of all sample points, where the random variable takes the value n , which is abbreviated this way.

The event that $X = n$.

This is the probabilistic meaning.

You just don't think in terms of sample points.

If X is a random variable, the events $(X = n)$ are disjoint and $\bigcup_n (X = n) = \Omega$

it is understood that this is an integer random variable.

This is intuitively obvious, because if X takes the value 3, it can not take the value 2.

This is just common sense.

The sequence $p_n = P(X = n)$ is called the probability distribution of the random variable X .

notice that we don't use 2 sets of parentheses and we ought to. We will only use one. This is, strictly speaking:

$P(X = n)$ — the probability of an event (i.e., the event $(X = n)$).

It satisfies $p_n \geq 0$ and

$$\sum_n p_n = \sum_n P(X = n) = P\left(\bigcup_n (X = n)\right) = P(\Omega) = 1$$

The notion of a probability distribution is historically much older than the notion of a random variable.

For a long time, people could conceive of probability distributions, but they couldn't conceive of random variables.

The idea of random variables is a 20th century idea. There were stalwarts who refused to accept it and said only distributions exist and not random variables.

This is the bare definition. In order to understand the definition, let's enrich it by typical examples. If they don't tell you the typical examples, then you don't know.

Let's see some typical examples. We already have seen them.

Example in Ω_{MB}

We talked about the number of balls in the 1st box.

The occupancy number.

But what is the occupancy number?

The occupancy number is a random variable.

So, we were talking about random variables w/o realizing it.

Θ_1 is a random variable

For once, let's use sample points:

$\Theta_1(\omega)$

ω is a placement of the balls into the boxes.

So, for that placement, you count how many balls are in this 1st box.

So it's a function.

There are various numbers of balls in the 1st box for the placements $\omega \in \Omega$.

This function Θ_1 is from the placement of balls into boxes to the number of balls in box 1.

It is a function.

However, the spirit of probability is to use sample points as little as possible. As the great mathematician John von Neumann used to say:

"The spirit of probability is to be pointless"

If you have to use sample points, you haven't really thought probability. Of course, von Neumann had an ulterior motive. He was developing quantum probability. And in quantum probability, there aren't any points.

In order to understand quantum probability, you have to understand all of classical probability without ever mentioning the sample points.



What is the probability distribution of the random variable θ_1 ?
We already have this.

$$p_i = P(\theta_1 = i) = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k}$$

And we did verify that: [2/11/98.7-8]

$$\sum_{i=0}^n p_i = \sum_{i=0}^n P(\theta_1 = i) = \sum_{i=0}^n \frac{\binom{k}{i} (n-1)^{k-i}}{n^k} = 1$$

$$= P(\hat{U}(\theta_1 = i)) = P(\Omega)$$

Thus p_i is the probability distribution of the random variable θ_1 .

Now, what if I said at this point, "Let's do the same for θ_2 ."
You would say "That's kind of stupid."
Why?

This stupidity conceals deep concepts.

So now I say: " θ_2 is also a random variable"

And θ_1 is different from θ_2 , since the number of balls in the 1st box is not the same as the number of balls in the 2nd box.

They are related, but that doesn't mean anything. Wisky-washy words.

$$P(\theta_2 = i) = P(\theta_1 = i) \quad \leftarrow \text{however, the probability of the event } (\theta_2 = i) \text{ is the same as the event } (\theta_1 = i).$$

It took people 200 years to unscramble what's the same and what's not the same here.

For a long time, people would say they are the same. $(\theta_1 = i)$ and $(\theta_2 = i)$

But they are not the same.

They have the same distributions

But that does not entitle you to say they are the same random variables.

We reap the fruits of 200 years of misunderstandings.

We will, therefore, say that:

Two random variables are identically distributed when they have the same probability distributions.

Two random variables X and Y are said to be identically distributed when $P(X=n) = P(Y=n)$ for all n .

NB: being identically distributed is not the same as being the same.



Example: in Ω_{MB} ,

Θ_1 and Θ_2 are identically distributed,

since $P(\Theta_1 = i) = P(\Theta_2 = i)$ for all i

but random variables Θ_1 and Θ_2 are different

By the way, Θ_1 and Θ_2 are identically distributed in the other statistics (namely: Ω_D , Ω_{BE} , Ω_{FD}).

So it's an intuitive concept one has to distinguish. identically distributed
-vs-
same random variable
This is strictly a probabilistic concept.
It should guard you from using loosely the term "the same"

Next, we say:

Two random variables X and Y are independent when

$$P((X=n) \cap (Y=k)) = P(X=n)P(Y=k) \text{ for all integers } n \text{ and } k$$

In other words, the values of one variable are totally unrelated, to use sloppy language, to the values of the other variable.

Similarly for 3 or more random variables.

NB: pairwise independence is not the same as independence

(similar to what we discussed with independence of events [2/6/98.4-5])

You must have that, for all possible subsets, the probability of the intersections of the random variable events = product of probabilities of the individual random variable events - for all values of the random variables.

For X, Y, Z to be independent, ^{random variables} we must have:

$$P((X=n) \cap (Y=k)) = P(X=n)P(Y=k)$$

$$P((X=n) \cap (Z=j)) = P(X=n)P(Z=j)$$

$$P((Y=k) \cap (Z=j)) = P(Y=k)P(Z=j)$$

} all pairs

and

$$P((X=n) \cap (Y=k) \cap (Z=j)) = P(X=n)P(Y=k)P(Z=j)$$

} and
triple

for all n, k, j

} over all
 n, k, j



2/20/98.5

Research Problem:

Again, similar to independence of events, as discussed in [2/6/98.5]:

If you have a lot of random variables and you assume certain subsets to be independent, then you find that the other conditions are automatically satisfied.

No one has ever really figured it out. It's a nice research problem to work on.

If you give me \$1000, I'll work on it. I'll go off for a weekend.

- Example: θ_1 and θ_2 in MB are not independent random variables

$$P(\theta_1=i) \cap (\theta_2=j) = \frac{\binom{k}{i} \binom{k-i}{j} (n-2)^{k-i-j}}{n^k} \quad \leftarrow \text{We already discussed this [2/10/98.9]}$$

$$\frac{k!}{i!(k-i)!} \frac{(k-i)!}{j!(k-i-j)!} = \binom{k}{i, j, k-i-j} = \frac{\binom{k}{i, j, k-i-j} (n-2)^{k-i-j}}{n^k}$$

$\uparrow$ 1st box $\uparrow$ 2nd box $\uparrow$ elsewhere

but

$$P(\theta_1=i) P(\theta_2=j) = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k} \cdot \frac{\binom{k}{j} (n-1)^{k-j}}{n^k}$$

$$P(\theta_1=i) \cap (\theta_2=j) \neq P(\theta_1=i) P(\theta_2=j)$$

Therefore, they are not independent.

But, intuitively, of course they are not independentIf you put more balls into the 1st box, then the 2nd box will have fewer balls.

So, we didn't have to do this.

It's just a matter of common sense that $\theta_1 + \theta_2$ are not independent.

I did it once just to show you that the math checks your common sense.



Let's do some examples of random variables that are independent.
 Example - again in Maxwell-Boltzmann statistics

Let X_1 = position of the first ball

This is Maxwell-Boltzmann.

That means we label the boxes $1, 2, \dots, n$.

$$U = \{1, 2, \dots, n\}$$

Let X_2 = position of the second ball

etc.

I claim that X_1 and X_2 are independent.
 That should be intuitively clear.

Balls do not think.

So, where the first ball goes, has no influence on where the second ball goes.

Let's verify this mathematically.

Let's see that the formal notion of independence does coincide with our intuitive notion of independence.

The probability distribution:

$$P(X_1 = i) = \frac{1}{n}$$

the event that
the 1st ball
goes into the
box labelled i .

← all the boxes are equally possible.
Or we can compute this directly as:

$$P(X_1 = i) = \frac{n^{k-1}}{n^k} = \frac{1}{n}$$

Similarly:

$$P(X_2 = j) = \frac{1}{n}$$

Random variables X_1 and X_2 are identically distributed

$$P(X_1 = i) = P(X_2 = i) \text{ for all } i$$

"If you don't see that, I can not explain it." ←

Do you know how Leibniz used to say the famous philosopher?

"That's because of the Principle of Sufficient Reason"

$$P((X_1 = i) \cap (X_2 = j)) = \frac{n^{k-2}}{n^k}$$

$$= \frac{1}{n^2}$$

$$= P(X_1 = i) P(X_2 = j) \text{ for all } i, j$$

Therefore random variables X_1 and X_2 are independent.

← the 1st ball goes into the i th box. Period.
 The 2nd ball goes into the j th box. Period.
 There is only one way to do that.
 The remaining $k-2$ balls can go into any
 of the n boxes.

Similarly, random variables X_1, X_2 , and X_3 are independent.

- Now we can backtrack.
We defined the Maxwell-Boltzmann sample space by saying we have balls into boxes and so forth.
But suppose we wanted a purely abstract definition of the Maxwell-Boltzmann sample space, as you might find in some probability books.

Abstract Definition of Maxwell-Boltzmann Statistics

You have k random variables $X_1, X_2, \dots, X_k$, which are independent, identically distributed, and $P(X_i = i) = \frac{1}{n}$, $i = 1, 2, \dots, n$.

So the Theory of Maxwell-Boltzmann statistics is the theory of k independent, random variables taking the distribution $p_i = \frac{1}{n}$.

This takes away some of the fun, because you don't visualize things. It's a good, rigorous definition, but if you see a definition like this, you should immediately pick up a piece of paper and start drawing balls and boxes. Otherwise, you can't visualize it.

Such a definition, by the way, does not work in the case of the other statistics, where you have indistinguishable balls.

- Let's now see some fancier examples of random variables from the theory of the Bernoulli Process:

Famous random variables in the Bernoulli Process

$$\Omega = \{ \omega = (\omega_1, \omega_2, \dots), \omega_n \in \{0, 1\} \}$$

all infinite sequences of 0's and 1's.

X_n = outcome of the n^{th} toss

Let's look at it from the wrong point of view - namely, as a function.
This says that:

$$X_n(\omega) = \omega_n \quad \leftarrow \text{Sort of tautological}$$

ω is a sample point, which is an infinite sequence, you pick the n^{th} entry of it.



The probability distribution of the random variable X_n :

$$P(X_n=1) = p$$

$$P(X_n=0) = q = 1-p$$

$$P(X_n=j) = 0 \text{ if } j \notin \{0,1\}$$

We already verified, in different terms, that the random variables X_1 and X_2 are independent.

In other words, the outcome of the 1st toss is independent of the outcome of the 2nd toss.

We verified that explicitly at the beginning of the course.

We can now restate our verification in terms of random variables

It's much more efficient.

More generally, any finite set $X_1, X_2, \dots, X_k$ is independent

So, if you take any distinct set of tosses, they are independent random variables.

Now, let's take this random variable:

$$X_1 + X_2 + \dots + X_n = S_n$$

↑ this random variable is given the unfortunate name S_n , which we used in a different sense last time.

It's not my fault. All the books do it. So you better get used to the ambiguity.

Random variables being functions, you can add them.
And multiply them (not divide them)

What does S_n measure? The number of Heads (1's) in the first n tosses.

Let's find the probability distribution of S_n :

This is Fermi-Dirac statistics. We have n boxes, which are the tosses, and k indistinguishable balls, which are the k Heads.
the win

$$P(S_n=k) = \binom{n}{k} p^k q^{n-k}$$

$\binom{n}{k}$ ways of choosing k Heads.
 $\binom{n}{k}$ wins

For each win (i.e., k heads, $n-k$ tails), the probability is:

$$p^k q^{n-k}$$

$P(S_n = k) = \binom{n}{k} p^k q^{n-k}$ is called the binomial distribution

Since:

$$\bigcup_k (S_n = k) = \Omega \Rightarrow P\left(\bigcup_k (S_n = k)\right) = P(\Omega) = 1$$

Each of these events is disjoint,
(n is fixed)

it follows that:

$$\sum_k P(S_n = k) = P\left(\bigcup_k (S_n = k)\right)$$

$$\sum_k \binom{n}{k} p^k q^{n-k} = 1$$

and since $1 = p + q$, we have:

$$\boxed{\sum_k \binom{n}{k} p^k q^{n-k} = (p+q)^n}$$

The 3rd proof of the binomial theorem

↳ hardly a "proof" since one could put $(p+q)^n$ on R.H.S too

Now, let's jazz this up a little.

These random variables are kind of boring.

Now let's look at random variables that are the hard core of probability.

What random variables are the hard core of probability?

It's not a secret.

They are:

Waiting Times

Most probability problems boil down to computing waiting times.
You will eventually learn.

Let's define a random variable connected w/ this Bernoulli Process:

W_1 = waiting time for the first Head (1)

Notice how artificial it would be to think of this as a function:

~~$W_1(\omega)$~~

↑ a sample pt ω is a sequence of 1's and 0's.
You look for where the first 1 occurs in ω .
Don't look at it this way.

Look at it pointlessly

What is $P(W_1 = n)$?

How does the event $(W_1 = n)$ occur?

It can occur in only one way:

The first $n-1$ tosses must be Tails (0)
the n^{th} toss must be Heads (1)

$$P(W_1 = n) = q^{n-1}p$$

Let's note that:

$$P(W_1 = 1) = p$$

$$\text{Since } \underbrace{\bigcup_n (W_1 = n)} = \Omega$$

That's a lie.

Where is the lie? why did I glibly write that?

Because, except for an event with probability 0, the above equality is true.

True, except for the sequence which is all 0's.

$$P(0, 0, 0, \dots) = q^\infty$$

the event where
all of the
infinite tosses
are 0.

$$\text{Assuming } 0 < p < 1 \\ 1 > 1-p > 0 \\ \underbrace{1-p}_q > 0$$

$$\Rightarrow q^\infty = 0$$

$$P(0, 0, 0, \dots) = 0$$

This event, all tosses are 0, has probability 0.

We will often write equalities this way:

$$\bigcup_n (W_1 = n) \stackrel{?}{=} \Omega$$

meaning that the two events are equal and they differ, at most by an event of probability 0.

This is a very useful distinction from the ordinary notion of equality.

Q/A: The probability of any sample point is 0 in the Bernoulli process, if $0 < p < 1$.

$$P(\omega) = 0$$

↑ single sample pt

An infinite product of numbers between 0 and 1 is 0.

Q: Can you remove a countable number of events?

A: We'll do that. That's a good idea.

A countable number of events with probability 0.

You see things like this, for example, when we deal w/ distances bt natural numbers.

With this notion of equality, from this identity:

$$\bigcup_n \underbrace{(W_1 = n)}_{\text{disjoint events}} = \Omega$$

we infer, by taking probabilities of both sides:

$$\sum_n P(W_1 = n) = 1$$

where, again, $P(W_1 = n) = q^{n-1} p$

and we have proved the following identity:

$$p + pq + pq^2 + \dots = 1$$

You can now, of course, a posteriori, realize that this is:

$$\begin{aligned} p(1 + q + q^2 + \dots) &= p \frac{1}{1-q} \\ &= \frac{p}{p} \\ &= 1 \end{aligned}$$



2/20/98.12

So, the probability distribution of the waiting time for the first success is called the geometric distribution.

$$P(W_1 = n) = q^{n-1} p \text{ is called the } \underline{\text{geometric distribution}}$$

Don't let this go to your head.
These are purely historical names
(in this case, because of the geometric series, you see.)

Let's jazz this up, even more.

Let's now consider the waiting time for the k^{th} success (k^{th} head)

A reasonable thing to ask.

You'll see we jazz this up to the point that you won't believe it.

I've got something up my sleeve to give you in about 3 weeks, which has to do w/ Heads and Tails, that you wouldn't believe.


Absolutely, one of the deepest computations ever made. You will see a paradox. We are working up to it.

One of the most extraordinary paradoxes in probability.

W_k = waiting time for k^{th} 1

Let's find the probability distribution for this random variable:

$$P(W_k = n) = \begin{cases} 0 & \text{if } n < k \\ ? & \end{cases} \quad \leftarrow \text{it doesn't come out. You have to see this.}$$



We have n tosses.
The last one must be a 1, since it's a waiting time for the k^{th} 1.

In the preceding $n-1$ tosses, we must have $k-1$ 1's.
That's kind of Fermi-Diracish.

So this gives:

$$\binom{n-1}{k-1} p^k q^{n-k}$$

from $n-1$ tosses, we must have $k-1$ 1's. the last toss has to be a 1. You have $k-1$ 1's somewhere in between the rest have to be 0

Therefore, we can rewrite this as:

$$P(W_k = n) = \binom{n-1}{k-1} p^k q^{n-k}, \quad n \geq k$$

This is called the Negative Binomial Distribution.

And we have proved, by probabilistic methods:

$$\sum_{n \geq k} \binom{n-1}{k-1} p^k q^{n-k} = 1$$

$$\begin{cases} \sum_{n \geq k} P(W_k = n) = 1 \\ \downarrow \\ \sum_{n \geq k} P(W_k = n) = 1 \end{cases}$$

This is a non-trivial identity.

These are the easy ones.

These are absolutely universal distributions.

Even the notation is universally adopted - by everyone.

Next time, we'll do some non-trivial ones.

The easy probability distributions of random variables for the Bernoulli Process:

$$P(S_n = k) = \binom{n}{k} p^k q^{n-k}$$

of Heads in
1st n tosses

Binomial Distribution

$$P(W_1 = n) = q^{n-1} p$$

tosses until
1st Head

Geometric Distribution

$$P(W_k = n) = \binom{n-1}{k-1} p^k q^{n-k}$$

waiting time
for kth head

Negative Binomial Distribution

Random Variables (cont'd) (integer)

Last time, we saw the first examples of random variables.

Let's discuss some more.

But let's review the concept in light of the examples we have covered.

Again, I remind you, we are dealing w/ a special kind of random variable, which is integer random variable.

Meaning they are integer value functions.

Functions that take on integer values.

We consider those in advance, because it is easy to develop the theory.

Later we will consider the general case.

(it is understood that we are talking about integer random variables)

Given a probability space Ω, P , a random variable X on Ω is a function from Ω to the integers, with the property that for every integer n , $\{\omega : X(\omega) = n\} = \{X = n\}$ is an event.

The set of all sample points ω for which $X(\omega) = n$.

This is the last time I'm going to write it this way.

From now on, we use this notation. The event that $X = n$.

Pointlessly

I want to fill in some details that I glossed over last time.

$\bigcup_n \{X = n\} = \Omega$ except, possibly, for an event of null probability

I didn't tell you this last time, because I didn't want to overdo it

What does this mean?

It means that the random variable need not be defined on the entire sample space. It may be undefined on an event of null probability.

For example, last time we considered the waiting time for the first head in the Bernoulli Process. [2/20/98.10-12]

Strictly speaking, the waiting time is infinity for the one sample point that consists only of all tails.

The waiting time for the first head is not very well defined.

* This usually happens on infinite probability spaces.

- The sequence $p_n = P(X=n)$ is the probability distribution of the random variable X .
- X and Y are independent random variables when:
intuitively, when every event definable in terms of X is independent of every event definable in terms of Y .
This is the intuitive idea.
It is rendered by saying:

$$P(X=n) \cap (Y=k) = P(X=n)P(Y=k) \text{ for all integers } k, n$$

- X and Y are identically distributed (i.i.d.)

when:

universal notation,
used by everybody

$$P(X=n) = P(Y=n) \text{ for all } n$$

We've seen examples of random variables that are not the same, but are identically distributed.

- In the literature, you will see the following notation extremely frequently:
i.i.d. = independent and identically distributed

Examples:

in Ω_{MB}

occupation numbers θ_1 and θ_2 are i.i.d. but not i.i.d. [2/20/98.5]

in Ω_B

X_1 and X_2 (outcome of 1st and 2nd tosses, respectively) are i.i.d.

Bernoulli

In fact, the whole sequence $X_1, X_2, \dots$ is i.i.d.
(X_n being the outcome of the n th toss)

Q: Why can an infinite set of events be independent?

A: Good point. You get a chocolate bar.

Every finite subset of events $X_1, X_2, \dots$ is independent.

Let's see some more examples of random variables.

Further examples

in Ω_B , W_1 = waiting time for first Head (1)

Last time, we saw the probability distribution for W_1 [2/20/98.10-12]:

$$P(W_1 = n) = q^{n-1}p = \text{geometric distribution}$$

Let's consider the problem of the following random variable:

T_2 = the gap between the first and second Heads (1)

²(you'll see the reason for the subscript 2)
in a minute

{exclude toss of first Heads
include toss of second Heads}

All right. Let's see if we can compute the probability distribution of T_2 .

$$P(T_2 = n) = ?$$

Well, we have:

$$\bigcup_k (W_1 = k) = \Omega \quad \leftarrow \text{except for an event of probability 0}$$

Therefore, the event $(T_2 = n)$ can be written as:

$$\begin{aligned} (T_2 = n) &= (T_2 = n) \cap \Omega \\ &= (T_2 = n) \cap \left(\bigcup_k (W_1 = k) \right) \end{aligned}$$

now use the distributive law for events:

$$(T_2 = n) = \bigcup_k \left((T_2 = n) \cap (W_1 = k) \right)$$

the union of these is disjoint.

Because if $W_1 = 31$ then W_1 can not be 32.

Therefore, since $P(\text{LHS}) = P(\text{RHS})$, and $P(\text{RHS}) = \text{sum of probabilities}$:


$$(*) \quad P(T_2 = n) = \sum_k P((T_2 = n) \cap (W_1 = k))$$

So now we are confronted w/ computing the probability:

$$P((T_2=n) \cap (W_1=k)) = ?$$

That's not too hard.

Let's draw a picture. We have our tosses:



the event
($W_1=k$)

the event ($T_2=n$)

The next 1 after the first 1
takes n tosses.

The event $((T_2=n) \cap (W_1=k))$ can occur only in this way.
So we have here:

$$P((T_2=n) \cap (W_1=k)) = q^{k-1} p q^{n-1} p$$

Now that we have this, we can rewrite (*):

$$\begin{aligned} P(T_2=n) &= \sum_k q^{k-1} p q^{n-1} p \\ &= q^{n-1} p \sum_k q^{k-1} p \end{aligned}$$

This is the sum of
the geometric distribution

$$P(W_1=k) = q^{k-1} p$$

W_1 = waiting time for 1st Head

and we showed [2/20/98, 11-12] that

$$\sum_k q^{k-1} p = 1$$

Therefore:

$$P(T_2=n) = q^{n-1} p$$

We have obtained the following result:

The probability distribution of the gaps between the 1st and the 2nd Heads
is identically distributed with that of the waiting time for the 1st Heads

T_2 and W_1 are i.i.d.

But you have no doubt noticed that the preceding argument proves a lot more.
 Ha. Ha. Ha.
 What we have really proved is this,
 By the same argument:

$$\underbrace{P(W_1=k) \wedge (T_2=n)}_{q^{k-1} p q^{n-1} p} = \underbrace{P(W_1=k)}_{q^{k-1} p} \underbrace{P(T_2=n)}_{q^{n-1} p}$$

Obviously.
If you don't see it, I can't explain it.

That means the random variables W_1 and T_2 are independent
 W_1 and T_2 are i.i.d.

Now we readjust our intuition.

We say "Of course they have to be independent. No matter where the first Head occurs, we just start all over again."

A hand waving argument which is best rendered by the preceding mathematical argument.

Similarly, we now rename $W_1 \Rightarrow T_1$ and
 we give W_1 two names

$T_3 =$ gap between the 2nd + 3rd Heads (1)

$T_4 =$ " " " 3rd + 4th " "

etc.

We define all the gaps:

$T_1, T_2, T_3, \dots =$ gaps between successive Heads (1's)

An extension of the above argument, which is pretty routine, shows that the random variables $T_1, T_2, \dots$ are independent and identically distributed

$T_1, T_2, \dots$ are i.i.d.

↑
 { identically distributed with
 the geometric distribution }

We will see later that we can reconstruct the Bernoulli Process with these random variables alone.

The entire Bernoulli Process can be reconstructed from these data alone.



OK, Let's see some more random variables.


Sampling

That's a major business, in statistics anyway.

This is time to have the exact, clear definition of what sampling is about.

There are a number of basic sample spaces, not all of them finite.

Sampling is just balls into boxes.




Intuitively, the function from the balls into the boxes means that, at various times, you sample a member of the population.

Then you have a choice of 2 things. Either:

- 1) you throw it away $\leftarrow$ sampling w/o replacement
 -or-
 2) you put it back $\leftarrow$ sampling w/ replacement

Sampling with replacement is just Ω_{MB} , except for one feature that comes in in sampling.

What makes sampling slightly different from ordinary Maxwell-Boltzmann statistics is that the population is usually of 2 or more colors:



Note only are the boxes distinguishable, but they are also endowed w/ certain features (e.g., red or black).

You sampling has to do w/ determining how many reds or blacks there are and also computations pertaining thereto.

So sampling is more than just Maxwell-Boltzmann, in that you have a partition of the population into several characteristics.

We will just take a partition of 2, since it's easier to compute. But you can easily imagine any number of characteristics.

Sampling w/o replacement

Once you sample a member of the population, you can not replace it.
For example, you sample fruit flies and they die, because you kill them. That's sampling w/o replacement.

very often, you have to
do that

The sample space changes.

The sample space is a variant of Maxwell-Boltzmann where the occupation numbers can only be 0 or 1.

We have seen this (occupation numbers only 0 or 1) in Fermi-Dirac.
Now, I'm going to show you something very remarkable that is going to happen.

w/o replacement:

* Both samples and population are distinguishable, but occupation numbers are 0 or 1.

Q/A: Population is distinguishable. Different and distinct members.

You can not have an indistinguishable population.

Ex: There are people in this room, some of whom wear red socks some of whom wear black.
They are still distinct.

In real life, we don't have indistinguishable things.
(You are, in part, spoiling my act)
Population is distinguishable, until further notice.

What is indistinguishable — you have to prove that.
Things in the real world are always distinguishable.

There is a famous paper by Galileo where he is trying to convince someone that when you toss 2 dice, they are not indistinguishable. They are distinguishable. The probability has to be computed on the basis of the dice being distinguishable.
This is one of the dialogues by Galileo. It takes 13 pages.

Q: Could you give an example where the samples are distinguishable?

A: They are distinguishable. At a time, the first is picked out. Then the second is picked out. Etc.

Or John picks out the first. Mary picks out the second.

There is always some criteria that you can find to make them distinguishable, in real life.



So now you can ask: "Where does indistinguishable happen? Does it happen?"
That's where you talk about subsets.

When you pick a subset - that's where you get indistinguishable.
But we are not picking subsets here.

We are sampling.

And yet, watch what happens. Now this is interesting.

Let's consider now, in each of these 2 cases, the random variable

$R = \#$ of red members in a sample

Let's compute, in each case, what the probability distribution of R is.

Case 1 Ω_{MB} (Sampling w/ replacement)

$$P(R=i) = \frac{\binom{k}{i} r^i b^{k-i}}{n^k}$$

we have i red

from the sample of k , we choose i red. There are $\binom{k}{i}$ ways to do this.

For each $\binom{k}{i}$, there are $r^i b^{k-i}$ ways to put i balls in red boxes and $k-i$ balls in black boxes.

numerator is # of ways you can form i reds

denominator is # of sample points.

What have we done?

We've given the 5th proof of the binomial theorem.

$$\underbrace{\sum_i P(R=i)}_{\text{all disjoint}} = \Omega$$

$$\sum_i P(R=i) = 1$$

$$\sum_i \frac{\binom{k}{i} r^i b^{k-i}}{n^k} = 1$$

and since $r+b=n$

$$\sum_i \binom{k}{i} r^i b^{k-i} = (r+b)^k$$

That's the binomial theorem.



2/23/98.9

Now let's do it w/ feeding.
We kill the samples, Cruelly.

Case 2 (Sampling w/o replacement)

I know a biologist once who worked w/ mice all the time.
I won't say who it is, but it's someone very close to me.
And if she wasn't careful, she would kill the mouse.
She was so used to breaking their necks.
She sampled mice. And when she was asked to do samples with replacement, she had to be extremely careful. Because normally, she would do sampling w/o replacement.

This is a new sample space.
It doesn't have a name - no standard naming.

Ω_{MBwOR}

Maxwell-Boltzmann without Replacement

How many sample points are there?

You have k distinguishable balls going into n boxes.
And at most one ball per box.
We've done that.

$$|\Omega_{MBwOR}| = (n)_k \leftarrow \text{good old lower factorial}$$

We now consider the random variable R .

$$P_{MBwOR}(R=i) = \frac{\binom{k}{i} \binom{n}{i} (n-i)^{k-i}}{(n)_k}$$

out of the k balls, I pick i balls

I place the i balls into the red boxes. How many ways can I place i balls into the red boxes w/ occupation numbers 0 or 1?

$$\binom{n}{i}$$

The remaining $k-i$ go into black boxes w/ occupation numbers 0 or 1.

2/23/98.10

Now, let's do it the WRONG way.

Many books do it the wrong way.

They imagine that the balls are indistinguishable.

They say, "Let's do this with Fermi-Dirac statistics"

In Ω_{FD} , we can ask the same question.

We have r red boxes and b black boxes.

In Ω_{FD} , the random variable R has the probability distribution:

$$P_{FD}(R=i) = \frac{\binom{r}{i} \binom{b}{k-i}}{\binom{n}{k}}$$

the balls are now indistinguishable.

i balls go into the r red boxes.
That means I pick a subset of i elements from the set of red boxes.
This can be done in $\binom{r}{i}$ ways.

The remaining $k-i$ balls go into the b black boxes.

$$\text{Remember that } \binom{n}{k} = \frac{(n)_k}{k!}$$

$$= \frac{\frac{(r)_i}{i!} \cdot \frac{(b)_{k-i}}{(k-i)!}}{\frac{(n)_k}{k!}}$$

$$= \frac{\binom{k}{i} \binom{r}{i} \binom{b}{k-i}}{(n)_k}$$

Thus:

$$P_{MBWR}(R=i) = P_{FD}(R=i)$$

} why?

Now, my question to you is: Why?

I thought about this last night before going to bed for a half an hour - how best to say it.

I couldn't find a good explanation.

I give you 30 points if you find a good way of justifying it, that is completely convincing.

I can give you a mathematical interpretation.

But I'd like an intuitive, convincing, justification.



Warning - most of the books assume a priori that the balls are indistinguishable.

Which they have no right to do!

This distribution is called the hypergeometric distribution, usually written in the form:

$$P(R=i) = \frac{\binom{r}{i} \binom{b}{k-i}}{\binom{n}{k}}$$

Research Problem

Now develop an analogous coincidence for Bose-Einstein.

Something similar happens.

If you relate Bose-Einstein to Dispositions, you get again, if the balls go into either red or black, you get a similar coincidence, whether the balls are distinguishable or indistinguishable.

And explain. It's easy to work out. But it would be nice to have an explanation. So it's understandable to the public. Other than math.

Relate similarly Dispositions and Bose-Einstein.

So these are some examples of sampling.

Let's introduce a 3rd sample space associated w/ sampling.

Ω_s

↑ sampling

Now, you remember when we discussed the random variables associated w/ Bose-Einstein statistics, you could do it in 2 ways:

- 1) you could associate it w/ placing balls into boxes
- or- 2) by taking the random variable X_1 to be the position of the 1st ball
 X_2 " " " " " 2nd ball
 etc.

In this way, we get random variables that are independent

Now let's do this again, with an infinite number of balls.

The sample space is this.

I have an infinite number of balls, which are called samples.

The way you represent the population is one big urn.

The boxes become balls.



Urn

boxes become balls, which are red or black



This is an occupancy situation

I pick a ball and throw it in here.

The balls are the letters of the alphabet.

Remember, there are distribution and occupancy points of view.

This is the occupancy point of view:

Pull the letter of the alphabet and place it in its place

Pull an infinite number of samples with replacement

- • pull out a ball
- write the color of the ball
- put back the ball

For example: r, b, b, r, r, b, r, b, ...

$$X_i = \begin{cases} 1 = \text{red} & \text{with probability } \frac{r}{r+b} \\ 0 = \text{black} & \text{" " " } \frac{b}{r+b} \end{cases}$$

We have that $X_1, X_2, \dots$ are independent random variables

This is equivalent to the Bernoulli Process except that we could have 3 or more colors and then we would have a slight generalization of the Bernoulli Process.

Q: Within each color are the balls distinguishable?

A: Yes. They are distinguishable.

In real life, they are distinguishable.

But you have paradoxes where the answer comes out the same.

However, you have to be careful. There are some goofy mistakes made all the time.

These are the basic sampling spaces.

Armed w/ these examples, we can go back to the theory and develop further the theory of random variables.

We introduce the notion of joint distribution and expectation.

You see, it's all balls into boxes when all is said and done. Most problems in probability can be reduced to balls into boxes.

Joint Distribution

Joint Distribution of two ^{integer} random variables X and Y is the probability:

$$P((X=n) \cap (Y=k))$$

In common parlance, this is called a correlation. It tells you how much the occurrence of the event $(X=n)$ affects the occurrence of the event $(Y=k)$.

When the random variables X and Y are independent, the probability is the product $(P(X=n)P(Y=k))$. In general, it's not the product, as the random variables are generally not independent.

This probability is indicated by a double indexed sequence

$$P((X=n) \cap (Y=k)) = p_{nk} \leftarrow$$

What are the properties of this doubly indexed sequence? Very simple, Two properties.

Properties

$$1) p_{nk} \geq 0$$

$$2) \sum_k p_{nk} = \underbrace{P(X=n)}_{\text{namely the probability distribution of } X}, \quad \sum_n p_{nk} = \underbrace{P(Y=k)}_{\text{namely the probability distribution of } Y}$$

Let's verify this and then we quit.

$$\sum_k p_{nk} = \sum_k P((X=n) \cap (Y=k))$$

$$= P\left(\bigcup_k ((X=n) \cap (Y=k))\right)$$

these are disjoint sets for each value k .

using the distributive law for events:

$$= P\left((X=n) \cap \left(\bigcup_k (Y=k)\right)\right)$$

this is the whole space
 Ω

$$\sum_k p_{nk} = P(X=n)$$

As desired.

It's all the distributive law and Boolean algebra that does this trick.

By the way - this is where quantum probability parts company with classical probability.

In quantum probability the distributive law of events does not hold.

Therefore, having joint distributions are a pain in the neck in quantum probability.

So far, everything I've said works in quantum probability too.

But here, that's where it stops.

That's where you start having headaches like the uncertainty principle.

We'll talk about this later.

Random Variables: joint distribution and expectation

X and Y are ^{integer} random variables

Associated w/ each random variable, we have a probability distribution:

$$P(X=n) = p_n$$

where:

$$p_n \geq 0$$

$$\sum_n p_n = 1$$

$$P(Y=k) = r_k$$

where:

$$r_k \geq 0$$

$$\sum_k r_k = 1$$

At the end of last time, we defined the joint distribution of two random variables:

$$P((X=n) \cap (Y=k)) = p_{nk} \leftarrow \text{sequence w/ 2 indices}$$

And, thereafter, we showed that:

$$\sum_k p_{nk} = p_n$$

$$\sum_n p_{nk} = r_k$$

The standard language is that when you are given a joint distribution (p_{nk}) , and you compute from the joint distribution, the single distributions, the single distributions are called: the marginals

The sequence p_n , as well as the sequence r_k , are the 2 marginals of the double sequence p_{nk} .

It's just the way people speak.

The joint distribution of 2 random variables completely determines the statistical behavior of those random variables.

If you know the joint distribution, then you know also the individual distributions, by these formulas,

$$p_{nk} \begin{cases} \sum_k p_{nk} = p_n \\ \sum_n p_{nk} = r_k \end{cases}$$

In principle, knowing the joint distribution is knowing everything.

2/25/98.2

In the same way, of course, you can define the joint distribution of 3 random variables.

Similarly,

$$P((X=n) \cap (Y=k) \cap (Z=t)) = p_{nkt} \leftarrow \text{sequence w/ 3 indices}$$

Completely determines the statistical behavior of 3 random variables

And you have similar identities for the marginals.
For example:

$$\sum_k p_{nkt} = p_{nt} = P((X=n) \cap (Z=t))$$

you sum all the values of $(Y=k)$

Intuitively, this is clear. You make the value of Y irrelevant by summing over all possible values of Y .

And you have these fancy identities. They look fancy, but they are really trivial.

- If X and Y are independent, then $p_{nk} = p_{nk}$ and conversely. This is just another restatement of the definition of independence in fancy language. Technically speaking, 2 random variables are independent when their joint distribution is the product of the marginals.
which makes it sound very high floutent.

Since we are on this definition kick, let's define a stochastic process:

- A stochastic process is a probability space Ω, P , together with a family of random variables and the joint distributions given of any finite subset.

Needless to say, these joint distributions must be consistent with one another. In a sense that I will not even write. Because you know perfectly well. If you write up the marginal probabilities, automatically they are consistent with the random variables defined on a probability space.

2/25/98.3

For example, the Bernoulli Process is a stochastic process.

Because it is a probability space, together with the random variables $X_1, X_2, \dots$, which are the outcomes of the tosses.

So a stochastic process is very general.
It can be almost anything.

Now let's see some examples of joint distributions.

And this is a typical situation, if I may be allowed to philosophize here, where we see just how negative definitions are.

The moment you see the example, you see that there is something missing from that definition where we can not really render it clear.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



• Example $\Omega_{MB} \leftarrow$ a fairly familiar object by now

Let's take the joint distributions of occupation numbers.
We've already computed 3 or 4 times the distributions of the occupation numbers.

$$P(\theta_1 = i) = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k} \quad \left. \vphantom{\frac{\binom{k}{i} (n-1)^{k-i}}{n^k}} \right\} \begin{array}{l} \text{distribution of} \\ \text{the first occupation} \\ \text{number} \end{array}$$

It's also the distribution of any occupation number, as all the occupation are identically distributed (i.i.d.), for reasons that I can not explain, if you don't see that, [2/20/98.5]

To miss one of those completely obvious things, and there are many in probability, sometimes you just have to click - you have to see that two occupation numbers are i.i.d. Imagine someone who can not see that. What can we do about it?

I'm saying that because I have something up my sleeve where it's not quite easy to see that.

Now, let's compute the joint distribution of the 1st and 2nd occupation numbers:

$$P((\theta_1 = i) \cap (\theta_2 = j)) = \frac{\binom{k}{i, j, k-i-j} (n-2)^{k-i-j}}{n^k}$$

We've done this already twice.
Let's do it this time w/ multinomial coefficients

2/25/98.4

k balls, you pick i and put them in box 1.
From the remaining, you pick j and put them in box 2.

$$\binom{k}{i} \binom{k-i}{j} = \frac{k!}{i! (k-i)!} \frac{(k-i)!}{j! (k-i-j)!}$$

$$= \binom{k}{i, j, k-i-j} \leftarrow \text{multinomial coefficient}$$

The remaining balls go elsewhere (i.e., neither box 1 nor 2):

$$(n-2)^{k-i-j}$$

Now comes the fun.

From this we obtain, using the marginal identity $\sum_j p_{ij} = p_i$:

$$\sum_j p_{ij} = p_i$$

$$\sum_j \frac{\binom{k}{i, j, k-i-j} (n-2)^{k-i-j}}{n^k} = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k}$$

← The distribution of θ_1

$$\sum_j p((\theta_1=i) \cap (\theta_2=j)) = P(\theta_1=i)$$

Summing over all possible values of θ_2

$$\text{The identity } \sum_j \frac{\binom{k}{i, j, k-i-j} (n-2)^{k-i-j}}{n^k} = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k}$$

is a special case of the identity $\sum_j p_{ij} = p_i$.

The innocent little marginal identity becomes this.

That is what the definition does not tell you.

So, in principle, even though we could have computed the distribution of θ_1 by summing the joint distributions (over all values for θ_2), in practice,

$$p_i = \sum_j p_{ij}$$

it's easier to compute it directly and get a non-trivial binomial identity. And that's easily the case.

2/25/98.5

Let's jazz it up.

Let's take the joint distribution of all occupation numbers.

$$P(\theta_1=i) \cap (\theta_2=j) \cap \dots \cap (\theta_n=r) = \frac{\binom{k}{i,j,\dots,r} \cdot 1}{n^k} \leftarrow \text{they're all placed}$$

That's it. That's the joint distribution of all the occupation numbers.
And now, by adding over $j, \dots, r$, this sum of the joint distribution is the marginal p_i :

$$\sum_{j,\dots,r} p_{i,j,\dots,r} = p_i$$

$$\sum_{j,\dots,r} \frac{\binom{k}{i,j,\dots,r} \cdot 1}{n^k} = \frac{\binom{k}{i} (n-1)^{k-i}}{n^k}$$

We've added this over all the occupation numbers except the i th

Try to prove this directly!

You can, but it's work.

So you see, that definition conceals more than it displays

Ok, Let's do another one.

Let's do a new kind of distribution.

Namely:

Example - Multivariate hypergeometric distribution

A nice mouthful. You can impress your friends when they hear you say this. When you apply for a job.


Now what's the most rational way of doing this?
What do you think this is?

Balls into boxes.

We have a population of n members. These are the boxes.
And we take a sample of k elements.

The population is a population of people, like in this room, and on your neck is branded a letter of the alphabet. Some people branded $a, b, \dots, c$.
They are distinct people.

Therefore, you have the partition of the population, called boxes:



$a, b, \dots, c$ are a number of letters (or colors, etc)

Altogether $a + b + \dots + c = n$

{ There are always n boxes + today they are called population }



We pick a sample and then we ask "How many people branded a are in the sample?" "How many branded b?" etc.
This generalizes the hypergeometric distribution we discussed last time [2/23/98.9-11]

Let $A = \#$ of a-people in a sample of k balls into n boxes

$B = \#$ of b-people " " " " " " " " " "

etc.

Random variables.

Let's find their distribution.

$$P(A=i) = \frac{\binom{a}{i} \binom{n-a}{k-i}}{\binom{n}{k}}$$

pick i that are branded a .
from branding other than a ($n-a$), pick remaining $k-i$.

This is the hypergeometric distribution we discussed last time.
No multivariate yet.

Now, let's take the joint distribution of A and B :

$$P((A=i) \cap (B=j)) = \frac{\binom{a}{i} \binom{b}{j} \binom{n-a-b}{k-i-j}}{\binom{n}{k}}$$

From a , you pick i
From b , you pick j
From the rest ($n-a-b$), you pick $k-i-j$

Again, as we discussed last time, this is the same as if the balls were indistinguishable.
(I hope some of you have come up w/ a nice explanation).

This is a bivariate hypergeometric distribution

You may ask why do we call this a bivariate hypergeometric distribution.
The answer will be given later when we do randomization.
Because I want to give you the real answer, not some phony one.

The marginal identity gives us:

$$\sum_j P_{ij} = P_i$$

$$\sum_j \frac{\binom{a}{i} \binom{b}{j} \binom{n-a-b}{k-i-j}}{\binom{n}{k}} = \frac{\binom{a}{i} \binom{n-a}{k-i}}{\binom{n}{k}}$$

We've proved this.

We don't need to prove it in a complicated way

By the way, this is not as impressive as it looks.
A more interesting identity is:

Since $\sum_i p_i = 1$

$$\sum_i P(A=i) = 1$$

$$\sum_i \frac{\binom{a}{i} \binom{n-a}{k-i}}{\binom{n}{k}} = 1$$

$$\boxed{\sum_i \binom{a}{i} \binom{n-a}{k-i} = \binom{n}{k}}$$

Vandermonde's identity

We've proved this.
It doesn't need any proof

Now let's go all the way.
Let's take the joint distribution of all the brands (colors, etc)

$$P((A=i) \cap (B=j) \cap \dots \cap (C=r)) = \frac{\binom{a}{i} \binom{b}{j} \dots \binom{c}{r}}{\binom{n}{k}}$$

Multivariate
hypergeometric
distribution

automatically $i+j+\dots+r = n$, the number of boxes

From this, you can get fantastic identities by using marginals.
Statisticians would spend 2 or 3 weeks just writing down all these identities individually.

{ I worked out also the Bose-Einstein.
You can do the same stuff w/ Bose-Einstein.
Big deal. }



Let's go into expectation.

2/25/98.8

Expectation - a fundamental concept

X = integer random variable

$E(X)$ = average value it takes

expectation
of X

How do we write down correctly the average value it takes?
This is simple,

$$E(X) = \sum_n n P(X=n)$$

That's the average value.

You weigh each value of X (i.e., n) by the probability that X takes that value (i.e., $P(X=n)$).

This is provided that the series converges.

If the series doesn't converge, then we say that the random variable does not have an expectation.

It only has an infinite expectation if all the terms $n P(X=n) > 0$.

Now comes the one fact in probability that has kept probability in business.

As a matter of fact, if not for this, probability would not be in business.

It's true, even for quantum probability.

It's really a god-given fact.

Namely - the expectation of a sum of random variables is always equal to the sum of the expectations.

Even when the random variables are dependent.

Always

Basic fact

For any two (or more) random variables X and Y , we have

$$E(X+Y) = E(X) + E(Y)$$

This is a very powerful identity, because in many circumstances, it is the only way we have to work with dependent random variables.

Because in many cases, computing the joint distribution is a pain in the neck. So you have to use other techniques. And this is what you use.

[I hope you have started the problem set, by the way.
Yesterday I spent 6 hours on the next problem set.
In vicious glee.]

Proof

Let's take the random variable $X+Y$.Let's take the event $(X+Y=n)$.

Can we simplify this event? Well, yes.

Intuitively, it's like this:

$$(X+Y=n) = \bigcup_{i+j=n} ((X=i) \cap (Y=j))$$

The event $(X+Y=n)$ if and only if X takes some value i and Y takes some value j and $i+j=n$. This is the only way $(X+Y=n)$. And these are disjoint ways.

We have the above set theoretic identity.

By definition, the expectation of the random variable $X+Y$ is:

$$E(X+Y) = \sum_n n P(X+Y=n)$$

$$= \sum_n n \left[\bigcup_{i+j=n} ((X=i) \cap (Y=j)) \right] \quad \leftarrow \text{using above identity}$$

$$= \sum_n n \sum_{i+j=n} P((X=i) \cap (Y=j)) \quad \leftarrow \text{probability of disjoint events equals sum of probabilities}$$

Combine these sums into one double sum:

you can't simplify this because the random variables X and Y may not be independent

($\sum_{i+j=n}$ over all n is the same as the sum over all i and over all j independently)

$$= \sum_{i,j} (i+j) P((X=i) \cap (Y=j))$$

$$= \sum_{i,j} i P((X=i) \cap (Y=j)) + \sum_{i,j} j P((X=i) \cap (Y=j))$$

$$= \sum_i i \underbrace{\sum_j P((X=i) \cap (Y=j))}_{\text{this is the marginal identity}} + \sum_j j \underbrace{\sum_i P((X=i) \cap (Y=j))}_{\text{ditto}}$$

$$= \sum_i i P(X=i) + \sum_j j P(Y=j)$$

$$E(X+Y) = E(X) + E(Y)$$

Taking joint probability and adding over one variable gives the marginal

When we went over marginal identities, I was training you to see this.
Expectation is always additive.

Now, let's use this fact.

What shall we compute next?

The expectation of the occupation number in Maxwell-Boltzmann.
We always start w/ that.

• Example in $\mathcal{L}_{MB}$

$$E(\theta_i) = \sum_i i P(\theta_i = i) \quad \text{by definition}$$

We've computed $P(\theta_i = i)$ several times.
It is:

$$= \sum_i i \frac{\binom{k}{i} (n-1)^{k-i}}{n^k} = ?$$

What is this equal to?
It looks like a mess.

Therefore, we ring the alarm bell and say, there must be an easier way of doing this.

Actually, if you think carefully, I've already done this - in a different notation. You have it in your notes.

Now let's try to figure this out in some clever way.

$\theta_1, \theta_2, \dots, \theta_n$ are i.i.d. (identically distributed)

Hence,

$$E(\theta_1) = E(\theta_2) = \dots = E(\theta_n)$$

Obviously, since they all have the same probability distribution $P(\theta_i = i)$

$$E(\theta_1 + \theta_2 + \dots + \theta_n) = k$$

I take the random variable $\theta_1 + \theta_2 + \dots + \theta_n$.

What's that?

The number of balls (k).

Furthermore, this random variable takes the value k with probability 1.

$$P(\theta_1 + \theta_2 + \dots + \theta_n = k) = 1$$

Therefore, its expectation is:

$$E(\theta_1 + \theta_2 + \dots + \theta_n) = \sum_i i P(\theta_1 + \theta_2 + \dots + \theta_n = i)$$

$$= k \cdot 1$$

$$E(\theta_1 + \theta_2 + \dots + \theta_n) = k$$

By the main theorem:

$$E(\theta_1 + \theta_2 + \dots + \theta_n) = \underbrace{E(\theta_1) + E(\theta_2) + \dots + E(\theta_n)}_{\text{since } E(\theta_1) = E(\theta_2) = \dots = E(\theta_n)}$$

$$= n E(\theta_1)$$

Therefore,

$$k = n E(\theta_1)$$

$$\Rightarrow E(\theta_1) = \frac{k}{n}$$

which also gives us the identity:

$$E(\theta_1) = \sum_i i \frac{\binom{k}{i} (n-1)^{k-i}}{n^k} = \frac{k}{n}$$

Another fancy binomial identity proved cheapo.




Let's do another one.
A real fancy one.

- Example Ω hyper
 $\uparrow$ multivariate hypergeometric distribution
 with $a, b, \dots, c$ colors (brands)

What is the expectation of A ?

$$E(A) = ?$$


So we have a population of n people whose necks are branded $a, b, \dots, c$.
 And we take a sample of k people.
 What is the expected number of people who are branded a ?



$$\begin{aligned} E(A) &= \sum_i i P(A=i) \\ &= \sum_i i \frac{\binom{a}{i} \binom{n-a}{k-i}}{\binom{n}{k}} \end{aligned}$$

This looks like a mess.
 You can do it.
 Let's see if we can do it
 another way.

We look at the problem from the following point of view:



We are picking a sample of k .
 Let's look at the 1st box (i.e., the 1st person Mr. Torres).
 Either the 1st box (i.e., Mr. Torres) is picked or is not picked.

2/25/98.13

What is the probability Mr. Torres gets picked?

It's $\frac{k}{n}$. ← If you don't see that, I can't explain it.

Let's introduce a random variable Z_i for the i^{th} box:

$$Z_i = \begin{cases} 1 & \text{if } i^{\text{th}} \text{ box is picked} \\ 0 & \text{if } i^{\text{th}} \text{ box is not picked} \end{cases}$$

$$\text{Prob}(Z_i = 0) = \frac{n-1}{n} \cdot \frac{n-2}{n-1} \cdot \frac{n-3}{n-2} \cdots \frac{n-k}{n-k+1} = \frac{n-k}{n}$$

$$\text{so } \text{Prob}(Z_i = 1) = 1 - \frac{n-k}{n} = \frac{k}{n}$$

The probability distribution of Z_i is:

$$P(Z_i = 1) = \frac{k}{n}$$

$$P(Z_i = 0) = 1 - \frac{k}{n}$$

The number of elements picked out of people who are branded a is simply:

$$Z_1 + Z_2 + \dots + Z_a = A$$

What's the problem?

The random variables Z_i are not independent.

However, the expectation is additive - even if they are not independent.

Therefore, we can still compute the expectation.

$$\begin{aligned} E(A) &= E(Z_1 + Z_2 + \dots + Z_a) \\ &= E(Z_1) + E(Z_2) + \dots + E(Z_a) \end{aligned}$$

← {whether the Z_i are independent or not.}

and since $Z_1, Z_2, \dots, Z_n$ are i.i.d.

$$E(A) = a \frac{k}{n} \leftarrow \text{a result that takes people 3 pages to work out, if they evaluate the sum.}$$

And we have another binomial identity:

$$\sum_i i \frac{\binom{a}{i} \binom{n-a}{k-i}}{\binom{n}{k}} = a \frac{k}{n}$$

2/25/98.14

Now let's do a couple of examples from the Bernoulli Process.

Ah, I forgot.

There's a nice trick to compute the expectation of the waiting time for the 1st Heads.

Done in a clever way.

I'll do it the hard way.

• Example Bernoulli

$W_1 = T_1 =$ waiting time for the 1st Heads

$$E(T_1) = \sum_{n \geq 0} n P(T_1 = n)$$

I'm ashamed of what I am about to do.

You don't remember how it goes?

There's a really simple trick.

You could just write down the answer and call it a day.

$$= \sum_{n \geq 1} n q^{n-1} p \quad \leftarrow P(T_1 = n) = \underbrace{q^{n-1}}_{n-1 \text{ Tails}} \underbrace{p}_{\text{Head}}$$

$$= \frac{d}{dq} \sum_{n \geq 0} q^n p$$

$$= \frac{d}{dq} \left(\frac{1}{1-q} \right) p \quad \leftarrow \begin{array}{l} \text{telescoping sum:} \\ 1 + q + q^2 + \dots = \frac{1}{1-q}, \quad q < 1 \end{array}$$

$$= \frac{1}{(1-q)^2} p$$

$$= \frac{1}{1 - (1-p)^2} p$$

$$E(T_1) = \frac{1}{p}$$

This is not the right way to do it.
But it works.



2/25/98.15

The point I was coming to is this:
The waiting time for the k^{th} success (k^{th} Heads) is:

$$W_k = \underbrace{T_1}_{\substack{\text{waiting time} \\ \text{for 1st Heads} \\ \text{(success)}}} + \underbrace{T_2}_{2^{\text{nd}}} + \underbrace{T_3}_{3^{\text{rd}}} + \dots + \underbrace{T_k}_{k^{\text{th}}}$$

As we've seen [2/23/98.3-5], $T_1, T_2, \dots, T_k$ are identically distributed (i.i.d.)

$$\begin{aligned} E(W_k) &= E(T_1 + T_2 + \dots + T_k) \\ &= E(T_1) + E(T_2) + \dots + E(T_k) \\ &= k \frac{1}{p} \end{aligned}$$

$$E(W_k) = \frac{k}{p}$$

Doing it the hard way gives you:

$$\begin{aligned} E(W_k) &= \sum_n n P(W_k = n) \\ &= \sum_n n \binom{n-1}{k-1} q^{n-k} p^k \end{aligned}$$

negative binomial distribution
[2/20/98.12-13]

We have yet another binomial identity:

$$\sum_n n \binom{n-1}{k-1} q^{n-k} p^k = \frac{k}{p}$$

[Next time I will figure out a better way to compute $E(W_k)$]

The moral is that expectation is directly computed from the sum of expectations of very simple random variables, which are not always independent.

Conditional Probability

The basic technique for computing probabilities

What is conditional probability?

Let's take a sample space Ω , with sample point ω .

On this sample space, you have events which have a probability.



The idea is this.

Suppose that you know that event B has occurred. Then the probability of an event is altered by your knowledge that the event B has occurred. How?

Intuitively, it's easy to answer "How?"

You have an event A, which has some probability.

You know that event B has happened.

So only the part of A that is inside B can happen.

Therefore, the probability of A, given that B has happened will depend only on the intersection of A and B ($A \cap B$).

- Q: When you say that "B has happened," do you know that nothing other than B has happened?
- A: Yes. Perhaps something more than B has happened, but my knowledge of it stops there. The probabilities all say something about what we don't know. If we knew everything, we'd do this by the sample points.

Let's repeat this in set theoretic terms.

Instead of taking the whole set Ω , we take only the set of sample points that are in the set B.

Then we want the probability of sample pts of A that are in the set B.

This probability, for events in the set B, has to be computed in terms of the given probability of Ω .

So the temptation is to say:

If we know that the event B with $P(B) > 0$ has happened, then the "new" probability of any event A will depend only on $A \cap B$ (e.g., the "new" probability of the whole sample space Ω will depend only on $\Omega \cap B = B$)

However, the "new" probability can not be $P(A \cap B)$. The "new" probability of an event can not be $P(A \cap B)$. Because then we would have:

$$P(\Omega \cap B) = P(B) \neq 1$$

"new" probability of whole sample space is not 1

If we have to renormalize our probability in the knowledge that the event B has happened, then after renormalization, we have to end up w/ a probability. Therefore, the only way we have to renormalize a probability of A in the knowledge that B has happened is to set:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

} probability of event A, given that event B has happened.

read: Conditional probability of A given B

For example:

$$P(\Omega|B) = \frac{P(\Omega \cap B)}{P(B)}$$

(probability of whole sample space given event B has happened.)


$$= \frac{P(B)}{P(B)}$$

$$= 1$$

That kind of checks.

This notation, of universal use, is somewhat misleading.
Because of the following:

We have a sample space with events and we are given a probability. With conditional probability, we keep the same events, but we change the probability.



We also write: $P_B(A) = P(A|B)$

to indicate that the conditional probability is a new probability.

P_B satisfies the three probability axioms:

$$1) P_B(\Omega) = 1$$

$$2) P_B(A^c) = 1 - P_B(A)$$

$$3) \text{ if } A_1, A_2, \dots \text{ is a disjoint sequence of events}$$

$$\text{then } P_B(A_1 \cup A_2 \cup \dots) = P_B(A_1) + P_B(A_2) + \dots$$

In other words, P_B has full rights to be called a probability.



There is no reason to assume that given a family of events, that there should be a unique probability for that family.

There are lots of probabilities,

Now we see some of them - the conditional probabilities

Now the verification of these 3 axioms: you just write down the definition and it checks. I leave it to you,

$$(1) P_B(\Omega) = \frac{P(\Omega \cap B)}{P(B)}$$


$$= \frac{P(B)}{P(B)}$$

$$P_B(\Omega) = 1 \quad \checkmark$$

$$(2) P_B(A^c) \stackrel{?}{=} 1 - P_B(A)$$

$$\frac{P(A^c \cap B)}{P(B)} \stackrel{?}{=} \frac{P(B)}{P(B)} - \frac{P(A \cap B)}{P(B)}$$

$$P(B) \stackrel{?}{=} P(A^c \cap B) + P(A \cap B)$$



Since B is the union of disjoint sets:

$$B = (A^c \cap B) \cup (A \cap B)$$

$$\therefore P(B) = P(A^c \cap B) + P(A \cap B)$$

We can, thus, erase the "?" and we have:

$$P_B(A^c) = 1 - P_B(A) \quad \checkmark$$

$$(3) P_B(A_1 \cup A_2 \cup \dots) = \frac{P((A_1 \cup A_2 \cup \dots) \cap B)}{P(B)}$$

$$= \frac{P((A_1 \cap B) \cup (A_2 \cap B) \cup \dots)}{P(B)}$$

$$= \frac{P(A_1 \cap B) + P(A_2 \cap B) + \dots}{P(B)}$$

$$P_B(A_1 + A_2 + \dots) = P_B(A_1) + P_B(A_2) + \dots \quad \checkmark$$

each of the A_i are disjoint, thus each of $(A_i \cap B)$ are too.

What's interesting is to understand what conditional probability really means.
In practice. When you see what it really means, you will realize that this is the greatest discovery in the theory of probability.

Because, as we will see from the examples, it really renders our intuitive notion of how the probability of an event A changes, in the knowledge that some other event B has happened.

So let's see now the significant examples.

As with all of this, you don't understand anything w/o significant examples.

It's always best to start w/ the most Mickey Mouse possible example.

Because those Mickey Mouse examples are likely to give away the secrets more than the fancy ones.

So let's start w/ a super Mickey Mouse example.

Super Mickey Mouse.

What do you think I'm going to do?

Balls into boxes

Which one? Maxwell-Boltzmann - ordinary balls into ordinary boxes

Example in Ω_{MB}

Let's compute the following conditional probability.

What's the probability that the 2nd box has j balls given that the 1st box has no balls?

$$P(\theta_2 = j | \theta_1 = 0)$$

Now you can work it out intuitively.
Because intuitively, if you assume the 1st box has no balls, it is as if you don't see that, I'm sorry.
though you had $n-1$ boxes.

Let's make this come out automatically from the definitions of conditional probability.

Let's pretend we don't know what our intuition tells us and let's see what happens.

Use the definition:

$$P(\theta_2 = j | \theta_1 = 0) = \frac{P((\theta_2 = j) \cap (\theta_1 = 0))}{P(\theta_1 = 0)} \quad \leftarrow \text{The scriptures say}$$

That's what the scriptures say. Let's compute this.

We've computed it umpteen times. Till you're sick of it.

What's the numerator?

Well, it's the joint distribution of 2 random variables

The numerator is:

We take j balls out of K and place them in box 2.

None of them go into box 1.

You're left w/ $K-j$ balls which go into $n-2$ boxes (i.e., neither box 1 nor 2)

$$P((\theta_2=j) \cap (\theta_1=0)) = \frac{\binom{k}{j} (n-2)^{k-j}}{n^k}$$

The denominator:

That's easy. All k balls go into the boxes other than box 1.
This can happen $(n-1)^k$ ways.

$$P(\theta_1=0) = \frac{(n-1)^k}{n^k}$$

So we have:

$$\begin{aligned} P(\theta_2=j | \theta_1=0) &= \frac{\binom{k}{j} (n-2)^{k-j}}{\cancel{n^k}} \\ &= \frac{\binom{k}{j} (n-2)^{k-j}}{(n-1)^k} \end{aligned}$$

This is precisely what we got from the intuitive idea.
So this begins to convince you that maybe there is something to this.
Contrary to your initial skepticism.

OK - now let's do a non-intuitive example.

• Example - Bernoulli Process Ω_B

Q/A: Is there a bias.
No. Coin is fair

Let event $A = (X_2 = 1)$

namely, the 2nd toss is a Heads (1)

event $B =$ at least one of X_1, X_2 equals Heads (1)

Let's compute:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

$A \cap B =$ event that:
 $(X_2 = 1)$ and at least
one of X_1, X_2
equals 1

So this consists of the following events:

$$A \cap B = ((X_1 = 1) \cap (X_2 = 1)) \cup ((X_1 = 0) \cap (X_2 = 1))$$

disjoint events

$$P(A \cap B) = P((X_1 = 1) \cap (X_2 = 1)) + P((X_1 = 0) \cap (X_2 = 1))$$

$$= \frac{1}{4} + \frac{1}{4}$$

$$P(A \cap B) = \frac{1}{2}$$

$$\text{Note that } P(B) = \frac{3}{4}$$

So:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}$$

$$= \frac{\frac{1}{2}}{\frac{3}{4}}$$

$$P(A|B) = \frac{2}{3}$$

It's not what we expected

Q/A: Is this the Monte Hall problem?

You can't go very far w/ just the definition.
You now have to develop the basic techniques for working with conditional probabilities. And we'll do very fancy problems.

Basic Techniques for working w/ conditional probabilities

There are 3 basic rules, which will be used over and over again and which can solve extraordinary problems, if properly used.

Rule 1 - Rule of successive probabilities

Remember that $P(A|B) = \frac{P(A \cap B)}{P(B)}$

Therefore, we can consider:

$P_B(A|C)$ because P_B is a probability

Lets work it out:

$$P_B(A|C) = \frac{P_B(A \cap C)}{P_B(C)}$$

and this simplifies to:

$$= \frac{\frac{P(A \cap C \cap B)}{P(B)}}{\frac{P(C \cap B)}{P(B)}}$$

since $P(B)$ is assumed to be > 0 .

$$= \frac{P(A \cap C \cap B)}{P(C \cap B)}$$

This can be viewed as another conditional probability

$$\left\{ \begin{array}{l} \text{This is the rule of} \\ \text{successive probabilities} \end{array} \right\} = P(A|C \cap B) = \underbrace{P((A|C)|B)}$$

we rewrite

$P_B(A|C)$ in standard terms

That's Rule 1:

iterated conditioning

Rule 2 is the most important:

2/27/98.8

Rule 2 - Law of alternatives

a partition is a set of sets

You must grasp the notion of this on a partition.

A partition π of a probability space Ω , P is a family of events $B \in \pi$ with the following properties:

(1) if $B \in \pi$ then $P(B) > 0$


$P(B)$ is positive

$B, B' \in \pi, B \neq B'$

(2) if $B \neq B' \in \pi$ then $B \cap B' = \emptyset$ except for an event of zero probability
(any 2 members of π are disjoint)

(3) $\bigcup_{B \in \pi} B = \Omega$ (except for an event of zero probability)

The intuitive picture of a partition is this:



You have your sample space Ω
And you cut it up.

The members of the partition are called blocks
if $B \in \pi$, B is called a block of π

So it's an intuitive concept - cutting up the sample space into blocks.

Example of a partition

we've only seen integer random variables to this point


if X is any (integer) random variable, the events $(X=n)$ for which $P(X=n) > 0$ are a partition of the sample space Ω .
positive

In other words, every block consists of all sample points where the random variable takes the value n .

Two sample points ω_1, ω_2 belong to the same block iff X of those sample points is the same (i.e., $X(\omega_1) = X(\omega_2)$).

This is an important concept.

So to every random variable, you can associate a partition.



Most partitions come from random variables in this way.
Not all, but most.

There is a beautiful algebra of partitions, which is analogous to the algebra of sets, which we might go over in the super-class.

Partitions are duals to sets in a deep philosophical sense.

A partition is a basic notion for classification.

It is the dual notion of the event.

Later on, we will see that, just as you assign probabilities to events, there is also a number you assign to partitions, which is called the entropy. The analogy is Entropy is to partitions what probability is to events.

Q/A: Note that property (2) says that the intersection of 2 distinct blocks B and B' is null, except for an event of zero probability.

$$B \cap B' = \emptyset \cup \text{possibly an event of zero probability}$$

Property (3) says that:

$$\bigcup_{B \in \pi} B = \Omega - \text{possibly an event of zero probability}$$

The intersection of two distinct blocks is disjoint (subject to the above proviso). They can't overlap.

In other words, partitions are classifications. You can view the random variable as the classification criteria.

Q/A: It is much more elegant to state properties (2) + (3) as:

$$(2) \quad P(B \cap B') = 0$$

$$(3) \quad P(\Omega - \bigcup_{B \in \pi} B) = 0$$

} this way you avoid having to mention the proviso "possibly an event of zero probability"

By the way, if you want to be really fancy, you can prove that in a partition, there is at most a countable number of blocks.
This follows automatically.

Q: Does having a countable number of blocks then imply that there be some random variable X that forms the partition?

A: You can always invent a random variable.

You can always find a random variable to choose the π with. Many of them. The point is that the partition abstracts the blocks of the random variable that determines the partition.

In other words, the blocks of a partition are distinct, but they are not distinguished by specifics because of the random variable.

I leave this for next hour.

We now return to the Law of Alternatives:

if π is a partition of $\underbrace{\Omega, P}_{\substack{\text{sample space } \Omega \\ \text{with probability } P}}$
then

$$P(A) = \sum_{\substack{B \in \pi \\ \text{sum over all} \\ \text{blocks in the} \\ \text{partition}}} P(A|B)P(B) \quad \leftarrow \text{That's the Law of Alternatives}$$

Proof

$$\text{RHS} = \sum_{B \in \pi} \frac{P(A \cap B)}{P(B)} \cancel{P(B)}$$

$$= \sum_{B \in \pi} P(A \cap B)$$

The blocks B are disjoint.

Therefore $A \cap B$ are even more disjoint.

$$= P\left(\bigcup_{B \in \pi} (B \cap A)\right)$$

Now, we use the distributive law of $\cup, \cap$.

$$= P\left(\underbrace{\left(\bigcup_{B \in \pi} B\right) \cap A}_{\Omega \text{ (from property (3))}}\right)$$

$$= P(\Omega \cap A)$$

$$P(A) = P(A)$$

as desired

Q: If we use the restated property (3) $P(\Omega - (\bigcup_{B \in \pi} B)) = 0$,
how do we carry out this proof?

A: Instead of "=" say "=", except for an event of probability zero"

By the way, this is what failed in quantum probability, because in quantum probability, the distributive law doesn't hold. So the whole idea of conditional probability still exists, but it gets much more sophisticated, as Mr. Feynman showed when he was taking this course in 1939. Why don't you show this? He could do it in 1939. You can do it too. Actually it was 1938. At that time the course was taught by Norbert Wiener. Sorry.



Rule 3 - Law of successive conditioning

Last night, just as I went to bed I thought "Gee, the examples are too trivial." I thought - maybe we can do the problem of the matching w/ conditional probability. And since I couldn't sleep, I worked out the problem of the matching using conditional probability. That's one of the examples coming up. It's quite interesting.

If events B_1 and B_2 are independent, then

$$P(B_1 | B_2) = P(B_1) \quad \leftarrow \left\{ \begin{array}{l} \text{The idea is that event } B_2 \\ \text{doesn't influence } B_1. \text{ Therefore} \\ \text{it shouldn't change.} \end{array} \right\}$$

and conversely.

This is immediately checked:

$$\begin{aligned} P(B_1 | B_2) &= \frac{P(B_1 \cap B_2)}{P(B_2)} \\ &= \frac{P(B_1) \cancel{P(B_2)}}{\cancel{P(B_2)}} \\ &= P(B_1) \quad \checkmark \end{aligned}$$

Since B_1 & B_2 are independent,
 $P(B_1 \cap B_2) = P(B_1) P(B_2)$

The law of successive conditioning is based on the following imaginary reasoning:

If B_1 and B_2 are independent, then $P(B_1 \cap B_2) = P(B_1) P(B_2)$

Suppose B_1 and B_2 are not independent. Can we fix this so we have something like it? Yes.

In general:

$$P(B_1 \cap B_2) = P(B_1) P(B_2 | B_1)$$

Trivial. From the definition.

So this gives you a new way of looking at it - Conditional Probability. You can compute the probability of the intersection of any two events as the probability of the 1st event times the probability of the 2nd event, given that the 1st has happened.

This is kind of nice.

So nice, that you can generalize to any number of events.

Let's see.

$$\begin{aligned}
 P(B_1 \cap B_2 \cap B_3) &= P(B_1) P(B_2 | B_1) P(B_3 | B_1 \cap B_2) \\
 &= \cancel{P(B_1)} \frac{P(\cancel{B_2} \cap B_1)}{\cancel{P(B_1)}} \frac{P(B_3 \cap B_1 \cap B_2)}{\cancel{P(B_1 \cap B_2)}} \\
 &= P(B_1 \cap B_2 \cap B_3) \checkmark
 \end{aligned}$$

It's very nice.

It tells us how to immediately generalize this for n events.
Exactly the same argument works for n events.

$$\begin{aligned}
 P(B_1 \cap B_2 \cap \dots \cap B_n) &= P(B_1) P(B_2 | B_1) P(B_3 | B_1 \cap B_2) \dots \\
 &\quad P(B_n | B_1 \cap B_2 \cap \dots \cap B_{n-1})
 \end{aligned}$$

That's the law of successive conditioning.

Now you are wondering: "This looks so abstruse, it can not have any practical application."

Next time we'll spend the whole just doing applications of these laws.
And you'll see how extraordinarily they simplify computations of probability.

Conditional Probability (cont'd) Probability trees

We continue our discussion of conditional probability, which, as we said, is the most important technique for the computation of probability.
Let's review some fundamental facts:

In a sample space Ω , the probability of A , given that B has happened, is defined as:

$$P(A|B) = \frac{P(A \cap B)}{P(B)} \quad \left. \vphantom{\frac{P(A \cap B)}{P(B)}} \right\} P(A \cap B) \text{ normalized by } P(B).$$

This, of course, requires that the probability of the hypothesis B is positive:

$$P(B) > 0$$

We have seen already some examples (and we will keep convincing ourselves) that this corresponds indeed to the intuitive notion of the event B happening. How the probability of A changes in the knowledge that the event B happens.

The fundamental properties of traditional probability are 3, but for the purposes of today's lecture, let's recall 2.

1) Law of alternatives

You are given $\Pi \equiv$ a partition of Ω

As you know, Π is a family of disjoint events called blocks, which cover the entire sample space, except for probability zero.

Each block has positive probability.


Blocks of probability 0 are not blocks.

We saw that every integer random variable defines a partition.

Let me write a formula that we didn't see last time:

$$(*) \quad P(A) = \sum_{B \in \Pi} P(A \cap B)$$

B ranges over
all the blocks
of the partition
 Π .



This is proved by the picture:

The probability of A is the sum of the probabilities of the intersection of A with the blocks of the partition.

That's obvious.

From (*) we get the Law of Alternatives, which is a more sophisticated version, which states:

$$P(A) = \sum_{B \in \pi} P(A|B) P(B)$$

The Law of Alternatives

The Law of Alternatives follows from formula (*) immediately, since:

$$P(A|B) = \frac{P(A \cap B)}{P(B)} \Rightarrow P(A \cap B) = P(A|B) P(B)$$

So the two formulas are really equivalent, although they have different interpretations.

2) Law of successive conditioning

$B_1, B_2, \dots, B_n$ - finite set of events

The probability of all these events happening - if they were independent, that would be the product of the probabilities. They are not independent, so we fix that.

How do we fix it?

Using conditional probabilities.

$$P(B_1 \cap B_2 \cap \dots \cap B_n) = P(B_1) P(B_2|B_1) P(B_3|B_1 \cap B_2) \dots P(B_n|B_1 \cap B_2 \cap \dots \cap B_{n-1})$$

Important special cases:

- if A and B are independent then $P(A|B) = P(A)$

This follows immediately from the definition, but intuitively, this confirms our interpretation of conditional probability.


Because this formula says that whether or not B happens does not influence the probability of A .

$$\begin{aligned} P(A|B) &= \frac{P(A \cap B)}{P(B)} \\ &= \frac{P(A) P(B)}{P(B)} \\ &= P(A) \end{aligned}$$

- if $\overbrace{A \subseteq B}^{A \text{ contained in } B}$ then $P(A|B) = \frac{P(A)}{P(B)}$

This further confirms our intuitive interpretation. Let's draw a picture:

If you know that B has happened, then probability of A must be normalized to 1.



so we renormalize the probability of A .

$$\begin{aligned} P(A|B) &= \frac{P(A \cap B)}{P(B)} \\ &= \frac{P(A)}{P(B)} \end{aligned} \quad A \cap B = A$$

- if $B_1 \supseteq B_2 \supseteq \dots \supseteq B_n$ then the Law of Successive conditioning becomes:

$$P(B_1 \cap B_2 \cap \dots \cap B_n) = P(B_1) P(B_2|B_1) \underbrace{P(B_3|B_2) P(B_4|B_3) \dots P(B_n|B_{n-1})}_{P(B_3|B_1 \cap B_2)}$$

$$P(B_3|B_1 \cap B_2)$$

since $B_1 \supseteq B_2$,

$$B_1 \cap B_2 = B_2.$$

Thus:

$$P(B_3|B_1 \cap B_2) = P(B_3|B_2)$$

This is trivial.


None the less, the formula is useful, as you will see.

Now we want to get a really good feeling for these formulas.
So we want to look at them from different perspectives.
These are the crucial formulas for the computation of probabilities.

I can give you the theory first and the examples later, or the examples first.
I'll give you the theory first.
Slightly more sadistic.

- A very important type of sample space is a sample space which is constructed on trees.
We want to discuss them in great detail because they occur in practice.

Sample spaces on trees (none of this is in the book, by the way)



this is a sample point w_2 .

We mark it just at the end, but we really mean the unique branch starting at the root and ending here is the sample point.

A probability tree means that you give probabilities to each of these sample points that add up to 1.

$$\sum_{i=1}^{19} P(w_i) = 1$$

The sample spaces that are associated w/ trees have for a sample point, a complete branch of the tree.

But watch the following:

Every vertex of the probability tree is an event,

Take, for example, the vertex labelled A_i .

A_i is the event containing all the sample points passing through A_i .




if vertex j is closer to the root than vertex i ,

$$A_j \supseteq A_i$$

event A_j contains A_i

In the example above, the event A_i contains the 4 sample points $w_{11}, w_{12}, w_{13}, w_{14}$.
The event A_j contains all these sample points, plus extras.
The farther up the tree you go, the bigger the event.

If you go down, from the root of the tree, you get a set of smaller events, until you get to the sample points.



if A_i are all vertices immediately below B (say),
then the events A_i are a partition of B .

This is obvious.

So you say "So what?"

Now you can do a wonderful thing.


We can give an interpretation to the edges of the tree

By conditional probability.

Watch.


Basic fact - The sample points are the complete path
Everything else follows from this basic fact.

If we have A immediately above B , so that $A \supseteq B$, then we
assign to the edge (A, B) the conditional probability $P(B|A)$.



Why is that good?

Watch - from the Law of successive conditioning:



← I want the probability of this sample point.
 $P(\omega)$

$P(A_1 \cap A_2 \cap \dots)$ We can express this in terms of conditional probability.

$$P(\omega) = \underbrace{P(A_1 | \Omega)}_{\text{this is the same as } P(A_1)} P(A_2 | A_1) P(A_3 | A_2) \dots$$

$$P(\omega) = P(A_1) P(A_2 | A_1) P(A_3 | A_2) \dots$$

Notice that:

$$A_1 \supseteq A_2 \supseteq A_3 \supseteq \dots$$

So we are in the special case of the law of successive conditioning where the events are contained in one another [3/2/98.3]

The point is, in real life, you are not given the probabilities of the sample points. You are only given the conditional probabilities. And you build the probabilities of the sample points from the conditional probabilities.

That's called Bayesian Theory.

If you are given the probabilities on the path, then you can determine the probabilities of all the vertices.

The vertices correspond to events.


The edges correspond to conditional probabilities.

Conversely:

If you are given conditional probabilities to the edges, which are consistent, then you have a sample space. And you define the probabilities on all the paths.

This is the way a lot of sample spaces are defined.
Just like this.
Even infinite sample spaces.

For example



You are admitted to MIT and you take 18.01.
If you pass, you take 18.03.
If you flunk, you flunk out.

If you pass 18.03, you take either 18.06 or 18.313.
If you flunk it, you switch to either Course 15 or Course 21.

All you are given are the conditional probabilities (e.g., that you pass 18.313, given that you've passed 18.03).
That's the natural data.

You are not given the probability of the sample points (i.e., your history at MIT).

The event that you graduate is a union of sample points.
All the sample points which end up w/ you graduating.

Flunking out is another event.

So you do computations this way.

Example The test problem - a classical example

In a hospital, people have a disease (the black plague) with probability:

$$\begin{aligned}
 p \text{ for plague} &\rightarrow p = .1 \\
 p^c &= 1 - .1 \\
 &= .9
 \end{aligned}$$

There is a test for the plague that is 80% accurate:

$$\begin{aligned}
 \text{accurate} & a = .8 \\
 \text{inaccurate} & i = 1 - .8 \\
 & = .2
 \end{aligned}$$

3/2/98.7

You go to the hospital to get tested and the test turns out positive.
What is the probability that you have the plague?

Let's build a tree.


Either you have the plague or you don't have the plague.

If you have the plague and the test is accurate, then it will come out as your having the plague. Or the test is inaccurate, then it will come out as your not having the plague.

If you don't have the plague and the test is accurate, test result is you don't have the plague. Or the test is inaccurate, then it will come out as your having the plague.

Let P = event that you have the plague

A = event that test result says you have the plague



We want the conditional probability of the event that you have the plague, given that the test result says you have the plague.
Namely:

What is $P(P|A)$?


$$P(P|A) = \frac{P(P \cap A)}{P(A)}$$

(1) The event A = union of the two sample points w_1 and w_4

$$A = \{w_1\} \cup \{w_4\}$$

$$\text{Thus } P(A) = P(w_1) + P(w_4)$$

(2) In this first approximation, let's assume that the accuracy/inaccuracy of the test results are independent of whether you have the plague or not.
That is:



$$P(P|A) = \frac{P(P \cap A)}{P(\omega_1) + P(\omega_4)} \leftarrow \text{from (1)}$$

$$= \frac{P(\omega_1)}{P(\omega_1) + P(\omega_4)} \leftarrow P \cap A = \{\omega_1, \omega_2\} \cap \{\omega_1, \omega_4\} = \{\omega_1\}$$

$$= \frac{(.1)(.8)}{(.1)(.8) + (.9)(.2)} = .308$$


How comforting.

By the way, now let's assume the test accuracy $a = .9$ and inaccuracy $i = .1$.

$$\text{Then } P(P|A) = \frac{(.1)(.9)}{(.1)(.9) + (.9)(.1)}$$

$$P(P|A) = .5 \leftarrow \text{The probability that you have the plague, given that your test result says you have it is a fair coin toss.}$$

The probability of the sample point ω_1 is the conditional probability times the conditional probability



$$\begin{aligned} P(\omega_1) &= P(P \cap A) \\ &= P(P|A) P(A|P) \\ &= P(P) P(A|P) \\ &= (.1)(.8) \end{aligned}$$

Doctors make mistakes 80% of the time on this and have to hire statisticians at \$100,000/year. Because they can't do this computation.


How can we jazz it up?

We can say the test has different accuracy according to whether you administer it to a patient with the plague, or w/o the plague.

How would we set up the tree in that case?

We set up the tree in much the same way.

In this case, we no longer have independent events.



Setting it up as a tree helps a lot.

Just as the probability of your flunking out can be set up as a probability tree, with all the conditional probabilities.

You can compute, for example, "What's the probability that I flunk out, given that I've taken 18.313?"

Or I can ask "Given a student that has flunked out, what is the probability he has taken 18.313?"

It's all in the tree. You can work it out. Once you've set up the tree, you can answer all these questions.

Now what do you think comes next?

Balls into boxes.

We jazz up Maxwell-Boltzmann this way:

I have an indefinite number of boxes. It might even be infinite.

Each box is an urn.

Every box (urn) has the same number of balls (n).


The 1st box has a_1 good balls, $n-a_1$ bad balls

" 2nd " " a_2 " " , $n-a_2$ " "

etc.

I would be grateful
for a good interpretation
of this.
A gripping interpretation.

Let's restrict ourselves to a finite number of boxes/urns (r):



all boxes have
 n balls

Each urn has a different number of good balls.


My strategy is the following:

1) pick an urn at random

then 2) pick k balls out of that urn


Let random variable A = number of good balls I pick.
What is the probability distribution of A ?

Strictly speaking, I'm deviating from my usual terminology.



This should be a population.
The balls are members of a population.
The balls are in the form of boxes in
Maxwell-Boltzmann.
The balls are really boxes.

Let's set up a tree for the sample space.
We have one edge for each urn. So,
you choose one of $U_1, \dots, U_r$.




Having chosen an urn, we are in
the ordinary Maxwell-Boltzmann
situation. We put k balls into
 n boxes.
The boxes are labelled good or
not good

You can do this w/ replacement or w/o replacement.

Let's do both cases.

First:

With replacement



The sample points $w_1, \dots, w_k$ correspond to all the possible sets of k distinguishable balls that can be extracted from urn U_1 .

Similarly, for each subtree under urn U_2 , the sample points w correspond to all the possible sets of k distinguishable balls that can be extracted from that urn.

The sample points w of our sample space Ω are these sets.

This is very important - there are different sample spaces Ω_{ms} , since we pick urns first.

In the absence of any further information, we give the choice of urns the same probability ($\frac{1}{r}$).

We will see later what happens if you alter these probabilities - if you give an uneven probability for the choice of urns.

We want the probability distribution $P(A=i)$.


This can be computed in one, and only one, way - with the law of alternatives.

$$P(A=i) = \sum_{j=1}^r P(A=i | U_j) P(U_j)$$

probability of the event
that we extract i
good balls, given that
we choose urn j .

$$\left\{ \begin{array}{l} U_j = \text{event that we choose} \\ \text{urn } j \\ P(U_j) = \frac{1}{r} \end{array} \right\}$$

We are adding all the sample points
as per theory. We didn't even need
to know the law of alternatives!



$\bigcirc$ = sample points w/ good
sets of k extracting
that have i good balls.

$$= \frac{1}{r} \sum_{j=1}^r P(A=i | U_j)$$

$$P(A=i) = \frac{1}{r} \sum_{j=1}^r \binom{k}{i} \left(\frac{a_j}{n}\right)^i \left(\frac{n-a_j}{n}\right)^{k-i}$$

probability ball
is a good one.

probability ball
is not a good one.

3/2/98.11

You mentally check that this is a probability distribution:

$$\sum_{i=0}^k P(A=i) = 1$$

$$\sum_{i=0}^k \frac{1}{r} \sum_{j=1}^r \binom{k}{i} \left(\frac{a_j}{n}\right)^i \left(\frac{n-a_j}{n}\right)^{k-i} = 1$$

$$\sum_{i=0}^k \binom{k}{i} \underbrace{\frac{1}{r} \sum_{j=1}^r \left(\frac{a_j}{n}\right)^i \left(\frac{n-a_j}{n}\right)^{k-i}} = 1$$

Since, from the binomial theorem:

$$\sum_{i=0}^k \binom{k}{i} \underbrace{p^i q^{k-i}} = 1, \text{ where } p+q=1$$

That's pretty easy.

Now, in the same sample space, let's ask another question.

What is the probability I've chosen the 1st urn, given that I have extracted i good balls?

$$P(U_1 | A=i) = \frac{P(U_1 \cap (A=i))}{P(A=i)}$$

We can rewrite the numerator as:

$$\begin{aligned} P(U_1 \cap (A=i)) &= P((A=i) \cap U_1) \\ &= P((A=i) | U_1) P(U_1) \end{aligned}$$

by definition of conditional probability.

$$= \frac{P((A=i) | U_1) P(U_1)}{P(A=i)}$$

We know these data:

$$P((A=i) | U_1) = \binom{k}{i} \left(\frac{a_1}{n}\right)^i \left(\frac{n-a_1}{n}\right)^{k-i}$$

$$P(U_1) = \frac{1}{r}$$

$$P(A=i) = \frac{1}{r} \sum_{j=1}^r \binom{k}{i} \left(\frac{a_j}{n}\right)^i \left(\frac{n-a_j}{n}\right)^{k-i}$$

$$= \frac{\cancel{\binom{k}{i}} \left(\frac{a_1}{n}\right)^i \left(\frac{n-a_1}{n}\right)^{k-i} \cdot \cancel{\frac{1}{r}}}{\cancel{\frac{1}{r}} \sum_{j=1}^r \cancel{\binom{k}{i}} \left(\frac{a_j}{n}\right)^i \left(\frac{n-a_j}{n}\right)^{k-i}}$$

$$P(U_1 | A=i) = \frac{\left(\frac{a_1}{n}\right)^i \left(\frac{n-a_1}{n}\right)^{k-i}}{\sum_{j=1}^r \left(\frac{a_j}{n}\right)^i \left(\frac{n-a_j}{n}\right)^{k-i}}$$

The idea is this. You don't know which urn the balls are extracted from. And you want an estimate of the probability that you've picked a given urn on the basis of how many good balls you have. That estimate is $P(U_1 | A=i)$.

↳ the probability that you've picked urn 1, given that you've extracted i good balls.

This is called the Bayesian probability.

We will discuss, next time, Bayes' Law.

All this nonsense about Bayes Law disappears when you visualize everything in terms of trees.



Quiz 1 average = 31
median = 60

Conditional Probability (cont'd)


The main technique for computing probabilities.

$$P(A|B)$$

We're going to examine this from all sorts of different viewpoints.
Let's start with a couple of easy problems.

An urn has 2 balls, colored a or b. The balls are distinguishable.
An a-ball is extracted (a ball is extracted from the urn + it turns out to be an a-ball).
What is the probability that both balls are a-balls?

This is a typical problem of the kind that we have begun to name last time as Bayesian, after the writer Sir John Bayes, who, by accident, wrote two pages on this.



We don't know what the distribution of the balls in an urn is.
Our sample space consists of first imagining 3 possibilities for the urn.
Then the extraction.

Let event A = number of a-balls extracted from the single extraction
event U = number of a-balls in the urn before extraction

All together, you have 4 sample points ($\omega_1, \omega_2, \omega_3, \omega_4$)

In the absence of any further information, we assume we choose one of the 3 urn configurations with equal probability.

This is something we will discuss later.

Under what circumstances should we make a different assumption?

This is something that is extremely hard to justify.

That is, why is it, in the absence of any further information, that you assume all the possibilities have equal probability.

Q: By symmetry?

A: No. That's circular.

Reams of words have been written on this.

There are people in the philosophy department who are working full-time on this.

Let's first consider $P(A=1)$.

Evaluated by the law of alternatives:

$$P(A=1) = \sum_{(U=i) \in \Pi} P(A=1|U=i) P(U=i)$$

$$= \underbrace{P(A=1|U=0)}_0 P(U=0) + P(A=1|U=1)P(U=1) + P(A=1|U=2)P(U=2)$$

This is 0. Since the urn has no a-ball, the probability of extracting an a-ball is 0.

The event $(A=1|U=1)$ has the sample point ω_2

$$P(A=1|U=1) = \frac{1}{2}$$

The event $(A=1|U=2)$ has the sample point ω_1 .

$$P(A=1|U=2) = 1$$

$$\begin{aligned} P(A=1) &= \frac{1}{2} \cdot \frac{1}{3} + 1 \cdot \frac{1}{3} \\ &= \frac{1}{2} \end{aligned}$$

Now to the question being asked,
What's being asked is $P(U=2|A=1)$.

if you extract one a-ball, what's the probability that the urn had two a-balls?

We have the whole sample space on our tree diagram.
So you can compute this w/ conditional probability.

$$P(U=2|A=1) = \frac{P((U=2) \cap (A=1))}{P(A=1)} \quad \text{by definition of conditional probability.}$$

This is one of the most important tricks in probability.

What we are given are the conditional probabilities coming down the tree.



So we rewrite the numerator:

3/6/98.3

$$\begin{aligned} P((U=2) \cap (A=1)) &= P((A=1) \cap (U=2)) \\ &= P(A=1 | U=2) P(U=2) \\ &= \frac{P(A=1 | U=2) P(U=2)}{P(A=1)} \end{aligned}$$

Note that: $P(A=1 | U=2) = 1$

Because this is certain

$$P(U=2) = \frac{1}{3}$$

$$P(A=1) = \frac{|\{\omega_1, \omega_2\}|}{|\Omega|} = \frac{2}{4} = \frac{1}{2}$$

$$\begin{aligned} P(U=2 | A=1) &= \frac{1 \cdot \frac{1}{3}}{\frac{1}{2}} \\ &= \frac{2}{3} \end{aligned}$$

Now you think you all know that.
Therefore I pose a harder problem.
Now we do this all over with feeling.

↑ do you know where this expression comes from?
It comes from a New Yorker cartoon of the '40's.
Which you have never seen, which portrays the scene of a
film being filmed in Hollywood. A love scene - a woman +
a man on a sofa. The camera around, the director.
They just finished shooting the love scene and the
director says to the two actors:
"Now let's do again all over with, with feeling."

This is a tough problem.

3/6/98.4

Example - Bernoulli Process (the first run, of course)

Let A = event that a run of h ones occurs before a run of t zeros.

Let's compute $P(A)$.

The way it's done in the book is lousy.

For the sake of this argument, say $h=4$, $t=3$

The method will be completely apparent from this example.

Consider $P(A | (X_1=0) \cap (X_2=1))$

This is the conditional probability that A occurs (namely, a run of 4 heads occurs before a run of 3 tails) given that the 1st toss is a tail and the 2nd toss is a head.

I claim that this $(X_1=0)$ is irrelevant.

Namely that:

$$P(A | (X_1=0) \cap (X_2=1)) = P(A | X_2=1)$$

The probability of event A , given that $(X_2=1)$ is independent of $(X_1=0)$.

Let's write this in a pedantic way:

$$P(A | (X_1=0) \cap (X_2=1)) = P((A | X_1=0) | X_2=1) \quad \text{I'm just rubbing it in.}$$

$$= P_{(X_2=1)}(A | X_1=0)$$

we write the subscript here to stress the fact that conditional probability is a probability.

If you have tosses 0, 1, then a number of tosses and then the event A occurs:

→ 0 1 ... A

(Note: if instead we considered the case $X_1=1$, then for event A , $(X_1=1)$ would be relevant.)

It doesn't influence the event A ever occurring. It's shit.

$$\text{So, } P_{(X_2=1)}(A | X_1=0) = P_{(X_2=1)}(A)$$

The occurrence of $(X_2=1)$ makes $(X_1=0)$ irrelevant.

If $X_1=0$, it can not play any role. It's independent.

This 0 can not play any role in the event A occurring.

The event that the 1st toss is 0 is independent of the event A occurring.

If you don't see that, I can't explain it any more.


$$= P_{(X_2=1)}(A)$$

$$P(A | (X_1=0) \cap (X_2=1)) = P(A | X_2=1)$$

As claimed.

Probability Tree

This is a partition of the sample space.



$$P(A) = P(\underbrace{1111...}_{\text{event A}}) + P(\underbrace{1110...}_{\text{event A}}) + P(110...) + P(10...) + P(01...) + P(001...)$$

The tree displays it right away.

Now you say, what's $P(10...)$?

This is the probability that the event A occurs, given that $(X_1=1)$ and $(X_2=0)$.

$$P(A | X_1=1, X_2=0) = P(A | X_2=0) \leftarrow \text{for exactly the same argument at end of previous page. The leading 1 has no influence on the occurrence of event A.}$$

sometimes we write commas instead of Ω .

$$= P(A | X_1=0) \leftarrow \text{since what happened before is independent, probability of A occurring after 2nd toss is 0 is the same as the probability of A occurring given the 1st toss is 0.}$$

Similarly, $P(A | X_1=1, X_2=1, X_3=0) = P(A | X_3=0)$

It's only X_3 that matters.

X_1 and X_2 can not make A occur at all.

The event A is independent of $(X_1=1)$ and $(X_2=1)$; if you have $(X_3=0)$

Similarly, $P(A|X_1=1, X_2=1, X_3=1, X_4=0) = P(A|X_1=0)$

Exactly the same argument holds.

None of X_1, X_2, X_3 can influence event A .

Events $(X_1=1), (X_2=1), (X_3=1)$ are independent of A , given that $(X_4=0)$

Same idea gives:

$$P(A|X_1=0, X_2=1) = P(A|X_2=1) = P(A|X_1=1)$$

The event $(X_1=0)$ is independent of A , given that $(X_2=1)$

So you can cross out $X_1=0$ and say probability of A , given that 1 occurs after the second toss.

But this is the same as probability of A , given that 1 occurs after the 1st toss.

Similarly,

$$P(A|X_1=0, X_2=0, X_3=1) = P(A|X_3=1)$$

$(X_1=0)$ and $(X_2=0)$ are independent of A , given that $(X_3=1)$

So you might as well start at the 3rd toss w/ a 1 and say "What's the probability of A occurring after the 3rd toss, given that the 3rd toss is a 1?"

But $P(A|X_3=1)$ does not know that two tosses have already gone off.
probability that A occurs after the 3rd toss, given that the 3rd toss is 1.

So you relabel the tosses, as if the 3rd was the 1st.

$$\begin{array}{l} \text{related} \\ \text{tosses} \end{array} \left\{ \begin{array}{l} P(A|X_3=1) \Rightarrow \begin{array}{ccccccc} ? & ? & 1 & \dots & A \\ 1^{\text{st}} & 2^{\text{nd}} & 3^{\text{rd}} & & \end{array} \\ P(A|X_1=1) \Rightarrow \begin{array}{ccccccc} 1 & \dots & A \\ 1^{\text{st}} & & \end{array} \end{array} \right.$$

Now we are done.

$$P(1111\dots) = P((X_1=1) \cap (X_2=1) \cap (X_3=1) \cap (X_4=1))$$

$$= p * p * p * p$$

$$\boxed{P(1111\dots) = p^4}$$

$$P(1110...) = P(A \cap (X_1=1) \cap (X_2=1) \cap (X_3=1) \cap (X_4=0))$$

from definition of conditional probability:

$$= P(A | X_1=1, X_2=1, X_3=1, X_4=0) P((X_1=1) \cap (X_2=1) \cap (X_3=1) \cap (X_4=0))$$

we've already shown that
this is:

$$P(A | X_1=0)$$

expand using law of successive conditioning

$$= P(A | X_1=0) P(X_1=1) P(X_2=1 | X_1=1) P(X_3=1 | X_1=1, X_2=1) P(X_4=0 | X_1=1, X_2=1, X_3=1)$$

Note that the coin tosses are independent

$$P(X_i = 0/1 \mid \text{any other tosses}) = P(X_i = 0/1)$$

$$= P(A | X_1=0) P(X_1=1) P(X_2=1) P(X_3=1) P(X_4=0)$$

$$P(1110...) = P(A | X_1=0) p^3 q$$

$$P(110...) = P(A \cap (X_1=1) \cap (X_2=1) \cap (X_3=0))$$

$$= P(A | X_1=1, X_2=1, X_3=0) P((X_1=1) \cap (X_2=1) \cap (X_3=0))$$

$$P(110...) = P(A | X_1=0) p^2 q$$

$$P(10...) = P(A \cap (X_1=1) \cap (X_2=0))$$

$$= P(A | X_1=1, X_2=0) P((X_1=1) \cap (X_2=0))$$

$$P(10...) = P(A | X_1=0) p q$$

$$P(01...) = \cancel{P(A \cap (X_1=0) \cap (X_2=1))}$$

$$P(A \cap (X_1=0) \cap (X_2=1))$$

$$= P(A | X_1=0, X_2=1) P((X_1=0) \cap (X_2=1))$$

$$P(01...) = P(A | X_1=1) q p$$

3/6/98.8

$$\begin{aligned}
 P(001\dots) &= P(A \cap (X_1=0) \cap (X_2=0) \cap (X_3=1)) \\
 &= P(A | X_1=0, X_2=0, X_3=1) P((X_1=0) \cap (X_2=0) \cap (X_3=1))
 \end{aligned}$$

$$P(001\dots) = P(A | X_1=1) q^2 p$$

Q/A: You don't include $P(000\dots)$, as this means a run of 3 tails occurs before any heads - thus the event A (a run of 4 heads before a run of 3 tails) can not happen. $P(000\dots) = 0$

Therefore:

$$P(A) = P(1111\dots) + P(1110\dots) + P(110\dots) + P(10\dots) + P(01\dots) + P(001\dots)$$

$$P(A) = p^4 + P(A|X_1=0)p^3q + P(A|X_1=0)p^2q + P(A|X_1=0)pq + P(A|X_1=1)qp + P(A|X_1=1)q^2p$$

And now we are behind the 8-ball.

Because, where are we going to get probabilities $P(A|X_1=0)$ and $P(A|X_1=1)$?

Where did I go wrong? (On purpose, I must say)

(consider $P(A|X_1=1)$)

$$P(A|X_1=1) = \frac{P(A \cap (X_1=1))}{P(X_1=1)} \quad \text{counts only events under } (X_1=1), \text{ that is}$$

the entire left subtree on [3/6/98.5], normalized by $P(X_1=1)$.
That's what conditional probability is all about.

$$P(A|X_1=1) = \frac{P(1111\dots) + P(1110\dots) + P(110\dots) + P(10\dots)}{P(X_1=1)}$$

We've already computed the terms in the numerator.
For the denominator, we know that:

$$P(X_1=1) = p$$



$$P(A|X_1=1) = p^3 + P(A|X_1=0)p^2q + P(A|X_1=0)pq + P(A|X_1=0)q$$



rearranging and reordering:

3/6/98.9

$$= (q + pq + p^2q) P(A|X_i=0) + p^3$$

$$= q(1 + p + p^2) P(A|X_i=0) + p^3$$

$$= q \left(\frac{1-p^3}{1-p} \right) P(A|X_i=0) + p^3$$

$$1-p=q$$

$$P(A|X_i=1) = (1-p^3) P(A|X_i=0) + p^3$$

Why is this good?

Because this expresses $P(A|X_i=1)$ in terms of $P(A|X_i=0)$.

We do the same on the right subtree and we get an expression where $P(A|X_i=0)$ is in terms of $P(A|X_i=1)$.

We end up w/ 2 linear equations with 2 unknowns.

We can solve this.

Now consider $P(A|X_i=0)$

Now we consider the right subtree, counting events under $(X_i=0)$, that is the entire right subtree on [3/6/98.5], normalized by $P(X_i=0)$.

$$P(A|X_i=0) = \frac{P(A \cap (X_i=0))}{P(X_i=0)}$$

$$P(A|X_i=0) = \frac{P(01\dots) + P(001\dots)}{P(X_i=0)} \leftarrow \text{computed in previous pages}$$

$$= P(A|X_i=1)p + P(A|X_i=1)qp$$

$$= p(1+q) P(A|X_i=1)$$

$$= p \left(\frac{1-q^2}{1-q} \right) P(A|X_i=1)$$

$$1-q=p$$

$$P(A|X_i=0) = (1-q^2) P(A|X_i=1)$$


Now, We can solve the two simultaneous linear equations and obtain:

$$P(A|X_i=1)$$

$$P(A|X_i=0)$$

3/6/98.10

Consider again our probability tree of [3/6/98.5] - only now just the first branching.



The Law of Alternatives gives:

$$P(A) = \underbrace{P(A|X_i=1)} P(X_i=1) + \underbrace{P(A|X_i=0)} P(X_i=0)$$

we just showed
how to compute
these.

$$P(A) = P(A|X_i=1) p + P(A|X_i=0) q$$

If you can find a better way, I'll give you 100 points.
The answer is given in the book.


In all these probability problems, you have to get a recursion for the probability that you are trying to solve.
In essence, what we got is a recurrence difference equation for the probability that we solved.

You have 2 problems on this, for next week, which are fairly tough.


By this method we can solve any problem, whatsoever, like:
What's the probability that the pattern 011100 occurs before pattern 100011 (the complement)?

Just draw the tree and do the possibilities.
It comes out.

Summarizing what we've observed



We have a sample space comprised of a number of events (B_i).



A

we have an event A here.

At the end of the branches, we have the sample points,

Now you can compute the probability of A by conditioning on the probabilities of the B_i 's (the Law of Alternatives):

$$P(A) = \sum_i P(A|B_i)P(B_i)$$

However, these B_i are used in another sense, as well.

For example - you are given several B_i 's, one of which is the good one.

You perform experiment A .

You find that $A=32$.

On the basis of these data, what's the probability that the first B_i (i.e., B_1) is the good one?

$$P(B_1|A=32) = ?$$

where A = random variable

B_i = event that B_i is good

$$P(B_1|A=32) = \frac{P(B_1 \cap (A=32))}{P(A=32)}$$

definition of conditional probability

Bayes' Law is sometimes written like this, w/ the denominator expanded by the Law of Alternatives:

$$P(B_1|A=32) = \frac{P(A=32|B_1)P(B_1)}{P(A=32)} = \frac{P(A=32|B_1)P(B_1)}{\sum_i P(A=32|B_i)P(B_i)}$$

This is called Bayes' Law.

You decide among several hypotheses ($B_1, \dots, B_n$) on the basis of performing an experiment. The result of the experiment determines the probability of each of the possible hypotheses.

We are going to see several more examples of this, because it's very important.

Conditional Probability and Bayes' Law (cont'd)

The Law of Alternatives:

$$P(A) = \sum_{B \in \pi} P(A|B) P(B)$$

where π is a partition of the sample space Ω .
I.e., π is a family of events such that:

(1) for $B \in \pi$, $P(B) > 0$

(2) if $B, B' \in \pi$, $B \neq B'$, then $B \cap B' = \emptyset$

(3) $P(\bigcup_{B \in \pi} B) = 1$

probability of union of all the blocks of $\pi = 1$.

The Law of Alternatives, along with the Law of Successive Conditioning, is the main means for computing probabilities.

We are going to go over some more examples. There are so many angles to it.

There are so many ways to do things.

You never finish learning.

When you think you've learned all about it, you discover some new twist, where by you apply things in a new way.

Here's a new twist. Suppose we compute the conditional probability of A , given event C .

$$P(A|C) = P_C(A)$$

Now, we've said many times that $P(A|C)$ is itself a probability.

Namely, defining a conditional probability is taking the same event and renormalizing the given probability — replacing it with another probability $P_C(A)$, which has all the same properties of all probabilities.

Therefore, the Law of Alternatives for the conditional probability P_C reads as follows:

$$P_C(A) = \sum_{B \in \pi} P_C(A|B) P_C(B)$$

The Law of Alternatives is valid for every probability.
Therefore, it is valid for the probability P_C .

Now let's unscramble this and write it in the "/" (given) notation.

3/9/98.2

Let's see what it looks like.

Sounds like child's play. Watch.

$$P(A|C) = \sum_{B \in \Pi} P((A|B)|C) P(B|C)$$

From the Law of Successive Probabilities, we have [2/27/98.7]:

$$P((A|B)|C) = P(A|B \cap C)$$

$$P(A|C) = \sum_{B \in \Pi} P(A|B \cap C) P(B|C)$$

This is the Law of Alternatives, for conditioning relative to C .

It's the same Law, but when applied, you have to be very careful about the conditioning event.

- Let's do the same problem I did last time w/ probability trees [3/6/98.4-10], but this time, w/o probability trees.

Example - Bernoulli Process

Let A = event that the first run of h ones precedes the first run of t zeros.

Let's just fool around a little bit.

What is $P(A|(X_1=0) \cap (X_2=0) \cap \dots \cap (X_t=0))$?

$$P(A|(X_1=0) \cap (X_2=0) \cap \dots \cap (X_t=0)) = 0$$

Because if the first t tosses are 0, then the first run of t zeros occurs before any ones. A run of t zeros occurs certainly before a run of h ones.

Therefore, this probability is 0.

If you don't see it, I can't explain it.

Let's jazz this up a little bit.

What is $P(A|(X_1=0) \cap (X_2=1))$?

This is the probability of the event A taken on all sequences of zeros and ones that start with 01.

You just restrict to that sample space of all sequences that start w/ 01.



I claim that:

$$P(A | (X_1=0) \cap (X_2=1)) = P(A | (X_1=1))$$

↑

because the 1st zero can not possibly influence the occurrence of A - in view of the fact $(X_2=1)$, so long as $h, t \geq 2$.

This again - if you don't see it, I can't explain it.

$P(A | (X_1=0) \cap (X_2=1))$ is independent of $(X_1=0)$, because the event is between, i.e., $(X_2=1)$ breaks any sequence that could satisfy A.

It's a crucial fact.

This is a very simple application of the Law of Alternatives.

Let's take as a partition $\pi = \{ \underbrace{(X_1=1)}_{B_1}, \underbrace{(X_1=0)}_{B_2} \}$

$$P(A) = \sum_{B \in \pi} P(A|B) P(B)$$

←

$B_1 + B_2$ are disjoint blocks
Their union is the entire sample space
 $P(\cup_{B \in \pi} B) = 1$.
Therefore, the Law of Alternatives applies.

$$= P(A | (X_1=1)) \underbrace{P(X_1=1)}_p + P(A | (X_1=0)) \underbrace{P(X_1=0)}_q$$

$$(***) \quad P(A) = P(A | X_1=1) p + P(A | X_1=0) q$$

Let W^0 = waiting time for the first zero

To solve $P(A)$, I need to compute $P(A | X_1=1)$ and $P(A | X_1=0)$.

First, we compute $P_{(X_1=1)}(A)$, by conditioning on the following partition:

$$\pi = \{ (W^0 \leq h), (W^0 > h) \}$$

- 1) These two events are disjoint (i.e., blocks)
- 2) Their union is the whole sample space.

Therefore, we can apply the Law of Alternatives

3/9/98.4

We apply the Law of Alternatives, but w/ a kick.

We apply it for the probability $P_{(X_1=1)}$:

$$P_{(X_1=1)}(A) = \sum_{B \in \Pi} P_{(X_1=1)}(A/B) P_{(X_1=1)}(B)$$

$$= P_{(X_1=1)}(A/W^0 \leq h) P_{(X_1=1)}(W^0 \leq h) + P_{(X_1=1)}(A/W^0 > h) P_{(X_1=1)}(W^0 > h)$$

Let's try to simplify this.
First consider:

$$P_{(X_1=1)}(A/W^0 > h)$$

This is the probability of event A, given that the waiting time for the first zero is strictly greater than h.

That means at least the first h tosses are ones.

Thus event A happens absolutely.

The first h tosses are one.

Therefore:

$$P_{(X_1=1)}(A/W^0 > h) = 1$$


Next, we consider $P_{(X_1=1)}(W^0 > h)$:

$(X_1=1)$ means the first toss is one.

$(W^0 > h)$ means that at least the first h tosses are one.

One of them, the first has already happened. So the next h-1 tosses have to be one.

① by probability tree:



$$P_{(X_1=1)}(W^0 > h) = p^{h-1}$$

{ ②/A: Note that getting a run of more than h heads includes getting a run of h heads. }

② Alternatively - by the Law of Alternatives:

$$P_{(X_1=1)}(W^0 > h) = P_{(X_1=1)}(W^0 > h \mid \text{tosses } 2 \dots h \text{ all } 1) P_{(X_1=1)}(\text{tosses } 2 \dots h \text{ all } 1)$$

$(X_1=1)$ and $(\text{tosses } 2 \dots h \text{ all } 1)$ mean the first h tosses are 1. Thus $(W^0 > h)$ is a certainty: 1

+ $P_{(X_1=1)}(W^0 > h \mid (\text{tosses } 2 \dots h \text{ all } 1)^c) P_{(X_1=1)}((\text{tosses } 2 \dots h \text{ all } 1)^c)$
 $(\text{tosses } 2 \dots h \text{ all } 1)^c$ means one of the first h tosses is 0. Thus $(W^0 > h)$ can never happen: 0

Now let's evaluate the first term.

First, we have:

$$\begin{aligned} P_{(X_i=1)}(W^0 \leq h) &= 1 - P_{(X_i=1)}(W^0 > h) \\ &= 1 - p^{h-1} \end{aligned}$$

Finally, we are left with the toughest part:


$$P_{(X_i=1)}(A/W^0 \leq h)$$

Let's think of what this is.

This is the probability $P_{(X_i=1)}$ that some zeros occur ($W^0 \leq h$) before you have a run of h ones.

So the zeros associated with the event ($W^0 \leq h$) break the run.

Therefore, any run that has occurred before is useless.



The only thing that matters is that given event ($W^0 \leq h$), there are some zeros that break any sequence, therefore you might as well assume you start with a zero.

$$P_{(X_i=1)}(A/W^0 \leq h) = P(A/X_i=0)$$

So we therefore obtain, using these terms to rewrite the above equation:

$$P_{(X_i=1)}(A) \rightarrow P(A/X_i=1) = P(A/X_i=0)(1-p^{h-1}) + 1 \cdot p^{h-1}$$

(*)

$$P(A/X_i=1) = P(A/X_i=0)(1-p^{h-1}) + p^{h-1}$$

This is a linear equation linking the conditional probability of the event given ($X_i=1$) with the conditional probability of the event, given ($X_i=0$). Notice we get both conditional probabilities.

We need another linear equation linking $P(A/X_i=0)$ to $P(A/X_i=1)$. Let's do that.

Let W' = waiting time for first one

Consider the partition:

$$\pi = \{(W' > t), (W' \leq t)\}$$

These two events form a partition

They are disjoint.

And their union is the whole sample space.

We therefore condition, by the Law of Alternatives:

$$P_{(X_1=0)}(A) = \sum_{B \in \pi} P_{(X_1=0)}(A|B) P_{(X_1=0)}(B)$$

$$= P_{(X_1=0)}(A|W' > t) P_{(X_1=0)}(W' > t) + P_{(X_1=0)}(A|W' \leq t) P_{(X_1=0)}(W' \leq t)$$

Let's simplify this.

Consider the first term:

$$P_{(X_1=0)}(A|W' > t)$$

The waiting time for the first one is greater than t .

This means that at least the first t tosses are zeros.

Therefore you've lost.

You've lost completely, if this happens.

Event A (event that first run of h ones precedes first run of t zeros) can never happen.


Since, if $(W' > t)$, A can not happen

$$P_{(X_1=0)}(A|W' > t) = 0$$

Now let's do the second term:

$$P_{(X_1=0)}(A|W' \leq t)$$

$(W' \leq t)$ means that the first 1 occurs before t zeros occur.



therefore, it's as if you were starting with a one

$$P_{(X_1=0)}(A|W' \leq t) = P(A|X_1 = 1)$$



3/4/98.7

Finally:

$$P_{(X_1=0)}(W' \leq t) = 1 - P_{(X_1=0)}(W' > t) \\ = 1 - q^{t-1}$$

We therefore obtain, using these terms to rewrite the above equation:

$$P_{(X_1=0)}(A) = 0 + P(A|X_1=1)(1-q^{t-1})$$

$$P(A|X_1=0) = P(A|X_1=1)(1-q^{t-1})$$

Solve (*) and (**) for $P(A|X_1=0)$ and $P(A|X_1=1)$
Then substitute into (***) to obtain the desired $P(A)$.

$$P(A) = \frac{p^{h-1}(1-q^t)}{p^{h-1} + q^{t-1} - p^{h-1}q^{t-1}}$$

I confess that I do not see this result intuitively.
It's not exactly intuitively clear.


One remark

This reasoning is exactly the same as the reasoning we went through before [3/6/98.4-10].
These are 2 ways of saying the same thing.
I see it more clearly by drawing trees.
I would never have guessed this conditioning if I hadn't drawn the tree.

You have 3 problems in the problem set that are hard, where you really have to draw the tree.
I suggest you do the last problems in the problem set first, because they are very easy.

BTW - You might compare this treatment of the problem w/ that in the book.
Notice that the book has a mistake.
Although the result works out.

$P_{(X_1=0)}(W' > t)$
a) by probability tree:



b) by Law of Alternatives:

$$P_{(X_1=0)}(W' > t) = P_{(X_1=0)}(W' > t | \text{tosses } 2 \dots t \text{ all } 0) P_{(X_1=0)}(\text{tosses } 2 \dots t \text{ all } 0) \\ + P_{(X_1=0)}(W' > t | \text{tosses } 2 \dots t \text{ not all } 0) P_{(X_1=0)}(\text{tosses } 2 \dots t \text{ not all } 0) \\ \text{1 - since first 1 after } t \text{ tosses } 0 \dots 0 \dots \\ 0 - \text{since at least one of tosses } 2 \dots t \text{ is } 1$$

• Example - this typical Bayes' Law

This is the Problem of Scientific Induction

You have:

Population of n members,
 a = good members

n = known
 a = unknown } We know that the population has n members, but we don't know how many good guys there are.

We are looking at a sample of size k
w/o replacement.

Each sample guy is killed.

On the basis of this sample, we have the random variable A :

Let random variable A = number of a -guys in the sample

The Problem of Scientific Induction, or Bayes' Law, is this.
 On the basis of the number of a -guys in the sample of size k ,
 what's the most probable inference of the number of a -guys in the
population?

This is a problem of universal occurrence.

I state it in these Mickey Mouse terms.

But you can state it in terms that are very highfalutin.

Life and death situations where people will pay millions of dollars.

I leave it to you to be applied mathematicians, which is to recognize when you walk into a problem that is just this.

I can't teach you that.

Let's set it up.

We want to set up a sample space where all possible distributions of numbers of a -guys (good guys) and bad guys exist.

And then we sample one of the population

We set up $n+1$ urns.

The urn ($U=j$) contains $\begin{cases} j & a\text{-guys (good guys)} \\ n-j & \text{bad guys} \end{cases}$

We set up the experiment as follows:

We first pick one of the urns.


Then from that urn, we sample k members.

Then we look for a -guys.

That's our sample space.



Sample space Ω is:



What's missing?

In order for this to be a probability space, we have to assign probabilities.

The sample points are the branches of the tree. {NB: branches, not just the leaves}
 You have to assign probabilities in such a way that the probabilities added up over all the branches add up to 1.

Let's take urn j.

Without knowing the sample space, we can compute the conditional probability:

$$P(A=i | U=j) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}} \leftarrow \begin{cases} \text{From the } j \text{ good guys, pick } i \\ \text{From the } n-j \text{ bad guys, pick } k-i \end{cases}$$

$\leftarrow$ because you have a sample of size k

probability of event $(A=i)$,
 given that you've chosen
 the population w/ j a-guys
 (good guys).
 We've done this before
 [2/23/98.8-11]

That is, the hypergeometric distribution


Remember - there was something very funny happening here (The population is the urns. The sample is the balls).
 We obtained the same result whether the balls are distinguishable or indistinguishable.
 This is an important fact - physically.
 Namely, if you sample w/o replacement, it doesn't matter whether the balls are distinguishable or not.

If we knew the probability of choosing urn j $P(U=j)$, then we could determine the probability distribution of A on this entire sample space, immediately, from the Law of Alternatives.

$$P(A=i) = \sum_{j=0}^n P(A=i|U=j) P(U=j)$$

Unfortunately, we have no information as to which distribution of populations $P(U=j)$ is more popular than which.

So we have $p_j = P(U=j)$.



$$\text{where } p_0 + p_1 + \dots + p_n = 1$$

Q: If you have absolutely no way of knowing which urn is more probable than another, why can't you just say the probability of each of the $n+1$ urns is just $\frac{1}{n+1}$?

$$p_i = \frac{1}{n+1}$$

A: You're spoiling my act.

Someone may have slipped us some information:

You can't have more than 15 good guys

-or-

Good guys come in multiples of three.

But in the absolutely most general case, this is what we get.


Therefore, the probability distribution of the random variable A is the following:

$$P(A=i) = \sum_{j=0}^n P(A=i|U=j) P(U=j) \leftarrow \text{Very standard application of the Law of Alternatives}$$

$$= \sum_{j=0}^n \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}} p_j$$



Now the sample space is completely defined.
The sample point is the complete branch.



The edges, as we've said, are the conditional probabilities.
Since this edge ends at ω_{24} , we have:

$$P(\omega_{24}) = \underbrace{P(U=1)}_{p_i} P(A=i | U=1)$$

So now we can compute everything about this sample space.
For example, we can compute:

$$P(U=j | A=i)$$

This is the process of induction.

If you draw a sample of k guys and i of them are good (i.e., the event $(A=i)$), which $(U=j)$ is the most probable distribution?

$$P(U=j | A=i) = \frac{P((U=j) \cap (A=i))}{P(A=i)} \quad \text{definition of conditional probability}$$

$$(*) \quad P(U=j | A=i) = \frac{P(A=i | U=j) \overbrace{P(U=j)}^{p_j}}{P(A=i)} \quad \text{rewrite numerator}$$

Now some jargon. In expression (*):

$p_j = P(U=j)$ is called the prior

$\frac{P(A=i | U=j)}{P(A=i)}$ is called the likelihood

$P(U=j | A=i)$ is called the posterior

These are weird terms, but everybody uses them. They were invented by business people, so they are really weird.

Begin Digression

Some philosophical remarks.

This is something I should have told you before, but I missed it.

Point 1

Expectation of a random variable:

If X is an integer random variable,
then

$$E(X) = \sum_n n P(X=n)$$

sum over all
integers n

Then, on the basis of this definition, we verified [2/25/98, 8-11]:

$$E(X+Y) = E(X) + E(Y)$$

by an argument that is not at all trivial.

You should have asked, at this point, where does this come from?
Let's digress for a minute and motivate the definition.Suppose sample space Ω is finite.And $X =$ random variable.I want to define the average value of that random variable.What would be the most natural definition of average value for that random variable?

The most natural definition of average would be:

Intuitive
Point of View

$$E(X) = \sum_{\omega \in \Omega} \underbrace{X(\omega)}_{\substack{\text{value of the random variable } X \\ \text{at the sample point } \omega}} \underbrace{P(\omega)}_{\substack{\text{times} \\ \omega \text{ probability of}}}$$

sum over all sample
points

This is as natural a definition of the average of a random variable as one can have.
Intuitive.

Add up the values, each one weighted by its probability.

In fact, if you use this definition, then the main formula $[E(X+Y) = E(X) + E(Y)]$
is obvious. Completely obvious.

Because you have:

$$\begin{aligned} E(X+Y) &= \sum_{\omega \in \Omega} (X(\omega) + Y(\omega)) P(\omega) \\ &\quad \text{which is obviously:} \\ &= \sum_{\omega \in \Omega} X(\omega) P(\omega) + \sum_{\omega \in \Omega} Y(\omega) P(\omega) \\ &= E(X) + E(Y) \end{aligned}$$



3/11/98.2

From this point of view, the linearity of expectation is something completely trivial.
Of course, we assume that our sample space Ω is finite.

Now let's see that the intuitive definition reduces to the definition we gave before:

$$E(X) = \sum_{\omega \in \Omega} X(\omega) P(\omega)$$

$$E(X) = \sum_n n P(X=n)$$

$$\sum_{\omega \in \Omega} X(\omega) P(\omega) = \sum_n \sum_{\{\omega: X(\omega)=n\}} X(\omega) P(\omega)$$

split this sum into sums,
one for each value of
 $X(\omega)$

First, I sum over all ω for which $X(\omega)=n$.
Then, I sum over all n .

That gives me the sum over all Ω .

The whole sample space partitions according to the value
of the random variable.

Now we have the Principle of Who is Buried in Grant's Tomb.
What is:

$$\sum_{\{\omega: X(\omega)=n\}} X(\omega) P(\omega) \quad ?$$

Since $X(\omega)=n$ then, then you have n here.
If you don't see that, I leave.

$$= \sum_n \sum_{\{\omega: X(\omega)=n\}} n P(\omega)$$

$$= \sum_n n \sum_{\{\omega: X(\omega)=n\}} P(\omega)$$

$$= \sum_n n P(X=n)$$

{And this is precisely the
other formula.}

This is another of "Who's buried in
Grant's Tomb" arguments.

$\sum_{\{\omega: X(\omega)=n\}} P(\omega)$, sum of $P(\omega)$
where $X(\omega)=n$

is the probability of the event
that $(X=n)$, namely:

$P(X=n)$. Again, if you don't
see that, I quit.

$$\sum_{\omega \in \Omega} X(\omega) P(\omega) = \sum_n n P(X=n)$$

intuitive
definition

unintuitive
definition

Therefore, for a finite sample space, the intuitive definition is an equal to the unintuitive one.

Unfortunately, most sample spaces of this world are not finite.
And, therefore, the intuitive definition can not be used.
However, the unintuitive definition still works.

And so, in the knowledge that the intuitive and the unintuitive are equivalent in the simple case of the finite sample space, we introduce the unintuitive definition and we pawn that off as the expectation value.

That's the motivation that I didn't give you before.

This is one of the greatest arguments of mathematics - don't kid yourselves.
It took hundreds of years for this to be discovered.
Now we work it out.

It took a long time for people to realize this very simple fact.

So that's the computation of expectation.
Intuitively, when we are given:

$$\sum_{\omega \in \Omega} X(\omega) P(\omega)$$

we think of the sum of infinitesimal probabilities over $X(\omega) P(\omega)$.

That can be done by integrating the Lebesgue integral.

But we don't want to do that now.

Since we don't want to do the fancy integration, we use the unintuitive definition:

$$= \sum_n n P(X=n)$$


That's the first point.



Point 2

Another philosophical point:

Two descriptions of the classical statistics.

The classical statistics, which are the statistics of k balls into n boxes.The sample points are functions from a set B to a set U (or something akin to functions, like equivalence classes, when the balls are indistinguishable).

We said that there are two points of view, whereby you can visualize the stochastic processes.

1) Distribution point of view

where you imagine placing the balls into the boxes.

Mostly, you consider the occupation number

 $B = \text{balls}$ $U = \text{boxes}$ the other
is:2) Occupancy point of view $U = \text{population}$ $B = \text{sample of that population}$ The formulas are the same, but sometimes the logic becomes clear if you look at it one way, or the other.

Let's make this distinction more rigorous.

Distribution:We define the random variable capital theta Θ : $\Theta_j = \text{number of balls in the first } j \text{ boxes}$
(we assume the boxes are numbered)

$$= \underbrace{\Theta_1 + \Theta_2 + \dots + \Theta_j}_{\text{sum of occupation numbers for boxes 1 to } j}$$

sum of occupation numbers for boxes 1 to j The sequence of random variables $\Theta_1, \Theta_2, \dots, \Theta_n$ completely determines any phenomenon pertaining to the occupation numbers of the boxes.For example, if you want the occupation number Θ_j , you take:


$$\Theta_j = \Theta_j - \Theta_{j-1}$$

I am introducing this term to motivate the next discussion.

Occupancy: This point of view is more sophisticated.

Lets do first Ω_{MB}

From an occupancy point of view, that's the situation where we: [2/13/98.2]
 sample, put it in position 1,
 sample, put it in position 2,
 $\vdots$
 sample, put it in position k .



random variable $X_1 =$ first sample

For the occupancy point of view, the stochastic process is completely described by giving k independent random variables $X_1, X_2, \dots, X_k$ that are independent and identically distributed (i.i.d.) and with the probability that:

$$P(X_1 = j) = \frac{1}{n} \quad \left. \vphantom{P(X_1 = j)} \right\} \text{Because the probability of a ball coming from any box is } \frac{1}{n}.$$

The Maxwell-Boltzmann process can be completely described by saying I have a sequence of k i.i.d. random variables with this distribution.

The occupancy description carries more information than the distribution description. Because, for example, in the distribution description, you only count how many balls are in each box. You can not say whether ball 3 has gone into box 5. Θ_j doesn't provide this information.

Whereas the occupancy description provides the entire information. For example, $(X_3 = 5)$ is the event that the 3rd sample was picked from the 5th box.

Let's do the occupancy description for the other stochastic processes. Namely, for Fermi-Dirac and Bose-Einstein. As you see, this is purely philosophical.

Let's take Fermi-Dirac [2/17/98.4]

We'd like to give an occupancy description in terms of a sequence of random variables.

But that's impossible, because the balls are indistinguishable.

What do we do?

We take another sample space, which is Maxwell-Boltzmann conditioned on the boxes having 0 or 1 ball.

$$\Omega_{FD} = \Omega_{MB} \text{ conditioned on } \Theta_j = \begin{cases} 0 \\ 1 \end{cases} \text{ for all } j$$

Then divide by $k!$ and you get Fermi-Dirac.

Q: Why not use Ω_{BE} ← Bose-Einstein?

A: Because I want the balls to be distinguishable.

In Bose-Einstein, the balls are indistinguishable, so I can not talk about X_1 = first sample, X_2 = second sample, etc.

In both Fermi-Dirac and Bose-Einstein, what you are doing is:
distribution description: placing marks on boxes

— $\times$ $\times$... $\times$

— or —
occupancy description: taking a sample and forgetting in which order you took it in.
This is why we divide by $k!$

If you want the order of the sample — every hour on the hour, we sample one, then they are distinguishable because we can tell which one has been sampled which hour.

So, we are forced to use Maxwell-Boltzmann. We condition it this way and divide by $k!$ to get Fermi-Dirac.

$X_1, X_2, \dots, X_k$ = successive balls extracted
(the samples)

These are no longer independent

$$P(X_1 = j) = \frac{1}{n}$$

What is $P(X_2 = j)$?

Let's work it out.

$$P(X_2=j) = \sum_{B \in \pi} P(X_2=j | B) P(B)$$

Law of Alternatives.

For the blocks B , choose the disjoint $(X_1=i)$, $i=\{1, \dots, n\}$

$$= \sum_{i=1}^n \underbrace{P(X_2=j | X_1=i)}_{\text{given that the 1st sample chosen is } i} \underbrace{P(X_1=i)}_{\frac{1}{n}}$$

given that the 1st sample chosen is i ($X_1=i$), the 2nd sample can not be i , since this is Fermi-Dirac (Pauli exclusion). So instead of n , there are $n-1$ possible samples for the 2nd choice. Thus:

$$P(X_2=j | X_1=i) = \frac{1}{n-1}$$

Finally, note that $P(X_2=j | X_1=j) = 0$, as the 2nd choice can not choose the same as the 1st choice.

$$= \sum_{\substack{i=1 \\ i \neq j}}^n P(X_2=j | X_1=i) P(X_1=i)$$

$$= (n-1) \left(\frac{1}{n-1} \right) \left(\frac{1}{n} \right)$$

$$P(X_2=j) = \frac{1}{n}$$

$X_1, X_2, \dots, X_K$ are identically distributed (i.i.d.)

What about the joint distribution?

$$P((X_1=i) \cap (X_2=j)) = \underbrace{P(X_2=j | X_1=i)}_{\frac{1}{n-1}} \underbrace{P(X_1=i)}_{\frac{1}{n}}$$

as discussed above.



$$P((X_1=i) \cap (X_2=j)) = \left(\frac{1}{n-1}\right) \left(\frac{1}{n}\right) \neq \underbrace{P(X_1=i) P(X_2=j)}_{\left(\frac{1}{n}\right) \left(\frac{1}{n}\right)}$$

$$P((X_1=i) \cap (X_2=j)) \neq P(X_1=i) P(X_2=j)$$

Therefore X_1 and X_2 are not independent

$X_1, X_2, \dots, X_K$ are not independent

So for this Fermi-Dirac case, you have these identically distributed random variables, which are not independent.

This sequence of random variables, together w/ their joint distributions, completely determines the process.

In fact, instead of the joint distribution, you might think of describing the process by the conditional probability.

{ For Bose-Einstein, the occupancy description is a similar situation,
I leave it to you as an exercise. }

Remember for Disposition (Flags), we discussed a distribution interpretation, but we did not discuss an occupancy interpretation.
That's very easy to work out. I leave it to you.

3rd Last bit

3/11/98.9

Bayes' Law

A = event

π = partition of sample space Ω

The Law of Alternatives gives:

$$P(A) = \sum_{B \in \pi} P(A|B) P(B)$$


Bayes' Law is a jazzing up of the above.

Bayes' Law:

$$P(B|A) = \frac{P(A|B) P(B)}{P(A)} = \frac{P(A|B) P(B)}{\sum_{B \in \pi} P(A|B) P(B)}$$

We are beginning to see some applications of this and we're coming to this in a minute.

This is interpreted as (let me do this in terms of sets):



You have the partition π with blocks B_i .

Then you have the event A .

Then you want to estimate the probability of each block B_i of the partition on the basis of the intersection of A with that block.

In order to estimate $P(B|A)$, you need to know conditional probabilities $P(A|B)$.

$$P(B|A) = \frac{P(A|B) P(B)}{\sum_{B \in \pi} P(A|B) P(B)}$$

posterior
likelihood
prior

It is impossible to understand this w/o examples.

End: Digression

Example 1

Given a population of n members, a of which are good. $a = \text{unknown}$
 We are allowed to sample, w/o replacement, k numbers
 And we find $A = \text{number of good members in the sample.}$

$\uparrow$
 clearly A is a random variable


What is the best estimate of the number of good guys in the population,
 given the number of good guys in a sample of size k ?

A

That's the basic problem.

We want the solution to the problem: Best guess of a based on random variable A .

Set up a sample space Ω :



consider the various number of possibilities for the number of a 's in the population. We are sampling from one of these possibilities, k members. But we don't know from which possibility

How do we set up the sample space?

We need to add probabilities for each branch.

Let $U = \text{number of } a\text{'s in the population}$

$\uparrow$ another random variable.

$$P(A=i/U=j) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}}$$

we've computed this a few times.
 For example, see [2/9/98.9].
 From j , choose i good guys.
 From remaining $n-j$, choose $k-i$ bad guys.

the hypergeometric distribution

Watch it, because I am preparing a trap!

This seems like something elementary, but you should be on guard.

From the Law of Alternatives:

$$P(A=i) = \sum_{j=0}^n P(A=i|U=j) P(U=j)$$

By Bayes' Law:

$$P(U=j|A=i) = \frac{P(A=i|U=j) P(U=j)}{\sum_{j=0}^n P(A=i|U=j) P(U=j)}$$

These are different j .
Strictly speaking, this is correct, as in denominator, j refers to summation variable. Although it is slightly deceiving.

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i} P(U=j)}{\sum_{j=0}^n \binom{j}{i} \binom{n-j}{k-i} P(U=j)}$$

← { the $\binom{n}{k}$ in the denominator of the hypergeometric distributions cancel out. }

We can not determine this until we have the probability $P(U=j)$.
We have to choose a prior probability, namely:

$$P(U=j)$$

It comes naturally to say that if we know absolutely nothing, then we assume that all the possibilities have the same probability.

But we will see that that is not the most natural thing to assume.

Case 1 - Uniform prior (each of $\overbrace{0a's, 1a, 2a's, \dots, na's}^{n+1}$ is equally probable)

$$P(U=j) = \frac{1}{n+1}, \text{ for all } j$$

Then (*) simplifies as follows:

$$\begin{aligned} P(U=j|A=i) &= \frac{\binom{j}{i} \binom{n-j}{k-i} \frac{1}{n+1}}{\sum_{j=0}^n \binom{j}{i} \binom{n-j}{k-i} \frac{1}{n+1}} \\ &= \frac{\binom{j}{i} \binom{n-j}{k-i}}{\sum_{j=0}^n \binom{j}{i} \binom{n-j}{k-i}} \end{aligned}$$

3/11/98.12

We do not have a good answer, unless we are able to give a closed form for the denominator.

Simplify the denominator as follows:
We want a closed form for:

$$\sum_{j=0}^n \binom{j}{i} \binom{n-j}{k-i}$$


It's not that easy.

You can do it directly.

But we'll do the following side argument:

In Fermi-Dirac statistics, take sample space Ω_{FD} with $k+1$ balls,
 $n+1$ boxes.

Let random variable C_{i+1} = the position of the $i+1^{st}$ box checked.



$$P(C_{i+1} = j+1) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n+1}{k+1}}$$

There are i checks in the first j boxes.

$(C_{i+1} = j+1)$ means the $j+1^{st}$ box has a check, by definition.

The remaining $k-i$ checks go into the following $n-j$ boxes.

total number of possibilities.

$$P(C_{i+1} = j+1) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n+1}{k+1}}$$

3/11/98.13

Since $P(C_{i+1} = j+1)$ is a probability distribution, its sum over all j must equal 1.

$$\sum_{j=0}^n P(C_{i+1} = j+1) = 1 \quad \left. \vphantom{\sum_{j=0}^n} \right\} \text{by definition of probability distribution}$$

$$\sum_{j=0}^n \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n+1}{k+1}} = 1$$

Which gives the identity:

$$\sum_{j=0}^n \binom{j}{i} \binom{n-j}{k-i} = \binom{n+1}{k+1}$$

This completes the side argument.

Now we can reduce $P(U=j|A=i)$ to:

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\sum_{j=0}^n \binom{j}{i} \binom{n-j}{k-i}} \quad \leftarrow \text{substitute the above identity.}$$

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n+1}{k+1}}$$



Let's verify this binomial identity with a probabilistic model.

Model - Fermi-Dirac with $k+1$ checks, $n+1$ boxes


$k+1$ checks, at most one check per box

$\frac{\checkmark}{1} \quad \frac{\checkmark}{2} \quad \frac{\checkmark}{3} \quad \frac{\checkmark}{4} \quad \dots \quad \frac{\checkmark}{n+1}$ boxes are numbered by position
positions

Let random variable $C_1 =$ position of 1st check

Ex: In the above, $C_1 = 3$

Let's consider the probability distribution of the event $(C_1 = j+1)$, i.e., that the first check occurs at position $j+1$:



$$P(C_1 = j+1) = \frac{\binom{j}{0} \binom{n-j}{k}}{\binom{n+1}{k+1}}$$

$$P(C_1 = j+1) = \frac{\binom{n-j}{k}}{\binom{n+1}{k+1}}$$

Since C_1 is a random variable, the sum over all j of its probability distribution must equal 1. We have:

$$\sum_j P(C_1 = j+1) = 1 \quad \left. \vphantom{\sum_j} \right\} \text{by definition}$$

$$\sum_j \frac{\binom{n-j}{k}}{\binom{n+1}{k+1}} = 1$$

$$\sum_j \binom{n-j}{k} = \binom{n+1}{k+1}$$

$$\binom{n}{k} + \binom{n-1}{k} + \dots = \binom{n+1}{k+1}$$

which verifies the case for $i=0$, as desired.

Case $i=1$ We want to show that the binomial identity holds for $i=1$:


$$\sum_j \binom{j}{1} \binom{n-j}{k-1} \stackrel{?}{=} \binom{n+1}{k+1}$$

$$\binom{j}{1} = \frac{j!}{1!(j-1)!} = j = \text{number of ways to pick one element from } j$$

$$\sum_j j \binom{n-j}{k-1} \stackrel{?}{=} \binom{n+1}{k+1}$$

$$\binom{n-1}{k-1} + 2 \binom{n-2}{k-1} + 3 \binom{n-3}{k-1} + \dots \stackrel{?}{=} \binom{n+1}{k+1}$$

Another interesting binomial identity.

We use the same Fermi-Dirac probabilistic model as before ($k+1$ checks, $n+1$ boxes), but now we consider the random variable: $C_2 =$ position of the 2nd checkThe probability distribution for C_2 is:

$$P(C_2 = j+1) = \frac{\binom{j}{1} \binom{n-j}{k-1}}{\binom{n+1}{k+1}}$$

$$P(C_2 = j+1) = \frac{j \binom{n-j}{k-1}}{\binom{n+1}{k+1}}$$

And since the sum over all possible values of a probability distribution equals one, we have:

$$\sum_j C(z=j+1) = 1$$

$$\sum_j \frac{j \binom{n-j}{k-1}}{\binom{n+1}{k+1}} = 1$$

$$\sum_j j \binom{n-j}{k-1} = \binom{n+1}{k+1}$$

$$\binom{n-1}{k-1} + 2 \binom{n-2}{k-1} + 3 \binom{n-3}{k-1} + \dots = \binom{n+1}{k+1}$$

which verifies the case for $k=1$, as desired.

- You can get additional interesting binomial identity proofs, in exactly the same way, for different k .
- Almost every binomial identity has a probabilistic model proof. Experience, for example taking 18.313, will teach you how to prove those.
- Note: instead of binomial identities, we can get multiset identities in the same fashion, but instead of Fermi-Dirac statistics, we use Bose-Einstein statistics for the probabilistic model proof.

identity	probabilistic model proof
binomial	Fermi-Dirac statistics
multiset	Bose-Einstein statistics



(i.e., w/ feeling. We kill the sampled members.)

3/13/98.5

Bayes' Theorem for Sampling w/o Replacement

population: n members (not random)
sample: k

Let's define the random variables:

U = number of good guys (in the population)

A = number of good guys in the sample

Bayes' Law, compliments of Reverend Sir Thomas Bayes, tells us:

$$P(U=j|A=i) = \frac{P(A=i|U=j) P(U=j)}{P(A=i)}$$

The meaning of this is that we perform experiment A. Given that we observe ($A=i$), $P(U=j|A=i)$ gives the probability that the population has j good guys.
We are estimating the number of good guys in the population, based on the observed number of good guys in the sample.

$$= \frac{P(A=i|U=j) P(U=j)}{\sum_j P(A=i|U=j) P(U=j)} \leftarrow \begin{array}{l} \text{rewrite } P(A=i) \text{ by} \\ \text{Law of Alternatives} \end{array}$$

As we discussed last time [3/11/98.10], $P(A=i|U=j)$ is the hypergeometric distribution:

$$P(A=i|U=j) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}}$$

$$(*) \quad P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i} \underbrace{P(U=j)}_{\text{prior probability}}}{\sum_j \binom{j}{i} \binom{n-j}{k-i} P(U=j)} \leftarrow \text{denominator}$$

Typically, we are provided with the conditional probabilities $P(A=i|U=j)$ and $\underbrace{P(A=i)}_{\text{prior probability}}$. That is, we can readily compute the likelihood probability.

In the absence of any further information, what do we assign to the prior?



Case 1 - Uniform Prior

This summarizes what we discussed last time [3/11/98.10-13].

Assume $P(U=j) = \frac{1}{n+1}$, for all j

Then equation (*) simplifies to:

$$P(U=j|A=L) = \frac{\binom{j}{L} \binom{n-j}{k-L} \frac{1}{n+1}}{\sum_j \binom{j}{L} \binom{n-j}{k-L} \frac{1}{n+1}}$$

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n+1}{k+1}} \quad \leftarrow \text{from binomial identity proved probabilistically [3/11/98.12-13]}$$

Now we ask "Is the assumption of a uniform prior sensible?"
Is there another prior that should be considered?

↑ Surprisingly, the answer is yes.
The conjugate prior

Case 2 - Conjugate Prior

members of the population are Bernoulli random variables ← (each equal to 0 or 1)

$$\underbrace{X_1 \quad X_2 \quad X_3 \quad \dots \quad X_n}_{n \text{ members}} \quad \text{population}$$

As we recall from [2/20/98.8]:

$$P(X_i=1) = p$$

$$P(X_i=0) = q = 1-p$$

random variable U = number of ^{it's} good guys in the sample
 $= X_1 + X_2 + \dots + X_n$

We've shown the probability distribution of U is the binomial distribution:

$$P(U=j) = \binom{n}{j} p^j q^{n-j}$$

this is the conjugate prior



With this prior $P(U=j)$, equation (*) gives the posterior probability as:

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i} P(U=j)}{\sum_j \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}} P(U=j)}$$

$$= \frac{\frac{\binom{j}{i} \binom{n-j}{k-i} \binom{n}{j} p^j q^{n-j}}{\binom{n}{k}}}{\sum_j \frac{\binom{j}{i} \binom{n-j}{k-i} \binom{n}{j} p^j q^{n-j}}{\binom{n}{k}}} \quad \text{A mess!}$$

the denominator works out as follows.
Trust me, I worked it out.
You, too, are welcome to.

$$\binom{k}{i} p^i q^{k-i}$$

$$= \frac{\frac{\binom{j}{i} \binom{n-j}{k-i} \binom{n}{j} p^j q^{n-j}}{\binom{n}{k}}}{\binom{k}{i} p^i q^{k-i}}$$


$$P(U=j|A=i) = \binom{n-k}{j-i} p^{j-i} q^{n-k-j+i}$$

Alternatively

But I didn't have to calculate this mess.
I can reason as follows:

Since all members of the population are random,
choosing the first k random members as the sample is the same as
choosing randomly k members.


So, for a sample of size k, we might as well choose
 $X_1, X_2, \dots, X_k$.



Therefore, the random variable A is defined as:

$$A = \text{number of } \overbrace{\text{good guys}}^{1's} \text{ in the sample} \\ = X_1 + X_2 + \dots + X_k \quad \leftarrow \text{first } k \text{ random variables}$$

We next consider the conditional probability $P(U=j|A=i)$ directly, from the Bernoulli Process point of view.



$P(U=j|A=i)$ is just the binomial distribution:

$$P(U=j|A=i) = \binom{n-k}{j-i} p^{j-i} q^{n-k-j+1}$$

given event $(A=i)$, number of ways of choosing remaining $j-i$ good guys from remaining $n-k$ random variables,
 probability of good guys p
 probability of bad guys q

As desired, there is agreement in the posterior probability $P(U=j|A=i)$, whether calculated analytically (i.e., evaluation of the mess on [3/13/98.7]) or viewed from the Bernoulli process point of view.

Note that with the prior given as a binomial distribution, the posterior was also a binomial distribution:


$$\begin{array}{ll} \text{prior} & P(U=j) = \binom{n}{j} p^j q^{n-j} \\ \text{posterior} & P(U=j|A=i) = \binom{n-k}{j-i} p^{j-i} q^{n-k-j+1} \end{array} \quad \leftarrow \text{both binomial distributions}$$

This is an example of Conjugate Prior Theory.

In practice, the prior used depends on the application. Many different priors are plausible (we've shown two). A statistician's sense and experience are involved in choosing the appropriate prior for a particular problem.

Now, let's do the same w/o feeling - sampling w/ replacement:

• Bayes' Theorem for Sampling w/ Replacement



n = population size

j = number of good guys in population

Sample of size k with replacement

Let U = number of good guys (in the population)

A = number of good guys in a random sample of k with repetition


Computation of the conditional probability $P(A=i|U=j)$ is straightforward (for example, see [2/23/98.8]):

$$P(A=i|U=j) = \frac{\binom{k}{i} j^i (n-j)^{k-i}}{n^k}$$

In a sample of size k , there are $\binom{k}{i}$ ways to have i good guys.

Each of the i good guys can come from any of the j good members of the population $\Rightarrow j^i$

Each of the $k-i$ bad guys can come from any of the $n-j$ bad members of the population $\Rightarrow (n-j)^{k-i}$



$$P(A=i|U=j) = \binom{k}{i} \left(\frac{j}{n}\right)^i \left(\frac{n-j}{n}\right)^{k-i} = \binom{k}{i} \left(\frac{j}{n}\right)^i \left(1 - \frac{j}{n}\right)^{k-i}$$

Since the above is a probability distribution for the random variable A , it satisfies:

$$\sum_i P_{(U=j)}(A=i) = 1, \text{ for all } j$$

which gives the binomial identity (our friend again, the binomial theorem):

$$\sum_i \binom{k}{i} \left(\frac{j}{n}\right)^i \left(\frac{n-j}{n}\right)^{k-i} = 1, \text{ for all } j$$

Bayes' Law gives:

$$P(U=j|A=L) = \frac{P(A=L|U=j) P(U=j)}{P(A=L)}$$

$$= \frac{P(A=L|U=j) P(U=j)}{\sum_j P(A=L|U=j) P(U=j)} \quad \leftarrow \text{Law of Alternatives}$$

Substituting our recently derived

$$P(A=L|U=j) = \binom{k}{L} \left(\frac{j}{n}\right)^L \left(\frac{n-j}{n}\right)^{k-L}$$

$$= \frac{\binom{k}{L} \left(\frac{j}{n}\right)^L \left(\frac{n-j}{n}\right)^{k-L} P(U=j)}{\sum_j \binom{k}{L} \left(\frac{j}{n}\right)^L \left(\frac{n-j}{n}\right)^{k-L} P(U=j)}$$

$$P(U=j|A=L) = \frac{\left(\frac{j}{n}\right)^L \left(\frac{n-j}{n}\right)^{k-L} P(U=j)}{\sum_j \left(\frac{j}{n}\right)^L \left(\frac{n-j}{n}\right)^{k-L} P(U=j)}$$

Let's consider the posterior probability with different prior probabilities :

(1) Uniform Prior (Sampling w/ Replacement)

Assume $P(U=j) = \frac{1}{n+1}$

which gives the posterior as:

$$P(U=j|A=L) = \frac{\binom{k}{L} \left(\frac{j}{n}\right)^L \left(\frac{n-j}{n}\right)^{k-L} \cancel{\left(\frac{1}{n+1}\right)}}{\sum_j \binom{k}{L} \left(\frac{j}{n}\right)^L \left(\frac{n-j}{n}\right)^{k-L} \cancel{\left(\frac{1}{n+1}\right)}}$$

this is the discrete beta function,
denoted as :

$$C_{L,k,n}$$

3/13/98.11

$$P(U=j|A=i) = C_{i,k,n}^{-1} \left(\frac{j}{n}\right)^i \left(\frac{n-j}{n}\right)^{k-i}$$

Coverage of Sampling w/ Replacement

Case 1 - uniform prior

Case 2 - conjugate prior

continues on AfterNotes [AN 3/16/98.6-9]
lecture 16

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



Conjugate PriorsSampling w/o Replacement

population : n
sample : k

Random Variables:

U = number of good in population

A = number of good in sample

$$P(A=i|U=j) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}}$$

from the j good, we choose i . from the $n-j$ bad, we choose $k-i$.
number of points in sample space,

$$P(U=j|A=i) = \frac{P(A=i|U=j) P(U=j)}{P(A=i)}$$

$$= \frac{P(A=i|U=j) P(U=j)}{\sum_j P(A=i|U=j) P(U=j)} \quad \leftarrow \text{Law of Alternatives}$$

$$= \frac{\frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}} P(U=j)}{\sum_j \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}} P(U=j)}$$

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i} P(U=j)}{\sum_j \binom{j}{i} \binom{n-j}{k-i} P(U=j)}$$

this is the posterior given arbitrary priors

Now we have the following cases:

• Case 1 - Uniform prior

$$\text{prior } P(U=j) = \frac{1}{n+1}$$

This gives:

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\sum_j \binom{j}{i} \binom{n-j}{k-i}}$$

← we proved binomial identity probabilistically [3/11/98, 12-13]:

$$\sum_j \binom{j}{i} \binom{n-j}{k-i} = \binom{n+1}{k+1}$$

$$P(U=j|A=i) = \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n+1}{k+1}}$$

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



• Case 2 - Conjugate Prior

$$\text{prior } P(U=j) = \binom{n}{j} p^j q^{n-j}$$

$$P(A=i) = \sum_j P(A=i|U=j) P(U=j)$$

$$= \sum_j \frac{\binom{j}{i} \binom{n-j}{k-i}}{\binom{n}{k}} \binom{n}{j} p^j q^{n-j}$$

$$= \sum_j \frac{k! (n-k)!}{i! (k-i)!} \frac{j!}{i! (j-i)!} \frac{(n-j)!}{(k-i)! (n-j-k+i)!} \frac{1}{(n-j)!} p^j q^{n-j}$$

$$= \sum_j \frac{k!}{i! (k-i)!} \frac{(n-k)!}{(j-i)! (n-k-j+i)!} p^j q^{n-j}$$

$$= \sum_j \binom{k}{i} \binom{n-k}{j-i} p^j q^{n-j}$$

$$= \binom{k}{i} \sum_j \binom{n-k}{j-i} p^j q^{n-j}$$

Let $t = j - i$

$$= \binom{k}{i} \sum_t \binom{n-k}{t} \underbrace{p^i p^t}_{p^t} \underbrace{q^{k-i} q^{n-k-t}}_{q^{n-j}}$$

$$= \binom{k}{i} p^i q^{k-i} \sum_t \binom{n-k}{t} p^t q^{n-k-t}$$

1, by the binomial theorem

$$P(A=i) = \binom{k}{i} p^i q^{k-i}$$

Now we figure our posterior probability:

$$P(U=j|A=i) = \frac{P(A=i|U=j) P(U=j)}{P(A=i)}$$

$$= \frac{\binom{j}{i} \binom{n-j}{k-i} \binom{n}{j} p^i q^{n-j}}{\binom{k}{i} p^i q^{k-i}}$$

$$= \frac{\cancel{j!} (n-k)!}{\cancel{j!} \cancel{(k-i)!}} \frac{\cancel{j!} \cancel{(k-i)!}}{\cancel{j!} \cancel{(k-i)!}} \frac{(n-j)!}{(n-j-k+i)!} \frac{\cancel{j!}}{\cancel{j!} \cancel{(n-j)!}} p^{j-i} q^{n-j-k+i}$$

$$= \frac{(n-k)!}{(j-i)! (n-k-j+i)!} p^{j-i} q^{n-j-k+i}$$

$$P(U=j|A=i) = \binom{n-k}{j-i} p^{j-i} q^{n-j-k+i}$$

$$P(A=i)(U=j) = \binom{n-k}{j-i} p^{j-i} q^{n-j-k+i}$$



Since this is a probability, if you fix i and add up over j , the result must equal 1:

$$\sum_j P_{(A=i)}(U=j) = 1$$

$$\sum_j \binom{n-k}{j-i} p^{j-i} q^{n-j-k+i} = 1$$

More binomial identities!

Conjugate Prior

$$P(U=j) = \binom{n}{j} p^j q^{n-j}$$

$$\sum_j \binom{n}{j} p^j q^{n-j} = 1$$

Posterior

$$P_{(A=i)}(U=j) = \binom{n-k}{j-i} p^{j-i} q^{n-j-k+i}$$

$$\sum_j \binom{n-k}{j-i} p^{j-i} q^{n-j-k+i} = 1$$

{close to, but not quite,
a binomial distribution}

Q/A: We start w/ the conjugate prior distribution, which is concentrated from 0 to n .
Then we end up with a posterior distribution, which is concentrated from i to n
(since event $(A=i)$ is given).
So the posterior distribution is the distribution for some other family.
So our family for this posterior distribution is not closed under sampling, which
you want from conjugate priors - unless you define a broader family in general.
But if you define a broader family, you will need to do this computation
for all representatives of the broader family.
It's unclear what the broader family should be like.
It is not a well-defined concept, this notion of conjugate priors.

In this case, the conjugate prior and posterior are roughly something similar, but not quite
the same.

In class, we discussed a direct probabilistic interpretation to find the conditional probability
 $P(U=j|A=i)$ [3/13/98.7-8].



Uniform Prior

$$P(U=j) = \frac{1}{n+1}$$

Posterior

$$P_{(A=i)}(U=j) = \frac{\binom{i}{j} \binom{n-j}{k-i}}{\binom{n+1}{k+1}}$$

Note how very different the prior + the posterior are in this case. The posterior is very close to a hypergeometric distribution.

The conjugate prior and its posterior are much closer than the uniform prior and its posterior.

Q: In looking for prior probabilities that give posteriors in the same family, what do we look for?

A: This happens mostly when you take samples w/ replacement. Then it's no problem.

But if you take samples w/o replacement, the combinatorics get in the way.

- I don't think anybody worked out the Pólya case.
I just put it in.
I have no idea what it is.

- I have also not seen anywhere worked out (in this sense) the posterior multivariate hypergeometric distribution.

I think for the multivariate geometric, it is something very similar to this. For Pólya, I really don't know.

- In practice, most people use sampling w/ replacement. But that's not very good w/ small samples.
- I suspect that except for this posterior from ^{conjugate prior} multivariate hypergeometric, and Pólya, these are the only 3 that have this weird so-called prior, which is computable. No other family has this computable so-called conjugate

$$P(U=j | A=i) = \frac{P(A=i | U=j) \underbrace{P(U=j)}_{\text{Conjugate Priors}}}{P(A=i)}$$

Posterior

1) approx binomial
2) multivariate hypergeometric
3) Pólya

Posterior + Conjugate Prior
similar distributions



- The problem is that for a long time, people didn't believe in Bayes' Law. And now, there is a big turn towards Bayesian statistics. If you look at all the classical statistics books, the word Bayes is forbidden. And all this stuff is left out.

- I would love to have an argument that gives the binomial distribution as posterior.

Sampling w/ Replacement

$$P(A=i|U=j) = \frac{\binom{k}{i} j^i (n-j)^{k-i}}{n^k}$$

from k , pick i good ones

given one of $\binom{k}{i}$ ways to pick i good:
 number of ways to have i good guys
 number of ways to have i bad guys

number of sample points

$$\begin{aligned} P(U=j|A=i) &= \frac{P(A=i|U=j) P(U=j)}{P(A=i)} \\ &= \frac{P(A=i|U=j) P(U=j)}{\sum_j P(A=i|U=j) P(U=j)} \\ &= \frac{\frac{\binom{k}{i} j^i (n-j)^{k-i}}{n^k} P(U=j)}{\sum_j \frac{\binom{k}{i} j^i (n-j)^{k-i}}{n^k} P(U=j)} \end{aligned}$$

$$P(U=j|A=i) = \frac{\left(\frac{j}{n}\right)^i \left(\frac{n-j}{n}\right)^{k-i} P(U=j)}{\sum_j \left(\frac{j}{n}\right)^i \left(\frac{n-j}{n}\right)^{k-i} P(U=j)}$$

Just need to plug in the priors.

Case 1 - Uniform prior

Assume $P(U=j) = \text{same}$, for all j

So,

$$P(U=j|A=i) = \frac{\left(\frac{i}{n}\right)^i \left(\frac{n-i}{n}\right)^{k-i} P(U=j)}{\sum_j \left(\frac{i}{n}\right)^i \left(\frac{n-i}{n}\right)^{k-i} P(U=j)}$$

$$P(U=j|A=i) = \frac{\left(\frac{i}{n}\right)^i \left(\frac{n-i}{n}\right)^{k-i}}{\underbrace{\sum_j \left(\frac{i}{n}\right)^i \left(\frac{n-i}{n}\right)^{k-i}}}$$

The denominator is the discrete beta function

You can evaluate it, but it's pointless.

You sum over j .

Note: if you look in the books, they tend to do the integral, because the beta function integral has been well described.

We refer to the discrete beta function, which depends on i, k, n , as:

$$C_{i,k,n}$$

$$P(U=j|A=i) = C_{i,k,n}^{-1} \left(\frac{i}{n}\right)^i \left(\frac{n-i}{n}\right)^{k-i}$$

Posterior for uniform prior
with replacement

$$\text{We know that } \sum_j P_{(A=i)}(U=j) = 1$$

Thus it is understood that the $C_{i,k,n}$ constant will be such that the above sum equals 1.

Statisticians often ignore such constants, with the understanding that probabilities add up to 1, and write such probabilities as:

$$P(U=j|A=i) \propto \left(\frac{i}{n}\right)^i \left(\frac{n-i}{n}\right)^{k-i}$$

↑
proportional



Case 2 - Conjugate Prior

Assume $P(U=j) = C_{t,u,n}^{-1} \left(\frac{j}{n}\right)^t \left(\frac{n-j}{n}\right)^{u-t}$

So,

$$P(U=j|A=i) = \frac{\left(\frac{j}{n}\right)^i \left(\frac{n-j}{n}\right)^{k-i} C_{t,u,n}^{-1} \left(\frac{j}{n}\right)^t \left(\frac{n-j}{n}\right)^{u-t}}{\sum_j \left(\frac{j}{n}\right)^i \left(\frac{n-j}{n}\right)^{k-i} C_{t,u,n}^{-1} \left(\frac{j}{n}\right)^t \left(\frac{n-j}{n}\right)^{u-t}}$$

$$= \frac{\left(\frac{j}{n}\right)^{i+t} \left(\frac{n-j}{n}\right)^{k+u-i-t}}{\sum_j \left(\frac{j}{n}\right)^{i+t} \left(\frac{n-j}{n}\right)^{k+u-i-t}}$$

This denominator is the discrete beta function denoted as:

$$C_{i+t, k+u, n}$$

$$P(U=j|A=i) = C_{i+t, k+u, n}^{-1} \left(\frac{j}{n}\right)^{i+t} \left(\frac{n-j}{n}\right)^{k+u-i-t}$$

$$P(U=j|A=i) \propto \left(\frac{j}{n}\right)^{i+t} \left(\frac{n-j}{n}\right)^{k+u-i-t}$$

Conjugate Prior

$$P(U=j) = C_{t,u,n}^{-1} \left(\frac{j}{n}\right)^t \left(\frac{n-j}{n}\right)^{u-t}$$

Posterior

$$P_{(A=i)}(U=j) = C_{i+t, k+u, n}^{-1} \left(\frac{j}{n}\right)^{i+t} \left(\frac{n-j}{n}\right)^{k+u-i-t}$$

Since these are probabilities, then have to add up to 1.

So you also have the discrete beta function identity series:

$$\sum_j P(U=j) = 1$$

$$\Rightarrow \sum_j \left(\frac{j}{n}\right)^t \left(\frac{n-j}{n}\right)^{u-t} = C_{t,u,n}$$

$$\sum_j P_{(A=i)}(U=j) = 1$$

$$\Rightarrow \sum_j \left(\frac{j}{n}\right)^{i+t} \left(\frac{n-j}{n}\right)^{k+u-i-t} = C_{i+t, k+u, n}$$

what does it mean to take $P(U=j) = C_{t,u,n}^{-1} \left(\frac{j}{n}\right)^t \left(\frac{n-j}{n}\right)^{u-t}$ as prior?

We already did this experiment and got the result $P(U=j|A=t)$.

We did the experiment with a sample of size u and got t good.

Then, having done this experiment, we use that as prior for a second experiment, pretending that we haven't done the first experiment.

So, for the 2nd experiment, the prior is really:

$$\underbrace{P(U=j)}_{\text{prior for 2nd experiment}} = \underbrace{P(U=j|A=t)}_{\text{result from 1st experiment}}$$

$$= C_{t,u,n}^{-1} \left(\frac{j}{n}\right)^t \left(\frac{n-j}{n}\right)^{u-t}$$

This means we have previously done the experiment.
We use that information as our prior.
Plug it in and we get the posterior.

Q/A: Start the experiment and take a sample of size t .
Then I pretend I have not done the experiment.
I say "Let's use the result of this experiment as prior."

Then I do the actual experiment. Taking a sample of k , I find l good.
The result of this experiment is as if you had a single experiment with a sample size of $u+k$ and found $t+l$ good.

1st pretend ↑ actual experiment 1st pretend ↑ actual experiment

$$\underbrace{P(U=j|A=l)}_{\text{posterior}} = C_{t+l,k+u,n} \left(\frac{j}{n}\right)^{t+l} \left(\frac{n-j}{n}\right)^{k+u-t}$$

- There is a fancy book on this by Sir Harold Jeffreys entitled "Theory of Probability".
Some statisticians have made an entire career rewriting chapters of that book.
700 pages. No table of contents. No index.
It's all streams of consciousness — you have to know where you are.
This is sort of the bible of Bayesian Probability.
I'm almost sure this is in Sir Harold Jeffreys.
The sampling w/o replacement examples are not.



I did all this because I want to get to De Finetti's Theorem.

↑ one of the few honest theorems
in this business.

Given an infinite sequence of exchangeable random variables that take values 0 + 1.
That's the same as:

- 1) Taking a number p between 0 and 1
with your favorite distribution.
- 2) After you do that, you play Bernoulli Trials.

So the prior can be any infinite sequence of 0 and 1, which is exchangeable,
i.e., position independent.

This means that there is a posterior of the Bernoulli Process with an arbitrary prior
on the probability of the bias of the coin.


It's a really nice theorem.
I'll do it in class.

In some sense you can not have infinite sequences of exchangeable 0-1 random
variables. They are very close to coin tosses.

Continuous Random Variables (Beginning)

Let me skip the motivation and just start with the definitions, then go into the theory, then we'll give motivational examples.

So far, we have considered random variables that are integer valued. A random variable X on the probability space Ω, P that takes integer values. Guess what's coming next?



Random variable X , whose values are arbitrary real numbers.

General Definition of a Random Variable ← no restrictions

A random variable X on a probability space is a function from Ω to the real numbers such that for every real number a and b ,

$$\{\omega : a < X(\omega) \leq b\} = (a < X \leq b) \text{ is an event.}$$

the set of all ω for which $X(\omega)$ lies in $(a, b]$

which we call, from now on, the event that X takes values in $(a, b]$.


This is an almost vacuous definition.

event $X \leq t$

To this end, it suffices to assume that $(X \leq t)$ is an event for all real numbers t , since

$$\underbrace{(a < X \leq b)}_{\text{The event that } X \text{ is in } (a, b]} = \underbrace{(X \leq b)}_{\text{The event that } X \leq b} - \underbrace{(X \leq a)}_{\text{The event that } X \leq a}$$

And now you realize why we say $(a < X \leq b)$ instead of $(a \leq X \leq b)$ ^{strictly}.





- The function $F(t) = \overbrace{P(X \leq t)}^{\text{probability of the event } X \leq t}$ is called the cumulative distribution function of the random variable X .

Example:

$$\underbrace{P(a < X \leq b)}_{\substack{\text{Probability that } X \text{ lies} \\ \text{in } (a, b]}} = \underbrace{F(b)}_{P(X \leq b)} - \underbrace{F(a)}_{P(X \leq a)}$$

- The random variable X takes values between a and b when

$$P(a \leq X \leq b) = 1$$

$\left\{ \begin{array}{l} \text{Here, we want } \leq \\ \text{Because } X \text{ could take value } a \text{ w/ positive probability} \end{array} \right\}$

i.e., $P(X > b) = 0$

and $P(X < a) = 0$

- Two random variables X and Y are independent when:

$$P((a < X \leq b) \cap (c < Y \leq d)) = P(a < X \leq b) P(c < Y \leq d) \text{ for all } a, b, c, d$$

- The cumulative probability distribution $F(t)$, also written as $F_X(t)$, satisfies:

$$\lim_{t \rightarrow \infty} F(t) = 1$$

$$\lim_{t \rightarrow -\infty} F(t) = 0$$

$\left\{ \begin{array}{l} \text{When you want to} \\ \text{identify which random} \\ \text{variable this is the} \\ \text{cumulative probability} \\ \text{distribution of.} \end{array} \right\}$

and $F(t) \leq F(t')$ if $t < t'$

In other words, $F(t)$ is increasing, and it tends to zero at $-\infty$ and 1 at $+\infty$.
As follows immediately, from the definition:

$$F(t) = P(X \leq t)$$

- If the random variable X takes the values between a and b , we say that X is a continuous random variable when there is a function $f(t)$ for which:

$$P(c < X \leq d) = \int_c^d f(t) dt, \text{ for } a \leq c \leq d < b$$

Probability of the event
that X takes values
in $(c, d]$

This is equivalent to:

$$F(d) - F(c) = \int_c^d f(t) dt$$

$$\left\{ P(c < X \leq d) = F(d) - F(c) \right\}$$

and too:

$$F(t) = \int_{-\infty}^t f(s) ds$$

$$\left\{ \begin{array}{l} P(-\infty < X \leq t) = \int_{-\infty}^t f(s) ds \\ F(t) - \underbrace{F(-\infty)}_0 = \int_{-\infty}^t f(s) ds \\ F(t) = \int_{-\infty}^t f(s) ds \end{array} \right\}$$

and too:

$$f(t) = \frac{d}{dt} F(t) = \frac{d}{dt} P(X \leq t)$$

for t between a and b

- Sometimes one writes:

$$\text{dens}(X=t) = f(t) \leftarrow \left\{ \text{dens}(X=t) \text{ instead of } f(t) \right\}$$

and calls it the density of the continuous random variable X (density is 0 outside the interval $[a, b]$).

- Density is always non-negative
 $\text{dens}(X=t) = f(t) \geq 0$

This is the continuous extension of the notion of random variables that we have been dealing with.

But now, shifting from the integers to the real numbers.

And we are going to see that a lot of the phenomena that we've already studied (e.g., balls into boxes) have continuous analogues.

Q: We have that: $F(t) = \int_{-\infty}^t f(s) ds$

and $f(t) = \frac{d}{dt} F(t)$

What if $f(t)$ is not a continuous function? In which case $F(t)$ will not be differentiable everywhere.

For example, what if $f(t)$ is a step function?

A: Then X is not a continuous random variable.

It's not true that every random variable is either integer or continuous.

There are other random variables, which we will consider later.

Continuous random variables are a special kind of random variable, just as integer random variables are a special kind of random variable.

This is the basic definition. It is vacuous, like all definitions.

But, gradually you will see what it means.

Q/A: if a random variable is continuous, we say that these integrals

$$\int_c^d f(t) dt, \text{ for } a \leq c \leq d \leq b$$

and $\int_{-\infty}^t f(s) ds$

automatically converge.

Let me write another formula:

$$\int_{-\infty}^{\infty} f(t) dt = 1$$

$$P(-\infty < X \leq \infty)$$

[Q/A: $f(t)$ has to be a function]

Now let's see examples.

As usual, the examples are not what one expects.

You are used to seeing functions like functions and integers like integers.


Now you see that this relationship between the cumulative distribution of a random variable and the density of a random variable, which are the probabilistic analogues of the fundamental theorem of the Calculus

Now the density of a random variable is not the probability.
It's something like the mass density.

By the way, I did not say what it means for two random variables to be identically distributed. It's obvious what it means. They have the same cumulative probability distribution functions.

Example 1 The uniform or Dirichlet process

Take an interval $[0, a]$ and pick a point at random.



pick a point
at random

} The simplest possible random variable
you can think of.

The sample space $\Omega = [0, a]$
is the interval
 $[0, a]$

sample point $\omega \in [0, a]$

The events are all intervals $(c, d]$ and any events that can be obtained on these intervals by ^{taking} union, intersection, complement, and infinite countable unions of disjoint ~~intervals~~ intervals.

Q: How do you get an event of 0?

A: You can easily get, from half open intervals, closed intervals, using these operations. So you can easily get the closed interval $[c, d]$. Thus this is also an event.

Then you can subtract the two:

$$[c, d] - (c, d] = c \leftarrow \text{to get a point}$$

So any sensible set is in the interval $[0, a]$.

These are sometimes called Borel sets.

~~These are~~
We define the probability as:

$$P((c, d]) = \frac{d-c}{a}$$



Then, by a theorem that we stated, but did not prove, this probability extends, in a well-defined way, to all events

↑ all Borel sets


By the fundamental theorem we stated at the beginning of the course, this gives a well-defined probability on all events.

Let $X = \text{random point in } [0, a]$

↑ a random variable

$$P(c < X \leq d) = \frac{d-c}{a} \quad \leftarrow \text{we define it that way}$$

probability that X
takes a value in
 $(c, d]$



$$\underbrace{\text{dens}(X=t)}_{\text{density of this random variable}} = \begin{cases} \frac{1}{a} & \text{if } 0 \leq t \leq a \\ 0 & \text{elsewhere} \end{cases}$$

X is a continuous random variable on the interval $[0, a]$

$$P(c < X \leq d) = \int_c^d \underbrace{\frac{1}{a}}_{\text{dens}(X=t) = f(t) = \frac{1}{a}} dt$$

$$\text{dens}(X=t) = f(t) = \frac{1}{a}$$



- Let's pick n points independently.

Dirichlet Process - Take $X_1, X_2, \dots, X_n$ independent random variables, identically distributed with X_1 .

we'll define the distribution in a moment.


What is the Dirichlet Process really?

Really, what we are doing is - guess what?

We take the interval $[0, a]$ and we drop n balls on that interval

So it's balls into boxes.

Except the boxes are now continuous.



Now we can ask questions. The git of the questions are those that we asked for Maxwell-Boltzmann statistics. Because these are distinguishable balls.


Q/A: Unfortunately in the book, I wrote X_n all over the place. If we kept the original notation (i.e., k balls into n boxes), we'd instead have independent random variables $X_1, X_2, \dots, X_k$ for the k balls and the interval (i.e., the boxes) would be $[0, n]$.

I debated last night whether to change it everywhere. Let's keep it as written.

So let's ask a suitable question.

Occupation Numbers


We take the interval $[0, a]$ and we divide it into intervals:



Now, we have genuine boxes, except they are of different sizes.

So the probability of a ball going into a box is different from the probability of a ball going into another box, because it could be bigger or smaller.

Let's compute the probability that the intervals have the following occupation numbers:



Let's introduce a convenient and efficient notation for this type of computation. From now on I talk about balls. The X_i refer to balls.

Let $N(t)$ = number of X_i in the interval $[0, t)$
(i.e., number of balls falling in $[0, t)$)

$$P(N(t)=k) = \binom{n}{k} \left(\frac{t}{a}\right)^k \left(\frac{a-t}{a}\right)^{n-k}$$

from the n balls,
you pick k .

The X_i are independent, so their probabilities multiply.

$\left(\frac{t}{a}\right)^k$ is the probability that they drop in $[0, t)$

$\left(\frac{a-t}{a}\right)^{n-k}$ is the probability that they drop outside $[0, t)$

$P(N(t)=k)$ = probability that there are k balls in $[0, t)$

Now let's go back to the event illustrated at the top of the page:

Event:

$$A = (N(b)=i) \cap (N(c)-N(b)=j) \cap (N(d)-N(c)=k) \cap (N(a)-N(d)=n-i-j-k)$$

We use multinomial coefficients to compute $P(A)$.

We have n balls, of which i go into the first interval

j " " " second "

k " " " third "

$n-i-j-k$ " " " remaining interval

$$\left\{ \begin{matrix} n \\ i, j, k, n-i-j-k \end{matrix} \right\}$$

read "multinomial n
 $i, j, k, n-i-j-k$ "

This is what multinomial coefficients are for.

$$P(A) = \binom{n}{i, j, k, n-i-j-k} \left(\frac{b}{a}\right)^i \left(\frac{c-b}{a}\right)^j \left(\frac{d-c}{a}\right)^k \left(\frac{a-d}{a}\right)^{n-i-j-k}$$

you pick the balls
that go in each
box.

Now they go independently, so the
probabilities multiply.
 i go independently into the first interval
 j " " " second "
etc.

And that's the way you compute probabilities.

3/16/98.9

Let's repeat what we just did in high school terms. Really high class.
What we've done is compute the joint distribution of a random function.
Watch.

What is $N(t)$, really?

$N(t)$ is infinitely many random variables, one for each t .

And these random variables are not independent.

(As the formula for $P(A)$ clearly indicates.)

In some sense $N(t)$ tells you everything about occupation numbers.

$N(t)$ is a random variable for each t

One way of saying this (i.e., that you have a random variable for each t) is to say you have a random function.

$N(t)$ is a random function

{The word function is somewhat misleading here, but not too much.}

The random function is a function of t , which is a random variable for each t .
The random function $N(t)$, associated w/ the Dirichlet Process, contains all the information about occupation numbers for the Dirichlet Process.


Anything about the balls dropping in the interval $[0, a]$, you can infer by studying joint distributions of the random function.

You think that's easy? Just wait.

• Example

You can view the dropping of the balls on the interval $[0, a]$ as putting checks on the interval.

Once the ball is dropped, you can forget about which ball is which.
We'll talk later about indistinguishable balls.



The probability that two balls will fall to the same place is 0.

So we may disregard it from now on.

We can ask the question:

When you drop all n balls, how close to the origin is the ball that is closest to the origin?

There is a standard name for that, used by everybody.

not my fault

$$X_{(1)} = \min(X_1, X_2, \dots, X_n)$$

read: "X order 1"

What is the probability distribution of $\widetilde{X_{(1)}}$?


$$P(X_{(1)} > t) = ?$$

The event $(X_{(1)} > t)$ is the same as the event $(N(t) = 0)$.
If you don't see that, I can't explain it.

$$(X_{(1)} > t) = (N(t) = 0)$$

Therefore,

$$P(X_{(1)} > t) = \left(\frac{a-t}{a}\right)^n$$



$$\begin{aligned} P(X_{(1)} > t) &= P(N(t) = 0) \\ &= \binom{n}{0} \left(\frac{t}{a}\right)^0 \left(\frac{a-t}{a}\right)^n \\ &= \left(\frac{a-t}{a}\right)^n \end{aligned}$$

And therefore,

the cumulative distribution of the random variable $X_{(1)}$ is:

$$\begin{aligned} F(t) &= P(X_{(1)} \leq t) \\ &= 1 - P((X_{(1)} \leq t)^c) \\ &= 1 - P(X_{(1)} > t) \\ &= 1 - \left(\frac{a-t}{a}\right)^n \end{aligned}$$

$$F(t) = 1 - \frac{(a-t)^n}{a^n}$$

Next, we compute the density of the continuous random variable.
Here we have the first non-trivial example of a continuous random variable, whose density is not obvious.

We have that:

$$\int_{-\infty}^t f(s) ds = F(t) \implies f(t) = \frac{d}{dt} F(t)$$

So the density $f(t)$ is:

$$f(t) = \text{dens}(X_{(1)} = t)$$

$$= \frac{d}{dt} \left(1 - \frac{(a-t)^n}{a^n} \right)$$

$$f(t) = \frac{n(a-t)^{n-1}}{a^n}$$

Observe that:

$$\int_0^a \frac{n(a-t)^{n-1}}{a^n} dt = 1$$


Automatically!

Since the integral of the density is always 1.

$$\begin{aligned} 1 &= \int_{-\infty}^{\infty} f(t) dt = \int_{-\infty}^{\infty} \frac{n(a-t)^{n-1}}{a^n} dt \\ &= \int_0^a \frac{n(a-t)^{n-1}}{a^n} dt \end{aligned}$$

Why is this result significant?

If you plot these curves, you'll find they are more and more condensed towards the origin.



graph of $\text{dens}(X_{(1)} = t)$

as n increases (i.e., more balls dropped), more mass is concentrated to the origin. That is, for any arbitrarily small ϵ , $\lim_{n \rightarrow \infty} \int_0^\epsilon f(t) dt \rightarrow 1$

You get:

$$P(X_{(1)} \leq \varepsilon) = F(\varepsilon) \quad \text{cumulative distribution}$$

$$= \int_{-\infty}^{\varepsilon} f(t) dt$$

$$P(X_{(1)} \leq \varepsilon) = \int_0^{\varepsilon} \frac{n(a-t)^{n-1}}{a^n} dt$$

{ for fixed ε , the larger n is, the more the mass of the density curve moves towards the origin. }

So the more balls you drop, the bigger the chance that the closest ball will go to the origin!

This gives the exact expression, by integrating the density.

So you have to get used to the fact that besides probability, there is this other measure, which is density.

density : probability :: mass density : mass

mass density means nothing in itself.
It means something only when integrated.

The same happens now with density.

Probability density means nothing in itself.
It means something when you integrate it over an interval, in which case, it is the probability of a continuous random variable taking values within that interval.

The probability of a continuous random variable taking one value is 0.

$$P(X=x) = \lim_{h \rightarrow 0} P(x-h < X \leq x)$$

$$= \lim_{h \rightarrow 0} \int_{x-h}^x f(t) dt$$

$$P(X=x) = 0$$

This is what the calculus is all about.

Your chickens are coming to roost now, if you don't know calculus.

From now on, we use calculus.

• Next comes the question:

$X_{(2)}$ is the next ball after $X_{(1)}$
reads: "X order 2"

Then you have:

$$X_{(1)} < X_{(2)} < X_{(3)} < \dots < X_{(n)}$$

We are going to compute the densities of these random variables, which are not independent, of course.

The order in which you find balls as you go down the interval.
These are called the:

Order Statistics

Tremendously important in statistics.

With these, you can answer questions such as how big these intervals are.

To do so, you need to know the facts:

gaps
biggest
smallest
median

[On the problem set that I gave you, there are 2 problems that are unsolved.
You are MIT students, you're not students at Oshkosh College.
It would be insulting to you if I gave you only solved problems all the time.]

Continuous Random Variables (cont'd)

Last time we saw the definition of Random Variable in general.

if X is a random variable and $P(a < X \leq b) = 1$, then X is said to be continuous when it has a density $f(t) = \text{dens}(X=t)$, i.e., when:

$$P(c < X \leq d) = \int_c^d f(t) dt = \int_c^d \text{dens}(X=t) dt \quad a \leq c \leq d \leq b$$


probability that X
is in the interval
 $(c, d]$

This is the continuous analogue of what we have been doing so far.

Given a random variable taking its values in the interval $(a, b]$, this random variable is said to be continuous when there is a density function

(the probability that a random variable takes a given value is 0. ($P(X=\pi) = 0$).

All we can talk about is that the random variable takes a value between c and d .



Exactly the analogue of what we call calculus.
That's what this is all about.

A density is characterized by:

(1) $f(t) \geq 0$ for t between a and b

(2) $\int_a^b f(t) dt = 1$

since $\int_a^b f(t) dt = P(a < X \leq b) = 1$

by assumption (see above)

This is the continuous analogue of the probability distribution of an integer random variable. Integer random variables have probability distributions, which are sequences and if you sum all the sequences, they add up to 1.

Here we have functions integrating to 1. It's exactly parallel.

In addition to the density, we have the cumulative distribution.

We saw last time that cumulative distribution is defined for all random variables, not just continuous random variables.

We will see later examples of random variables that are neither continuous nor integer.

You can immediately imagine one - take the sum of a continuous random variable and an integer random variable. The resulting random variable is neither continuous nor random.

$$\underbrace{X}_{\text{continuous random variable}} + \underbrace{Y}_{\text{integer random variable}}$$

Cumulative Distribution

$$F(t) = F_X(t) = P(X \leq t)$$

↑ we use this notation when we want to be explicit about which random variable we are talking about.

For a continuous random variable, you have certain important relationships between the cumulative distribution and the density, which are essential to manipulate.

Let's review them:

$$P(c < X \leq d) = F(d) - F(c)$$

By an obvious argument, which we made last time.

Also:

$$F(t) = \int_{-\infty}^t \text{dens}(X=s) ds$$


I put $-\infty$ instead of a , since $\text{dens}(X=s)$ will be 0 when $s < a$. So it doesn't matter.

And from this, by the Fundamental Theorem of the Calculus, you obtain the fundamental relation:

$$\frac{d}{dt} F(t) = \text{dens}(X=t) = f(t)$$

The cumulative distribution is the integral of the density.
The density is the derivative of the cumulative distribution.

What does the cumulative distribution look like?



$F(t)$ is increasing. Why? Because of the definition:

$$F(t) = P(X \leq t)$$

Since probability is non-zero, if you increase t , then the probability of the event $(X \leq t)$ either stays the same, or increases for the larger event.

$$\lim_{t \rightarrow -\infty} F(t) = 0$$

$$\lim_{t \rightarrow \infty} F(t) = 1$$

If the random variable is continuous then, of course, the cumulative distribution will also be continuous - in fact, differentiable.

Let's add a few tidbits to what we said last time:

if X is continuous, then there is a number m , for which $F(m) = \frac{1}{2}$.

Because $F(m)$ goes from 0 to 1 continuously.

So there is one number where $F(m) = \frac{1}{2}$.

This number m is the median of the random variable X .

Similarly, you can define the quartiles, as they say in statistics.

$m_{1/4}$ with the property that $F(m_{1/4}) = \frac{1}{4}$ ← that's called the first quartile


$m_{3/4}$ " " " " $F(m_{3/4}) = \frac{3}{4}$

This is the precise definition of the common sense, everywhere used, notion.

You can interpret these numbers explicitly in terms of probabilities.
Let's redefine the median w/o using the cumulative distribution.

Q: Is there only one median m ?

A: No. Assuming we are talking about continuous random variables. If it's not continuous, all bets are off.
Consider the following density:



$$F(t) = \int_{-\infty}^t f(s) ds$$

$F(t) = \frac{1}{2} \forall t$ in this interval

$m \in$ this interval


So there can be many medians m

If $f(t) > 0$ $\uparrow$ strictly, then there can be only one median m

$$P(X < m) = \frac{1}{2} = P(X \geq m)$$

This is just repeating the definition.
This is for continuous random variables.

Let's add one more concept.
That's the expectation.



For the integer random variables, the idea of expectation was clear, if you had a finite sample space. There, you took the value of the random variable at each point, multiplied it by the probability at that point and added it all over:

$$E(X) = \sum_n n P(X=n)$$

Now, we are in a similar situation - even worse, since we have a continuum. Points have probability 0.

We define the expectation for a continuous random variable X :

$$E(X) = \int_a^b t \underbrace{f(t)}_{\text{density}} dt$$

The expectation is the average of the density, which is the continuous analogue of the notion of expectation that we had for integer random variables.

And we'll see later that the expectation of the sum of continuous random variables is the sum of the expectations, even though the random variables may not be independent.

$$E(X_1 + X_2 + \dots + X_n) = E(X_1) + E(X_2) + \dots + E(X_n)$$

Example Uniform (Dirichlet) Distribution

This is one of the fundamental stochastic processes.
To put it in really blunt terms:

Take n measurements of some quantity.
Then you want all the statistical facts about these n measurements.

Really rock bottom phenomena.

A measurement of the quantity that is known to lie between 0 and a .
Beyond which, we know nothing whatsoever.

The uniform process is characterized by being given n random variables $X_1, \dots, X_n$ which are independent and identically distributed.


The probability distribution:

$$P(c < X_1 \leq d) = \int_c^d \frac{1}{a} dt \quad \text{if } 0 \leq c \leq d \leq a$$

$$= \frac{d-c}{a}$$

In other words, if I have an interval $[0, a]$ and I pick a point at random, I can not tell anything about the probability that this point is some specific point - that probability is 0.

$$P(X_1 = x) = 0$$



But we can tell the probability that the picked point lies in an interval.
For example:

$$P(x-\epsilon < X_1 \leq x+\epsilon) = \int_{x-\epsilon}^{x+\epsilon} \frac{1}{a} dt$$

$$= \frac{2\epsilon}{a}$$

The density $f(t) = \frac{1}{a}$

The density is not the probability.

It's just like mass density and mass. Same analogy. It's extremely similar.
As a matter of fact there is a concept relevant to math + physics that, as you will see, is very clearly visualized in probability.

The variables are all independent and they all have this obvious uniform distribution.
 $P(c < X_1 \leq d)$

What can we say about them?

We associate w/ the Dirichlet Process the integer random variable:

$U(t) = \text{number of } X_i\text{'s that are } \leq t$

We computed the probability distribution of this last time [3/16/98.8]:

$$P(U(t) = k) = \binom{n}{k} \left(\frac{t}{a}\right)^k \left(\frac{a-t}{a}\right)^{n-k}$$

We saw that $U(t)$ is a random variable for each t .
And this situation is called random functions.

A random function is a random variable for each value of t . Observe that to get a random function is not just to describe one random variable for each t . All the joint distributions for any finite set of t , are needed to completely describe the random function.

That's what we did last time. We subdivided the interval into a number of boxes and then we computed the probability that 1st box had i X 's, 2nd j X 's, etc. [3/16/98.7-9].
This gives you joint distributions of the integer random variable $U(t)$.

$U(t)$ is an integer random variable associated w/ the Dirichlet Process.

On Friday, Problem Set w/ 20 problems. A regular problem set and training problems for the quiz.

There is a very nice collection of problems made up by a British Reverend Whitworth who include some beautiful problems. I'll crib from those. They're all stated in British terms. Some Bayes problems that are really super.

By the way, if you look up in the literature on Bayesian stuff, the philosophers have written many books w/ names such as "Paradoxes of Rationality and Cooperation" Several hundred pages.



On the other hand, the order statistics, which we began to discuss last time, are continuous random variables.

These are probably the most important single sequence of random variables there is. Perhaps 20 books have been written on the subject of order statistics.

$$\underbrace{X_{(1)}}_{\text{read: "X order 1"}} = \min(X_1, X_2, \dots, X_n)$$

$$X_{(2)} = \text{next after the minimum}$$

$$\vdots$$

$$\vdots$$

$$X_{(n)} = \max(X_1, X_2, \dots, X_n)$$

These are the
order statistics

The order statistics are the X_i rearranged in increasing order.

In statistical situations, in problems of distribution, that's what happens. We will later consider problems of occupancy.

So let's now study the order statistics

These are not independent random variables.

I have to curb my tendencies to philosophize.

But, philosophically, the order statistics correspond to indistinguishable balls.

I'm not going to explain that.

If you don't see it, forget it.

Shifting from $X_1, X_2, \dots, X_n$ to $X_{(1)}, X_{(2)}, \dots, X_{(n)}$ means that you are shifting from distinguishable balls to indistinguishable.

Q: Can we assume that no two X_i are the same?

A: Always. Probability is 0. These are continuous random variables.

We don't even mention it anymore.

Our first job is to compute the probability distribution of the order statistics. We'll do it two ways.


First we do it brutal. Then we do it using calculus.

$$P(X_{(k)} \leq t) = ?$$

The cumulative distribution
of the k^{th} order statistic

Well, what do you do to compute the probability?

Let's write the event $(X_{(k)} \leq t)$ in terms of events that we can compute.



What does it mean for $(X_{(k)} \leq t)$?

At least k X 's fall in interval $(0, t]$.

This kind of trick is the only trick there is. Once you've mastered that trick, you can do everything.

So we have:

$$(X_{(k)} \leq t) = (U(t)=k) \cup (U(t)=k+1) \cup \dots \cup (U(t)=n)$$

Therefore, these events being disjoint, we are saved from doing mutual inclusion-exclusion.

$$P(X_{(k)} \leq t) = P(U(t)=k) + P(U(t)=k+1) + \dots + P(U(t)=n)$$

$$= \binom{n}{k} \frac{t^k (a-t)^{n-k}}{a^n} + \binom{n}{k+1} \frac{t^{k+1} (a-t)^{n-k-1}}{a^n} + \dots + \binom{n}{n} \frac{t^n (a-t)^{n-n}}{a^n}$$

i.e., all of them fall in the interval $(0, t]$.

Now, I have an unpleasant announcement.
Let's compute the density of this random variable.
That's the derivative of the mess above.

1st - Density - The Hard Way

$$\text{dens}(X_{(k)} = t) = \frac{d}{dt} F(t)$$

$$= \frac{d}{dt} P(X_{(k)} \leq t)$$

Miracles of cancellation occur!
 $(fg)' = fg' + gf'$

$$= \binom{n}{k} \frac{k t^{k-1} (a-t)^{n-k}}{a^n}$$

$$+ \binom{n}{k} \frac{(n-k) t^k (a-t)^{n-k-1}}{a^n}$$

} derivative of 1st term

$$+ \binom{n}{k+1} \frac{(k+1) t^k (a-t)^{n-k-1}}{a^n}$$

$$- \binom{n}{k+1} \frac{(n-k-1) t^{k+1} (a-t)^{n-k-2}}{a^n}$$

} of 2nd term

$$+ \dots$$

$$- \dots$$

} :

$$+ \frac{n t^{n-1}}{a^n}$$

$$- \dots$$

} of last term

Miracles of cancellation occur because alternating $-/+$ terms have the same factors for t and $(a-t)$ and the coefficients cancel:

$$-\binom{n}{k+i}(n-k-i) + \binom{n}{k+i+1}(k+i+1) \stackrel{?}{=} 0$$

$$-\frac{\cancel{n!} (n-k-i)}{(k+i)!(n-k-i)!} + \frac{\cancel{n!} (k+i+1)}{(k+i+1)!(n-k-i-1)!} \stackrel{?}{=} 0$$

$$-\frac{1}{(k+i)!(n-k-i-1)!} + \frac{1}{(k+i)!(n-k-i-1)!} = 0 \quad \checkmark$$

So only the first term remains and we have:

(*)

$$\text{dens}(X_{(k)}=t) = \binom{n}{k} \frac{k t^{k-1} (a-t)^{n-k}}{a^n}$$

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



2nd - Density - The Easy Way

Now let's rederive this result directly using calculus, without going through the cumulative distribution.

Again, with feeling:

By definition, we have:

$$P(t < X_{(k)} \leq t+dt) = \int_t^{t+dt} \text{dens}(X_{(k)}=s) ds$$


$$\lim_{dt \rightarrow 0} P(t < X_{(k)} \leq t+dt) = \lim_{dt \rightarrow 0} \int_t^{t+dt} \text{dens}(X_{(k)}=s) ds$$

$$= \lim_{dt \rightarrow 0} \text{dens}(X_{(k)}=t) dt$$

which gives:

$$\lim_{dt \rightarrow 0} \frac{P(t < X_{(k)} \leq t+dt)}{dt} = \text{dens}(X_{(k)}=t)$$

This is calculus.



The event that the k^{th} order statistic is in the interval $(t, t+dt]$.
 $(t < X_{(k)} \leq t+dt)$

multinomial coefficient


probability of all events where more than 1 X (ball) falls in $(t, t+dt]$

$$P(t < X_{(k)} \leq t+dt) = \binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{dt}{a}\right) \left(\frac{a-t-dt}{a}\right)^{n-k} + O(dt^2)$$

probability of event that $k-1$ X 's (balls) are in $[0, t]$ exactly 1 in $(t, t+dt]$ exactly $n-k$ in $(t+dt, a]$

$$\sum_{\substack{u+v+w=n \\ u \geq k-1 \\ v \geq 1 \\ w \geq 0}} \binom{n}{u, v, w} \left(\frac{t}{a}\right)^u \left(\frac{dt}{a}\right)^v \left(\frac{a-t-dt}{a}\right)^w$$

Note that $v \geq 2$



$$\begin{aligned} \lim_{dt \rightarrow 0} \frac{P(t < X_{(k)} \leq t+dt)}{dt} &= \lim_{dt \rightarrow 0} \left\{ \binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{1}{a}\right) \left(\frac{a-t-dt}{a}\right)^{n-k} + O(dt) \right\} \\ &= \binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{1}{a}\right) \left(\frac{a-t}{a}\right)^{n-k} \end{aligned}$$

which gives:

$$\text{dens}(X_{(k)} = t) = \binom{n}{k-1, 1, n-k} \frac{t^{k-1} (a-t)^{n-k}}{a^n}$$

Done. Directly. No computation necessary.

This is the same as equation (8), obtained from differentiating the cumulative density as:

$$\binom{n}{k} k = \binom{n}{k-1, 1, n-k}$$

$$\begin{aligned} \frac{n!}{k!(n-k)!} k &= \frac{n!}{(k-1)! 1! (n-k)!} \\ &= \binom{n}{k-1, 1, n-k} \checkmark \end{aligned}$$



Now let's reap something from our labors.
As a consequence, we obtain the following:

$$\int_0^a \text{dens}(X_{(k)} = t) dt = 1 \quad \leftarrow \left\{ \begin{array}{l} \text{We know that the integral of the} \\ \text{density is 1} \end{array} \right\}$$

This gives us an extraordinary integral:

$$\int_0^a \binom{n}{k} \frac{k t^{k-1} (a-t)^{n-k}}{a^n} dt = 1$$

Automatically, as they say
Give that to your friends to
evaluate directly.

$$\int_0^a \frac{n!}{(k-1)!(n-k)!} \frac{t^{k-1} (a-t)^{n-k}}{a^n} dt = 1$$

Let's rewrite this in a way that makes
where it comes from.

Namely:

$$\int_0^a t^{k-1} (a-t)^{n-k} dt = \frac{a^n (k-1)!(n-k)!}{n!}$$

This is the Beta Function Integral

The distribution of the k^{th} order statistic is called the Beta density distribution:

$$P(X_{(k)} \leq t)$$


Not bad.
Let's do a little more.



This is one of the cutest arguments in this business.
Really cute. Super cute.
You have to really use probability on this one.
Watch.

We now have the probability distribution of all the order statistics.
Where the 1st point goes, the 2nd point, etc.
Now we want the gaps between successive points.

If you drop n points on the interval $[0, a]$, there are $n+1$ gaps



It is our duty to compute the probability densities of these gaps.
How far apart are the points?


Here we have the main result.
These gaps are identically distributed. They are not independent, of course.

L_i is the gap between $X_{(i-1)}$ and $X_{(i)}$ ← where $X_{(0)} = 0$
 $X_{(n+1)} = a$

I claim that random variables $L_1, \dots, L_{n+1}$ are identically distributed.

It doesn't seem obvious until I show you the trick,
which is the following:

Instead of dropping n points on the interval $(0, a]$, let's take a circle w/ circumference a and drop, on this circumference, $n+1$ points.



If I drop $n+1$ points at random, then you have $n+1$ gaps on this circumference.

Now I say - in the case of a circumference of the circle, any 2 gaps are identically distributed. If you don't see that, I can not explain it. In fact, I can not prove it.


By the Principle of Sufficient Reason.

So by dropping $n+1$ points on the circumference of a circle, it is obvious that the gaps are i.i.d.

3/18/98.13

Now, instead of dropping n points on the interval $(0, a]$, I do the following:

I drop $n+1$ points on the circumference. Then I see where the first one lands. Then with a scissors, I cut the circle. The result is the same as dropping n points on the interval $(0, a]$.



That's a purely probabilistic argument.

In other words, the experiment of dropping n points is exactly the same as dropping $n+1$ points on the circle.

This gives you an intuitive grasp that the gaps are identically distributed. We will later prove it again, using conditional probability.

This is fairly convincing.

And this leads immediately to a way of computing the expectation.

$$E(L_1) = E(L_2) = \dots = E(L_{n+1}) \leftarrow \text{Because these random variables are identically distributed.}$$

$$\text{Note that: } L_1 + L_2 + \dots + L_{n+1} = a$$

$$E(L_1 + L_2 + \dots + L_{n+1}) = a$$

$$\text{So, } E(L_1) + E(L_2) + \dots + E(L_{n+1}) = a$$

{ Here we anticipate that the expectation of the sum is the sum of the expectations, which we have not yet proved. }

This gives:

$$E(L_1) = \frac{a}{n+1}$$

No computation necessary.

What's more, we even get the expectation of the order statistics.

$$X_{(k)} = L_1 + L_2 + \dots + L_k$$

We easily get the expectation:

$$\begin{aligned} E(X_{(k)}) &= E(L_1 + L_2 + \dots + L_k) \\ &= \underbrace{E(L_1) + E(L_2) + \dots + E(L_k)}_k \text{ and since } E(L_1) = E(L_2) = \dots = E(L_k) \end{aligned}$$

$$E(X_{(k)}) = \frac{ka}{n+1}$$

You can also compute this expectation by evaluating the integral [3/18/98.5]:

$$\begin{aligned} E(X_{(k)}) &= \int_{-\infty}^{\infty} t \underbrace{f(t)}_{\text{we determined the density earlier [3/18/98.8-10]}} dt \\ &= \int_0^a t \binom{n}{k} \frac{k t^{k-1} (a-t)^{n-k}}{a^n} dt \end{aligned}$$

Same result. So we also have another identity:

$$\int_0^a \binom{n}{k} \frac{k t^k (a-t)^{n-k}}{a^n} dt = \frac{ka}{n+1}$$

This is the easy stuff.

The hard stuff is to find the maximum length of a gap.

To solve this, we have to use all the machinery of conditional probability, in the continuous case.



Problem 6b of PS #6 is an open question.
In the Pólya Urn process, no conjugate prior is known for Bayes Theory.

Continuous Random Variables (cont'd)

I want to really rub in order statistics

Order statistics (cont'd)

Let X = continuous random variable

Then:


$$P(a < X \leq b) = \int_a^b \text{dens}(X=t) dt$$

probability that X
lies bt any two
numbers $a + b$

integrating the density gives the probability

density itself doesn't mean anything.

It is not a probability.




This discussion went on for 200 years after the discovery of the calculus.

People would say this is meaningless.
And it is meaningless.

Nowadays there is something called non-standard analysis, whereby you can make $\text{dens}(X=t)$ meaningful if you talk about infinitesimal probability.
It comes at a cost of enlarging the real numbers, adding infinitesimals, which can be done in a very simple way.
It hasn't taken on yet, but eventually, people will catch on.

Meanwhile, we have to get used to the density.
The order statistics we saw comes from the Dirichlet process, where we drop n points on the interval $(0, a]$, by looking at the points in the order in which they appear on the interval



↑ the closest point to 0, after dropping n points, is $X_{(1)}$ (read, "X order 1"),
next is $X_{(2)}$, and so on,

3/20/98.2

The uniform process, together with the order statistics, is the basic statistical gadget we work with.

We take n measurements in a given interval.

We want to really study what probability can say about this basic process, taking n measurements.

We saw, last time, how to compute the density of the order statistics:

$$\text{dens}(X_{(n)} = t)$$

Let's do it again - with less feeling.

We've done this twice already.

Once - by computing the cumulative distribution

Remember, the cumulative distribution is:

$$P(X_{(n)} \leq t) = F_{X_{(n)}}(t)$$

This is a genuine probability

From this, we get the density by taking the derivative of the cumulative distributions:

$$\text{dens}(X_{(n)} = t) = \frac{d}{dt} F_{X_{(n)}}(t)$$

Which, by the Fundamental Theorem of the calculus, is the same as saying:

$$F_{X_{(n)}}(t) = \int_{-\infty}^t \text{dens}(X_{(n)} = s) ds$$

We tend to poo-poo this, but it's highly non-trivial.


This is the whole calculus. 200 years ago, people would shout at each other about this.

We really see the calculus in action when dealing with density.

So, I want to do it again to really rub it in.

Unless you unanimously rise and say "No."





$X_{(k)}$ lies in $(t, t+h]$

Let $u(t) = \text{number of } X_i \text{ in } [0, t)$

Computations w/ the Dirichlet Process are performed by reducing it to the random function that is associated w/ the Dirichlet Process.

event $(t < X_{(k)} \leq t+h) =$

$$\begin{aligned} & (u(t) = k-1) \cap ((u(t+h) = k) \cup (u(t+h) = k+1) \cup \dots \cup (u(t+h) = n)) \\ & \cup (u(t) = k-2) \cap ((u(t+h) = k) \cup (u(t+h) = k+1) \cup \dots \cup (u(t+h) = n)) \\ & \cup \vdots \\ & \cup (u(t) = 0) \cap ((u(t+h) = k) \cup (u(t+h) = k+1) \cup \dots \cup (u(t+h) = n)) \end{aligned}$$


We can simplify by the distributive law:

$$\begin{aligned} & ((u(t) = k-1) \cap (u(t+h) - u(t) = 1)) \cup ((u(t) = k-1) \cap (u(t+h) - u(t) = 2)) \cup \dots \\ & \cup ((u(t) = k-2) \cap (u(t+h) - u(t) = 2)) \cup ((u(t) = k-2) \cap (u(t+h) - u(t) = 3)) \cup \dots \\ & \cup \vdots \\ & \cup ((u(t) = 0) \cap (u(t+h) - u(t) = k)) \cup ((u(t) = 0) \cap (u(t+h) - u(t) = k+1)) \cup \dots \end{aligned}$$

So we have that:

$$\text{event } (t < X_{(k)} \leq t+h) = \underbrace{((u(t) = k-1) \cap (u(t+h) - u(t) = 1))}_{\text{disjoint events that have 2 or more points in } (t, t+h]} \cup \dots$$

(i.e., all but first term (boxed) disjunction above)



Therefore, the probability is:

$$P(t < X_{(k)} \leq t+h) = \binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{h}{a}\right)^1 \left(\frac{a-t-h}{a}\right)^{n-k} + \underbrace{O\left(\left(\frac{h}{a}\right)^2\right)}_{\text{from at least 2 points in } (t, t+h]}$$

In terms of cumulative distributions, we have:

$$P(t < X_{(k)} \leq t+h) = \underbrace{F_{X_{(k)}}(t+h)}_{P(X_{(k)} \leq t+h)} - \underbrace{F_{X_{(k)}}(t)}_{P(X_{(k)} \leq t)}$$

which gives:

$$F_{X_{(k)}}(t+h) - F_{X_{(k)}}(t) = \binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{h}{a}\right) \left(\frac{a-t-h}{a}\right)^{n-k} + O\left(\left(\frac{h}{a}\right)^2\right)$$


divide by h , and take limit as $h \rightarrow 0$:

$$\lim_{h \rightarrow 0} \frac{F_{X_{(k)}}(t+h) - F_{X_{(k)}}(t)}{h} = \binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{1}{a}\right) \left(\frac{a-t}{a}\right)^{n-k} + \underbrace{O(h)}_{\substack{\text{as they said 200 years ago,} \\ \text{you drop higher order} \\ \text{infinitesimals.}}} \rightarrow 0$$

This is precisely the derivative of the cumulative distribution, which is exactly the density.

$$\text{dens}(X_{(k)}=t) = \frac{\binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{1}{a}\right) \left(\frac{a-t}{a}\right)^{n-k}}{\frac{n!}{(k-1)!(n-k)!}}$$

- Note that we could have argued, from the start, that no more than 1 point could fall in $(t, t+h]$, as h is infinitesimal.



From now on, we're going to argue that way.
I'm not going to go through this again.
The whole idea of the calculus is to simplify.

- Plotting the curves of the density, $\text{dens}(X_{(k)}=t)$, tells you how the probability is distributed.

Is the density lumped on one side? Then it's more probable that event occurs on that side. You understand the probability of a continuous random variable by looking at the density, not by looking at the cumulative distribution.

If you plot the density, you see the lumps.


Ex: unimodal density — one lump
bimodal " — two lumps

{ It's very instructive to plot $\text{dens}(X_{(k)}=t)$ to see how various order statistics fit.

- Range $\triangleq X_{(n)} - X_{(1)}$
- Midpoint, when n is odd $\triangleq X_{(\frac{n+1}{2})}$

You can go on forever w/ some very delicate computations of this kind.

- Let's take now two order statistics, $X_{(j)}$ and $X_{(k)}$:



We ask: How are $X_{(j)}$ and $X_{(k)}$ correlated?

These are clearly not independent random variables.
Therefore, we have to develop the analogue for the joint distribution of integer random variables.

Now we need the continuous analogue of the joint distribution.

Then we have to do a complete computation of the correlation of all the order statistics.

All possible ones.

Sorry - that's the way it is.

This requires that we first develop the theory of joint distribution of continuous random variables. This is farraginous.
Don't let yourself be distracted by this. } consisting of a confused mixture.
(see comments at end of [3/20/98.7]) } A half-page

- Joint distribution of 2 continuous random variables

First, the bad news: you have to use partial derivatives, ∂ and multiple integrals, $\iint$ } the definitions involve these.

The good news - in practice, there are all these tricks to avoid them.

Let X, Y be two continuous random variables.

$$\underbrace{P(X \leq t \cap Y \leq s)}_{\text{a genuine, garden variety probability.}} = \int_{-\infty}^s \int_{-\infty}^t \overbrace{\text{dens}(X=u, Y=v)}^{\text{joint density}} du dv$$

joint cumulative distribution

NB: this is not the intersection of 2 sets
These are not events

Random variables $X + Y$ are jointly continuous when their joint cumulative distribution function is obtained by integrating a function of these variables that is called the joint density.

Q/A: everything is non-negative, so the order of integration doesn't matter.

$$\iint \text{dens}(X=u, Y=v) du dv = \iint \text{dens}(X=u, Y=v) dv du$$

$\text{dens}(X=t, Y=s)$ is called the joint density of X and Y

Let's fool around a little bit.

The marginal density $\text{dens}(X=t)$ is obtained by integrating out the Y variable:

$$\text{dens}(X=t) = \int_{-\infty}^{\infty} \text{dens}(X=t, Y=s) ds$$

and similarly:

$$\text{dens}(Y=s) = \int_{-\infty}^{\infty} \text{dens}(X=t, Y=s) dt$$

The probability $P(X \leq t)$ is:

$$P(X \leq t) = \int_{-\infty}^t \int_{-\infty}^{\infty} \text{dens}(X=u, Y=v) dv du$$

the Y variable is integrated out

and then integrating w/ these limits gives the cumulative distribution $P(X \leq t)$.

The basic formula to remember is:

$$P((X \leq t) \cap (Y \leq s)) = \int_{-\infty}^s \int_{-\infty}^t \text{dens}(X=u, Y=v) du dv$$

Rewriting this in a slightly different form:

$$P((a < X \leq b) \cap (c < Y \leq d)) = \int_{t=a}^b \int_{s=c}^d \text{dens}(X=t, Y=s) ds dt$$

useful notation to keep track of limits and variables of integration.

This tells you that any event involving two continuous random variables can be computed, and the probability of any such event can be computed by integrating the joint density.

Lastly, by the Fundamental Theorem of the Calculus, the partial derivative of the joint cumulative distribution is the joint density:

$$\underbrace{\text{dens}(X=t, Y=s)}_{\text{joint density}} = \frac{\partial^2}{\partial t \partial s} \underbrace{P((X \leq t, Y \leq s))}_{\text{joint cumulative distribution}}$$

You can go on forever w/ these formulas and in some statistics books, you have 20 pages doing this stuff w/ $i=2, \dots, n$ random variables.
Your eyes can't see anymore.

It's clear that once you understand this for $i=2$ that you can see what happens for $i=3, \dots$, any number of random variables.

Today : finish Chapt 3
joint distributions/densities
expectation (again)

Wed : in-class review

Fri : Quiz

it doesn't matter if they are
continuous, integer, etc.

Def The joint distribution of any two random variables X, Y is
$$P(X \leq x, Y \leq y)$$


Caution: In the literature, it's a totally standard technique to use
lower-case x, y for the variables
upper-case X, Y for the random variables
this does get people into trouble, as you're using a letter in
two ways.

Def If X, Y are continuous random variables, then their joint density
is:

$$\text{dens}(X=x, Y=y) = f(x, y) = \frac{\partial^2}{\partial x \partial y} P(X \leq x, Y \leq y)$$

derivative of the probability
function

Example - Drop a point uniformly on triangle enclosed by $x \geq 0, y \leq 1, x=y$:




Let continuous r.v. X, Y be the coordinates
of the drop point.

What does it mean to say that a point is dropped uniformly
on this triangle?
It's a statement about what the density function should
look like.

uniformly $\Rightarrow$ $\text{dens}(X=x, Y=y) = c$
implies
density function
is a constant $\hat{=}$ a constant

We want to figure out what this constant c is.

To get this joint density function, we know that if we integrate it over the whole space, we get 1.
The whole space here is over a certain area.
What's the area?



We know that:

$$P(X \leq 1, Y \leq 1) = 1$$

joint distribution

$$P(X \leq 1, Y \leq 1) = \int_0^1 \int_0^y f(x, y) dx dy$$

We've just argued that:

$$f(x, y) = \text{dens}(X=x, Y=y) = c \quad \text{when } 0 \leq x \leq y \leq 1$$


$$= \begin{cases} \int_0^1 \int_0^y c dx dy & \text{when } 0 \leq x \leq y \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

$$= \begin{cases} c * \overbrace{\int_0^1 \int_0^y dx dy}^{\text{area of triangle} = \frac{1}{2}} & \text{when } 0 \leq x \leq y \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

$$P(X \leq 1, Y \leq 1) = \begin{cases} \frac{c}{2} & \text{when } 0 \leq x \leq y \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

This gives:

$$1 = \frac{c}{2} \Rightarrow c = 2$$




$$\text{dens}(X=x, Y=y) = \begin{cases} 2 & 0 \leq x \leq y \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

Once you have the density function, we can get anything we want.
From this, we can compute $P(X \leq x, Y \leq y)$

$$P(X \leq x, Y \leq y)$$

• case 1 - $0 \leq y \leq x \leq 1$




we've just shown this is 2 within the triangle

$$P(X \leq x, Y \leq y) = \iint_{y \leq x} \text{dens}(X=x, Y=y) dx dy$$

To keep our sanity, let's switch the variables of integration:

$$\begin{aligned} x &\rightarrow s \\ y &\rightarrow t \end{aligned}$$




$$= \iint_{t \leq s} 2 ds dt$$

$$= \int_0^y \int_0^t 2 ds dt$$

$$= 2 \int_0^y t dt$$

$$P(X \leq x, Y \leq y) = y^2 \quad \text{when } 0 \leq y \leq x \leq 1$$

• case 2 - $0 \leq x \leq y \leq 1$



2 everywhere w/in triangle


$$P(X \leq x, Y \leq y) = \iint_{y \leq x} \text{dens}(X=x, Y=y) dx dy$$

$$= (2 * \text{area triangle}) + (2 * \text{area rectangle})$$

$$= 2 \left(\frac{y^2}{2} \right) + 2(y-x)x$$

$$P(X \leq x, Y \leq y) = y^2 + 2x(y-x) \quad \text{when } 0 \leq x \leq y \leq 1$$

• case 3 - otherwise



$$P(X \leq x, Y \leq y) = \iint_{yx} \overbrace{\text{dens}(X=x, Y=y)}^0 dx dy$$

$$P(X \leq x, Y \leq y) = 0$$

And we have:

$$P(X \leq x, Y \leq y) = \begin{cases} y^2 + 2x(y-x) & \text{when } 0 \leq x \leq y \leq 1 \\ y^2 & 0 \leq y \leq x \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



• Are the continuous random variables X and Y of this example uniformly distributed?


We look at the marginal density functions, which are obtained from the joint density function.

If the marginal density functions are constants, then X and Y are uniformly distributed.

$$f_X(x) = \int_{-\infty}^{\infty} f(x,y) dy$$

integrate away the y component.

Integrate over all values of y allowed by a particular x .



$$= \int_x^1 2 dy \quad \left\{ \begin{array}{l} \text{we need only concern ourselves w/ limits of integration } [x, 1], \\ \text{because outside of that region the density function is 0.} \\ (f(x,y)=0) \end{array} \right.$$


$$f_X(x) = 2(1-x)$$

X is not uniformly distributed, as $f_X(x)$ is not constant.

Similarly:

$$f_Y(y) = \int_{-\infty}^{\infty} f(x,y) dx$$

$$= \int_0^y 2 dx$$



$[0, y]$ all values of x allowed by a particular y

$$f_Y(y) = 2y$$

Y is not uniformly distributed, as $f_Y(y)$ is not constant.

• Are the continuous random variables X and Y of this example independent?

How can we test for this?

$$\text{Test 1 : } P(X \leq x, Y \leq y) \stackrel{?}{=} P(X \leq x) P(Y \leq y)$$

alternatively:

$$\text{Test 2 : } \text{dens}(X=x, Y=y) \stackrel{?}{=} \text{product of the marginals} \\ f_X(x) f_Y(y)$$

These tests and independence are all equivalent.

In this example:

$$\text{dens}(X=x, Y=y) \stackrel{?}{=} f_X(x) f_Y(y)$$

$$2 \stackrel{?}{=} 2(1-x) 2y$$

↑
No!

Product of marginals does not equal 2.

Therefore X and Y are not independent r.v.

Let's do a fancier example:

3/30/98.6


Joint densities and joint distributions for the order statistics for the uniform process

Uniform process - dropping n points on interval $[0, a]$

$X_{(i)}$ = position of the i^{th} largest point

Recall the density function for the i^{th} order statistic:


$$\text{dens}(X_{(i)} = x) = \binom{n}{i-1, 1, n-i} \left(\frac{x}{a}\right)^{i-1} \left(\frac{1}{a}\right) \left(\frac{a-x}{a}\right)^{n-i}$$



$$\text{dens}(X_{(i)} = x) = n \binom{n-1}{i-1} \left(\frac{x}{a}\right)^{i-1} \left(\frac{1}{a}\right) \left(\frac{a-x}{a}\right)^{n-i} \quad \text{where } 0 \leq x \leq a$$

Let $X_{(j)}$ and $X_{(k)}$ be random variables such that $j < k$.
We want to compute the joint distribution:

$$P(X_{(j)} \leq x, X_{(k)} \leq y) = \binom{n}{j, k-j, n-k} \left(\frac{x}{a}\right)^j \left(\frac{y-x}{a}\right)^{k-j} \left(\frac{a-y}{a}\right)^{n-k}$$



exactly j fall in $(0, x]$
exactly $k-j$ fall in $(x, y]$
exactly $n-k$ fall in $(y, a]$

+ ...

all terms where more than j fall in $(0, x]$
and where more than k fall in $(0, y]$

So $P(X_{(j)} \leq x, X_{(k)} \leq y)$ is a big expression, with lots of terms.

Recall that the density is:


$$\text{dens}(X_{(j)} = x, X_{(k)} = y) = \frac{\partial}{\partial x} \frac{\partial}{\partial y} P(X_{(j)} \leq x, X_{(k)} \leq y)$$

to get the density, we take the derivative of the probability.

3/30/98.7

It turns out that when we take the derivative that the ... terms for $P(X_{(j)} \leq x, X_{(k)} \leq y)$ are zero. So we could simply take the derivative.

Alternatively, what's the probability of dropping points with arbitrarily small intervals $[x, x+h]$, $[y, y+g]$?



Let $A_{h,g}$ = event that $X_{(j)}$ (i.e., the j^{th} order statistic) falls in $[x, x+h]$ and $X_{(k)}$ falls in $[y, y+g]$

$$\text{dens}(X_{(j)}=x, X_{(k)}=y) = \lim_{h \rightarrow 0} \lim_{g \rightarrow 0} \frac{P(A_{h,g})}{hg}$$

probability of event $A_{h,g}$ is not hard to figure out, as we have exact occupation numbers in each interval.

$$= \lim_{h \rightarrow 0} \lim_{g \rightarrow 0} \frac{\binom{n}{j-1, 1, k-j-1, 1, n-k} \left(\frac{x}{a}\right)^{j-1} \left(\frac{h}{a}\right)^1 \left(\frac{y-x-h}{a}\right)^{k-j-1} \left(\frac{g}{a}\right)^1 \left(\frac{a-y-g}{a}\right)^{n-k}}{hg}$$

$$\text{dens}(X_{(j)}=x, X_{(k)}=y) = \begin{cases} \binom{n}{j-1, 1, k-j-1, 1, n-k} \left(\frac{x}{a}\right)^{j-1} \left(\frac{1}{a}\right)^1 \left(\frac{y-x}{a}\right)^{k-j-1} \left(\frac{1}{a}\right)^1 \left(\frac{a-y}{a}\right)^{n-k} & \text{when } 0 \leq x \leq y \leq a \\ 0 & \text{otherwise} \end{cases}$$

Warning:

Note that $\text{dens}(X_{(j)}=x, X_{(k)}=y)$ looks a lot like $P(X_{(j)} \leq x, X_{(k)} \leq y)$ for this problem. Do not get confused by this. Density function and probability distribution function are different beasts.


Remember: Density can be bigger than 1, unlike probability distribution function. There's no upper bound on density function. (density does have to be non-negative)



What if we want the joint density for all of the order statistics?

What is $\text{dens}(X_{(1)}=x_1, \dots, X_{(n)}=x_n)$?

Divide the interval $[0, a]$ into intervals which include arbitrarily small intervals about the order statistics:



$$\text{dens}(X_{(1)}=x_1, \dots, X_{(n)}=x_n) = \lim_{h_1, h_2, \dots, h_n \rightarrow 0}$$

$$\frac{P\left(\bigcap_{i=1}^n (x_i \leq X_{(i)} \leq x_i + h_i)\right)}{h_1 \cdot h_2 \cdot \dots \cdot h_n}$$

$$= \lim_{h_1, h_2, \dots, h_n \rightarrow 0} \frac{\binom{n}{1, 1, \dots, 1} \prod_{i=1}^n \left(\frac{h_i}{a}\right)}{\prod_{i=1}^n h_i}$$

$$\binom{n}{1, 1, \dots, 1} = \frac{n!}{1! 1! \dots 1!} = n!$$

$$\text{dens}(X_{(1)}=x_1, \dots, X_{(n)}=x_n) = \frac{n!}{a^n} \quad \text{when } 0 < x_1 < x_2 < \dots < x_n < a$$

The joint density function for all the order statistics.

There are whole books on order statistics.

For example, there are a lot of nice tests for determining whether the random number generator on a computer actually provides random numbers. You can do this by checking the order statistics.

If the distribution for the order statistics of your random number generator doesn't come out right, perhaps you have a bug in your codes.

- Integration of the density function over the whole sample space gives 1.

$$\iint \dots \int \text{dens}(X_{(1)}=x_1, \dots, X_{(n)}=x_n) dx_1 dx_2 \dots dx_n = 1$$

$$\int_0^a \int_0^{x_{n-1}} \dots \int_0^{x_2} \frac{n!}{a^n} dx_1 dx_2 \dots dx_n = 1$$



the reason the integration limits are like this is because, for example,
 the 1st order statistic lies between 0 and x_1
 the 2nd " " " " 0 and x_2
 etc.

which gives the cool integral identity (try proving this directly from calculus):

$$\int_0^a \int_0^{x_{n-1}} \dots \int_0^{x_2} dx_1 dx_2 \dots dx_n = \frac{a^n}{n!}, \quad 0 < x_1 < x_2 < \dots < x_n < a$$

Expectation

For a continuous random variable, we have:

$$E(X) = \int_{-\infty}^{\infty} x \text{dens}(X=x) dx$$

Back to our example [3/30/98.4]:

given the density function: $f_Y(t) = 2t$, $0 \leq t \leq 1$
 the expected value of the random variable Y is:

$$E(Y) = \int_{-\infty}^{\infty} t f_Y(t) dt$$

$$= \int_0^1 t (2t) dt$$

$$= \int_0^1 2t^2 dt$$

$$= \left. \frac{2t^3}{3} \right|_0^1 = \frac{2}{3}$$

$$E(Y) = \frac{2}{3}$$

- The main trick w/ expectations is the following basic fact:

For any continuous random variables X, Y ,

$$E(X+Y) = E(X) + E(Y)$$

Proof

The proof practically jumps right out from the definition:

$$E(X+Y) = \iint (x+y) \underbrace{f(x,y)}_{\text{joint density}} dx dy$$

By linearity of integrals, this can be separated into:

$$= \iint x f(x,y) dx dy + \iint y f(x,y) dx dy$$

$$= \int x \underbrace{\int f(x,y) dy}_{f_x(x)} dx + \int y \underbrace{\int f(x,y) dx}_{f_y(y)} dy$$

$f_x(x) \leftarrow \text{marginals} \rightarrow f_y(y)$

$$= \underbrace{\int x f_x(x) dx}_{E(X)} + \underbrace{\int y f_y(y) dy}_{E(Y)}$$

$$E(X+Y) = E(X) + E(Y) \quad \checkmark$$

- As a nice example of this:
Expectation of order statistics

Anytime someone says "expectation," think Basic Fact $\Rightarrow E(X+Y) = E(X) + E(Y)$
How can I break it up into the expectation of a sum of random variables?
This should be a knee jerk reaction.

We want $E(X_{(i)})$

The way we do this is to look at the gaps

$$L_1 = X_{(1)}$$

the first
gap

$$L_i = X_{(i)} - X_{(i-1)}$$

} gaps in a uniform process

If we can figure out the expectation for the gaps $E(L_i)$, then we can clearly figure out $E(X_{(i)})$

$$E(L_1 + \dots + L_{n+1}) = a$$

← The expectation of the sum of all the gaps — we already know this, as the sum of all the gaps must cover the whole interval $[0, a]$.

Also, by the basic fact of expectations:

$$E(L_1 + \dots + L_{n+1}) = E(L_1) + \dots + E(L_{n+1})$$

Previously ([3/18/98.12-13]), we argued that the gaps are all identically distributed (equidistributed).



The intuitive circle argument.

$$E(L_1 + \dots + L_{n+1}) = (n+1)E(L_1)$$

$$\text{Thus: } E(L_1) = \frac{a}{n+1}$$

From the definition of gap, we have: $L_i = X_{(i)} - X_{(i-1)}$

$$\Rightarrow X_{(i)} = L_i + X_{(i-1)}$$

Successive applications of gap definitions gives:

$$\begin{aligned} X_{(i)} &= L_i + \underbrace{X_{(i-1)}}_{L_{i-1} + \underbrace{X_{(i-2)}}_{L_{i-2} + \underbrace{X_{(i-3)}}_{\vdots}} + L_1 \end{aligned}$$

$$X_{(i)} = L_1 + L_2 + \dots + L_i$$

Taking expectations:

$$\begin{aligned} E(X_{(i)}) &= E(L_1 + L_2 + \dots + L_i) \\ &= E(L_1) + E(L_2) + \dots + E(L_i) \end{aligned} \quad \left. \vphantom{\begin{aligned} E(X_{(i)}) &= E(L_1 + L_2 + \dots + L_i) \\ &= E(L_1) + E(L_2) + \dots + E(L_i) \end{aligned}} \right\} \text{from basic fact of expectations}$$

We just showed that:

$$E(L_1) = E(L_2) = \dots = E(L_i) = \frac{q}{n+1}$$

$$= i E(L_1)$$

$$E(X_{(i)}) = i \frac{q}{n+1}$$

John Guidi
guidi@mit.edu

18.313

Lecture 21 (Sara Billey)
4/1/99.1

I did not attend this lecture and no tape of the lecture is available.

The lecture was an in-class review of the material to be tested by the Quiz on Friday.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



Within the text, pagination is of the form *lecture_date.page*. Note that a different convention is used in both the table of contents and the index, where pagination is of the form *lecture_number.page*. The table of contents provides the mapping between *lecture_number* and *lecture_date*. For example, page 1.1 in the table of contents and index corresponds to page 2/4/98.1 in the text, where the topic anomaly detection is discussed.

Index

- Abel polynomials, SC2.8
- Alternatives
 - law of, 11.8, 13.11, 14.1
 - law of (continuous analogue), 23.5
- Anomaly detection, 1.1
- Average measurements, 33.15
- Average number of blips, 29.6
- Axioms of probability, 1.3
- Balls into boxes, 3.11, 4.10, SC2.3
 - forest of rooted trees, SC2.8
- Bayes' estimate, 25.4
 - using conjugate prior, 25.7
- Bayes' law, 12.5, 13.11, 14.8, 15.9
 - (continuous analogue), 24.10, 25.1
 - densities, 33.18
- Bayes' theorem
 - sampling without Replacement, 16.5
- Bell numbers, SC2.9
- Bernoulli process, 1.7, 2.6, 28.1, 28.2
 - conditional probability, 11.5, 13.4, 14.2
 - density plots, 24.13
 - random variables, 8.7, 8.8, 8.10, 9.3
 - sample space, 1.7
- Bernoulli run, infinitely often, 2.11, 3.2
- Beta function, 25.1
 - discrete, AN16.7, AN16.8
 - integral, 18.11, 24.11, 24.12, 25.2
- Binomial
 - coefficients, 4.2, 6.5
 - distribution, 8.9
 - identity, 6.5, 6.6, 6.10, 7.7, 10.5, 10.11, 10.13, 10.15, 15.13, 16.1, 18.14
 - theorem, 4.4, 4.8, 8.9, 9.8
- Bits, SC4.2
- Blips, 28.12, 29.1
 - colored, 30.17, 31.1
- Blocks, 11.8
- Borel sets, 17.5
- Borel-Cantelli lemma, 3.2
- Bose-Einstein statistics, 5.7, 16.4
- Buffon needle problem, SC3.2, SC3.8
- Cauchy's functional equation, 26.14, 28.5
- Central limit theorem, 34.11
- Circuit theoretic interpretation
 - Boolean algebra, SC1.7
- Cluster analysis, 1.1
- Coin
 - fair, 34.9


- tossing, 1.7
 - continuous, 29.8
 - gaps, 9.3
- Computing
 - π , SC3.6
 - perimeters, SC3.12
- Conditional probability, 11.1, 11.7
 - continuous, 22.1
 - density, 22.11
 - law of alternatives, 14.1
- Conditioning event of probability
 - zero, 23.1
- Confidence intervals, 24.15, 25.3
- Continuity property, 2.3
- Convex set, SC3.7
- Convolution, 26.8
- Countable additivity, 2.2
- Cumulative distribution, 17.2, 18.2
- DeFinetti's theorem, AN16.10
- Density, 17.3, 18.1
 - algebra, 26.1
 - function
 - using inverse function, 25.13
 - joint, independent random variables, 26.7
 - plots, 19.4, 24.13
- Difference operator, 4.3, SC2.11
 - backwards, 4.4
- Dirichlet distribution, 18.5
- Dirichlet process, 17.5, 22.2
 - random function, 19.3
- Disjoint intervals
 - Poisson process, 29.4
- Dispositions, 4.11, 5.3, SC2.8
- Distribution interpretation, 15.6
 - Maxwell-Boltzmann, 4.6, 6.1, 15.4
- Dropping points
 - on a circle, 18.12
 - on a triangle, 20.1
 - on an interval, 23.1
- Entropy
 - independent partitions, SC4.3
 - partitions, SC4.2, SC4.5
- Event, 1.3, 2.1
 - not certain, 3.10
 - not impossible, 3.10
- Expectation, 10.8, 15.1
 - random variable plus constant, 27.4
 - sum of random variables, 10.8
- Exponential density, 29.10
- Exponential distribution, 28.8, 28.15
- Factorial
 - lower, 4.3
 - rising, 4.4, 5.4
- Fair coin, 24.7
- Fermi-Dirac statistics, 6.4, 9.10, 15.12, 16.1
- Flags on poles, 4.11
- Gamma density, 29.12
- Gamma distribution, 29.11
- Gaps
 - coin tossing, 9.3
 - order statistics, 18.12
 - uniform process, 23.8
- Gaussian distribution, 35.1
- Geometric distribution, 8.12
- Hard spheres, 1.1
- Hypergeometric distribution, 9.11, 14.9
 - multivariate, 10.5
- Identically distributed, 8.3
- Inclusion-exclusion principle, 6.7, 7.1, 30.15, SC1.1

- Independence of
 - disjoint intervals, 29.5
 - disjoint intervals (Poisson process), 28.14
 - events, 1.5, 2.4
 - random variables, 8.4, 17.2
- Independent and identically distributed, 9.2
- Indicator random variables, 26.9
- Information Search, SC4.7
- Inspector's paradox, 27.11
- Intensity, 29.2, 31.3
- Joint
 - density, 19.6, 26.3
 - sum, 26.10
 - distribution, 9.13, 10.1
- Joint cumulative distribution
 - continuous random variables, 19.5
- Kolmogorov zero-one law, 3.8
- Laplace law of succession, 25.5
- Laplace transform, 31.4
- Large numbers
 - law of, 35.12
- Likelihood, 14.11, 33.17
- Marginal
 - density, 19.6, 20.4
 - distribution, 10.1
- Matching, 7.1
- Maximum entropy principle, SC4.6
- Maxwell-Boltzmann statistics, 4.1, 5.1, 8.7
 - conditional probability, 11.4
- Maxwell-Einstein
 - normal distribution derivation, 35.4
- Measure, 34.6, SC3.9
 - lines in the plane, SC3.11
- Measurements
 - assumptions, 33.16
 - average, 33.15
- Median, 18.3
- Memoryless property, 28.4
- Memorylessness
 - waiting time, 29.10
- Multinomial coefficient, 4.4, 4.10
- Multiset, 5.10
 - coefficients, 4.4, 5.5, 5.8, 6.5
 - identity, 16.4
- Needles on a stick, 24.1, 30.6, SC3.1, SC3.10
- Negative binomial distribution, 8.13
- Normal distribution, 33.14, 35.1
 - theory of, 28.1
- Normalize, 35.2
- Normalized random variable, 33.8
- Null probability, 9.1
- Occupancy interpretation, 15.6
 - Maxwell-Boltzmann, 4.6, 6.1, 15.5
- Occupation numbers, 30.10, 30.13
 - Bose-Einstein, 6.5, 7.6
 - Dirichlet process, 17.7
 - expectation, 10.10
 - Maxwell-Boltzmann, 4.7, 7.4
- Order statistics, 17.9, 18.7, 19.1, 20.6
 - conditional probability, 22.7
 - density, 18.8
 - expectation, 18.14, 20.11
 - joint density, 22.5
 - joint distribution, 22.4
- Origin
 - time to first return, 31.13
 - time to return, 31.7

- Pólya urn model, SC2.3
- Partitions, 11.8, SC2.9, SC4.1
 - algebra of, SC1.10
 - entropy, SC4.2, SC4.5
 - integer, SC2.5
- Perimeters, SC3.13
- Permutation, 4.2
- Pointless probability, 8.2, 26.9
- Poisson distribution, 29.4
- Poisson process, 28.1
 - events, 28.11
 - fundamental properties, 29.6
 - probability, 28.13, 29.2
 - sample space, 28.9
- Posterior, 14.11, AN16.4, 25.6
 - in statistics, 34.7
- Prior, 14.11
 - conjugate, 16.6, AN16.2, AN16.8, 24.16, 25.6
 - in statistics, 34.5
 - uniform, 15.11, 16.6, 16.10, AN16.2
 - uniform(continuous), 24.11
- Probability, 2.2
 - axioms, 1.3
 - distribution, 8.1
 - space, 1.4
- Probability distributions
 - algebra of, 25.9
- Problem of scales, SC4.9
- Quantum probability, 9.14
- Rényi's principle, 6.9, 7.8, SC1.5
- Rademacher's theorem, 3.9
- Random function, 17.9, 29.4, 35.10
- Random variable
 - continuous, 17.1
 - indicator, 26.9, SC1.1
 - integer, 8.1
 - interpretation of, 32.7
 - normalized, 33.8
 - standard normal, 32.5
 - standardized, 33.8
- Random walk, 31.5
 - as Bernoulli Process, 31.7
 - continuous analogue, 31.13
 - self avoiding, 1.2
 - symmetric, 32.1
- Rare events, 28.10
 - law of, 29.6
- Record values, 27.1
- Reluctant functions, SC2.7
- Run, 2.7
- Runners on a track, 27.8
- Sample
 - point, 1.7, 14.11
 - space, 1.3, 2.1
 - Poisson process, 28.9
 - space (infinite), 15.3
- Sampling
 - with replacement, 5.10, 9.6, 9.8, AN16.6
 - with replacement (infinite), 9.12
 - without replacement, 9.7, 9.9, AN16.1
- Schrödinger randomization, 30.9
- Scientific induction, 14.8
- Shannon coding theorem, SC4.12
- Standard deviation, 33.2
 - Bernoulli process, 33.3
- Standard normal
 - distribution, 32.2
 - random variable, 32.5, 33.10
 - random variable density, 34.1
 - table, 34.8
- Standardize, 35.2
- Standardized random variable, 33.8
- Statistics
 - assumptions, 33.15, 33.18, 34.3

- basic rule, 34.2
- Stirling number's of the second kind, SC2.11
- Stochastic process, 10.2
- Successive conditioning, 11.11, 12.5
- Successive probability, 11.7
- Sylvester's theorem, SC3.7
- Tail event, 3.8
- Test Problem, 12.6
- Tree
 - probability, 12.3, 13.5, SC2.1
- Uniform distribution, 18.5
- Uniform process, 17.5, 28.1
 - by conditioning Poisson process, 30.1
 - limit, 29.9
- Urn models, SC2.2
- Urn sampling
 - conditional probability, 13.1, 14.8
 - with replacement, 12.9
- Vandermonde's identity, 10.7
- Variance
 - addition formula, 32.8
 - Bernoulli process, 33.3
 - exponential random variable, 33.6
 - Poisson process, 33.5
 - properties of, 33.2
 - random variable, 32.6
- Waiting time, 14.3
 - Bernoulli process, 8.10, 10.14
 - memorylessness, 28.4
- Whitworth problems, 18.6
- Wiener process, 35.7
- z transform, 31.10
- Zero-one law, 3.8

Probability



John N. Guidi
 Lecture Notes - Spring 1998
 MIT Course 18.313
 Professor Gian-Carlo Rota
 Volume 2 of 2
 (Lectures 22–35, Super Classes 1–4)

John N. Guidi
Visiting Scholar in Applied Mathematics
Department of Mathematics
Room 2-392
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
Email: guidi@mit.edu

Copyright ©1998 John N. Guidi. All rights reserved.

These *Lecture Notes* originated from the lectures presented by Gian-Carlo Rota, Professor of Applied Mathematics, for course 18.313 - Probability, which he taught at MIT, during the Spring 1998 semester. The *Lecture Notes* were produced from notes I made during class, audio recordings I made of lectures, as well as clarifications and expansions I made of the material presented, after the fact. Although no attempt was made to provide a literal transcription of the lectures, an effort was made to present the material here in the same spirit and I hope to have introduced only a limited number of errors. The *Lecture Notes* are comprised of a two volume set. Volume 1 contains lectures 1-22, volume 2 contains lectures 22-35 and super class lectures 1-4.

Sara Billey, Assistant Professor of Applied Mathematics, who was the course director and whose recitations I attended, presented lectures 20, 21, and 27. The other recitation instructors were Alex Perlin and Ioana Popescu. Their efforts clearly contributed to the success of this course.

I was absent from lecture 25 and my thanks go to Jeff Lieberman, who taped the lecture for me and provided his notes, as well as to Ira Gerhardt and Rick Monte, who also provided their notes of this class.

During the course, I found that preparation of these *Lecture Notes* was a particularly useful way to really understand the material (or, as Professor Rota is fond of saying, "*to really rub it in*"). My gratitude goes to him for imparting his enthusiasm and sharing his profound knowledge of mathematics, as well as for presenting these superb lectures on probability.

John N. Guidi
October 15, 1998

Within the text, pagination is of the form *lecture_date.page*. Note that a different convention is used in both the table of contents and the index, where pagination is of the form *lecture_number.page*. The table of contents provides the mapping between *lecture_number* and *lecture_date*. For example, page 1.1 in the table of contents and index corresponds to page 2/4/98.1 in the text, where the topic anomaly detection is discussed.

Contents

Research Problems

Volume 1

Lecture 1 - [2/4/98]

1. Unsolved Problems of Probability	1.1
1.1 Cluster Analysis	1.1
1.2 Anomaly Detection	1.1
1.3 Hard Spheres	1.1
1.4 Self Avoiding Random Walk	1.2
2. Sample Space	1.3
2.1 Events	1.3
3. Probability	1.3
4. Consequences of the Axioms of Probability	1.4
5. Independence	1.5
6. Example - Coin Tossing (Bernoulli Process)	1.7

Lecture 2 - [2/6/98]

1. Sample Spaces - The Bernoulli Process (cont'd)	2.1
2. Review	2.1
3. Continuity Property of Probability	2.3
4. Independence - The Fundamental Notion of Probability	2.4
5. Bernoulli Process	2.6
5.1 Sample Points	2.6
5.2 Events	2.6
5.3 Probability	2.7
6. Example - Probability that a Run of n Heads ever occurs	2.9

Lecture 3 - [2/9/98]

1. The Bernoulli Process (cont'd)	3.1
2. Borel-Cantelli Lemma	3.2
3. Tail Event	3.8
4. Kolmogorov Zero-One Law	3.8
4.1 Example	3.9
5. Philosophical Interpretation of Probability	3.10

Lecture 4 - [2/11/98]

1. The Theory of Distribution and Occupancy (beg'g)	4.1
2. Maxwell-Boltzmann Sample Space	4.1
3. Quantities	4.2
3.1 Number of Permutations of a Set	4.2
3.2 Binomial Coefficient	4.2
3.3 Lower Factorial	4.3
3.4 Difference Operator	4.3
3.5 Rising Factorial	4.4
3.6 Backwards Difference Operator	4.4
3.7 Multiset Coefficients	4.4
3.8 Multinomial Coefficient	4.4
4. Binomial Theorem	4.4
5. Examples - Maxwell-Boltzmann Statistics	4.5
6. Distribution Interpretation	4.6
7. Occupancy Interpretation	4.6
8. Occupation Numbers	4.7
9. Proof of Binomial Theorem	4.7
10. Probability of Finding all Occupation Numbers	4.9
10.1 Multinomial Coefficient	4.9
11. Balls into Boxes	4.10
12. Dispositions	4.11

Lecture 5 - [2/13/98]

1. The Theory of Distribution and Occupancy (cont'd)	5.1
2. Maxwell-Boltzmann Statistics	5.1
2.1 Occupancy Interpretation	5.2
3. Occupation Number - Maxwell-Boltzmann Statistics	5.2
4. Dispositions	5.3
4.1 Occupation numbers!dispositions	5.4

4.2	Binomial identity	5.4
4.3	Probability Independent of Given Occupation Numbers . . .	5.5
5.	Bose-Einstein Statistics	5.7
5.1	Number of Sample Points - Proof 1	5.7
5.2	Number of Sample Points - Proof 2	5.8
5.3	Number of Sample Points - Proof 3	5.8
5.4	Interpretation	5.10
5.5	Sampling with Replacement	5.10
5.6	Multiset	5.10

Lecture 6 - [2/17/98]

1.	The Theory of Distribution and Occupancy (cont'd)	6.1
1.1	Maxwell-Boltzmann Statistics	6.1
1.2	Dispositions	6.1
1.3	Bose-Einstein Sample Space	6.2
2.	Fermi-Dirac Sample Space	6.4
3.	Probabilistic Identities	6.5
3.1	Binomial Coefficients	6.5
3.2	Multiset Coefficients	6.5
4.	Inclusion-Exclusion Principle	6.7
5.	Rényi's Principle	6.9
5.1	An Application	6.9
6.	Matching - Application of Rényi's Principle	6.12

Lecture 7 - [2/18/98]

1.	Inclusion/Exclusion Principle (cont'd)	7.1
2.	Example	7.1
2.1	Matching	7.1
2.2	All Boxes Occupied in Maxwell-Boltzmann Statistics	7.4
2.3	At Least One Box Occupied in Bose-Einstein Statistics . . .	7.6
3.	Probability that Exactly One Event Happens	7.8

Lecture 8 - [2/20/98]

1.	Random Variables (Integer)	8.1
2.	Probability Distribution	8.1
3.	Identically Distributed Random Variables	8.2
4.	Independent Random Variables	8.4

5. Maxwell-Boltzmann Statistics	8.5
5.1 Occupancy Numbers Not Independent	8.5
5.2 Position Random Variables	8.6
6. Abstract Definition of Maxwell-Boltzmann Statistics	8.7
7. Famous Random Variables in the Bernoulli Process	8.7
8. Outcome of n^{th} Toss	8.7
9. Number of Heads in First n Tosses	8.8
9.1 Proof of the Binomial Theorem	8.8
10. Waiting Time for First Head	8.9
10.1 Geometric Distribution	8.12
11. Waiting Time for k^{th} Head	8.12
11.1 Negative Binomial Distribution	8.13

Lecture 9 - [2/23/98]

1. Random Variables (Integer) (cont'd)	9.1
2. Event of Null Probability	9.1
3. Review	9.2
4. Gaps between Successive Heads	9.3
5. Sampling	9.6
5.1 Sampling with Replacement	9.6
5.2 Sampling without Replacement	9.7
6. Sampling with Replacement	9.8
6.1 Binomial Theorem	9.8
7. Sampling without Replacement	9.9
7.1 The Wrong Way - with Fermi-Dirac Statistics	9.10
7.2 Hypergeometric Distribution	9.11
8. Infinite Sampling with Replacement	9.12
9. Joint Distribution	9.13

Lecture 10 - [2/25/98]

1. Random Variables: Joint Distribution and Expectation	10.1
2. Stochastic Process	10.2
3. Joint Distribution	10.3
3.1 Occupation Numbers in Maxwell-Boltzmann	10.3
3.2 Combinatorial Identity	10.5
3.3 Multivariate Hypergeometric Distribution	10.5
3.4 Vandermonde's Identity	10.7
4. Expectation - a Fundamental Concept	10.8

5. Additivity of Expectations	10.8
6. Expectation	10.10
6.1 Maxwell-Boltzmann	10.10
6.2 Multivariate Hypergeometric Distribution	10.12
6.3 Bernoulli	10.14

Lecture 11 - [2/27/98]

1. Conditional Probability	11.1
2. Conditional Probability is a Probability	11.2
3. Example	11.4
3.1 Maxwell-Boltzmann Statistics	11.4
3.2 Bernoulli Process	11.6
4. Basic Techniques for Working with Conditional Probabilities	11.7
4.1 Rule of Successive Probabilities	11.7
4.2 Law of Alternatives	11.8
4.3 Law of Successive Conditioning	11.11

Lecture 12 - [3/2/98]

1. Conditional Probability (cont'd)	12.1
2. Review	12.1
3. Sample Spaces on Probability Trees	12.3
3.1 Vertices for Events	12.4
3.2 Edges for Conditional Probabilities	12.4
4. Bayesian Theory	12.5
5. Classical Example - The Test Problem	12.6
6. Probability Tree - Maxwell-Boltzmann with Replacement	12.9

Lecture 13 - [3/6/98]

1. Conditional Probability (cont'd)	13.1
2. Motivating Examples	13.1
3. Probability Tree - Bernoulli Process	13.4
3.1 Examples	13.4

Lecture 14 - [3/9/98]

1. Law of Alternatives for Conditional Probability	14.1
1.1 Examples - Bernoulli Process	14.2

2. Problem of Scientific Induction	14.8
3. Bayes' Law	14.8

Lecture 15 - [3/11/98]

1. Digressions	15.1
1.1 Expectation with Finite and Infinite Sample Spaces	15.1
1.2 Occupancy -vs- Distribution Perspectives	15.4
1.3 Bayes' Law	15.9
2. Bayes' Law	15.10
3. Uniform Prior	15.11
3.1 Fermi-Dirac Statistics	15.11
3.2 Binomial Identity	15.13

Lecture 16 - [3/13/98]

1. Bayesian Theory (cont'd)	16.1
2. Binomial Identity	16.1
2.1 Fermi-Dirac Position of First Check	16.1
2.2 Fermi-Dirac Position of Second Check	16.3
3. Bayes' Theorem for Sampling without Replacement	16.5
3.1 Uniform Prior	16.6
3.2 Conjugate Prior	16.6
4. Bayes' theorem!sampling with replacement	16.9
4.1 Uniform Prior	16.10

After Notes Lecture AN16 - [3/16/98]

1. Conjugate Priors	AN16.1
2. Sampling without Replacement	AN16.1
2.1 Uniform Prior	AN16.2
2.2 Conjugate Prior	AN16.2
3. Posterior's for Given Conjugate Prior's	AN16.4
4. Sampling with Replacement	AN16.6
4.1 Uniform Prior	AN16.7
4.2 Conjugate Prior	AN16.8
4.3 Discrete Beta Function Identity Series	AN16.8

Lecture 17 - [3/16/98]

1. Continuous Random Variables	17.1
1.1 General Definition of a Random Variable	17.1
1.2 Independence of Random Variables	17.2
1.3 Cumulative Distribution Function	17.2
1.4 Density of Continuous Random Variable	17.3
2. Example - The Uniform (Dirichlet) Process	17.5
3. Given Occupation Numbers in Uniform Process Intervals	17.7
4. Order Statistics	17.9
4.1 Cumulative Distribution	17.10
4.2 Density	17.11
5. Continuous Random Variable Taking a Single Value	17.12

Lecture 18 - [3/18/98]

1. Continuous Random Variables (cont'd)	18.1
2. Characterization of Density	18.1
3. Cumulative Distribution	18.2
4. Median of a Random Variable	18.3
5. Expectation of a Continuous Random Variable	18.5
6. Example - Uniform (Dirichlet) Distribution	18.5
6.1 Density - the Hard Way	18.8
6.2 Density - the Easy Way	18.9
6.3 Beta Function Integral	18.11
7. Gaps between Points Dropped on a Line	18.12
7.1 A Purely Probabilistic Argument	18.12
8. Expectation of the Order Statistics	18.14
8.1 Binomial Identity	18.14

Lecture 19 - [3/20/98]

1. Order Statistics	19.1
2. The k^{th} Order Statistic of the Dirichlet Process	19.3
2.1 Events	19.3
2.2 Probability	19.3
2.3 Density Function	19.4
3. Density Curves	19.4
4. Joint Distribution of 2 Continuous Random Variables	19.5
5. Marginal Density	19.6
6. Join Density as Derivative of Join Cumulative Distribution	19.7

Lecture 20 (Sara Billey) - [3/30/98]

1. Join Distribution	20.1
2. Join Density	20.1
3. Example - Dropping Points Uniformly on a Triangle	20.1
4. Joint Density and Joint Distribution	20.6
4.1 Order Statistics of the Uniform Process	20.6
5. Expectation	20.9
6. Expectation of Order Statistics	20.11

Lecture 21 (Sara Billey) - [4/1/98]

1. In Class Review (no notes)	21.1

Volume 2

Lecture 22 - [4/6/98]

1. Continuous Conditional Probability	22.1
2. Dirichlet Process	22.2
2.1 Cumulative Distribution	22.3
2.2 Density	22.3
3. Example - Probability of Events with Order Statistics	22.4
3.1 Integrating the Joint Density	22.5
4. Conditional Probability	22.7
4.1 By the Book	22.7
4.2 By Intuition	22.7
5. Example - Using Inclusion-Exclusion Principle	22.9
6. Conditional Density	22.11

Lecture 23 - [4/8/98]

1. Dirichlet Process - Dropping Two Points	23.1
1.1 Intuition	23.1
1.2 Rigorous Computation	23.2
2. Dirichlet Process - Dropping Arbitrary Number of Points	23.3
2.1 Intuition	23.3
2.2 Rigorous Computation	23.4
3. Continuous Version of the Law of Alternatives	23.5
3.1 Proof	23.5

4. Gaps in the Uniform Process	23.8

Lecture 24 - [4/10/98]

1. Continuous Analogue of the Law of Alternatives	24.1
2. Example - Needles on a Stick	24.1
3. Example - Bernoulli Process: Is the Coin Fair?	24.7
3.1 Continuous Analogue of Bayes' Law	24.8
3.2 Uniform Prior gives Beta Function Integral	24.11
3.3 Density Plots	24.13
4. Confidence Intervals	24.14

Lecture 25 - [4/13/98]

1. Beta Function	25.1
2. Continuous Law of Alternatives	25.1
3. Continuous Analogue of Bayes' Law	25.1
4. Use of Uniform Density	25.2
5. Confidence Intervals	25.3
6. Bayes' Estimate	25.4
7. Application - the Laplace Law of Succession	25.5
8. Conjugate Prior	25.6
9. Bayes' Estimate using Conjugate Prior	25.7
10. The Algebra of Probability Distributions	25.9
11. Density of a Function in Terms of Inverse Function	25.13
12. Complicated Densities	25.14

Lecture 26 - [4/15/98]

1. The Algebra of Probability Densities (cont'd)	26.1
2. Densities	26.1
3. Joint Densities	26.3
3.1 Identity Analogous to Conditional Probability	26.3
4. Joint Density with Complicated Random Variables	26.5
5. Joint Density of Sum of Random Variables	26.6
6. Convolution - Joint Density of Sum of Independent Random Variables	26.7
7. Joint Density of Independent, Uniformly Distributed Random Variables	26.9
7.1 Indicator Random Variables	26.9

8. Cauchy's Functional Equation	26.14

Lecture 27 (Sara Billey) - [4/17/98]

1. Examples	27.1
1.1 Record Values	27.1
1.2 Expectation of Random Variable Plus a Constant	27.4
1.3 Runners on a Track	27.8
1.4 Inspector's Paradox	27.11

Lecture 28 - [4/22/98]

1. The Poisson Process	28.1
2. Back to the Bernoulli Process	28.2
3. Continuous Waiting Time	28.4
3.1 Memoryless Property	28.4
3.2 Cauchy's Functional Formula	28.5
3.3 Density of a Memoryless Waiting Time	28.6
3.4 Expectation of a Memoryless Waiting Time	28.7
4. Exponential Distribution	28.8
5. Sample Space of the Poisson Process	28.9
5.1 Physicists' Definition of the Poisson Process	28.10
5.2 Rare Events	28.10
5.3 Poisson Event	28.11
6. Probability of Poisson Process	28.12
6.1 Main Point	28.13
6.2 Axiom	28.14
7. Justification of Derivation of Exponential Distribution	28.15

Lecture 29 - [4/24/98]

1. Poisson Process - Motivation	29.1
2. The Four Fundamental Stochastic Processes	29.2
2.1 Bernoulli Process	29.2
2.2 Uniform Process	29.2
2.3 Poisson Process	29.2
2.4 Processes Pertaining to the Normal Distribution	29.2
3. Poisson Events	29.2
3.1 Poisson Probability	29.2
4. Random Function	29.3

5. Poisson Probability Distribution with Intensity α	29.4
6. Number of Blips in Disjoint Intervals is Independent	29.4
7. Expectation of a Poisson Process Random Function	29.5
8. The 7 Fundamental Properties of the Poisson Process	29.6
9. Property 1 - Law of Rare Events	29.6
10. Property 2 - Limit of Uniform Process	29.9
11. Property 3 - Memorylessness	29.10
11.1 Gamma Distribution	29.11

Lecture 30 - [4/29/98]

1. The 7 Fundamental Properties of the Poisson Process (cont'd) . . .	30.1
2. Property 4 - Uniform Process Obtained by Conditioning Poisson Process	30.4
2.1 Example - Needles on a Stick	30.6
3. Property 5 - Schrödinger Randomization	30.9
3.1 Maxwell-Boltzmann Statistics from Conditioning Poisson Pro- cess	30.10
3.2 Any Problem of Occupation Numbers	30.13
4. Property 6 - Poisson Blips in 2 Colors	30.17

Lecture 31 - [5/1/98]

1. Property 6 - Poisson Blips in 2 Colors (cont'd)	31.1
2. Property 7 - Laplace Transform	31.4
3. Random Walk	31.5
4. Time of First Return to Origin	31.7
4.1 z Transform	31.10
5. Continuous Analogue of Random Walk	31.13

Lecture 32 - [5/4/98]

1. Random Walk	32.1
2. Standard Normal Distribution	32.2
3. Standard Normal Distributed Random Variable	32.5
4. Variance of a Random Variable	32.6
5. Views of a Random Variable	32.7
5.1 A Random Phenomenon	32.7
5.2 Result of a Search	32.7
5.3 Measurement of an Imperfect Quantity	32.7

6. Variance Addition Formula	32.8

Lecture 33 - [5/6/98]

1. Variance (cont'd)	33.1
2. Density Plots	33.1
3. Properties of Variance	33.2
4. Examples	33.2
4.1 Bernoulli Process	33.2
4.2 Poisson Process	33.5
4.3 Exponential Random Variable	33.6
5. Expectation of a Linear Function of a Random Variable	33.8
6. Standard Deviation of a Linear Function of a Random Variable	33.8
7. Standardized (Normalized) Random Variable	33.8
8. Conformation of Properties of Normalized Random Variables	33.10
8.1 Expectation	33.10
8.2 Variance	33.10
9. Density of Linear Function of Normalized Random Variables	33.12
10. n Measurements of the Same Quantity	33.14
11. Statistics in One Easy Lesson	33.15
12. Likelihood	33.17

Lecture 34 - [5/8/98]

1. Normal Distribution (cont'd)	34.1
2. Basic Rule of Statistics	34.2
3. Bayes' Law for Densities	34.3
4. Priors	34.5
4.1 Honest Choice	34.5
4.2 Dishonest Choice	34.5
5. Measure	34.6
6. Prediction of Average of n Measurements	34.7
7. Probability Tables of Standard Normal Random Variable	34.8
8. Example - Fair Coin?	34.9
9. The Central Limit Theorem	34.11
9.1 Justification 1 - Psychological	34.11
9.2 Example - Coin Tosses to Choose Airline Seats	34.13

Lecture 35 - [5/11/98]

1. Gaussian (Normal) Distribution	35.1
2. Normalizing Random Variables	35.2
3. Gaussian (Normal) Distribution Justifications	35.3
3.1 The Central Limit Theorem (cont'd)	35.3
3.2 Maxwell-Einstein Derivation	35.4
3.3 Wiener's Characterization	35.7
4. The Law of Large Numbers	35.12
5. The Strong Law of Large Numbers	35.13

Super Class Lecture SC1 - [2/27/98]

1. Proof of Inclusion-Exclusion Principle	SC1.1
1.1 Indicator Random Variable	SC1.1
2. More on Rényi's Principle	SC1.5
3. Circuit Theoretic Interpretation of Boolean Operations	SC1.7
4. Algebra of Partitions	SC1.10

Super Class Lecture SC2 - [3/13/98]

1. Probability Trees	SC2.1
2. Urn Models	SC2.2
2.1 Pólya Urn Model	SC2.3
3. Balls Distinguishable, Boxes Indistinguishable	SC2.4
4. Balls Indistinguishable, Boxes Indistinguishable	SC2.5
4.1 Partition of an Integer into Summands	SC2.5
5. Reluctant Functions	SC2.7
6. Balls in the Box Form a Rooted Tree	SC2.8
6.1 Disposition	SC2.8
6.2 Abel Polynomials	SC2.8
7. Partitions	SC2.9
7.1 Bell Numbers	SC2.9
7.2 Stirling Number's of the Second Kind	SC2.10

Super Class Lecture SC3 - [4/10/98]

1. Needles on a Stick (again)	SC3.1
2. Buffon Needle Problem	SC3.2
2.1 From Needles to Curves	SC3.4
2.2 Computing π	SC3.5

3. Convex Set	SC3.7
4. Sylvester's Theorem	SC3.7
5. Measure	SC3.9
6. Measure of Lines in the Plane	SC3.11
7. Computing Perimeters	SC3.12

Super Class Lecture SC4 - [4/24/98]

1. Entropy and Information	SC4.1
2. Partition	SC4.1
3. Fundamental Properties of Entropy	SC4.3
3.1 Independent Partitions	SC4.3
3.2 Finer the Partition, Bigger the Entropy	SC4.5
3.3 Entropy of Meet	SC4.5
4. Maximum Entropy Principle	SC4.6
5. Information Search	SC4.7
6. Famous Problem of Scales	SC4.9
7. Shannon Coding Theorem	SC4.12

Index

Within the text, pagination is of the form *lecture_date.page*. Note that a different convention is used in both the table of contents and the index, where pagination is of the form *lecture_number.page*. The table of contents provides the mapping between *lecture_number* and *lecture_date*. For example, page 1.1 in the table of contents and index corresponds to page 2/4/98.1 in the text, where the topic anomaly detection is discussed.

Research Problems

1. Given a family of intersections of events, where the probability of each intersection is the product, determine if the events are independent 2.5
2. Show that the sample space and probability defined for the Bernoulli Process define a unique probability on all elements 3.1
3. Let A_n be event that all even number of tosses have equal numbers of 0's and 1's. Determine the probability that infinitely many of the A_n occur 3.7
4. Prove the Kolmogorov Zero-one Law 3.8
5. Provide an elegant simplification for the expression that gives the first occupation number in Bose-Einstein statistics 5.10
6. Provide a satisfactory explanation of the identity relating binomial and multiset coefficients 6.6
7. Given a number of random variables, of which certain subsets of intersections are independent, determine whether the random variables are independent 8.5
8. Relate Bose-Einstein statistics to dispositions and explain, with respect to sampling 9.11
9. Given a uniform process on the interval from zero to an unknown endpoint, compute the Bayes' estimate and determine which priors make sense 25.8
10. Find a nice formula for the density of n independent, uniformly distributed random variables 26.14

We're going to change the policy on problem sets from now on, because the material is getting very tough.

I'll give you problem sets that are very closely related to the material given in class. Instead of giving fancy, unsolved problems.

From now on, we're getting to the absolute heart of probability, as well as some of the toughest chapters of probability.

This is what you really have to learn. So far, balls into boxes and all that stuff you can visualize it.

When it comes to continuous probability, you really have to work on it to get it - to develop the concepts behind it.

I'm going to give you problems that are really lecture based from now on. No fancy stuff, no unsolved problems. Well, maybe unsolved problems.

The ideal thing is that at the end of this course, each one of you does a UROP with me and pulls out one of my unsolved problems and solves it & we write a joint paper.

You are young, you have lots of time.
Help me out.

Continuous conditional probability (beginning)

Now we are doing a very tough concept.

Let X = continuous random variable

Then we know that:

$$P(a < X \leq b) = \int_a^b \text{dens}(X=t) dt$$

In particular, we have that:

$$P(X \leq t) = \int_{-\infty}^t \text{dens}(X=s) ds$$

This is called the cumulative distribution.

The density in itself doesn't have a meaning as a probability.

It has a meaning only when integrated.

The density is always a non-negative quantity.

Our objective is to learn to work w/ these densities, even though they do not have a direct meaning.



4/6/98.2

- The toughest thing we are facing is to find the right definition, as well as interpretation, of this formula:

$$P(A|X=t) = ?$$

probability of event A , given that $X=t$, where X is a continuous random variable.

So it's a conditioning relative to a density.

This is what we have to work on.


Let's work up to it gradually, by increasing difficulty + working up to that.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS




- Let's start the whole discussion w/ one of the simplest examples we have seen. Let's take the interval $[0, a]$ and pick two points.

Dirichlet Process with $n=2$



} so we have two order statistics, $X_{(1)}$ and $X_{(2)}$

Let's remind ourselves that:




$$P(X_{(1)} > s) = \binom{2}{0,2} \left(\frac{a-s}{a}\right) \left(\frac{a-s}{a}\right)$$

$$= \frac{(a-s)^2}{a^2}$$

← both marbles fall between s and a , i.e., $(s, a]$.

and, as well:



$$P(X_{(2)} \leq t) = \binom{2}{2,0} \left(\frac{t}{a}\right)^2$$

$$= \frac{t^2}{a^2}$$

← both marbles fall between 0 and t , i.e., $(0, t]$.

4/6/98.3

Since we know $P(X_{(1)} > s)$, the cumulative distribution for the first order statistic is:

$$\begin{aligned} P(X_{(1)} \leq s) &= 1 - P(X_{(1)} > s) \\ &= 1 - \frac{(a-s)^2}{a^2} \end{aligned}$$

The densities are obtained by differentiating the cumulative distributions:

$$\begin{aligned} \text{dens}(X_{(1)} = s) &= \frac{d}{ds} \underbrace{P(X_{(1)} \leq s)}_{1 - \frac{(a-s)^2}{a^2}} \\ &= \frac{2(a-s)}{a^2} \end{aligned}$$

Similarly:

$$\begin{aligned} \text{dens}(X_{(2)} = t) &= \frac{d}{dt} \underbrace{P(X_{(2)} \leq t)}_{\frac{t^2}{a^2}} \\ &= \frac{2t}{a^2} \end{aligned}$$

This thing is tough. Let's not kid ourselves.

Inside the Dirichlet Process (Uniform Process), you find all of probability.

There are unsolved questions all over the place.

From such a simple thing as dropping n marbles, there is research going on.

So this is tough.

Genuinely tough.

In my days, this would be a graduate course.

In particular, we have:

$$\int_{-\infty}^{\infty} \text{dens}(X_{(1)} = s) ds = 1, \text{ by definition } \left\{ \begin{array}{l} \text{integral of density} \\ \text{is } 1 \end{array} \right\}$$

Thus:

$$\int_0^a \frac{2(a-s)}{a^2} ds = 1$$

You can verify this by hand,
but this is automatic.

We got a definite integral cheapo.

4/6/98.4

Similarity:

$$\int_{-\infty}^{\infty} \text{dens}(X_{(2)}=t) dt = 1$$

gives the definite integral:

$$\int_0^a \frac{2t}{a^2} dt = 1$$

Let's jazz this up.
We now ask:

$$P((X_{(1)} > s) \cap (X_{(2)} > t)) = ?$$

(always assume $s \leq t$)


I don't know.

Because to do a combinatorial count, I don't know how to fix it.

Hard to do directly.

But I do know how to compute:

$$P((X_{(1)} > s) \cap (X_{(2)} \leq t)) \leftarrow \text{This is easy}$$

the first marble
is $\geq s$ the second marble
is $\leq t$ Therefore both marbles fall in
the interval $[s, t]$.Two marbles falling in the same interval
is easy because the marbles are independent.


$$\begin{aligned} P((X_{(1)} > s) \cap (X_{(2)} \leq t)) &= \binom{2}{0,2,0} \left(\frac{t-s}{a}\right)^2 \\ &= \frac{(t-s)^2}{a^2} \end{aligned}$$

Now we can use this probability to get the desired $P((X_{(1)} > s) \cap (X_{(2)} > t))$.

By essential reasoning you've seen before, but don't mind seeing again:

$$\text{event } (X_{(1)} > s) \cap (X_{(2)} > t) = (X_{(1)} > s) - ((X_{(1)} > s) \cap (X_{(2)} \leq t))$$

from boolean algebra:



$$A \cap B = A - (A \cap B^c)$$

4/6/98.5

Therefore, because probabilities always subtract:

$$P((X_{(1)} > s) \cap (X_{(2)} > t)) = \underbrace{P(X_{(1)} > s)}_{\frac{(a-s)^2}{a^2} [4/6/98.2]} - \underbrace{P((X_{(1)} > s) \cap (X_{(2)} \leq t))}_{\frac{(t-s)^2}{a^2} \text{ we just computed this.}}$$

$$P((X_{(1)} > s) \cap (X_{(2)} > t)) = \frac{(a-s)^2}{a^2} - \frac{(t-s)^2}{a^2} = \frac{2}{a^2} \left(\frac{a^2}{2} - sa - \frac{t^2}{2} + st \right)$$

I don't know of an intuitive justification of this.

Now, let's do the same the hard way.

By integrating densities. Painful, but good.


Let's find the joint density $X_{(1)}$ and $X_{(2)}$, namely: $\text{dens}(X_{(1)}=s, X_{(2)}=t)$.

We then obtain the joint distribution by integrating the joint density - just like the scriptures say:

$$P((X_{(1)} > s) \cap (X_{(2)} > t)) = \iint_{\text{some set}} \text{dens}(X_{(1)}=u, X_{(2)}=v) du dv$$

NB: not the intersection of events

Now, I've told you, in the strictest confidence, that this kind of computation is avoided most of the time. We just avoided it, because we did it directly. But for once, let's make it come out.



where: $s < t$
 $s < X_{(1)} < X_{(2)}$
 $t < X_{(1)} \leq a$

Roll up our sleeves and let's do it.
 We have to find the limits of integration and do a double integral.

What are the limits of integration?


$$= \int_{u=t}^a \int_{v=s}^{X_{(1)}=u} \text{dens}(X_{(1)}=u, X_{(2)}=v) du dv$$

Labels in the diagram: $X_{(2)} \leq a$, $X_{(1)} < X_{(2)}$, $u=s$, $s < X_{(1)}$, $t < X_{(2)}$.

4/6/98.6

- Next, we need to compute the joint density.
We've done this several times.

$$\text{dens}(X_{(1)}=s, X_{(2)}=t) = ? \quad \text{We use the } \underline{\text{infinitesimal approach}}$$



$$\text{dens}(X_{(1)}=s, X_{(2)}=t) = \lim_{\substack{ds \rightarrow 0 \\ dt \rightarrow 0}} \frac{\binom{2}{0,1,0,1,0} \left(\frac{ds}{a}\right) \left(\frac{dt}{a}\right)}{ds dt}$$

This is, of course, an approximation argument by a limiting process.

$$\text{dens}(X_{(1)}=s, X_{(2)}=t) = \frac{2}{a^2}$$

↑ Recall, from Dr. Billy's lecture [3/30/98.8], that the joint density function for n order statistics, derived in exactly the same fashion, is:

$$\text{dens}(X_{(1)}=x_1, \dots, X_{(n)}=x_n) = \frac{n!}{a^n}$$

- Now, with the limits of integration and the joint density, we can go back to our double integral:

$$P((X_{(1)} > s) \cap (X_{(2)} > t)) = \int_{v=t}^a \int_{u=s}^v \underbrace{\frac{2}{a^2}}_{\text{dens}(X_{(1)}=u, X_{(2)}=v)} du dv$$

$$= \int_{v=t}^a \frac{2(v-s)}{a^2} dv$$

$$= \frac{2}{a^2} \left[\frac{v^2}{2} - sv \right]_{v=t}^a$$

$$= \frac{2}{a^2} \left(\frac{a^2}{2} - sa - \frac{t^2}{2} + st \right)$$

$$= \frac{(a-s)^2}{a^2} - \frac{(t-s)^2}{a^2}$$

So it checks out.

In 90% of the cases, we avoid evaluating double integrals.

✓ probability obtained by integration of joint density checks w/ directly obtained probability [4/6/98.4-5].

4/6/98.7

You think this is the end?
 Sorry - this is just the beginning.

- Now let's compute the conditional probability:

$$P(X_{(2)} > t | X_{(1)} > s) = ?$$

Probability that $X_{(2)} > t$
given that $X_{(1)} > s$.

Let's do it 2 ways.
 First, by the book. } This will justify our intuition, as often happens in life.
 Then, intuitively.

- 1) By the book:

$$P(X_{(2)} > t | X_{(1)} > s) = \frac{P((X_{(2)} > t) \cap (X_{(1)} > s))}{P(X_{(1)} > s)} \quad \left. \vphantom{\frac{P((X_{(2)} > t) \cap (X_{(1)} > s))}{P(X_{(1)} > s)}} \right\} \begin{array}{l} \text{by definition} \\ \text{of} \\ \text{conditional} \\ \text{probability} \end{array}$$

{ The numerator we just computed two different ways.
 The denominator we computed on [4/6/98.2]. }

$$= \frac{\frac{(a-s)^2}{a^2} - \frac{(t-s)^2}{a^2}}{\frac{(a-s)^2}{a^2}}$$

$$= 1 - \frac{(t-s)^2}{(a-s)^2}$$

- 2) By intuition:

What is $P(X_{(2)} > t | X_{(1)} > s)$ telling you?

You know that no marble has fallen in the interval $(0, s]$.

Since it is given that $X_{(1)} > s$ (i.e., first order statistic strictly $> s$).

It's as if you were dealing w/ a shorter interval.


Namely, the interval from s to a .

Therefore, you recalculate this notion on the interval $[s, a]$.

This is what it's telling you. $P(X_{(2)} > t | X_{(1)} > s)$ is the probability distribution of the second marble, given that both marbles fall in the interval $[s, a]$.



Consider $P_{(X_{(1)} > s)}(X_{(2)} \leq t)$



given that $X_{(1)} > s$, no balls drop in interval 0 to s , so we shorten interval to $[s, a]$.

$$\begin{aligned} P_{(X_{(1)} > s)}(X_{(2)} \leq t) &= \binom{2}{2, 0} \left(\frac{t-s}{a-s} \right)^2 \\ &= \frac{(t-s)^2}{(a-s)^2} \end{aligned}$$

And since we have:

$$P_{(X_{(1)} > s)}(X_{(2)} > t) = 1 - P_{(X_{(1)} > s)}(X_{(2)} \leq t)$$

Our intuition is confirmed.

$$P(X_{(2)} > t | X_{(1)} > s) = 1 - \frac{(t-s)^2}{(a-s)^2} \quad \checkmark$$



Now we try another one that is a little tougher.

$$P(X_{(2)} > t \mid X_{(1)} \leq s) = ?$$

By the definition of conditional probability:

$$P(X_{(2)} > t \mid X_{(1)} \leq s) = \frac{P((X_{(2)} > t) \cap (X_{(1)} \leq s))}{P(X_{(1)} \leq s)}$$

First- the numerator:

Recall that we have a nice probability that is almost that of the complement of the numerator's event [4/6/98.4], namely:

$$P((X_{(2)} \leq t) \cap (X_{(1)} > s)) = \frac{(t-s)^2}{a^2}$$

which gives:

$$P([(X_{(2)} \leq t) \cap (X_{(1)} > s)]^c) = 1 - P((X_{(2)} \leq t) \cap (X_{(1)} > s))$$

↓ by DeMorgan's Law

$$P((X_{(2)} > t) \cup (X_{(1)} \leq s)) = \boxed{1 - \frac{(t-s)^2}{a^2}}$$

↑ this probability differs from the numerator in that this operation is a union, not an intersection.

Let's use the inclusion-exclusion principle on the events of the numerator:

$$P((X_{(2)} > t) \cup (X_{(1)} \leq s)) = \underbrace{P(X_{(2)} > t)}_{\boxed{1 - \frac{(t-s)^2}{a^2}}} + \underbrace{P(X_{(1)} \leq s)}_{\substack{1 - P(X_{(1)} > s) \\ \frac{t^2}{a^2} \text{ see [4/6/98.2]}}} - \underbrace{P((X_{(2)} > t) \cap (X_{(1)} \leq s))}_{\substack{\text{the numerator,} \\ \text{what we desire.} \\ 1 - \frac{(a-s)^2}{a^2} \text{ see [4/6/98.3]}}}$$

Rearranging and simplifying gives:

$$P((X_{(2)} > t) \cap (X_{(1)} \leq s)) = 1 - \frac{t^2}{a^2} + \cancel{1 - \frac{(a-s)^2}{a^2}} - \cancel{1} + \frac{(t-s)^2}{a^2}$$

$$= 1 - \left(\frac{t^2 - (t-s)^2 + (a-s)^2}{a^2} \right)$$

Second - the denominator:

We already know this, we've done it a hundred times [4/6/98.3].

$$P(X_{(1)} \leq s) = 1 - \frac{(a-s)^2}{a^2}$$

Which gives the desired conditional probability:

$$P(X_{(2)} > t | X_{(1)} \leq s) = \frac{1 - \left(\frac{t^2 - (t-s)^2 + (a-s)^2}{a^2} \right)}{1 - \frac{(a-s)^2}{a^2}}$$

$$= \frac{\cancel{a^2} - \cancel{t^2} - (\cancel{t^2} - \cancel{2st} + \cancel{s^2}) + (\cancel{a^2} - \cancel{2as} + \cancel{s^2})}{\cancel{a^2} - (\cancel{a^2} - \cancel{2as} + \cancel{s^2})}$$

$$= \frac{2as - 2st}{2as - s^2}$$

$$P(X_{(2)} > t | X_{(1)} \leq s) = \frac{a-t}{a-\frac{s}{2}}$$

4/6/98.11

- Now we compute the density, given an event.

Conditional Density

$$\text{dens}(X_{(2)}=t | X_{(1)}>s) = \frac{d}{dt} P(X_{(2)} \leq t | X_{(1)}>s)$$

{ Take the conditional cumulative probability }
and differentiate.

$$= \frac{d}{dt} \frac{P((X_{(2)} \leq t) \cap (X_{(1)} > s))}{P(X_{(1)} > s)} \leftarrow \begin{array}{l} \text{numerator from [4/6/98.4]} \\ \text{probability both marbles fall in } [s, t] \end{array}$$


$$= \frac{d}{dt} \frac{\frac{(t-s)^2}{a^2}}{\frac{(a-s)^2}{a^2}}$$

$$= \frac{2(t-s)}{(a-s)^2}$$

Now, let's justify this intuitively.
Note that:

$$\text{dens}(X_{(2)}=t) = \frac{2t}{a^2}$$

$$\text{dens}(X_{(2)}=t | X_{(1)}>s) = \frac{2(t-s)}{(a-s)^2}$$



The condition is that $X_{(1)} > s$.

What does this mean in ordinary language?

It means that both marbles fall in the interval $(s, a]$.

Therefore, we have a uniform process on a shorter interval.

And we take the density in the shorter interval.

$\text{dens}(X_{(2)}=t | X_{(1)}>s)$ is exactly the formula for the density on the shorter interval:

$$\text{dens}(X_{(2)}=t) = \frac{2t}{a^2}, \text{ where } \begin{array}{l} t \Rightarrow t-s \\ a \Rightarrow a-s \end{array}$$

- So conditional densities are computed just as other densities. Either:
a) take the conditional cumulative distribution and differentiate
or- b) if you can, give an intuitive argument.



- Now comes the tough part.
What is:

$$P(X_{(2)} \leq t | X_{(1)} = s) = ?$$

random variable
 $X_{(1)}$ takes specific value s .

This is not yet defined.
Previously, we looked at probabilities and densities conditioned on events of continuous random variables, where the probability of the conditioning event was non-zero.
Here, the conditioning event, that $X_{(1)}$ takes a specific value s , has probability:

$$P(X_{(1)} = s) = 0$$

How will we define this?


- 1) Take the conditioning event that $X_{(1)}$ is in a small Δs neighborhood of s .
- 2) Then take the limit.

$$\lim_{\Delta s \rightarrow 0} P(X_{(2)} \leq t | s < X_{(1)} \leq s + \Delta s) = \lim_{\Delta s \rightarrow 0} \frac{P((X_{(2)} \leq t) \cap (s < X_{(1)} \leq s + \Delta s))}{P(s < X_{(1)} \leq s + \Delta s)}$$

This is the event of non-zero probability, that $X_{(1)}$ lies in the interval $(s, s + \Delta s]$.
So we rewrite it using the rule of conditional probability to get the RHS.

Let's compute this with probabilities, rather than with densities.

Numerator:



Could have $X_{(2)}$ maybe fall in $(s, s + \Delta s]$, but in limit $\Delta s \rightarrow 0$, only $X_{(2)}$ can be in $(s, s + \Delta s]$.

$$P((X_{(2)} \leq t) \cap (s < X_{(1)} \leq s + \Delta s)) = \binom{2}{0, 1, 1, 0} \left(\frac{\Delta s}{a}\right) \left(\frac{t - s - \Delta s}{a}\right) + (\Delta s)^2 (\text{stuff})$$

Denominator:

$$P(s < X_{(1)} \leq s + \Delta s)$$

We could do, cheapo. Is there something more elegant?
 $\lim_{\Delta s \rightarrow 0} P(s < X_{(1)} \leq s + \Delta s) = \lim_{\Delta s \rightarrow 0} \int_s^{s + \Delta s} \text{dens}(X_{(1)} = u) du$

{more next time}


$$= \lim_{\Delta s \rightarrow 0} \text{dens}(X = s) \Delta s$$

- Our motivation is the following:

Objective: $\text{dens}(X_{(1)} = s | X_{(2)} = t) = ?$

Continuous Conditional Probability (cont'd)

Let's continue from where we left off last time.
We take the Uniform or Dirichlet Process with $n=2$.
So we drop 2 points at random in an interval.



We're trying to find a rigorous definition of the following term:


$$P(X_{(2)} > t \mid X_{(1)} = s) = ?$$

in other words, the condition is an event of probability zero

$$P(X_{(1)} = s) = 0$$

We want to show that conditioning on an event of probability zero makes sense.

Let's first guess what it ought to be.
What should it be? Intuitively,



↑ we know that the first marble is placed at position s,
as the given conditioning event is $(X_{(1)} = s)$.


That's our knowledge.

The placement of $X_{(1)}$ is no longer a random event.
It's a certain event.

Given that we know that the first marble is placed at position s , what is the probability distribution for the second marble?

The answer intuitively is clear.

If the first marble is placed in position s , then you are left w/ a Dirichlet Process
with only one marble, on a smaller interval.



$\{s < t \text{ always}\}$

So it ought to be that the interval of the first order statistic $X_{(1)} (= X_{(1)})$ is $(s, a]$. So, you take one point in this interval.
So it ought to be:

$$P(X_{(2)} > t \mid X_{(1)} = s) = \frac{a-t}{a-s}$$

Now, let's confirm this by actual, rigorous computation.


The only possible definition of $P(X_{(2)} > t | X_{(1)} = s)$ is:

$$P(X_{(2)} > t | X_{(1)} = s) = \lim_{\Delta s \rightarrow 0} \underbrace{P(X_{(2)} > t | s < X_{(1)} \leq s + \Delta s)}$$

this is a genuine probability

$$(*) = \lim_{\Delta s \rightarrow 0} \frac{P((X_{(2)} > t) \cap (s < X_{(1)} \leq s + \Delta s))}{P(s < X_{(1)} \leq s + \Delta s)}$$

The numerator is:




{ In the limit, we need not concern ourselves w/ multiple balls in an interval. }

$$P((X_{(2)} > t) \cap (s < X_{(1)} \leq s + \Delta s)) = \binom{2}{0, 1, 0, 1} \left(\frac{\Delta s}{a}\right) \left(\frac{a-t}{a}\right)$$

The denominator is:

$$\underbrace{(s < X_{(1)} \leq s + \Delta s)}_{\text{event}} = (s < X_{(1)}) - (X_{(1)} > s + \Delta s)$$



Pure logic.
The limits of rationality

Probabilities always subtract, thus:

$$P(s < X_{(1)} \leq s + \Delta s) = \underbrace{P(s < X_{(1)})} - \underbrace{P(X_{(1)} > s + \Delta s)}$$

We've computed this a hundred times
[4/6/98.2]:

$$\frac{(a-s)^2}{a^2}$$

$$\frac{(a-s-\Delta s)^2}{a^2}$$

Therefore equation (*) is :

$$\begin{aligned}
 P(X_{(2)} > t \mid X_{(1)} = s) &= \lim_{\Delta s \rightarrow 0} \frac{2 \left(\frac{\Delta s}{a}\right) \left(\frac{a-t}{a}\right)}{\frac{(a-s)^2}{a^2} - \frac{(a-s-\Delta s)^2}{a^2}} \\
 &= \lim_{\Delta s \rightarrow 0} \frac{2 \Delta s (a-t)}{[\cancel{a-s} - (\cancel{a-s} - \Delta s)] [(a-s) + (a-s-\Delta s)]} \\
 &= \lim_{\Delta s \rightarrow 0} \frac{2 \cancel{a-s} (a-t)}{\Delta s [2(a-s) - \cancel{a-s}]} \cdot 0 \\
 &= \frac{a-t}{a-s} \quad \checkmark
 \end{aligned}$$

Lo and Behold, this confirms our intuition.

Let's do a fancy one.

Example: Uniform Process, n arbitrary
 $\nwarrow$ any number of marbles.

What is :

$$P(X_{(1)} > s \mid X_{(k)} = t) = ?$$

Let's first guess what it ought to be.



$\nwarrow$ we are given the data that the k^{th} marble falls exactly at position t .

Given these data, what is the probability distribution that $(X_{(1)} > s)$?

Well - let's think.

If you are given these data, then the positions of all marbles from $X_{(k+1)}$ to $X_{(n)}$ are irrelevant.

It is as if you dropped $k-1$ marbles on the interval $(0, t)$.

So it ought to be:

$$P(X_{(1)} > s \mid X_{(k)} = t) = \left(\frac{t-s}{t}\right)^{k-1}$$

4/8/98.4

Now, we justify it the hard way.
(Last time we do it the hard way)

$$P(X_{(1)} > s | X_{(k)} = t) = \lim_{\Delta t \rightarrow 0} P(X_{(1)} > s | t \leq X_{(k)} \leq t + \Delta t)$$

a genuine probability, as we are conditioning by an actual event of non-zero probability, not an event of zero probability.

$$(*) = \lim_{\Delta t \rightarrow 0} \frac{P((X_{(1)} > s) \cap (t < X_{(k)} \leq t + \Delta t))}{P(t < X_{(k)} \leq t + \Delta t)}$$


The numerator is:



{ We may assume WLOG that there is only 1 marble here, because Δt is small. We've gone through this crisis before. }

$$P((X_{(1)} > s) \cap (t < X_{(k)} \leq t + \Delta t)) = \binom{n}{0, k-1, 1, n-k} \left(\frac{t-s}{a}\right)^{k-1} \left(\frac{\Delta t}{a}\right) \left(\frac{a-t-\Delta t}{a}\right)^{n-k}$$

The denominator is:



{ again, WLOG, because of the calculus. }

$$P(t < X_{(k)} \leq t + \Delta t) = \binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{\Delta t}{a}\right) \left(\frac{a-t-\Delta t}{a}\right)^{n-k}$$

Note: a good check is that power of denominator adds up to n .

ex: that's how $k-1$ marbles fall independently in the interval $(0, t]$.

Therefore, $(*)$ equals:

$$\begin{aligned} P(X_{(1)} > s | X_{(k)} = t) &= \lim_{\Delta t \rightarrow 0} \frac{\binom{n}{0, k-1, 1, n-k} \left(\frac{t-s}{a}\right)^{k-1} \left(\frac{\Delta t}{a}\right) \left(\frac{a-t-\Delta t}{a}\right)^{n-k}}{\binom{n}{k-1, 1, n-k} \left(\frac{t}{a}\right)^{k-1} \left(\frac{\Delta t}{a}\right) \left(\frac{a-t-\Delta t}{a}\right)^{n-k}} \\ &= \frac{(t-s)^{k-1}}{t^{k-1}} \quad \checkmark \end{aligned}$$

It comes out. Our intuition is again confirmed.



In general:

$$P(A|X=t) = \lim_{\Delta t \rightarrow 0} P(A|t < X \leq t + \Delta t)$$

probability of event A,
given the event that
continuous random variable
X has value t

X is any continuous variable.
A is any event in the same
sample space.

Now you say "So what?"

Now we come to the deepest, and most important, formula in probability.
This is what I was leading up to.

Continuous Version of the Law of Alternatives

It is stated:

$$P(A) = \int_{-\infty}^{\infty} P(A|X=t) \text{dens}(X=t) dt$$

where X is a
continuous random variable

This is the continuous analogue of the Law of Alternatives.
It is amazing that it works.

Let's roll up our sleeves and prove it.

We prove it by a limiting process, as in the calculus. What did you expect?

Proof

... $t_1, t_2, t_3, \dots$

take events that X takes a value in the interval $[t_n, t_{n+1}]$.

event $(t_n < X \leq t_{n+1})$

These events are a partition of the sample space Ω

- ↑
- i) events disjoint
 - ii) events cover the sample space.

Hence, we can apply the Law of Alternatives. ← ordinary, garden variety.
Ordinary Law of Alternatives gives you:

$$P(A) = \sum_n P(A|t_n < X \leq t_{n+1}) P(t_n < X \leq t_{n+1})$$

4/8/98.6

Now we close our eyes and take the limit as all the intervals between successive $[t_n, t_{n+1}]$ simultaneously tend to 0.
Let's see what happens.

Recall that:

$$P(t_n < X \leq t_{n+1}) = \int_{t_n}^{t_{n+1}} \text{dens}(X=t) dt$$

Integral Mean Value Theorem

An Intro. to Numerical Analysis
Kendall E. Atkinson

Let $w(x)$ be non-negative + integrable on $[a, b]$ and let $f(x)$ be continuous on $[a, b]$. Then

$$\int_a^b w(x) f(x) dx = f(\xi) \int_a^b w(x) dx$$

for some $\xi \in [a, b]$

In this case, $w(x) = 1$,
 $f(x) = \text{dens}(X=x)$

$$\int_{t_n}^{t_{n+1}} \text{dens}(X=x) dx = \text{dens}(X=s_n) \times \int_{t_n}^{t_{n+1}} 1 dx$$

$s_n \in [t_n, t_{n+1}]$

Now we use something that you thought you'd never see again when you took 18.01.
The Mean Value Theorem

When you integrate a ^{continuous} non-negative quantity, it's the same as taking the length of the interval of integration times the value of the integrand at some intermediate value.

$$= (t_{n+1} - t_n) \text{dens}(X=s_n)$$

for some s_n between t_n and t_{n+1}

Therefore (**) can be simplified as follows:

$$P(A) = \sum_n P(A | t_n < X \leq t_{n+1}) (t_{n+1} - t_n) \text{dens}(X=s_n), \quad s_n \in [t_n, t_{n+1}]$$

This is the way you define an integral, as limit of all successive intervals $(t_n, t_{n+1}] \rightarrow 0$.

$$P(A) = \lim_{(t_n, t_{n+1}] \rightarrow 0} \sum_n \underbrace{P(A | t_n < X \leq t_{n+1})}_{\substack{\text{in the limit, tends to} \\ P(A|X=t)}} \text{dens}(X=s_n) (t_{n+1} - t_n), \quad s_n \in [t_n, t_{n+1}]$$

$$= \int_{-\infty}^{\infty} P(A|X=t) \text{dens}(X=t) dt$$

For the next couple of lectures, we will learn much more of this.
We will now apply this all over the place.
It is very important.

• Let's draw a consequence of: $P(A) = \int_{-\infty}^{\infty} P(A|X=t) \text{dens}(X=t) dt$

Applying this to the event:

$$A = (Y \leq s)$$

$\uparrow$ Y is another continuous random variable

you get:

$$P(Y \leq s) = \int_{-\infty}^{\infty} P(Y \leq s | X=t) \text{dens}(X=t) dt$$

Now take the derivative, relative to s , on both sides.

Then you get the density.

The derivative of the cumulative distribution is the density.

$$\frac{d}{ds} P(Y \leq s) = \text{dens}(Y=s)$$

$$\frac{d}{ds} P_{(X=t)}(Y \leq s) = \text{dens}(Y=s | X=t)$$

$$\frac{d}{ds} P(Y \leq s) = \frac{d}{ds} \left[\int_{-\infty}^{\infty} P(Y \leq s | X=t) \text{dens}(X=t) dt \right]$$


$$\text{dens}(Y=s) = \int_{-\infty}^{\infty} \frac{d}{ds} P(Y \leq s | X=t) \text{dens}(X=t) dt$$

$$(***) \quad \text{dens}(Y=s) = \int_{-\infty}^{\infty} \text{dens}(Y=s | X=t) \text{dens}(X=t) dt$$

4/8/98.8

Let's start w/ some easy applications.

Example: Uniform Process

The gaps are $L_1, L_2, \dots, L_{n+1}$
n+1 gapsBy an intuitive argument, we've convinced ourselves that the gaps are identically distributed (i.i.d.), even though they are not independent.


Let's now verify this rigorously:

Gaps in the uniform process are i.i.d.ProofLet's find the probability distribution of the 2nd gap (L_2).Let $A = (L_2 > t)$


By the Continuous Law of Alternatives:

$$P(L_2 > t) = \int_{s=0}^{a-t} P(L_2 > t | L_1 = s) \text{dens}(L_1 = s) ds$$

Limits of integration:

The first gap $L_1 = s$ ranges from $s = 0$ to $a - t$

upper bound:


because the 2nd gap $L_2 > t$
You can not go beyond that.
because otherwise the 2nd gap
would not be able to fit in
the interval $[0, a]$.

Therefore, integration limits are

$$\int_{s=0}^{a-t}$$

4/8/98.9

- Now we do $P(L_2 > t | L_1 = s)$ intuitively:



$$L_1 = s \Rightarrow 1^{\text{st}} \text{ gap} = \text{length } s$$

That means that the first marble falls at s .

Therefore the second marble, which falls after the first marble, falls in a smaller interval.

Since $L_2 > t$, the 2nd through n^{th} marble ($n-1$ total) fall on the smaller interval, as illustrated above. Therefore, we have the conditional probability:

$$P(L_2 > t | L_1 = s) = \left(\frac{a-s-t}{a-s} \right)^{n-1}$$

- Next, we need the density of the 1st gap - $\text{dens}(L_1 = s)$.

Note that the density of the 1st gap is the same as the density of the 1st order statistic $X_{(1)}$. The position of the 1st marble $X_{(1)}$ is the same as the gap of the 1st marble, $L_1 = X_{(1)}$.
Thus:

$$\text{dens}(L_1 = s) = \text{dens}(X_{(1)} = s)$$

We've computed density of order statistics two different ways - see [3/18/98. 8-10].

$$\text{dens}(X_{(k)} = s) = \binom{n}{k} \frac{k s^{k-1} (a-s)^{n-k}}{a^n}$$

$$= \frac{n(a-s)^{n-1}}{a^n}$$

- Now that we have $P(L_2 > t | L_1 = s)$ and $\text{dens}(L_1 = s)$, we can continue computing the integral for $P(L_2 > t)$:

$$P(L_2 > t) = \int_{s=0}^{a-t} P(L_2 > t | L_1 = s) \text{dens}(L_1 = s) ds$$

$$= \int_{s=0}^{a-t} \left(\frac{a-s-t}{a-s} \right)^{n-1} \frac{n(a-s)^{n-1}}{a^n} ds$$

4/8/98.10

$$= \int_{s=0}^{a-t} \frac{n (a-s-t)^{n-1}}{a^n} ds$$

$$= \frac{n}{a^n} \int_{s=0}^{a-t} (a-s-t)^{n-1} ds$$

$$= \frac{1}{a^n} \left[-(a-s-t)^n \right]_{s=0}^{s=a-t}$$

$$P(L_2 > t) = \frac{(a-t)^n}{a^n}$$

Probability that the 2nd
gap is greater than t

is exactly the same as Probability that the 1st gap
is greater than t .

$$P(L_2 > t) = P(L_1 > t) = \frac{(a-t)^n}{a^n}$$

This proves that L_1 & L_2 are identically distributed (i.i.d.)

What more do you want?

And similarly, you can show L_2 i.i.d. to L_3 , etc.



Continuous Conditional Probability (cont'd)

We begin to see applications of what is perhaps the fundamental identity of probability theory. Namely:

The continuous analogue of the Law of Alternatives:

$$P(A) = \int_{-\infty}^{\infty} P(A|X=t) \underbrace{\text{dens}(X=t)}_{\text{density of a continuous random variable,}} dt$$

probability of an event A

A concept which we took some time to define.

We saw one application of this last time, when we showed that if we take the first 2 gaps in the Dirichlet Process, then the 2 gaps L_1 and L_2 are identically distributed [4/8/98.8-10]. And, of course, from the computations we did last time, it's easy to infer that the computations can be further pursued to show that any 2 gaps (L_i and L_j) are identically distributed.


Now let's attack one of the toughest problems of probability, which is the problem of needles on a stick.

• Example Problem of needles on a stick

We are given a stick of length a :



We are given a supply of n identical needles of length h :



We drop the n needles at random on the stick.

The needle discipline is that the needles should fit entirely w/in the stick. They can not stick out.

So we drop the needles on the stick, subject to this needle discipline.

Problem: What is the probability that no 2 needles overlap?

Note - this problem is solved in the book, but there are several mistakes.

Let $B_{a,n}$ = event that no two needles overlap
out of n needles on interval of length a


4/10/98.2

Our objective is to compute the probability of the event $B_{a,n}$:


$$P(B_{a,n}) = ?$$

Step 1: rephrase the whole problem in terms of the Uniform Process
 In terms of the uniform process, what have we done?
 Let's decide what it means to drop a needle.
 You may use several tricks, which are equivalent.
 The one we will use is the following:


- ① We first pick the right hand endpoint of the needle.



- ② Then we drop the RHS onto the stick.
- ③ Then we put the needle's left hand endpoint onto the stick.



If we understand the dropping of the needles from this point of view, then the problem can be immediately reformulated in terms of the uniform process.



The dropping of the RHS of the needle closest to 0 is the same as choosing point $X_{(1)}$ of the Uniform Process, and so forth.

And therefore, the event that the needles shall not overlap (i.e., $B_{a,n}$) is equivalent to saying that the first n gaps must be $> h$.
 So, the event $B_{a,n}$ is the same as:

$$B_{a,n} = (L_1 > h) \cap (L_2 > h) \cap \dots \cap (L_n > h)$$

Nothing is said about the last gap, L_{n+1} , for which we don't care.

So that's the event whose probability we have to compute.

The way known to man of to compute this probability is by applying the Continuous Law of Alternatives, in a judicious way.

We condition relative to either a) the position of the 1st needle
or b) the position of the last (n^{th}) needle


Let's see what happens if we condition relative to the position of the 1st needle.
By the Continuous Law of Alternatives:

$$P(B_{a,n}) = \int_{-\infty}^{\infty} P(B_{a,n} | L_1 = t) \text{dens}(L_1 = t) dt$$

Now, let's put the limits of integration as carefully as we can.

Certainly, $L_1 > h$, otherwise the first needle doesn't fit,
in which case $P(B_{a,n} | L_1 = t) = 0$.

At the other limit, the largest 1st gap L_1 that is possible is:




$$= \int_h^{a-(n-1)h} P(B_{a,n} | L_1 = t) \text{dens}(L_1 = t) dt$$

Now we look at the conditional probability and say:

"What is this, really?"

Note that $L_1 = X_{(1)}$ ← {the 1st gap and the 1st order statistic} are the same.

If $L_1 = X_{(1)} = t$ (i.e., the 1st order statistic is a t), then this conditioning tells you that you are dropping $n-1$ needles on a stick of length $a-t$.



So,

$$P(B_{a,n} | L_1 = t) = P(B_{a-t, n-1})$$

$$(*) \quad = \int_h^{a-(n-1)h} P(B_{a-t, n-1}) \text{dens}(L_1=t) dt$$

Next, we will evaluate this. First by handwaving, then by rigorous mathematics.

By handwaving:

Since we have
$$P(B_{a,n}) = \int_h^{a-(n-1)h} P(B_{a-t, n-1}) \text{dens}(L_1=t) dt$$

we can use this equality to obtain an expression for $P(B_{a,n-1})$ in terms of $P(B_{a,n-2})$

We can repeatedly use this inequality and eventually get:

$$P(B_{a,n}) \text{ in terms of } P(B_{a,1})$$

The toilet trained way of saying this is:
By induction


By rigorous mathematics (i.e., by induction)

We want to show that:

$$P(B_{a,n}) = \frac{(a-nh)^n}{a^n}$$

At this point, I will assume that this probability has already been computed and shown to be correct up to $n-1$, which is the inductive hypothesis.

Basis: for $n=1$, the solution is trivial. One needle has to fall inside the stick. $(L_1 = X_{(1)}) > h$



$$P(B_{a,1}) = \frac{a-h}{a} \quad \checkmark$$

Inductive Step: By the inductive hypothesis, where I replace $\begin{cases} a \rightarrow a-t \\ n \rightarrow n-1 \end{cases}$, we have:

$$P(B_{a-t, n-1}) = \frac{((a-t) - (n-1)h)^{n-1}}{(a-t)^{n-1}}$$

Then we have $P(B_{a,n})$ from equation (*) by computing the integral:

$$P(B_{a,n}) = \int_h^{a-(n-1)h} \underbrace{\frac{((a-t) - (n-1)h)^{n-1}}{(a-t)^{n-1}}}_{\text{motus motandis}} \text{dens}(L_1=t) dt$$

motus motandis = Latin: You all would see, You all will have seen.

4/10/98.5

We've computed $\text{dens}(L_1=t)$ a hundred thousand times

$$\left\{ \begin{array}{l} [3/18/98.8-10] \\ [3/20/98.2-4] \end{array} \right\}$$

$$\begin{aligned} \text{dens}(L_1=t) &= \text{dens}(X_{(1)}=t) \\ &= \frac{n(a-t)^{n-1}}{a^n} \end{aligned}$$

Thus, we have:

mutatis mutandis

$$P(B_{a,n}) = \int_h^{a-(n-1)h} \frac{(a-t-(n-1)h)^{n-1}}{(a-t)^{n-1}} \cdot \frac{n(a-t)^{n-1}}{a^n} dt$$

$$= \int_h^{a-(n-1)h} \frac{n(a-t-(n-1)h)^{n-1}}{a^n} dt$$

$$= \frac{n}{a^n} \int_h^{a-(n-1)h} (a-t-(n-1)h)^{n-1} dt$$

Evaluate this integral by the substitution method (aka change of variables)

$$u = a-t-(n-1)h \quad du = -dt$$

$$\int_h^{a-(n-1)h} (a-t-(n-1)h)^{n-1} dt = - \int_{t=h}^{a-(n-1)h} u^{n-1} du$$

$$= - \frac{u^n}{n} \Big|_{t=h}^{a-(n-1)h}$$

$$= - \frac{(a-t-(n-1)h)^n}{n} \Big|_{t=h}^{a-(n-1)h}$$

$$= 0 - \left(- \frac{(a-h-(n-1)h)^n}{n} \right)$$

$$\int_h^{a-(n-1)h} (a-t-(n-1)h)^{n-1} dt = \frac{(a-nh)^n}{n}$$

So we have:



4/10/98.6

$$= \frac{n}{a^n} \frac{(a-nh)^n}{n}$$

$$P(B_{a,n}) = \frac{(a-nh)^n}{a^n}, \text{ as desired}$$

This completes the proof by induction of $P(B_{a,n})$, which was obtained using the Continuous Law of Alternatives.

That's it.

The problem of needles on a stick - I told you about this the 1st day of class.

The 2-Dimensional analogue of this problem is unsolved.

Want to get a Nobel Prize? Work this out.

(You'll get the Nobel Prize in Physics).

80% of solid-state physics is this problem.

I am not joking.

Q/A: I should have given the following as an exercise.

Work out all the configurations where the needles can fall off the right of the stick, but not the left. Or they can fall off the left of the stick.

Q/A: Another way of stating the problem is to say:

The right end of each needle must fall into the interval 0 to a.

This is a very famous problem. There are about 1,000 physics research papers written on this.

Now you get it in 10 minutes.

It's worth your tuition.

Some of you may help me do a UROP to do a search through the literature of all the physics papers that have been written on the uniform process w/o acknowledging it. That's called 1-dimensional physics.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



Lets go on to bigger + better things.
In fact, let's go back to our origins - as they say in philosophy.
What's the fundamental problem of probability?

A: Balls into boxes? $\leftarrow$ { No. You toss a coin.
Even more fundamental than balls into boxes. }

A: Tossing a coin

$\hookrightarrow$ the fundamental problem of probability

So let's go back to the Bernoulli Process and apply the Continuous Law of Alternatives to the Bernoulli Process.

Example - Bernoulli Process

I toss the coin 17 times and I get 5 heads.
Is the coin fair?

How do we answer that question?

We have the means to answer it. You are now in a position to answer this fundamental question of life.

Is the coin fair based on these data?

You don't know if the coin is fair. You are being asked the bias of the coin. You don't know the coin's bias.

So what do you do when you don't know the coin's bias?

You say the bias is a random variable

P = a continuous random variable w/ a density
= bias

Now, let's compute the probability that in n tosses, we get k heads.
The binomial distribution.

Do you remember? I hope so.

By the Continuous Law of Alternatives:

$$P(S_n = k) = \int_{-\infty}^{\infty} P(S_n = k | P = t) \text{dens}(P = t) dt$$

Same old
notation [2/10/98.8-9]

Random variable S_n is
number of heads in
 n tosses.

$\text{dens}(P = t) = 0$ for $t < 0, t > 1$

$$= \int_0^1 P(S_n = k | P = t) \text{dens}(P = t) dt$$

continuous random variable } bias

4/10/98.8

$$P(S_n=k | P=t) = \binom{n}{k} t^k (1-t)^{n-k}$$

probability of k heads,
given coin bias is t .
[2/20/98.8] } $\lim_{\Delta t \rightarrow 0} P(S_n=k | t \leq P \leq t+\Delta t)$

$$P(S_n=k) = \int_0^1 \binom{n}{k} t^k (1-t)^{n-k} \text{dens}(P=t) dt$$

Q: Isn't this the Beta Function?

A: Not quite, because the $\text{dens}(P=t)$ is an unknown function. You don't know what the density is. It can be anything. It could be bimodal.

Recall the Beta Function [3/18/98.11], which rewritten is:

$$\int_0^1 \binom{n}{k} t^{k-1} (1-t)^{n-k} dt = \frac{a^n}{k}$$

No, No. It's not Beta.

So, in order to solve the problem of whether the coin is biased on the basis of given data, we have to develop a continuous analogue of Bayes' Law.

Remember Bayes' Law:

$$P(A|B) = \frac{P(B|A) P(A)}{P(B)}$$

↑ posterior
↑ likelihood
← prior

} Trivial, as a mathematical identity.

NB: There is no common agreement among statisticians as to whether

$$P(B|A) \text{ --or-- } \frac{P(B|A)}{P(B)} \text{ is called the prior.}$$

Because of this lack of agreement, Bayes' Law is sometimes written as follows:

$$P(A|B) \propto P(B|A) P(A)$$

↑ proportional to

The proportionality factor is $\frac{1}{P(B)}$.

Why don't people like to write down the proportionality factor?
Because, you really can get it from the fact that:

$$\int_{a=-\infty}^{\infty} P(A=a|B) da = 1 = \int_{a=-\infty}^{\infty} \frac{P(B|A=a) P(A=a)}{P(B)} da = \frac{1}{P(B)} \int_{a=-\infty}^{\infty} P(B|A=a) P(A=a) da$$

4/10/98.9

You can get $\frac{1}{P(B)}$, the proportionality factor, from the fact that probabilities (in this case, the posterior probability) must integrate to 1.

The conditional probability $P(A/B) = P_{(B)}(A)$ over the whole sample space must be 1.

So, if you don't know the denominator in Bayes' Law, you can obtain it by a simple integration:

$$P(A/B) = \frac{P(B/A) P(A)}{P(B)} \leftarrow \int_{a=-\infty}^{\infty} P(B/A=a) P(A=a) da$$

This is why you will see, in the literature or when you get a summer job or so, the symbol $\propto$ appearing when people talk about Bayes' Law. Just to warn you.

- Let's see what the continuous analogue might be. The continuous analogue is obtained by a limiting process like we got when conditioning densities.

$$P(S_n=k | P=t) = \lim_{\Delta t \rightarrow 0} P(S_n=k | t < P \leq t + \Delta t)$$

this is the
likelihood,
analogous to
 $P(B/A)$

$$= \binom{n}{k} t^k (1-t)^{n-k}$$

The prior, analogous to $P(A)$, is:

$$\text{dens}(P=t)$$

↖ Note: this is a density, not a probability.

The denominator, analogous to $P(B)$, is:

$$P(S_n=k) = \int_{-\infty}^{\infty} P(S_n=k | P=t) \text{dens}(P=t) dt \leftarrow \left\{ \begin{array}{l} \text{from the continuous} \\ \text{Law of Alternatives,} \\ \text{as shown on [4/10/98.7]} \end{array} \right\}$$

The LHS, analogous to $P(A/B)$, is:

$$\text{dens}(P=t | S_n=k)$$

So the continuous analogue of Bayes' Law is:

$$\text{dens}(P=t|S_n=k) = \frac{P(S_n=k|P=t) \text{dens}(P=t)}{P(S_n=k)}$$

Instead of the event $(S_n=k)$, you can write an arbitrary event, of course.
The continuous analogue of Bayes' Law is an estimate of:

the density given that some event pertaining to that density has happened,
given the evidence

↑ (in this example, the evidence is that n tosses yields k heads)

So let's see how this works in practice.

From the previous pages, we have:

$$P(S_n=k|P=t) = \binom{n}{k} t^k (1-t)^{n-k}$$

$$P(S_n=k) = \int_0^1 \binom{n}{k} t^k (1-t)^{n-k} \text{dens}(P=t) dt$$

which gives:

$$(*) \quad \underbrace{\text{dens}(P=t|S_n=k)}_{\text{posterior}} = \frac{\overbrace{\binom{n}{k} t^k (1-t)^{n-k}}^{\text{likelihood}} \underbrace{\text{dens}(P=t)}_{\text{prior}}}{\int_0^1 \binom{n}{k} t^k (1-t)^{n-k} \text{dens}(P=t) dt}$$

$$\boxed{\text{dens}(P=t|S_n=k) \propto t^k (1-t)^{n-k} \text{dens}(P=t)}$$

This is a very important conclusion, from several points of view.

So, you toss the coin n times, and you get k heads. Given this evidence, then you make a guess of what you think the bias of the coin is.

← $\text{dens}(P=t)$

For example, you may have reason to believe that the coin is strongly biased around $P = \frac{1}{2}$. Or strongly biased with P between 0 and $\frac{1}{4}$. Or whatever.

So you plug $\text{dens}(P=t)$ in.

On the basis of your previous guess of the bias, you compute $\text{dens}(P=t|S_n=k)$, which in turn allows you to obtain the cumulative probability of the bias of the coin.

$$\int_{-\infty}^t \text{dens}(P=s|S_n=k) ds = P(P \leq t | S_n=k)$$

- What if you don't have any guess?
If you don't have any guess, then you assume that all possible values of p are equally probable. So continuous random variable P has a uniform distribution. That's called a uniform prior.

- With a uniform prior:


$$\text{dens}(P=t) = 1, \text{ for } 0 \leq t \leq 1$$

In this case, the Continuous analogue of Bayes' Law, equation (8), becomes:

$$\text{dens}(P=t | S_n=k) = \frac{t^k (1-t)^{n-k}}{\int_0^1 t^k (1-t)^{n-k} dt}$$

We can compute the denominator's integral in a similar fashion to that of [3/18/98.9-11], but this time we get the identity by using $n+1$ balls and focusing on $\text{dens}(X_{(k+1)}=t)$.

← the $(k+1)^{\text{st}}$ order statistic.



In the limit,

$$\text{dens}(X_{(k+1)}=t) = \binom{n+1}{k, 1, n-k} \left(\frac{t}{a}\right)^k \left(\frac{1}{a}\right) \left(\frac{a-t}{a}\right)^{n-k}$$

and, as a consequence of the fact that integral of density over the whole sample space equals one:

$$\int_{-\infty}^{\infty} \text{dens}(X_{(k+1)}=t) dt = 1$$


$$\int_0^1 \binom{n+1}{k, 1, n-k} t^k (a-t)^{n-k} dt = 1$$

(**) This is another version
of the Beta Function
Integral.

$$\int_0^1 t^k (a-t)^{n-k} dt = \frac{k! (n-k)!}{(n+1)!}$$

4/10/98.12

Precisely the same argument gives us the general Euler formula for the Beta Function Integral $B(a, b)$.



In the limit, we have:

$$\text{dens}(X_{(a)} = t) = \binom{a+b-1}{a-1, 1, b-1} t^{a-1} (1-t)^{b-1}$$

Then integrating the density to obtain probability:

$$\int_{-\infty}^{\infty} \text{dens}(X_{(a)} = t) dt = 1$$

$$\int_0^1 \binom{a+b-1}{a-1, 1, b-1} t^{a-1} (1-t)^{b-1} dt = 1$$

$$B(a, b) = \int_0^1 t^{a-1} (1-t)^{b-1} dt = \frac{(a-1)!(b-1)!}{(a+b-1)!}$$

Beta Function
Integral-general form

Now that we have computed the denominator, we have:

$$\text{dens}(P=t | S_n=k) = \frac{t^k (1-t)^{n-k}}{\left(\frac{k! (n-k)!}{(n+1)!} \right)}$$

Equation (**), the Beta Function Integral $\int_0^1 t^k (1-t)^{n-k}$ is exactly a proportionality constant required to ensure that:

$$\int_{-\infty}^{\infty} \text{dens}(P=t | S_n=k) dt = 1$$

Since $\left(\frac{k! (n-k)!}{(n+1)!} \right)$ is a constant wrt t , we write:

$$\text{dens}(P=t | S_n=k) \propto t^k (1-t)^{n-k}$$



4/10/98.13

- So now that we have $\text{dens}(P=t | S_n=k)$, what do we have philosophically? We have the answer to what the bias of the coin is, on the basis of observing k heads in n tosses is not given by a number. It is given by a probability distribution, as we might expect.

$$\int_{-\infty}^t \text{dens}(P=s | S_n=k) ds = P(P \leq t | S_n=k)$$

Let's see some cases of what happens if we observe $S_n=k$, for various n and k .

- Example - $n=2, k=0$ $\leftarrow$ we observe no heads in 2 tosses


What's the plot of $\text{dens}(P=t | S_n=k)$?

Clearly, we don't give a hoot about the proportionality constants here. We just plot:

$$\text{dens}(P=t | S_n=k) \propto t^k (1-t)^{n-k} \quad \text{--vs--} \quad t$$


Of course, the plot needs to be normalized, so that the area under the curve integrates to 1.

$$\text{dens}(P=t | S_2=0) \propto t^0 (1-t)^2$$



- Example - $n = \text{large}, k=0$ $\leftarrow$ let's say we toss it one billion times and get no heads.


$$\text{dens}(P=t | S_n=0) \propto (1-t)^n$$



This stresses the fact that the larger n is, and if you don't get any heads, the more the area of the density is close to 0. Which means most of the cumulative probability $P(P \leq t | S_n=0)$ is accounted for when the bias P is close to 0.

- Example - $n=2$, $k=1$ $\leftarrow$ 1 head in 2 tosses


$$\text{dens}(P=t | S_2=1) \propto t(1-t)$$



Symmetric function around $\frac{1}{2}$

- Example - $n=10$, $k=5$ $\leftarrow$ more tosses, but again, exactly half the tosses are heads.

$$\text{dens}(P=t | S_{10}=5) \propto t^5(1-t)^5$$



Symmetric function around $\frac{1}{2}$


{ Compared w/ above, it gets steeper and most of the area is around $\frac{1}{2}$. }

- For n even, with $k = \frac{n}{2}$ (i.e., half the tosses are heads), we have:

$$\text{dens}(P=t | S_n = \frac{n}{2}) \propto t^{\frac{n}{2}}(1-t)^{\frac{n}{2}}$$

This is a symmetric function around $\frac{1}{2}$

$$\begin{aligned} t \rightarrow 1-t &\Rightarrow (1-t)^{\frac{n}{2}} (1-(1-t))^{\frac{n}{2}} \\ 0 \leq t \leq 1 & \\ &= (1-t)^{\frac{n}{2}} t^{\frac{n}{2}} \quad \checkmark \end{aligned}$$




- As n gets large, you get bumps exactly where they should be. The bumps come out in exactly the best position. So it really works beautifully. Except for 1 kink.

It is my duty to inform you.
The boss will say: "I don't understand graphs. I want you to tell me whether the coin is fair, or not. Forget graphs."
What do you do?

You do confidence intervals.


Now, let me tell you that it comes out, in the Bernoulli Process, very simple. You simply take an interval that contains 95% of the area under the graph. And that interval is called a very good confidence interval. Or else, you take an interval that contains 99% of the area, and this interval is called extremely good.

interval w/ 95% of area - very good confidence interval
 " " 99% " " - extremely good confidence interval



Then you give the boss this. With 95% confidence, the bias is between this and that.

So, for example, in this example, as n becomes large and half the tosses are heads, the curve becomes very steep and the 95% confidence interval becomes smaller.



As n gets very large, the confidence interval gets smaller + smaller, shrinking about $\frac{1}{\sqrt{n}}$. You get more and more evidence that the bias is $\frac{1}{2}$.

This is what statistics is all about.

People don't accept the graphs, they want a number.

So you give them this 95% (or 99%) area.

Satisfy them.

- Now, these confidence intervals have to be computed numerically w/ tables and on the computer, because you want integration of the density function such that:

$$\int_a^b \text{dens}(P=t | S_n=k) dt = .95, \text{ where } [a, b] \text{ is the very good confidence interval}$$

$$\int_a^b \frac{t^k (1-t)^{n-k}}{\binom{n}{k}} dt = .95$$

Since $\int t^k (1-t)^{n-k} dt$ is not an elementary integral, this integral has to be computed by computer, or w/ tables.

4/10/98.16

I was going to do tables last night, but I was just too tired at 1 am.
So I didn't xerox the tables.
So I didn't give you any with this problem set. You'll get them next week.

I'll go over this again next time, because it's too important, even though elementary.

This is what's coming next:

Why must we go through all this rigamarol if we are going to take the uniform prior?

A: We don't always take the uniform prior.

There is another prior, which has physical significance.
This is called the conjugate prior, which is the following:

I take a coin and secretly toss it 100 times and keep a record.
Then I give you the coin w/o giving you my record.
You use the uniform prior.

I use the prior that corresponds to my previous tossings (from my record).
I get a better confidence interval than you do.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



I was absent from this class.
These Lecture 25 notes are based on the following, which were kindly provided by:

Lecture tape - Jeff Lieberman

Lecture notes - Ira Gerhardt
Jeff Lieberman
Rick Monte

These Lecture 25 notes are based on my interpretation of these materials.

Bayes Estimation (cont'd)

Last time, we derived the Beta function, which can be written in several ways:

$$\int_0^1 t^k (1-t)^{n-k} dt = \frac{k!(n-k)!}{(n+1)!}$$

And last time we studied the problem of estimating how fair a coin is on the basis of data obtained from tossing the coin a number of times.
We'll conclude the discussion today.

This discussion is, historically, one of the central chapters of statistics.
The discussion of the estimate of the fairness of the coin, using Bayes' Law - or rather, it's continuous analogue.

As we showed last time [4/10/98.7-8], by the Continuous Law of Alternatives:

$$P(S_n = k) = \int_0^1 \underbrace{P(S_n = k | P=t)}_{\substack{\text{probability that there} \\ \text{are } k \text{ heads in } n \\ \text{tosses.}}} \underbrace{\text{dens}(P=t)}_{\substack{\text{density of} \\ \text{the bias}}} dt$$

We also have, from the Continuous analogue of Bayes' Law:

Bayes' Law
for coin tossing

$$\text{dens}(P=t | S_n=k) = \frac{P(S_n=k | P=t) \text{dens}(P=t)}{P(S_n=k)} \quad \leftarrow P(S_n=k) \text{ from above}$$

Strictly speaking, we should find $\text{dens}(P=t | S_n=k)$ in two steps:

$$1) \quad P(P \leq t | S_n=k) = \frac{P((P \leq t) \cap (S_n=k))}{P(S_n=k)}$$

$$\text{then } 2) \quad \text{dens}(P=t | S_n=k) = \frac{d}{dt} P(P \leq t | S_n=k)$$

- If we have no previous information about the bias of the coin, then the density for the random bias P is the uniform density between 0 and 1
Therefore:

$$\text{dens}(P=t) = 1, \quad 0 \leq t \leq 1$$

And as we showed. [4/13/98.10-12]:

$$\text{dens}(P=t/S_n=k) = \frac{P(S_n=k|P=t) \text{dens}(P=t)}{P(S_n=k)}$$

$$= \frac{P(S_n=k|P=t) \text{dens}(P=t)}{\int_0^1 P(S_n=k|P=t) \text{dens}(P=t) dt}$$

{ w/ $\text{dens}(P=t)=1$,
this becomes the
Beta Function Integral. }

Q/A: Note that the binomial coefficients for $P(S_n=k|P=t)$ in the numerator + denominator will cancel.

$$(*) \quad \text{dens}(P=t/S_n=k) = \frac{t^k (1-t)^{n-k}}{\left(\frac{k! (n-k)!}{(n+1)!} \right)}$$

We're back to "Who's buried in Grant's Tomb?" We know that the integral of the probability density has to equal 1.
 S_0 :

$$1 = \int_{-\infty}^{\infty} \text{dens}(P=t/S_n=k) dt$$

$$= \int_0^1 \text{dens}(P=t/S_n=k) dt$$

$$= \int_0^1 \frac{t^k (1-t)^{n-k}}{\left(\frac{k! (n-k)!}{(n+1)!} \right)} dt \Rightarrow \int_0^1 t^k (1-t)^{n-k} dt = \frac{k! (n-k)!}{(n+1)!}$$


And this is precisely the Beta Function Integral, which we worked out a different way earlier [4/13/98.11] and was used in computing $\text{dens}(P=t/S_n=k)$ above.

- When people write densities, they often leave out the denominator.
For example:

$$\text{dens}(P=t/S_n=k) \propto t^k (1-t)^{n-k}$$

You can always recover the denominator, from the fact that the density has to integrate to 1.

- With the uniform prior, we saw last time how more + more evidence causes the graph to become steeper around the actual bias of the coin.



So if you toss the coin a large number of times, you are likely to get graphs that are very steep.
And this steep peak would be your guess of the bias.


Now, we have to repeat something we touched upon last time.
The integral under this curve is equal to 1, because it is a probability density.
But the woman in the street does not accept a curve as an answer.
To satisfy the man in the street, you do 2 things.
I told you one last time.

1) confidence intervals

What is a confidence interval? It's simply this.

The whole area under the curve is 1.

You take some subset of the area that includes 95% of the area.



Then you look at the 2 extremes of the interval.
This is called the confidence interval.

Last night, I was going to make up some problems for you, but I was just too tired to look up the tables.
I have some tables that are horribly, poorly done.
But you want to get some exercise, as trivial as they are, on computing confidence intervals. It's of practical value - you just look it up in a table, or you got a computer program to work it out.
If $K=32$ and $n=65$, you should be able to give a 95% confidence interval.

Q: What do you do if you have a bimodal distribution?

A: The notion of confidence intervals is not well defined.
In fact, it's not well defined when it's unimodal. You may not know exactly where the middle is.
The whole thing is handwaving, really.

In general, you take a sensible center and take the 95% confidence interval about that center.

Or maybe, maybe you take 95% of the area, starting from 0. It depends on the kind of question you are asking.

Confidence interval is not well defined.

What is well defined is the posterior
What you do w/ the posterior depends on who you work for.

That's only the first thing.
The second thing is:

The boss says: "I don't want all these confidence intervals. I want a number."

So what do you do?

You take the expectation of the posterior.

The average.

That's called the Bayes' Estimate

2) Bayes' Estimate

$E(P | S_n = k)$ ← The conditional expectation of P , given the event $(S_n = k)$.

the expected value
of the posterior

(This stuff is presented in statistics books as a thick mess. It's so simple. This is clearly Mickey Mouse. But in statistics books it takes 100 pages explaining this stuff. Besides, not everyone takes 18.313.)

The conditional expectation of the random variable P ← P = bias of the coin

If the boss wants a number, this is what you give her.

Let's see what it turns out to be.

This is a very interesting computation.

Q: Would you use this as the center for obtaining a confidence interval.

A: Again, often. In sensible situations, yes.

But if you have a weirdo prior, you might get a very sensitive situation, based on this weirdo prior.

By the definition of expectation:

$$E(P | S_n = k) = \int_{-\infty}^{\infty} t \text{ dens}(P=t | S_n=k) dt$$

$$= \int_0^1 t \text{ dens}(P=t | S_n=k) dt$$

This density is given by equation (*), under the uniform prior
We assume a uniform prior.

4/13/98.5

$$= \int_0^1 t \frac{t^k (1-t)^{n-k}}{\frac{k! (n-k)!}{(n+1)!}} dt$$

$$= \frac{(n+1)!}{k! (n-k)!} \int_0^1 t^{k+1} (1-t)^{n-k} dt$$

This is the Beta Function Integral

See [4/10/98.12];

$$a-1 = k+1 \Rightarrow a = k+2$$

$$b-1 = n-k \Rightarrow b = n-k+1$$

$$\int_0^1 t^{k+1} (1-t)^{n-k} dt = \frac{(k+1)! (n-k)!}{(n+2)!}$$

$$= \frac{\cancel{(n+1)!}}{\cancel{k! (n-k)!}} \cdot \frac{(k+1)! \cancel{(n-k)!}}{(n+2)!}$$

$$E(P | S_n = k) = \frac{k+1}{n+2}$$

This is the famous Bayes' Estimate.

If you toss a coin n times and you get k heads ($S_n = k$), the best estimate for the fairness of the coin is this number:

$$E(P | S_n = k) = \frac{k+1}{n+2}$$

Don't you ever forget that.

You that's what you give the boss, if she just wants a number.

• Application - the Laplace Law of Succession

It's most impressive what it says.

Assuming that the sun has risen n times, what is the probability that it will rise tomorrow again.

We view this as tossing a coin with a bias, which we don't know.

We use Bayes' and end up w/ a Bayes' Estimate with $K=n$ ^{event} ($S_n = n$) because it's always risen. Therefore:

$$E(\text{sun will rise tomorrow} | \text{sun has risen } n \text{ days}) = E(P | S_n = n) = \frac{n+1}{n+2}$$



4/13/98.6

Now you say: Excuse me. So far we have only worked with the uniform prior. I ran into a psychology professor the other day who said:

"All my life I have never used any other prior except uniform prior."

But that's the case.

Depending on the computations, you often use other priors.

Let's consider another prior - the conjugate prior.

↑ the term is justified from back in your previous problem sets.

Conjugate Prior

Before handing you the coin to test, I have privately tossed it j times and obtained i heads. I keep this information to myself.

Then you toss it, and you get k heads out of n tosses. All you can do is use the uniform prior, because you don't know what I've done w/ the coin.

But what I can do is use the extra information, from my private tossing, which will give me a different prior. What prior?

The posterior that comes from tossing the coin j times. That is the conjugate prior.

$$\text{Conjugate prior: } \text{dens}(P=t) \propto t^i (1-t)^{j-i}$$

I take a uniform prior, then perform my private tossings and based on the observed i heads in j tosses, obtain the posterior:

$$\text{dens}(P=t/S_j=i) \propto t^i (1-t)^{j-i}$$

So then, I use this private tossings posterior as my new prior, which is called a conjugate prior.

With this conjugate prior, I observe k heads in n tosses. The posterior of this conjugate prior has the same functional form. With this as conjugate prior, the posterior becomes:

$$\text{posterior: } \text{dens}(P=t/S_n=k) \propto t^k (1-t)^{n-k} \cdot \overbrace{t^i (1-t)^{j-i}}^{\text{dens}(P=t)}$$

This is really a lie, because I did not tell you I've already tossed the coin j times, obtaining i heads, before tossing n times & obtaining k heads.

$$\left\{ \begin{aligned} \text{dens}(P=t/S_n=k) &= \frac{P(S_n=k|P=t) \text{dens}(P=t)}{P(S_n=k)} \\ &\propto t^k (1-t)^{n-k} \text{dens}(P=t) \\ &\uparrow \text{ignore factors that integrate to 1} \end{aligned} \right.$$

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



$$\text{posterior} \Rightarrow \text{dens}(p=t|S_n=k) \propto t^{k+i} (1-t)^{n-k+j-i}$$

Therefore, the posterior will be much steeper, if I use this prior.
If you toss a coin - and, you take into account the information from previously tossing the coin, then your posterior is getting steeper (the powers are increased).

So, in the case of coin tossing, the conjugate prior is simply the results of the previous experiments of the same kind.

Again, why do I say proportional?

$$(*) \quad \text{dens}(p=t|S_n=k) \propto t^{k+i} (1-t)^{n-k+j-i}$$

Because I can always integrate the density between 0 and 1 and the result must be 1, so I can determine the proportionality factors.

That's the story of the conjugate prior.

The term conjugate prior arose from this example.

Then it was generalized for balls into boxes, etc.

Intuitively, it means the same thing has happened before (i.e., you are using information of the same kind).

- If you take the Bayes' Estimate of the posterior obtained w/ the conjugate prior, of course you obtain a much sharper estimate:

Bayes' Estimate using conjugate prior:

$$E(p|S_n=k) = \int_{-\infty}^{\infty} t \underbrace{\text{dens}(p=t|S_n=k)} dt$$

obtain proportionality factor by integrating density

posterior obtained from conjugate prior (equation (4.12))
From [4/10/98.12], proportionality constants for density Beta Function Integral:
 $a-1 = k+i \rightarrow a = k+i+1$
 $b-1 = n-k+j-i \rightarrow b = n-k+j-i+1$

$$\frac{(k+i)! (n-k+j-i)!}{(n+j+1)!} = \int_0^1 t^{k+i} (1-t)^{n-k+j-i} dt$$

$$= \int_0^1 \frac{(n+j)!}{(k+i)! (n-k+j-i)!} t^{k+i+1} (1-t)^{n-k+j-i} dt$$

$$= \frac{(n+j)!}{(k+i)! (n-k+j-i)!} \int_0^1 t^{k+i+1} (1-t)^{n-k+j-i} dt$$

This gives:

$$a-1 = k+i+1 \Rightarrow a = k+i+2$$

$$b-1 = n-k+j-i \Rightarrow b = n-k+j-i+1$$

$$\int_0^1 t^{k+i+1} (1-t)^{n-k+j-i} dt = \frac{(k+i+1)! (n-k+j-i)!}{(n+j+2)!}$$

$$= \frac{\cancel{(n+j)!}}{\cancel{(k+i)!} \cancel{(n-k+j-i)!}} \cdot \frac{\cancel{(k+i+1)!} \cancel{(n-k+j-i)!}}{(n+j+2)!}$$

$$E(P|S_n=k) = \frac{k+i+1}{n+j+2}$$

Bayes' Estimate w/ conjugate prior

This is more accurate than Bayes' Estimate w/ uniform prior

$$E(P|S_n=k) = \frac{k+1}{n+1}$$

compare

Let's conclude our discussion about Bayes', for the moment.
We now initiate a new chapter.

- An open question which I was going to give you as a research problem:
Suppose I have the uniform process in the interval $[0, a]$, but
I do not know a .
The research problem consists of 1) doing the Bayes' for a
2) finding out which priors make sense for a

It's a very interesting exercise.

You should try several priors. One of them, the prior is to know that the interval is no bigger than such + such. That's one sensible prior.

A second prior is based on the posterior, but based only on knowledge of the largest points.

The Algebra of Probability Distributions

This is one of the most counterintuitive chapters of probability.

Your intuition has to adjust to reality.

Your intuition goes wrong. Everybody's intuition goes wrong, at first, when dealing w/ this.

Suppose $X \geq 0$ is a continuous random variable that has

$$\text{dens}(X=t) = 0, \quad t < 0$$

(We are taking X non-negative for simplicity of exposition. Not for any other reason.)

Now we take the random variable X^2 .

And we ask, what is the density of the random variable X^2 ?

$$\text{dens}(X^2=t) = ?$$

The temptation to give a stupid answer is wrong.
For example:

$$\text{dens}(X^2=t) \neq [\text{dens}(X=t)]^2$$

Why is it not the density [↑]squared?

Because densities integrate to 1.

We know that:

$$\int_0^{\infty} \text{dens}(X=t) dt = 1$$

Thus:

$$\int_0^{\infty} [\text{dens}(X=t)]^2 dt \neq 1$$

Sorry.

There is one and only one way of doing this.
Let's take the event $(X^2 \leq t)$, where $t \geq 0$

This event is the same as the event:

$$(X^2 \leq t) = (X \leq \sqrt{t}) \quad \leftarrow \text{if you don't see it, I can't explain it.}$$

$$X^2 \leq t \Rightarrow X \leq \sqrt{t}.$$

We're just jazzing it up w/ events. It's the same fact.

Therefore:

$$P(X^2 \leq t) = P(X \leq \sqrt{t})$$

Cumulative probability distribution of X^2

But the cumulative probability distribution is the integral of the density.

$$= \int_0^{\sqrt{t}} \text{dens}(X=s) ds$$

But, my friends, the density is the derivative of the cumulative probability distribution.

$$\text{dens}(X^2=t) = \frac{d}{dt} P(X^2 \leq t)$$

$$= \frac{d}{dt} P(X \leq \sqrt{t})$$

$$= \frac{d}{dt} \int_0^{\sqrt{t}} \text{dens}(X=s) ds$$

Now, we use our finest resources from 18.01.

This is the derivative of a function of a function.

I will not insult you by reminding you of the formula for the derivative of a function of a function.

Right? You all know it inside and out.

$$h(t) = f(g(t)) = \int_0^{\sqrt{t}} \text{dens}(X=s) ds$$

$$\text{where } u = g(t) = \sqrt{t}$$

By the Chain Rule:

$$\frac{d}{dt} h(t) = \frac{d}{du} f(g(t)) \frac{d}{dt} g(t)$$

$$= \frac{df(u)}{du} \frac{du}{dt}$$

$$\textcircled{a} f(u) = \int_0^u \text{dens}(X=s) ds$$

$$= P(X \leq u)$$

$$\frac{df(u)}{du} = \frac{d}{du} P(X \leq u)$$

$$= \text{dens}(X=u)$$

$$= \text{dens}(X=\sqrt{t})$$

$$\textcircled{b} \frac{du}{dt} = \frac{d\sqrt{t}}{dt}$$

$$= \frac{1}{2\sqrt{t}}$$

$$\text{dens}(X^2=t) = \text{dens}(X=\sqrt{t}) \frac{1}{2\sqrt{t}}$$

4/13/98.11

- If we are not sure of our results, we double check it by verifying that it does indeed integrate out to 1.
This is an unnecessary computation, but we are insecure.

$$\int_{-\infty}^{\infty} \text{dens}(X^2=t) dt = \int_0^{\infty} \text{dens}(X^2=t) dt \leftarrow \begin{array}{l} \text{limits of integration 0 and } \infty \\ \text{since } \text{dens}(X^2=t) = 0 \text{ for } t < 0 \end{array}$$

$$= \int_0^{\infty} \text{dens}(X=\sqrt{t}) \frac{1}{2\sqrt{t}} dt$$

Let's substitute:

$$\sqrt{t} = s \Rightarrow t = s^2$$

$$dt = 2s ds$$

$$= \int_0^{\infty} \text{dens}(X=s) \frac{1}{2s} 2s ds$$

$$= \int_0^{\infty} \text{dens}(X=s) ds$$

This is 1, as discussed in [4/13/98.9], so it checks.

$$= 1 \checkmark$$

And that's the way it works.

- More generally:
Suppose we have a random variable X .
Let $g(x)$ be "any" function, in the ordinary sense of the word function.
 $\uparrow$ little x

Then $g(X)$ is a random variable.

You want the density of random variable $g(X)$:

$$\text{dens}(g(X)=t) = ?$$

That's no joke.

In general, it's pretty tough.

We'll work it out only for the following special case.



- Special Case - $X \geq 0$, continuous, and function g has the property:
- $$g'(x) \geq 0$$
- ↑ little x

Since the derivative is non-negative, the function g is strictly increasing.

Under these assumptions, we can repeat the reasoning we just performed for $g(X) = X^2$ word by word. But it's worth doing, because it's such an unusual computation that it takes a while to adjust to it.

$$P(g(X) \leq t) = P(X \leq g^{-1}(t))$$

↑
We always compute the cumulative distribution first, because that's the natural probability.

↑ the inverse function of g

This equality is because the event $\{g(X) \leq t\}$ is the same as the event that $\{X \leq g^{-1}(t)\}$ $g(x)$ is continuous, it is increasing, so there is only one inverse function.

$$= \int_0^{g^{-1}(t)} \text{dens}(X=s) ds \leftarrow \left\{ \begin{array}{l} \text{By definitions of cumulative} \\ \text{distribution and density.} \end{array} \right\}$$

(*)

{ One of the problems you have in the Problem Set due Friday involves this type of computation. }

- Now, may I remind you that inverse function means:

$$g(g^{-1}(t)) = t$$

Therefore, differentiating gives:

$$g'(g^{-1}(t)) \cdot g^{-1}'(t) = t^{t \rightarrow 1}$$

$$\Rightarrow \boxed{g^{-1}'(t) = \frac{1}{g'(g^{-1}(t))}}$$

An important formula

(**)

- In view of this formula, we can take equation (8)

$$P(g(X) \leq t) = \int_0^{g^{-1}(t)} \text{dens}(X=s) ds$$

and differentiate both sides:

$$\underbrace{\frac{d}{dt} P(g(X) \leq t)}_{\text{dens}(g(X)=t)} = \frac{d}{dt} \int_0^{g^{-1}(t)} \text{dens}(X=s) ds$$

Identical to [4/13/98.10], we use the Chain Rule:

$$h(t) = r(s(t)) = \int_0^{g^{-1}(t)} \text{dens}(X=s) ds$$

$$u = s(t) = g^{-1}(t)$$

$$\frac{d}{dt} h(t) = \frac{d r(s(t))}{d s(t)} \cdot \frac{d s(t)}{d t}$$

$$= \underbrace{\frac{d r(u)}{d u}}_{\textcircled{2}} \cdot \underbrace{\frac{d u}{d t}}_{\textcircled{1}}$$

$$\textcircled{2} \quad r(u) = \int_0^u \text{dens}(X=s) ds$$

$$= P(X \leq u)$$

$$\frac{d r(u)}{d u} = \text{dens}(X=u)$$

$$= \text{dens}(X=g^{-1}(t))$$

$$\textcircled{1} \quad \frac{d u}{d t} = g^{-1}'(t)$$

equation (8*)
gives:

$$= \frac{1}{g' g^{-1}(t)}$$

$$= \text{dens}(X=g^{-1}(t)) \frac{1}{g' g^{-1}(t)}$$

And we have:

$$\boxed{\text{dens}(g(X)=t) = \text{dens}(X=g^{-1}(t)) \cdot \frac{1}{g' g^{-1}(t)}}$$

That's the formula for the density of a function, under the given assumptions.

- You can check, just like in the previous example, that this integrates to 1. It's almost a miracle that this integrates to 1 automatically.

- Now, as a last word.
Suppose you have to find the density:

$$\text{dens}(X=s|Y^2=t) = ?$$

This implies a fairly delicate computation w/ the cumulative distribution.
You can not just close your eyes and square.

$$\text{dens}(X=s|Y^2=t) \neq [\text{dens}(X=s|Y=t)]^2$$

This is the kind of problem we have for Friday.
Don't wait until the last minute.

The only way to do this is to ^①rephrase everything in terms of probabilities.
For example, in terms of:

$$P(s < X \leq s+\Delta s | t < Y^2 \leq \Delta t)$$

} something like that

or in terms of cumulative distributions.

Then, once you have things in terms of probabilities, you are on safe ground.
Then, when the time is right, you differentiate

②

Such densities are highly non-trivial problems.

- { Next time, we'll continue w/ the Algebra of Probability Distributions.
We'll discuss the notion of convolution. }

The Algebra of Probability Densities (Cont'd)

X = continuous random variable w/ density $\text{dens}(X=t)$.
Then random variable $f(X)$ has a density that must be computed extremely delicately. $\uparrow$ {where f is a suitable function}

$$\text{dens}(f(X)=t) = ?$$

$\uparrow$ In particular, one has to guard against oversimplification.

What we saw is that $\text{dens}(f(X)=t)$ has to be computed by going back to the cumulative distribution and reasoning w/ the cumulative distribution.

We worked this out, last time, where

$X \geq 0$, X is a non-negative random variable
 $f(x) > 0$, $f(x)$ is a strictly increasing function

In more general cases, which we did not work out, working it out is kind of a mess.

It's one of the messiest computations in probability.

Fortunately, in special cases, there is always a way out, which is dictated by the example under consideration.

So, you can try to avoid such messy computations.

So let's do one last example.

Suppose we take the random variable $3X$.
What is:

$$\text{dens}(3X=t) \text{ in terms of } \text{dens}(X=t) = ?$$

It is not $\frac{1}{3} \text{dens}(X=t)$, because it doesn't integrate to 1

$$\int_{-\infty}^{\infty} \frac{1}{3} \text{dens}(X=t) = \frac{1}{3} \neq 1$$

And the density must integrate to 1.

So that's wrong.

This is what I call an oversimplification.

Well, the following events are the same.

$$(3X \leq t) = (X \leq \frac{t}{3})$$

If you don't see it, I can't explain it.

Therefore:

$$P(3X \leq t) = P(X \leq \frac{t}{3})$$

cumulative distribution
of the random variable
 $3X$.

$$= \int_{-\infty}^{\frac{t}{3}} \text{dens}(X=s) ds$$

Taking derivatives of both sides:

$$\frac{d}{dt} P(3X \leq t) = \frac{d}{dt} \int_{-\infty}^{\frac{t}{3}} \text{dens}(X=s) ds$$

The derivative of this
cumulative distribution
is $\text{dens}(3X=t)$, by
the fundamental theorem
of the calculus:

$$\frac{d}{dx} \int_a^x f(t) dt = f(x)$$

$$\frac{d}{dt} \int_{-\infty}^t \text{dens}(3X=s) ds = \text{dens}(3X=t)$$

$$h(t) = f(g(t)) = \int_{-\infty}^{\frac{t}{3}} \text{dens}(X=s) ds$$

$$\text{where } u = g(t) = \frac{t}{3}$$

$$h'(t) = \underbrace{f'(g(t))}_{(a)} + \underbrace{g'(t)}_{(b)}$$

$$(a) \quad f(u) = \int_{-\infty}^u \text{dens}(X=s) ds = P(X \leq u)$$

$$(b) \quad g(t) = \frac{t}{3} \\ g'(t) = \frac{1}{3}$$

$$\begin{aligned} \frac{d}{du} f(u) &= \frac{d}{du} P(X \leq u) \\ &= \text{dens}(X=u) \\ &= \text{dens}(X = \frac{t}{3}) \end{aligned}$$

By the Chain Rule:

$$\frac{d}{dt} \int_{-\infty}^{\frac{t}{3}} \text{dens}(X=s) ds = \text{dens}(X = \frac{t}{3}) \cdot \frac{1}{3}$$

$$\boxed{\text{dens}(3X=t) = \text{dens}(X = \frac{t}{3}) \cdot \frac{1}{3}}$$

And this integrates to 1.

$$\int_{-\infty}^{\infty} \text{dens}(3X=t) dt = \int_{-\infty}^{\infty} \text{dens}(X = \frac{t}{3}) \cdot \frac{1}{3} dt = 1$$

4/15/98.3

This is the density of $3X$.

I don't know of anyone who has an intuitive grasp of this computation w/ density.
It involves a very deep understanding of inverse functions.
There is something philosophically, extremely deep in this computation.

Don't let me get into it.
Because I'll never stop.

Let's do some other computations of this kind, so you get used to it.
Let's do a couple of computations w/ joint densities.

Joint Densities

Suppose we have continuous random variables X, Y , with given joint density: $\text{dens}(X=s, Y=t)$

$$\left\{ \begin{array}{l} \text{Recall, the marginal density:} \\ \text{dens}(X=s) = \int_{-\infty}^{\infty} \text{dens}(X=s, Y=t) dt \\ \text{in other words, if you wish one of the marginal densities,} \\ \text{you integrate out the other. You let the other random} \\ \text{variable take any value it could, which amounts to integrating.} \end{array} \right\}$$

Given this joint density, we want:

$$\text{dens}(X^2=u, Y=t) = ?$$

How do we compute that one?

Don't give me an over-simplified answer.

There is no way of guessing - actually some physicists guess and get it right,
How they do it, I don't know.

I guess the safest way to work this out is to reduce this to the case of one variable.

How do we do that?

Simple. We conditionalize it. That's easy.

Note that we have the following identity w/ densities, analogous to conditional probability:

$$\text{dens}(X=s | Y=t) = \frac{\text{dens}(X=s, Y=t)}{\text{dens}(Y=t)}, \quad \text{dens}(Y=t) \neq 0$$

by definition of density:

4/15/98.4

Proof

$$\text{dens}(X=s|Y=t) = \frac{d}{ds} P(X \leq s | Y=t)$$

In the limiting process, the following events are equal:

$$\lim_{\Delta t \rightarrow 0} (t < Y \leq t + \Delta t) = (Y=t)$$

$$= \frac{d}{ds} \lim_{\Delta t \rightarrow 0} P(X \leq s | t < Y \leq t + \Delta t)$$

this can be expressed as an ordinary conditional probability.

$$= \frac{d}{ds} \lim_{\Delta t \rightarrow 0} \frac{P((X \leq s) \cap (t < Y \leq t + \Delta t))}{P(t < Y \leq t + \Delta t)}$$

Only the numerator involves s.

$$= \lim_{\Delta t \rightarrow 0} \frac{\frac{d}{ds} P((X \leq s) \cap (t < Y \leq t + \Delta t))}{P(t < Y \leq t + \Delta t)}$$

Multiply by $1 = \frac{\frac{1}{\Delta t}}{\frac{1}{\Delta t}}$

$$= \lim_{\Delta t \rightarrow 0} \frac{\frac{d}{ds} \frac{P((X \leq s) \cap (t < Y \leq t + \Delta t))}{\Delta t}}{\frac{P(t < Y \leq t + \Delta t)}{\Delta t}}$$

$$= \lim_{\Delta t \rightarrow 0} \frac{\frac{\partial}{\partial s} \frac{P((X \leq s) \cap (Y \leq t + \Delta t)) - P((X \leq s) \cap (Y \leq t))}{\Delta t}}{\frac{P(Y \leq t + \Delta t) - P(Y \leq t)}{\Delta t}}$$

The numerator is precisely $\frac{\partial^2}{\partial s \partial t} P((X \leq s) \cap (Y \leq t))$

The denominator is precisely $\frac{d}{dt} P(Y \leq t)$

$$= \frac{\frac{\partial^2}{\partial s \partial t} P((X \leq s) \cap (Y \leq t))}{\frac{d}{dt} P(Y \leq t)} \left\{ \begin{array}{l} \text{definition of the joint} \\ \text{density } \text{dens}(X=s, Y=t) \end{array} \right\}$$

$$\left\{ \begin{array}{l} \text{definition of the} \\ \text{density } \text{dens}(Y=t) \end{array} \right\}$$

(*)

$$\text{dens}(X=s|Y=t) = \frac{\text{dens}(X=s, Y=t)}{\text{dens}(Y=t)}$$

QED.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



4/15/98.5

- With this identity in hand, we return to finding the joint density $\text{dens}(X^2=u, Y=t)$. This can be written as follows, using the previous identity.

$$\text{dens}(X^2=u, Y=t) = \text{dens}(X^2=u | Y=t) \text{dens}(Y=t)$$

assume $X \geq 0$

Density is an infinitesimal analogue of conditional probability, as we've discussed.

We've already computed $\text{dens}(X^2=u)$ [4/13/98.10]. Computing the conditional density follows exactly the same steps:

$$\begin{aligned} \text{dens}(X^2=u | Y=t) &= \frac{d}{du} P(X^2 \leq u | Y=t) \\ &= \frac{d}{du} P(X \leq \sqrt{u} | Y=t) \\ &= \frac{d}{du} \int_0^{\sqrt{u}} \text{dens}(X=s | Y=t) ds \\ &= \text{dens}(X=\sqrt{u} | Y=t) \cdot \frac{1}{2\sqrt{u}} \end{aligned}$$

$$\begin{aligned} &\text{dens}(X^2=u | Y=t) \\ &= \overbrace{\text{dens}(X=\sqrt{u} | Y=t)}^{\text{dens}(X^2=u | Y=t)} \cdot \frac{1}{2\sqrt{u}} \cdot \text{dens}(Y=t) \\ &= \frac{1}{2\sqrt{u}} \underbrace{\text{dens}(X=\sqrt{u} | Y=t)}_{\text{dens}(X^2=u | Y=t)} \cdot \text{dens}(Y=t) \end{aligned}$$

Using the general identity (*), we have:

$$\text{dens}(X=\sqrt{u} | Y=t) = \frac{\text{dens}(X=\sqrt{u}, Y=t)}{\text{dens}(Y=t)}$$

$$\text{dens}(X^2=u, Y=t) = \frac{1}{2\sqrt{u}} \text{dens}(X=\sqrt{u}, Y=t)$$

So we got it.

You have some problems like this in the problem set.

I can see no one has looked at the problem set yet.

You are waiting until Thursday night, as usual. Right?

Thursday night at 2:00am.

That's shocking.

Even though I did it too.

It's still shocking.



- Let's do another one.
Again, X and Y being as in the previous problem.
 X and Y are not independent.
Let's compute:

$$\text{dens}(X+Y=u) = ?$$

- First - an intuitive, physicist type argument.

The event $(X+Y=u)$ is meaningless, because it is an event of probability 0.
 $P(X+Y=u) = 0$.

$$(X+Y=u) \text{ iff } (X=s, Y=u-s) \text{ for some } s$$

$$\text{dens}(X+Y=u) = \int_{-\infty}^{\infty} \text{dens}(X=s, Y=u-s) ds \leftarrow \{\text{integrate out the } s\}$$

This is what we guess.
Now, let's check it by working out probabilities and then differentiating.

- Second - formally

by definition:

$$P(X+Y \leq u) = \iint \text{dens}(X=s, Y=t) ds dt$$

Now we are talking sets. The event $(X+Y \leq u)$ is meaningful.

Now, the painful part is to choose the domain of integration.

The domain of integration can be expressed set theoretically as:

$$\{s, t : s+t \leq u\}$$

Which doesn't help that much.
But at least you get some idea.

So we do a change of variables.
 $s+t \leq u \Rightarrow t \leq u-s$

$$P(X+Y \leq u) = \int_{t=-\infty}^u \int_{s=-\infty}^{\infty} \text{dens}(X=s, Y=t-s) ds dt$$

Therefore, we can take the derivative of both sides:

$$\frac{d}{du} P(X+Y \leq u) = \frac{d}{du} \int_{t=-\infty}^u \int_{s=-\infty}^{\infty} \text{dens}(X=s, Y=t-s) ds dt$$

this is:

$$\text{dens}(X+Y=u)$$

From the 2nd Fundamental Theorem of the Calculus:

$$\frac{d}{du} \int_a^u f(t) dt = f(u)$$

$$\text{dens}(X+Y=u) = \int_{-\infty}^{\infty} \text{dens}(X=s, Y=u-s) ds$$

which is the same as our intuitive argument, as desired.

As you suspect, this stuff can get pretty rough.
You can take one whole term of measure theory.
That's what this is all about.

Special Case - suppose X and Y are independent

By definition, if X and Y are independent:

$$P(X \leq s) \cap (Y \leq t) = P(X \leq s) \cdot P(Y \leq t)$$

$$\frac{\partial^2}{\partial s \partial t} P(X \leq s) \cap (Y \leq t) = \frac{\partial}{\partial t} \left[\frac{\partial}{\partial s} P(X \leq s) P(Y \leq t) \right]$$

$$\text{dens}(X=s, Y=t) = \frac{\partial}{\partial t} \left[P(Y \leq t) \underbrace{\frac{\partial}{\partial s} P(X \leq s)}_{\text{dens}(X=s)} \right]$$

$$= \text{dens}(X=s) \underbrace{\frac{\partial}{\partial t} P(Y \leq t)}_{\text{dens}(Y=t)}$$

$$\text{dens}(X=s, Y=t) = \text{dens}(X=s) \text{dens}(Y=t)$$



- Under this special case, where X and Y are independent, the formula for the density of the sum, equation (**), becomes something very interesting.

$$\begin{aligned} \text{dens}(X+Y=u) &= \int_{-\infty}^{\infty} \underbrace{\text{dens}(X=s, Y=u-s)}_{\text{independent}} ds \\ &= \int_{-\infty}^{\infty} \text{dens}(X=s) \text{dens}(Y=u-s) ds \end{aligned}$$

We have seen this integral in 18.03.

This is the convolution of two functions.

That's where convolution comes from.

We have therefore found a meaning to convolution.

$$\begin{aligned} \text{Let } h(u) &= \text{dens}(X+Y=u) \\ f(s) &= \text{dens}(X=s) \\ g(u-s) &= \text{dens}(Y=u-s) \end{aligned}$$

$$\text{Then: } h(u) = \int_{-\infty}^{\infty} f(s) g(u-s) ds$$

- In particular, if X and $Y \geq 0$, then:

$$(*) \quad \text{dens}(X+Y=t) = \int_0^t \text{dens}(X=s) \text{dens}(Y=t-s) ds$$

$\left\{ \begin{array}{l} X \text{ and } Y \geq 0 \\ \text{cuts down limits of} \\ \text{integration.} \\ \text{If } X < 0, \text{ dens}(X=s) = 0 \end{array} \right\}$

This is the convolution of two functions, as per 18.03.

The whole of the Laplace transform can be interpreted probabilistically.

And we will do that.

It's too bad we can't do that in 18.03 when we just had to dish out: This is convolution, take it home.

- Now you see that:

Convolution corresponds to taking the sum of two independent random variables.

$$\underbrace{\int_0^t \text{dens}(X=s) \text{dens}(Y=t-s) ds}_{\text{Convolution of their densities}} = \underbrace{\text{dens}(X+Y=t)}_{\text{density of their sums}}$$

4/15/98.9

This is very interesting, if you ask me.

• Tough Example

Let X, Y be independent random variables, uniformly distributed in $[0, a]$.

$$\text{dens}(X=s) = \begin{cases} \frac{1}{a} & \text{if } s \in [0, a] \\ 0 & \text{otherwise} \end{cases}$$

We will do this w/ indicator functions, which I discussed in SuperClass 1 [SC 2/27/98.1-4].

A short review.
(Open parentheses:

Indicator random variables

This is a pretty useful set theoretic device for computing probabilities.

A = an event

I_A = indicator random variable

$$= \begin{cases} 1 & \text{if } A \text{ happens} \\ 0 & \text{if } A \text{ does not happen} \end{cases}$$

} Probabilistic

How do you change this into set theoretic language?
This is just a probabilistic way of saying that:

$$I_A(\omega) = \begin{cases} 1 & \text{if } \omega \in A \\ 0 & \text{if } \omega \notin A \end{cases}$$

} Set Theoretic

read: "I sub A of sample point ω "

The probabilistic version is superior.

Because it gives you a better understanding of set theory of events.

As John von Neumann, who contributed greatly in developing probability of quantum mechanics, said:

"Probability must be pointless."

Whenever you use sample points, your argument is not perfect.

There is a mathematician who developed all of probability pointlessly. Very interesting reading. Some things are a little tougher to define - like convolution, w/o using points.



I_A = indicator random variable of the event A

You will find in the book the algebra of indicator random variables and the relationship bt. Boolean algebra and the algebra of indicator random variables. I leave it to you to look up.

For the present purpose, the most important thing to observe is this [Sc 2/27/98.2]:

$$\begin{aligned} E(I_A) &= \sum_n n P(I_A = n) \\ &= 1 \cdot P(I_A = 1) + 0 \cdot P(I_A = 0) \\ &= P(A) \quad \leftarrow \text{since } P(I_A = 1) = P(A) \end{aligned}$$

$E(I_A) = P(A)$

Namely, the probability of event A equals the expectation of the indicator random variable.

The use of indicator random variables is useful, because it reduces the computation of probability to the computation of expectation. And expectations are easier to deal w/ than probabilities.

So much for indicator random variables.
Close parentheses)

- Now back to our problem.
We want to write the density of X and of Y using indicator random variables.

$$\text{dens}(X=s) = \frac{1}{a} \underbrace{I_{[0,a]}(s)}_{\text{Indicator random variable on the interval } [0,a] \text{ of } s.}$$

$$\text{dens}(Y=t) = \frac{1}{a} I_{[0,a]}(t)$$

Therefore, from equation (*):

$$\text{dens}(X+Y=t) = \int_0^t \text{dens}(X=s) \text{dens}(Y=t-s) ds$$

Now we have to work in the above formulas for the densities,
Careful!

Watch me carefully.

$$= \int_0^t \frac{1}{a} I_{[0,a]}(s) \frac{1}{a} I_{[0,a]}(t-s) ds$$

$$= \int_0^t \frac{1}{a^2} I_{[0,a]}(s) I_{[0,a]}(t-s) ds$$

Now, there are 2 cases.

Another von Neumann story:

Everytime he went to a lecture and the lecturer said "Now there are 3 cases," von Neumann got up + left. He couldn't stand it.

We are lucky here. We have only 2.

• Case 1: $t \leq a$

$$I_{[0,a]}(s) = 1 \leftarrow \left\{ \begin{array}{l} s \text{ ranges from } 0 \text{ to } t. \\ \text{Since } t \leq a, \text{ we are assured} \\ \text{that the event } s \in [0,a] \text{ happens.} \end{array} \right\}$$

$$\text{dens}(X+Y=t) = \int_0^t \frac{1}{a^2} \cancel{I_{[0,a]}(s)}^1 I_{[0,a]}(t-s) ds$$

$$\text{dens}(X+Y=t) = \int_0^t \frac{1}{a^2} \underbrace{\cancel{I_{[0,a]}(t-s)}^1}_{I_{[0,a]}(t-s)} ds$$

$$I_{[0,a]}(t-s) = 1 \leftarrow \left\{ \begin{array}{l} s \in [0,t] \\ \text{so } t-s \in [0,t] \text{ and} \\ \text{since } t \leq a, \text{ we are} \\ \text{assured that the event} \\ t-s \in [0,a] \text{ happens.} \end{array} \right\}$$

$$= \int_0^t \frac{1}{a^2} ds$$

$$\boxed{\text{dens}(X+Y=t) = \frac{t}{a^2}, \quad t \leq a}$$



4/15/98.12

Case 2 : $a < t \leq 2a$ $I_{[0,a]}(s) = 1$, as long as $s \in [0, a]$, 0 otherwise

$$\text{dens}(X+Y=t) = \int_0^t \frac{1}{a^2} I_{[0,a]}(s) I_{[0,a]}(t-s) ds$$

$$= \int_0^a \frac{1}{a^2} I_{[0,a]}(t-s) ds$$

restrict limits of integration
to domain where $I_{[0,a]}(s) = 1$
It is 0 everywhere else.

Indicator function $I_{[0,a]}(t-s) = 1$ when: $0 \leq t-s \leq a$ $\Rightarrow$ Lower bound on s :

$$t-a \leq s$$

Upper bound on s :

$$s \leq t$$

$$I_{[0,a]}(t-s) = \begin{cases} 1 & \text{if } s \in [t-a, t] \\ 0 & \text{otherwise} \end{cases}$$

Note that the limits of integration are currently

 $\int_0^a$. Due to the behavior of $I_{[0,a]}(t-s)$,

we can change the limits:

$$\int_0^a \Rightarrow \int_{t-a}^a$$

$$I_{[0,a]}(s) = 1$$

$$I_{[0,a]}(t-s) = 1$$

$$I_{[0,a]}(s) = 1$$

AND

$$I_{[0,a]}(t-s) = 1$$

$$= \int_{t-a}^a \frac{1}{a^2} ds$$

$$= \frac{1}{a^2} (a - (t-a))$$



4/15/98.13

$$= \frac{2a - t}{a^2}$$

$$\text{dens}(X+Y=t) = \frac{2}{a} - \frac{t}{a^2}, \quad a < t \leq 2a$$

So we have:

$$\text{dens}(X+Y=t) = \begin{cases} \frac{t}{a^2} & , t \leq a \\ \frac{2}{a} - \frac{t}{a^2} & , a < t \leq 2a \\ 0 & , \text{o.w.} \end{cases}$$

That's the density.

Now where is the proof of the pudding?

Integrating the density has to equal 1.

If it doesn't equal 1, I've made a mistake.

Check:

$$\int_{-\infty}^{\infty} \text{dens}(X+Y=t) dt \stackrel{?}{=} 1$$

$$\int_{-\infty}^{\infty} \text{dens}(X+Y=t) dt = \int_0^{2a} \text{dens}(X+Y=t) dt$$

$$= \int_0^a \underbrace{\text{dens}(X+Y=t)}_{\frac{t}{a^2}} dt + \int_a^{2a} \underbrace{\text{dens}(X+Y=t)}_{\frac{2}{a} - \frac{t}{a^2}} dt$$

$$= \left. \frac{t^2}{2a^2} \right|_{t=0}^a + \left[\frac{2t}{a} - \frac{t^2}{2a^2} \right]_{t=a}^{2a}$$

$$= \frac{1}{2} + \cancel{4} - \cancel{2} - \cancel{2} + \frac{1}{2}$$

$$= 1 \quad \checkmark$$

So it is a density.



Now you see what you get into in this business.

- Research Problem:
Add n independent, uniformly distributed random variables.
Find a nice formula for:

$$\text{dens}(X_1 + X_2 + \dots + X_n = t) = ?$$

I want to do a seemingly disparate topic.

- Cauchy's functional equation

Suppose I have a function f that is continuous and satisfies the following functional equation:

$$f(x+y) = f(x)f(y)$$

for $x, y \geq 0$

What can you say about f ?
What you can say is this.

$$\text{Then } f(x) = c^x$$

↑ some constant c ,
raised to the x .

Why?

For the following reason.

First of all:

$$\begin{aligned} f(x+x) &= f(x)f(x) \\ &= c^x c^x \\ &= c^{2x} \end{aligned}$$

$$f(nx) = [f(x)]^n \quad \left. \begin{array}{l} \text{simply by iteratively applying} \\ f(x+y) = f(x)f(y) \end{array} \right\}$$

to:

$$f(\underbrace{x+x+\dots+x}_n) = \underbrace{f(x)f(x)\dots f(x)}_n$$

If we set $x = \frac{1}{n}$,

$$f(n \cdot \frac{1}{n}) = [f(\frac{1}{n})]^n$$

$$f(1) = [f(\frac{1}{n})]^n$$

And since $f(x) = c^x$,

$$f(1) = c$$

$$\text{Therefore: } f(1) = c = [f(\frac{1}{n})]^n \Rightarrow [f(1)]^{\frac{1}{n}} = c^{\frac{1}{n}} = f(\frac{1}{n})$$

Now, take:

$$f\left(\frac{k}{n}\right) = f\left(\underbrace{\frac{1}{n} + \frac{1}{n} + \dots + \frac{1}{n}}_k\right)$$

Therefore, you can apply the functional equation k times.

$$= \left[f\left(\frac{1}{n}\right)\right]^k$$

Using the previously proved identity;

$$\left[f(1)\right]^{\frac{1}{n}} = c^{\frac{1}{n}} = f\left(\frac{1}{n}\right)$$

$$= \left[f(1)\right]^{k/n} = c^{k/n}$$

$$\boxed{f\left(\frac{k}{n}\right) = c^{k/n}}$$

Therefore, for every rational number $\frac{k}{n}$, we have the conclusion that this f is assumed to be continuous.

Every number is a limit of rational numbers, so the above holds for all numbers.

$$\left. \begin{array}{l} f(x+y) = f(x)f(y) \\ x, y \geq 0 \end{array} \right\} \Rightarrow f(t) = c^t \text{ for all numbers } t$$

This is a very important fact.

The only function that satisfies the functional equation $f(x+y) = f(x)f(y)$ is $f(x+y) = c^{x+y}$

• An alternative proof ← a cheapo proof, as done by physicists

Assume, in addition to being continuous, that f is differentiable.

Take the functional equation and take 1) the derivative w.r.t x on both sides, and then 2) " " " y on both sides.

$$\left. \begin{array}{l} 1) f'(x+y) = f'(x)f(y) \\ 2) f'(x+y) = f(x)f'(y) \end{array} \right\} \frac{d}{dx} f(x+y) = \frac{d}{dy} f(x+y)$$

Combining gives:

$$f'(x)f(y) = f(x)f'(y)$$

Blindly assume $f(x) \neq 0$.

$$\frac{f'(x)}{f(x)} = \frac{f'(y)}{f(y)}$$

Now we apply the Principle of Ignorance.
 x and y do not know of each other's existence.

Since variables don't think, the only way to get the LHS = the RHS is
 if LHS = RHS = a constant:

$$\underbrace{\frac{f'(x)}{f(x)}}_{\text{only } x} = \underbrace{\frac{f'(y)}{f(y)}}_{\text{only } y} = k$$

Solve the differential equation:

$$f'(x) - kf(x) = 0$$

to obtain the exponential:

$$f(x) = C^x$$

$$\int_0^x \frac{f'(t)}{f(t)} dt = \int_0^x k dt$$

$$\text{Let } u = f(t)$$

$$\int_{t=0}^x \frac{1}{u} du = kt \Big|_0^x$$

$$\ln u \Big|_{t=0}^x = kx$$

$$\ln f(x) - \ln f(0) = kx$$

$$\ln \frac{f(x)}{f(0)} = kx$$

$$f(x) = \underbrace{f(0)} e^{kx}$$



Today: Some examples

Quiz: April 27 (Monday)

The examples are motivated by the fact that the marathon is coming up on Monday.

• Example 1 - "Record Values"

If you watched the Winter Olympics, it was kind of exciting.
In speed skating, they have these new kind of skates, the clap skates.
Every time they ran a race, they'd get a new world's record.
Then the next race, they'd get another new world's record.
And so on.

Let's try to analyze how often this happens, where you have a lot of turnover in world records.


Or, another way of saying this in the opposite direction, how long will the world record you just set last.

Let $X_1, X_2, \dots$ be an infinite number of independent and identically distributed (i.i.d.) continuous random variables
(we don't care about their distribution)
that represent (skater's time).

Let's assume that the X_i happen sequentially: X_1 happens, then X_2 , etc.

Let $N = n$, which is the first time $X_n > X_1$ (i.e., that X_n breaks X_1 's world record).

For example:



Here, $N=5$, since X_5 is the first $> X_1$.

The event $(N=n)$ is:


$$(N=n) = (X_2 < X_1) \cap (X_3 < X_1) \cap \dots \cap (X_{n-1} < X_1) \cap (X_n > X_1)$$

Rather than compute this probability directly, let's get at it indirectly, by considering an easier probability. Namely the probability of the event $(N > n)$.

The event $(N > n) = X_1$, the first racer (the first label), is larger than points at labels $X_2, \dots, X_n$.
You just need to consider the permutation of labels.

Event $(N > n)$:

4/17/98.2



$$P(N > n) = \frac{(n-1)!}{n!}$$

Total number of ways to permute labels X_2 through X_n , while keeping X_1 fixed.

$$= \frac{1}{n}$$

Total number of ways of permuting labels $X_1, \dots, X_n$.

Consider the event $(N > n-1)$. This is equivalent to:

$$(N > n-1) = \underbrace{(N=n) \cup (N > n)}_{\text{union of disjoint events}}$$

Taking probabilities:

$$P(N > n-1) = P(N=n) + P(N > n)$$

Which we can rearrange to give the probability of the event we first considered $(N=n)$:

$$\begin{aligned} P(N=n) &= P(N > n-1) - P(N > n) \\ &= \frac{1}{n-1} - \frac{1}{n} \end{aligned}$$

$$P(N=n) = \frac{1}{n(n-1)}$$

Once we have this, we can then calculate the expected value:

$$E(N) = \sum_{n=2}^{\infty} n P(N=n)$$

lower limit is 2, since you can't break a record unless a first skater sets it.

$$= \sum_{n=2}^{\infty} n \frac{1}{n(n-1)}$$

$$\approx \infty$$

4/17/98.3

$$E(N) \simeq \infty$$

X_1 holds the record and the expected amount of time that X_1 continues to hold the record is infinite.

$$(N = \infty) \Rightarrow X_{\infty} > X_1$$

{ there are an infinite number of points until we have the first occurrence of a point $> X_1$ }

So the fact that in the Olympics, these speed skaters kept skating new world records was actually quite a freak. There were people who had records all of 2 minutes. They should have expected their records to last forever.

{ Note that the differing athletic prowess, abilities, and performance capacities were not taken into consideration in arguing $P(N > n) = \frac{1}{n}$. We've blithely assumed all athletes are equal, which is a gross oversimplification. }



4/17/98.4


- Example 2 - " $E(X+C) \neq E(X)+C$ " $\leftarrow C$ is a constant
 $\uparrow$ this should shock you.

Take $a = \frac{3}{\pi^2}$

Let $X =$ continuous random variable with density:

$$\text{dens}(X=x) = f(x) = \begin{cases} a \frac{1}{|k|} (1 - |k-x|) & \text{for all } k = \pm 1^2, \pm 2^2, \dots \text{ and } |k-x| \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

k all squared integers



Just looking at this, you can guess the expected value of X , because of the symmetry about $x=0$.

$$E(X) = \int_{-\infty}^{\infty} x \overbrace{\text{dens}(X=x)}^{f(x)} dx = 0$$

This gives us that:

$$E(X) + C = 0 + C$$

$$\boxed{E(X) + C = C}$$

Next, I have to show you that the expected value of $X+C$ is something different.

First, we note that $f(x)$ is a valid density function.
 Check that integration equals 1.

$$\int_{-\infty}^{\infty} f(u) du \stackrel{?}{=} 1$$




$$\int_{-\infty}^{\infty} f(u) du = \text{area under the curve above}$$

Because $f(x)$ is symmetric about $x=0$,
we can focus on the curve to the right
of the origin.

$$= 2 \int_0^{\infty} f(u) du$$

Each triangle for $k=i^2$, $i=1, \dots, \infty$,
has area:



$$\begin{aligned} \text{Area} &= \frac{1}{2} \cdot 2 \cdot \frac{a}{i^2} \\ &= \frac{a}{i^2} \end{aligned}$$

$$= 2 \sum_{i=1}^{\infty} \frac{a}{i^2}$$

$$= 2a \underbrace{\sum_{i=1}^{\infty} \frac{1}{i^2}}$$

This is one of those beautiful identities
you can prove compliments of the Fourier Series expansion of $x+x^2$.

$$\sum_{i=1}^{\infty} \frac{1}{i^2} = \frac{\pi^2}{6}$$

$$= 2a \frac{\pi^2}{6}$$

And note that we had conveniently set

$$a = \frac{3}{\pi^2}$$

which ensures that the above equals 1.

$$\int_{-\infty}^{\infty} f(u) du = 1 \quad \checkmark$$

which confirms that $f(x)$ is a valid density function.

4/17/98.6

$$E(X+C) = \int_{-\infty}^{\infty} (x+C) f(x+C) d(x+C)$$

$$= \lim_{b \rightarrow \infty} \int_{-b-C}^{b-C} (u+C) f(u) du$$

We break this integral up into 3 parts:

$$= \lim_{b \rightarrow \infty} \left[\underbrace{\int_{-b-C}^{b+C} u f(u) du}_{\text{part 1}} - \underbrace{\int_{-b-C}^{b+C} u f(u) du}_{\text{part 2}} + \underbrace{\int_{-b-C}^{b-C} C f(u) du}_{\text{part 3}} \right]$$

Since $f(u)$ is symmetric about $x=0$, and the limits of integration are negatives of each other, this equals 0.
 $0 = E(X)$

$$\lim_{b \rightarrow \infty} C \int_{-b-C}^{b-C} f(u) du$$

This integral evaluates to 1 in the limit, as $f(u)$ is a valid density


= C

In order that $E(X+C) \neq E(X) + C$, this integral must be $\neq 0$.

The trick is to note the following:

$$\int_{k-1}^{k+1} u f(u) du \geq \underbrace{(k-1)}_{\text{lowest limit}} \int_{k-1}^{k+1} f(u) du \quad \Leftarrow \quad \left\{ \begin{array}{l} \int_{k-1}^{k+1} u f(u) du \geq \int_{k-1}^{k+1} (k-1) f(u) du \end{array} \right.$$

for any $k=2, 3, 4, \dots$, the above holds true



$$\begin{aligned} \int_{k-1}^{k+1} f(u) du &= \text{the area of the triangle} \\ &= \frac{1}{2} \geq \frac{a}{k} \\ &= \frac{a}{k} \end{aligned}$$

$$\boxed{\int_{k-1}^{k+1} u f(u) du \geq (k-1) \frac{a}{k} > 0}$$

for any $k=2, 3, 4, \dots$

4/17/98.7

With the knowledge that $\int_{k-1}^{k+1} u f(u) du > 0$ (i.e., the limits of integration are those of the triangle centered at $x = k$) for $k = 2^2, 3^2, \dots$,
 Let's consider, again, the segment:

$$\lim_{b \rightarrow \infty} \int_{b-C}^{b+C} u f(u) du$$

Let's impose that $C > 1$.

1) If, for some large value of b , there is a squared integer k , such that $[k-1, k+1] \subseteq [b-C, b+C]$,

then
$$\int_{b-C}^{b+C} u f(u) du > 0$$



In other words, the area of the triangle for k is included in the above integral.


Since this area is non-zero, this lends credence to the claim that:

$$E(X+C) \neq E(X) + C$$

2) BUT

If $[k-1, k+1] \not\subseteq [b-C, b+C]$, in other words $[b-C, b+C]$ does not cover any triangle, then:

$$\int_{b-C}^{b+C} u f(u) du = 0$$



So, in fact, the limit:


$$\lim_{b \rightarrow \infty} \int_{b-C}^{b+C} u f(u) du$$

does not exist. The above two cases are disparate and, as a result, the above limit is not defined.

• Example 3 — "Runners on a track"


You have $(n-1)$ runners on a circular track of circumference a .

Let $G_1, G_2, \dots, G_{n-1}$ be gaps between runners




Set up a camera near the track
What is the distribution of runners near the camera?
(i.e., what is the distribution of gaps nearest the camera?)

Let G_{n-1} be the gap nearest the camera



We can think of this as a uniform process, where we snip the circle at the point closest to the camera.


↓ unrolling



$G_1, G_2, \dots, G_{n-2}$ are equidistributed according to the same distribution as L_1 . $P(G_1 > t) = \dots = P(G_{n-2} > t) = P(L_1 > t)$.

↖ first gap in the uniform process

see [3/18/98, 12-13, 3/30/98, 11-12]



$$\text{event } (L_1 > t) = (X_{(1)} > t)$$

$$P(G_1 > t) = P(L_1 > t) = P(X_{(1)} > t)$$

$$= \left(\frac{a-t}{a} \right)^{n-1}$$

$P(L_1 > t)$, all $(n-1)$ points drop in here




4/17/98.9


$$P(G_1 > t) = P(G_2 > t) = \dots = P(G_{n-2} > t) = P(L_1 > t) = \left(\frac{a-t}{a}\right)^{n-1}$$

gaps G_1 through + including G_{n-2} are equidistributed, even though not independent

- However - gap G_{n-1} has a different distribution

G_{n-1} is the gap that was closest to the camera and was subsequently snipped.

race 

uniform process 

$$G_{n-1} = L_1 + L_n$$

$$\text{event } (G_{n-1} > t) = (L_1 + L_n > t)$$


Recall that gaps in the uniform process are exchangeable

$$= (L_1 + L_2 > t)$$

Taking the probabilities of these events:

$$P(G_{n-1} > t) = P(L_1 + L_2 > t)$$

$$X_{(2)} = L_1 + L_2$$



The second order statistic is precisely the sum $L_1 + L_2$.

$$= P(X_{(2)} > t) \quad \leftarrow \text{we can compute this de novo}$$

$$= 1 - P(X_{(2)} \leq t) \quad \leftarrow \text{since we already have computed } \text{dens}(X_{(2)} = t) \text{ [3/18/98.8-10], we can compute this readily, since:}$$

$$P(X_{(2)} \leq t) = \int_{-\infty}^{\infty} \text{dens}(X_{(2)} = s) ds$$



4/17/98, 10

@ Let's compute $P(X_{(2)} > t)$ de nuovo:

$$\text{event } (X_{(2)} > t) = ((X_{(1)} \leq t) \cap (X_{(2)} > t)) \cup ((X_{(1)} > t) \cap (X_{(2)} > t))$$

union of disjoint events

$$P(X_{(2)} > t) = P((X_{(1)} \leq t) \cap (X_{(2)} > t)) + P((X_{(1)} > t) \cap (X_{(2)} > t))$$



$$= \binom{n-1}{1, n-2} \frac{t}{a} \left(\frac{a-t}{a}\right)^{n-2} + \left(\frac{a-t}{a}\right)^{n-1}$$

$$P(X_{(2)} > t) = (n-1) \frac{t(a-t)^{n-2}}{a^{n-1}} + \left(\frac{a-t}{a}\right)^{n-1}$$

And we have:

$$P(G_{n-1} > t) = (n-1) \frac{t(a-t)^{n-2}}{a^{n-1}} + \left(\frac{a-t}{a}\right)^{n-1}$$

{ Explanation of
apparent contradiction. }

But probability calculated on
basis of uniform process w/
 $n-1$ points.

And we observe that:

$$P(G_1 > t) = \dots = P(G_{n-2} > t) = \left(\frac{a-t}{a}\right)^{n-1} \neq P(G_{n-1} > t) = \left(\frac{a-t}{a}\right)^{n-1} + (n-1) \frac{t(a-t)^{n-2}}{a^{n-1}}$$

uniform process
w/ $n-2$ points

On the surface, this seems kind of strange.

Simply by observing this process at a certain point, I changed the probability distribution of the observed gap. Now the gap nearest to the camera is always bigger:


$$P(G_{n-1} > t) > (P(G_1 > t) = \dots = P(G_{n-2} > t))$$

This is an apparent paradox.

How did the runners (i.e., gap G_{n-1}) know that the camera was there?

Well, they don't. There is no paradox.

Snipping the circle at the point in the gap G_{n-1} that the camera falls on is the same as dropping an extra point in the circle that breaks gap G_{n-1} .



There is an extra
point on this interval.

Example 4 - "Inspector's Paradox"

An inspector is sent to find out something about the process of milk being sold in a certain store.
Milk turns out to be sour every once in a while.

Any particular bottle of milk has probability p of being sour.
The inspector wants to know, when she arrives, what the expected amount of time between the last sour bottle of milk and the next sour bottle of milk is. We are assuming there is at least one sour bottle of milk.

Let ~~integer~~ random variable $T = I + J$

$$I = n - W_{N_n} \leftarrow N_n = \text{number of sour milk bottles in first } n \text{ bottles}$$


$\underbrace{\hspace{10em}}$
 waiting time for N_n^{th} sour bottle of milk.

At the $W_{N_n}^{\text{th}}$ time step, there was a bottle of sour milk

I = "time" (measured in bottles) since the last bottle of sour milk arrived. Includes both point W_{N_n} and point n .

$$J = W_{N_n+1} - n$$

J = "time" from n^{th} bottle until the next bottle of sour milk arrives. Does not include point n , but includes point W_{N_n+1} .



So the integer random variable $T = I + J$ is just a gap in the Bernoulli Process.

We want to figure out the distribution of the gaps in this process, given that the inspector has arrived.

The inspector takes data from various stores, noting the length of the gap from the point when she arrives for inspection.


The inspector claims that $P(T=k) \neq P(W_1=k)$

Implying that, somehow, gap T is special.

even that the 1st waiting time for a sour bottle of milk.


First, let's determine the distributions of I and J .

$$P(I=i) = p q^i$$




$$P(W_1 = i+1) = q^i p$$

$W_1 = 1^{\text{st}}$ waiting time for a bottle of sour milk




$$P(I=i) = P(W_1 = i+1) = p q^i$$

$$P(J=j) = q^{j-1} p$$



$$P(W_1 = j) = q^{j-1} p$$




$$P(J=j) = P(W_1 = j) = q^{j-1} p$$

So we have:

$$P(T=k) = P(I+J=k)$$

Event $(I+J=k)$



$$\text{Event } (I+J=k) = (W_2 = k+1)$$

{Event $(I+J=k)$ equivalent to
Waiting time for the 2^{nd}
sour bottle of milk at $k+1$.}

Hence:

$$P(I+J=k) = P(W_2 = k+1) \neq P(W_1 = k)$$

4/17/98.13

The apparent paradox is that someone comes and observes the gap T and, because of this, changes the probability distribution of $P(T=k)$ so that it is no longer:

$$P(T=k) \neq P(W_1=k)$$

The explanation of the apparent contradiction is to note that the manner in which we defined random variable $T = I + J$ requires the occurrence of 2 sour bottles of milk, one at each end of the gap T .

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



{ Thermodynamics is a sore thumb of mathematics.
It is a major failure of Mathematicians. }

The Poisson Process (beginning)

So far, we've studied basically 2 stochastic processes - the Bernoulli Process and the Uniform Process.

Both of these processes correspond to very intuitive experiments:

Bernoulli Process - corresponds to tossing the same coin repeatedly

Uniform Process - corresponds to picking n points on an interval

Now we have to bite the bullet and face up to the fact that there are, in nature, stochastic processes that we didn't know about.

And we have to adjust our intuition accordingly.

Not everything is intuitive at first, otherwise you wouldn't have anything to learn.

These two new processes we will cover from today through the end of the course. They are:

Poisson Process

Theory of the Normal Distribution

} They are the central core of both probability and statistics.

You could say that everything we've done so far was preparatory to studying these stochastic processes.

We begin w/ the Poisson Process.

The Poisson process exists in nature. It is extremely common. But it is something you have not heard about, probably.

It's good to develop it by developing an analogy with the Bernoulli process. So we have a table of analogies - what happens w/ Bernoulli what happens w/ Poisson

Philosophically, Poisson is continuous Bernoulli.

So let's go back to Bernoulli and do Bernoulli from a different point of view.

Back to Bernoulli

You remember that when we did Bernoulli, we introduced the waiting time W_1 [2/20/98, 10-13].

W_1 = waiting time for first heads.

As it's easy enough to show:

$$P(W_1 > n) = q^n$$

Because the event $(W_1 > n)$ means that the first n tosses are tails.

$$(W_1 > n) = \Omega - (W_1 = 1) - (W_1 = 2) - \dots - (W_1 = n)$$

$$P(W_1 > n) = P[\Omega - (W_1 = 1) - (W_1 = 2) - \dots - (W_1 = n)]$$

all disjoint events, and since probabilities subtract:

$$= P(\Omega) - P(W_1 = 1) - P(W_1 = 2) - \dots - P(W_1 = n)$$

$$= 1 - p - qp - \dots - q^{n-1}p$$

$$= 1 - p(1 + q + q^2 + \dots + q^{n-1})$$

$$\frac{1 - q^n}{1 - q}$$

$$= 1 - \frac{p(1 - q^n)}{1 - q} \leftarrow 1 - q = p$$

$$P(W_1 > n) = q^n \quad \checkmark$$

Lets compute this probability now:

$$P(W_1 > n+k | W_1 > n) = ?$$

We could call this the fools paradise probability.

Everybody who plays roulette believes,

if you waited n times for red ($W_1 > n$)

{event that first n spins are
not red}

$P(W_1 > n+k | W_1 > n)$ becomes more probable.

whats the probability that we wait
 k more times, given that you've
already waited n times.

By definition of conditional probability:

4/22/98.3

$$P(W_1 > n+k | W_1 > n) = \frac{P((W_1 > n+k) \cap (W_1 > n))}{P(W_1 > n)}$$

The event $(W_1 > n+k)$ contains $(W_1 > n)$
 $(W_1 > n+k) \supseteq (W_1 > n)$

So the intersection is:

$$(W_1 > n+k) \cap (W_1 > n) = (W_1 > n+k)$$

$$= \frac{P(W_1 > n+k)}{P(W_1 > n)}$$

$$= \frac{q^{n+k}}{q^n}$$

$$P(W_1 > n+k | W_1 > n) = q^k$$

$\uparrow$
 $P(W_1 > k)$

And we have that:

$$P(W_1 > n+k | W_1 > n) = P(W_1 > k)$$

In other words, the fact that you already waited n times
does not change the probability that you have to wait
 k more times.

Facts are facts.

- We could have developed the entire theory of the Bernoulli Process using the random variable W_1 .
Because we can define all other events pertaining to the Bernoulli Process in terms of the random variable W_1 .

For example, sums:

$$W_2 = W_1 + W_1$$

{waiting time
for second}

so forth and so on.

It's an interesting exercise.



That's what we are forced to do when we do the continuous analogue, for the Poisson process.

Buoyed by this example, we ask the following question:

Is there a continuous waiting time that has the same property:

$$P(W_1 > n+k | W_1 > n) = P(W_1 > k)$$

Let's give this property a name.
It's called memoryless property.

W_1 , the waiting time for the first head in the Bernoulli Process is memoryless.

A waiting time for something to happen is memoryless when knowledge that you have already waited so many units of time doesn't make the event more probable.

Now, we have the language to express that.
And we can ask the following question.

Is there a continuous analogue of this waiting time, where n is replaced by any positive number t , representing continuous time?

Is there a continuous random variable $T \geq 0$ that represents a memoryless waiting time? times are non-negative

The remarkable answer is Yes.
There is one, and only one.

↑ {The interesting situation is when there
is only one. We have that situation here.}

This is very remarkable that there is only one.
So, we will show there is only one.

Let's write precisely what we mean:

$$P(T > t+s | T > t) = P(T > s)$$

This is the memoryless property we want.

On the basis of this property alone, in a couple of lines, we derive the random variable.

By the definition of conditional probability:

$$P(T > t+s | T > t) = \frac{P((T > t+s) \cap (T > t))}{P(T > t)}$$

The event $(T > t+s) \supseteq (T > t)$.

If $(T > t+s)$ happens, then a posteriori $(T > t)$ happens.

$$(T > t+s) \cap (T > t) = (T > t+s)$$

$$P(T > t+s | T > t) = \frac{P(T > t+s)}{P(T > t)}$$

And since the memoryless property we want is (bottom of previous page):

$$P(T > t+s | T > t) = P(T > s)$$

Combining, we want:

$$\frac{P(T > t+s)}{P(T > t)} = P(T > s)$$

(*)

$$P(T > t+s) = P(T > t) P(T > s)$$

• Set $P(T > t) = f(t)$

Then equation (*) becomes:

$$f(t+s) = f(t)f(s)$$

We saw this two lectures ago [4/15/98.14-16].

It's Cauchy's functional equation, which we elaborately solved.

Assuming continuity alone, we proved that [4/15/98.15]:

$$f(t) = c^t \leftarrow \text{which can be written as: } e^{-\alpha t}$$

$$\text{Let } c = e^{-\alpha} \Rightarrow c^t = (e^{-\alpha})^t = e^{-\alpha t}$$

$$P(T > t) = f(t) = e^{-\alpha t}$$

this has to be a probability,

From Cauchy's functional equation, we have:

$$P(T > t) = e^{-\alpha t}$$

• Therefore, the density:

$$\text{dens}(T=t) = \frac{d}{dt} P(T \leq t)$$

$$= \frac{d}{dt} [1 - P(T > t)]$$

$$= \frac{d}{dt} (1 - e^{-\alpha t})$$

$$\text{dens}(T=t) = \alpha e^{-\alpha t}$$

• Now we can find the value of α .
We know that densities are required to integrate to 1.

$$1 = \int_{-\infty}^{\infty} \text{dens}(T=t) dt = \int_0^{\infty} \text{dens}(T=t) dt$$

$$= \int_0^{\infty} \alpha e^{-\alpha t} dt$$

$e^{-\alpha t} > 0$, which implies that
 α is positive.
 $\alpha > 0$

$$1 = \int_0^{\infty} \alpha e^{-\alpha t} dt$$

$$= \alpha \int_0^{\infty} e^{-\alpha t} dt$$

$$\text{Let } u = -\alpha t \\ du = -\alpha dt$$

$$= \cancel{\alpha} \int_{t=0}^{\infty} e^u \left(\frac{-1}{\cancel{\alpha}} \right) du = -e^{-\alpha t} \Big|_{t=0}^{\infty} = (-0) - (-1) = 1$$

{ Vacuously true,
Regardless of
 α . }

So: $\int_0^{\infty} \alpha e^{-\alpha t} dt = 1$, for any $\alpha > 0$.

Therefore, for every α positive, there is a random variable T , whose density is:

$$\text{dens}(T=t) = \alpha e^{-\alpha t}$$

and which represents a memoryless waiting time.

{ And we have seen that this is the only one.
From the definition, we've derived the only possible solution. }

Let's compute the expectation:

$$E(T) = \int_{-\infty}^{\infty} t \text{dens}(T=t) dt$$

$$= \int_0^{\infty} t \alpha e^{-\alpha t} dt$$

$$= \alpha \int_0^{\infty} t e^{-\alpha t} dt$$

$$\int \underbrace{t}_u \underbrace{e^{-\alpha t} dt}_{dv}$$

$$\boxed{u=t}$$

$$\boxed{du=dt}$$

$$\boxed{dv = e^{-\alpha t} dt}$$

$$v = \int e^{-\alpha t} dt$$

$$\text{let } w = -\alpha t$$

$$dw = -\alpha dt$$

$$v = \int e^w \left(\frac{-1}{\alpha} \right) dw$$

$$= \frac{-1}{\alpha} e^w$$

$$\boxed{v = \frac{-1}{\alpha} e^{-\alpha t}}$$

$$\int \underbrace{t}_u \underbrace{e^{-\alpha t} dt}_{dv} = \underbrace{t}_{u} \underbrace{\left(\frac{-1}{\alpha} \right) e^{-\alpha t}}_v - \int \underbrace{\left(\frac{-1}{\alpha} \right) e^{-\alpha t}}_v \underbrace{dt}_{du}$$

$$= -\frac{t}{\alpha} e^{-\alpha t} + \frac{1}{\alpha} \int e^{-\alpha t} dt$$

we've already
computed this

$$= -\frac{t}{\alpha} e^{-\alpha t} + \frac{1}{\alpha} \left(-\frac{1}{\alpha}\right) e^{-\alpha t}$$

$$\int t e^{-\alpha t} = \frac{1}{\alpha} \left(-t - \frac{1}{\alpha}\right) e^{-\alpha t}$$

Which gives:

$$= \left[\frac{1}{\alpha} \left(-t - \frac{1}{\alpha}\right) e^{-\alpha t} \right]_{t=0}^{\infty}$$

$$= 0 + \left(0 + \frac{1}{\alpha}\right)$$

$$E(T) = \frac{1}{\alpha}$$

A memoryless waiting time is completely determined by its expectation. For every expectation, there is a memoryless waiting time, which has density:

$$\text{dens}(T=t) = \alpha e^{-\alpha t}$$

and cumulative probability distribution:

$$P(T \leq t) = 1 - e^{-\alpha t}$$

as the one, and only one, density and cumulative probability distribution.

This distribution is called the exponential distribution.

- Now you ought to ask a question.
What question must you ask?

Q: What are the analogues of the second, and so forth, waiting times?

That's a good question.

It's not the one I had in mind, but it's good.

I'll answer your question, but let me ask the question you might have asked instead.

The question you must ask is this:

Q: Where on earth is the sample space?

Random variables exist on a sample space.

We worked out this w/o a sample space!

Where is it? Where are the sample points?

Where are the events?

What's the probability?

That's the question.

See, all this makes sense, but it assumes some sample space has been taken.

Making the sample space explicit is the theory of the Poisson Process.

There are many natural phenomena that obey the exponential distribution.
For example, you are at a telephone booth and someone is using the phone.
The waiting time for that person to finish is an exponential distribution.
The fact that you have already waited so many minutes does not make
more probable that the person will finish.

Although, the expected waiting time is consumed some by that person.

The confusion in people's minds, both in the case of winning at roulette and in this continuous case of the telephone booth, is the confusion between:

the expected waiting time - and - the conditional probability of
waiting more than you have
so far,

$$E(T)$$

$$P(T > t+s | T > t)$$

You have to make a careful distinction between these two.
That's what you learn probability for.
And you make money.

So any situation where something can terminate with a random time

You are waiting for something to happen.

You know, on the average, but you don't know exactly when it happens.

All the data you have is the average and it fits the exponential distribution.
Marvellous, and it works perfectly.

But the question is:

Where is the sample space?

So now we construct a sample space on which the exponential distribution exists.

For example, if we just had the waiting time for Bernoulli $P(W_1 > n) = q^n$,
then we would have known, sooner or later, that the sample space
consists of sequences of 0 and 1, and that W_1 is the waiting time for
the first 1.

$$\omega = 00101\dots$$

sample point is sequence of 0's and 1's

4/22/98.10


We reason backwards from the waiting time W_1 , and derive the full sample space of sequences of 0's and 1's, from just analysis of this waiting time.

Now we are in the situation where we are forced to do this backwards reasoning from random variable to derivation of sample space in spaces. Because we don't have anything - it's continuous.

Let me give you first the physicist's definition of the Poisson Process.

Physicist's Definition of the Poisson Process

You have the interval $[0, \infty)$



And, at any instant, there is an infinitesimal little person, who is tossing an infinitesimal coin.


The coin is infinitesimally expected to be 1, but mostly it is 0.

So mostly you get 0's, and every once in a while, you got a hit (1). The idea is that the hits are rare events.

Definition - a rare subset ω , of hits

$\omega \subseteq [0, \infty)$ is a subset of $[0, \infty)$, with the property that for every finite interval $[a, b]$, the set $\omega \cap [a, b]$ is finite.

Examples




ω is a rare subset

if you take any finite interval $[a, b]$, then only a finite number of hits are in it. Every finite interval has at most finitely many hits.

For a rare subset, you can not have any condensation (i.e., an infinite number of hits) anywhere. In any finite interval, there can not be an infinite number of hits.

Example:




$\omega \cap [a, b]$ is infinite

ω is not a rare subset



Sample space $\Omega = \underline{\text{all}}$ ^{set of} rare subsets of $[0, \infty)$

- A rare subset is analogous to the 1's in the Bernoulli Process. The elements not in the rare subset are the 0's. The idea is that we have phenomena that occur every once in a while. It occurs, then you wait, it occurs, you wait, it occurs, ... That's the idea and this is carried by the word rare.



{ The sample space Ω_{poisson} = set of all rare subsets of $[0, \infty)$.
Compare to sample space $\Omega_{\text{bernoulli}}$ = all infinite sequences of 0's and 1's. }

When you consider the rare subset in analogy to the Bernoulli Process, it's a sequence of 0's and 1's, but mostly 0's. Intuitively, You can't say that, of course.

- You should view a rare subset as a sequence of hits.
Or a time when an event happened, i.e., when the phone booth becomes available, or a signal comes in from outer space.
- So this is the analogue of all sequences of 0's and 1's in Bernoulli.
Now we have to define events and their probabilities.

The event was $(W_i = n) \leftarrow \{\text{you had the } i^{\text{th}} \text{ head at the } n^{\text{th}} \text{ toss.}\}$

You can't do that now, because here things are continuous.
So we have to do some sort of continuous analogue of that:

Poisson Event:

$$\left[\begin{array}{c} (a, b] \\ n \end{array} \right] = \left\{ \omega \in \Omega : |\omega \cap (a, b]| = n \right\}$$

This looks like a rather complicated definition.
Let's try to re-express it in probabilistic language.

Given $\omega =$ a rare set

Then you can always write ω as a sequence of hits (blips):

$$\omega = (\omega_1, \omega_2, \omega_3, \dots)$$

where $\omega_1 < \omega_2 < \omega_3 < \dots$

$$\omega_n \rightarrow \infty$$

{Where as in the Bernoulli Process, we had sequences of 0's and 1's, in the Poisson Process, here we just note the places where the 1's happen.}

$\left[\begin{smallmatrix} (a, b] \\ n \end{smallmatrix} \right]$ is the event that we get n blips between time interval a and time interval b .

Probabilistically, that's the event that we get n blips in the interval $(a, b]$.

These are not all of the events, of course. We start w/ these elementary events, as we've done w/ other processes, and obtain all possible events:

Any subset of Ω obtained from the elementary events $\left[\begin{smallmatrix} (a, b] \\ n \end{smallmatrix} \right]$ by taking unions, intersections, complements, and unions of countably disjoint events shall be called events.

{In practice, if you want to show that there are subsets of Ω that are not events, you have to use the axiom of choice.
But, for practical purposes, any subset of Ω is an event.}

Any property of rare sets (or blips) that you want to talk about is computed from these basic properties.

An event is the rare set where we get n blips in the interval $(a, b]$.

Now, I'll dish you out the probability.

Then, we'll spend an hour justifying the probability until you realize that this is what it must be.

I'll just tell you what it is. It shall be set, because I say so, to:

$$(*) \text{ Set } P_{\alpha} \left[\begin{smallmatrix} (a, b] \\ n \end{smallmatrix} \right] = \frac{(b-a)^n \alpha^n}{n!} e^{-(b-a)\alpha}$$

(probability depends on α)

probability that you get n blips in interval $(a, b]$

Now we have the following theorem:

Theorem - This definition makes sense

What does this theorem mean?

What does it mean to make sense?

Two things:

1) This definition is not contradictory

You can not manipulate it and get two different probabilities

2) This definition extends to all events

By the rules of boolean algebra, you can extend it to all events.

It is non-trivial to prove all the details.
But let's just check the main point.

Main Point:

The main point is this:

$$\left[\begin{array}{c} [a,b] \cup [b,c] \\ n \end{array} \right] = \left[\begin{array}{c} [a,c] \\ n \end{array} \right] \quad \left\{ \begin{array}{l} \text{These events are equal, as the intervals} \\ [a,b] \text{ and } [b,c] \text{ don't have any overlap} \\ \text{and } [a,b] \cup [b,c] \text{ has no gaps.} \end{array} \right\}$$

This, in turn, is equal to:

$$\left[\begin{array}{c} [a,c] \\ n \end{array} \right] = \bigcup_{i=0}^n \left(\left[\begin{array}{c} [a,b] \\ i \end{array} \right] \cap \left[\begin{array}{c} [b,c] \\ n-i \end{array} \right] \right)$$

Let's argue this probabilistically.

The event that you get n blips in the interval $(a,c]$ means that you get some blips in the interval $(a,b]$ and the rest of the blips in the interval $(b,c]$.

That's exactly what this says.

This is the set theoretic analogue of the binomial theorem.
It's really trivial.

This is a disjoint union.

Therefore, if our probability equation (*) is to be consistent, then we better have:

$$P_{\alpha}(\text{LHS}) = P_{\alpha}(\text{RHS})$$

Otherwise, we are in deep trouble.

• Proof $P_\alpha(\text{LHS}) = P_\alpha(\text{RHS})$

4/22/98.14

$$P_\alpha(\text{LHS}) = P_\alpha\left(\left[\begin{smallmatrix} [a, c] \\ n \end{smallmatrix}\right]\right) = \frac{(c-a)^n \alpha^n}{n!} e^{-(c-a)\alpha}$$

$$P_\alpha(\text{RHS}) = P_\alpha\left(\underbrace{\bigcup_{i=0}^n \left(\left[\begin{smallmatrix} [a, b] \\ i \end{smallmatrix}\right] \cap \left[\begin{smallmatrix} [b, c] \\ n-i \end{smallmatrix}\right]\right)}_{\text{disjoint union}}\right) = \sum_{i=0}^n P_\alpha\left(\left[\begin{smallmatrix} [a, b] \\ i \end{smallmatrix}\right] \cap \left[\begin{smallmatrix} [b, c] \\ n-i \end{smallmatrix}\right]\right)$$

There is no reason why these should be equal.

Because I've left out an axiom, for the purposes of drama.

The only way these can be equal is if we have the following axiom, in addition to (*).

Axiom

If $(a, b]$ and $(c, d]$ are disjoint intervals and k and n any non-negative integers, the events:

$$\left[\begin{smallmatrix} [a, b] \\ n \end{smallmatrix}\right] \quad \text{and} \quad \left[\begin{smallmatrix} [c, d] \\ k \end{smallmatrix}\right]$$

are independent.

In other words, the number of blips occurring in a given interval does not influence the number of blips occurring in a disjoint interval.

The blips are random.

• With this assumption of the axiom of independence, equation (**) simplifies to:

$$\left\{ \begin{array}{l} \text{events } A+B \text{ independent} \\ \text{has:} \\ P(A \cap B) = P(A)P(B) \end{array} \right\} \Rightarrow = \sum_{i=0}^n P_\alpha\left(\left[\begin{smallmatrix} [a, b] \\ i \end{smallmatrix}\right]\right) P_\alpha\left(\left[\begin{smallmatrix} [b, c] \\ n-i \end{smallmatrix}\right]\right)$$

Using the definition we have for P_α , i.e., equation (*):

$$= \sum_{i=0}^n \frac{(b-a)^i \alpha^i}{i!} e^{-(b-a)\alpha} \frac{(c-b)^{n-i} \alpha^{n-i}}{(n-i)!} e^{-(c-b)\alpha}$$

$\swarrow \quad \nwarrow$
 $e^{-(c-a)\alpha}$

$$= \sum_{i=0}^n \alpha^n \frac{1}{i!(n-i)!} (b-a)^i (c-b)^{n-i} e^{-(c-a)\alpha}$$

Multiply by $1 = \frac{n!}{n!}$





$$= \frac{\alpha^n}{n!} e^{-(c-a)\alpha} \sum_{i=0}^n \frac{n!}{i!(n-i)!} \binom{n}{i} (b-a)^i (c-b)^{n-i}$$

What's this?

The binomial theorem!

These are the binomial coefficients.

$$((b-a) + (c-b))^n = \sum_{i=0}^n \binom{n}{i} (b-a)^i (c-b)^{n-i}$$

$$(c-a)^n = \sum_{i=0}^n \binom{n}{i} (b-a)^i (c-b)^{n-i}$$

This is just:
 $(c-a)^n$

$$P_\alpha \left(\bigcup_{i=0}^n \left(\left[\frac{a}{i} \right] \cap \left[\frac{b}{n-i} \right] \right) \right) = \frac{(c-a)^n \alpha^n}{n!} e^{-(c-a)\alpha} \quad \checkmark$$

So, we have $P_\alpha(\text{LHS}) = P_\alpha(\text{RHS})$, as desired.
It checks.

So now you are beginning to believe there is something to this, after all.
So the Poisson Process is consistent.

- Let's now justify our derivation of the exponential distribution $P(T > t) = e^{-\alpha t}$.
The easiest thing we can do is this.
Let's compute the waiting time for the first blip in the Poisson Process.
You have random blips occurring. What's the waiting time for the first blip?
Give it a name.

T = waiting time for the first blip

$$P(T > t) = P(\text{no blips in interval } (0, t])$$

$$= P_\alpha([0, t])$$

The definition of P_α , i.e., equation (*), gives:

$$= \frac{(t-0)^0 \alpha^0}{0!} e^{-(t-0)\alpha}$$

$$P(T > t) = e^{-\alpha t} \quad \left. \vphantom{P(T > t)} \right\} \text{Precisely the exponential distribution, which we also derived using Cauchy's functional equation [4/22/98.4-6].}$$

4/22/98, 16

In conclusion, we have set up a sample space where the continuous waiting time makes sense.

And now, we're going to do all sorts of experiments on this sample space to convince ourselves that it is the only one possible.

Fact: The Poisson Process is more common than the Bernoulli Process in nature.

The Poisson Process is the process when you spread little seeds around a big field. Or you bomb someplace by dropping 18.313 students by parachute. Each 18.313 student you drop is a blip. And you are dropping over a big country. All you have is the average number of 18.313 students dropped per unit area.

That's the idea.

The idea is that all the data you have is:
number of blips per unit length

Everything else is completely random.

Under these circumstances,
there is only one way of
modeling this.

That's the Poisson Process.
It's a deep message.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



Seven Properties of The Poisson Process


The Poisson Process is very frequent in nature.
As a matter of fact, it occurs more frequently than the Bernoulli Process, or the Uniform Process.
But, it's something that we haven't seen.
So we have to get used to it.
We need to focus our intuition.

What's the idea?

The idea, roughly, speaking as a physicist is:

Instead of a finite interval, as we had in the Uniform Process,
we have a half infinite interval.

↑ {a half interval, because we want to see the interpretation in terms of waiting times, as one of several interpretations.



Then, we imagine flying over this interval, at a steady rate, w/ an indefinite supply of blips, and dropping blips at random.

The only decision is that the average number of blips per unit area is given by α .

α = average number of blips per unit area
↑ intensity

That's it.

This is the way you interpret picking infinitely many points.
Because if you pick infinitely many points, you can't divide by the random interval enclosing them.

And what turns out to be the case is that this is described by an exponential distribution.


We did the brutal, dry, mathematical definitions last time.
Now, we want to understand it from a different point of view.

Let's repeat our definition:

Our definition was based on the sample space consisting of rare subsets.

$$\Omega = \{ \omega = \text{rare subsets, where } \omega = (\omega_1, \omega_2, \dots) \}$$

↑ sample points



Rare means you can not have any accumulation points.

This is best stated by saying - Look at any finite interval on $[0, \infty)$ and there must be only a finite number of blips.

4/24/98.2

what we are studying is the appearance of these blips, randomly.
That's the idea.

- The Poisson Process is one of 4 fundamental stochastic processes that make up all of probability theory.

The 4 fundamental stochastic processes

1) Bernoulli Process

Tossing a coin.

2) Uniform Process

Picking points randomly from a finite interval.

3) Poisson Process

4) Process pertaining to the Normal Distribution

That's it.

From a conservative viewpoint, you can define probability theory as the study of these 4 stochastic processes and their degeneration.

Event - $\left[\begin{matrix} (a,b] \\ n \end{matrix} \right] =$ you have n blips in the interval $(a,b]$

We view this probabilistically as blips appear at random (someone drops them from an airplane).

And the event is that we look at the area in the interval $(a,b]$ and we see n blips (random blips).

You have to get used to the fact that what's random is:

a) it is random when the blips appear

b) there are infinitely many blips

It takes a little getting used to.

From these events, by the operations of boolean algebra, together with infinite unions intersections of sequences, we generate all possible events.

Main Result

We assign a probability to this basic event.

The main result of the theory of the Poisson Process states that this probability uniquely determines the probability of such events.


The probability we assign to such events is:

$$P_{\alpha} \left(\left[\begin{matrix} (a,b] \\ n \end{matrix} \right] \right) = \frac{\alpha^n (b-a)^n}{n!} e^{-(b-a)\alpha}$$

$\alpha =$ given intensity

Furthermore, the number of blips appearing in disjoint intervals is independent.

If you have two events:



This is a way of exactly rendering the fact that the number of blips that appear in a given interval has no influence, whatsoever, on the number of blips that appear in a disjoint interval.

It is amazing that this is the only consistent way of defining a probability. And that is what we have to convince ourselves of now.

- Let's introduce something analogous to what we have done for the Bernoulli and the Uniform Process. We introduce the random function $N_\alpha(t)$.

- 1) Bernoulli Process : random variable $S_n =$ number of heads in the first n tosses ^(1/2)
- 2) Uniform (Dirichlet) Process ^[3/18/98.6] : random function $U_\alpha(t) =$ number of points chosen in the interval $[0, t]$
- 3) Poisson Process : random function $N_\alpha(t) =$ number of blips in the interval $[0, t]$
depends on the intensity α of the Poisson Process.

NB: α is the only data on which the Poisson Process depends

The entire theory of the Poisson Process can be reformulated in terms of the random function $N_\alpha(t)$.

What's a random function?

A random function is a function of t , which has the property that for every t , $N_\alpha(t)$ is a random variable ^[3/18/98.6].

So, it's a random variable, depending on the parameter t .

- What is the probability distribution of this random function?

$$P(N_\alpha(t) = n) = ?$$

This is easily derived, once we realize that:

$$\text{event } (N_\alpha(t) = n) = \left[\begin{matrix} [0, t) \\ n \end{matrix} \right]$$

So, we really have 2 notations for the same event.

We could have used the notation $(N_\alpha(t) = n)$ from the start.

Therefore, we have, trivially:

$$P(N_\alpha(t) = n) = \frac{\alpha^n t^n}{n!} e^{-\alpha t}$$

The Poisson Distribution w/ intensity α

- Furthermore, the fact, that the number of blips in disjoint intervals are independent, can be expressed in terms of the random function, as follows:

if $[a, b)$ and $[c, d)$ are disjoint intervals, then the random variables

$$N_\alpha(b) - N_\alpha(a) \quad \text{and} \quad N_\alpha(d) - N_\alpha(c)$$

are independent.

This is a repetition of what we worked in to the definition of the Poisson Process.
Why?

Because $(N_\alpha(b) - N_\alpha(a) = m)$ the number of blips in the interval $[a, b)$

Similarly $(N_\alpha(d) - N_\alpha(c) = n)$ the number of blips in the interval $[c, d)$

$$= \left[\begin{matrix} [a, b) \\ m \end{matrix} \right] \quad \text{and} \quad \left[\begin{matrix} [c, d) \\ n \end{matrix} \right]$$

since $[a, b)$ and $[c, d)$ are disjoint, the originating random variables are independent.

Again, we could have defined the Poisson Process using random functions from the start.

- Observe, already, one decisive advantage of the Poisson Process over the Uniform Process.

The random function $U_\alpha(t)$ of the Uniform Process is not independent. If you take the number of points of 2 disjoint intervals of the Uniform Process, is not independent, because for have a fixed number of n total points.

But for the Poisson Process, the number of blips in disjoint intervals is independent.

This is very good, for all sorts of purposes.

So, you could say that the Poisson Process achieves something that we wanted for the Uniform Process, but couldn't get on a finite interval.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



- Let's compute the expectations:

$$E(N_\alpha(t)) = \sum_{n=0}^{\infty} n \underbrace{P(N_\alpha(t)=n)}_{\text{by definition:}}$$

$$= \sum_{n=0}^{\infty} n \frac{\alpha^n t^n}{n!} e^{-\alpha t}$$

[this probability given by the formula we just gave for the Poisson Distribution w/ intensity α .]

$$= \sum_{n=1}^{\infty} \frac{\alpha t \alpha^{n-1} t^{n-1}}{(n-1)!} e^{-\alpha t} + 0 \quad \leftarrow n=0$$

Set $j=n-1$ to obtain:

$$= \alpha t e^{-\alpha t} \sum_{j=0}^{\infty} \frac{\alpha^j t^j}{j!}$$

The Taylor Series expansion of $e^{\alpha t}$ is:

$$e^{\alpha t} = \sum_{j=0}^{\infty} \frac{f^{(j)}(0)}{j!} t^j$$

$$= \sum_{j=0}^{\infty} \frac{\alpha^j e^{\alpha \cdot 0}}{j!} t^j$$

$$= \sum_{j=0}^{\infty} \frac{\alpha^j t^j}{j!}$$

$$= \alpha t e^{-\alpha t} e^{\alpha t}$$

$$E(N_\alpha(t)) = \alpha t$$

A major result

To repeat, the expectation of the random variable $N_\alpha(t)$, for t fixed, equals αt .

$$E(N_\alpha(t)) = \alpha t$$

We can repeat this, as follows:

$$\frac{1}{t} E(N_\alpha(t)) = \alpha$$

This is exactly what I've said already 5 times.

The average number of blips per unit length is α .


$$\frac{E(N_\alpha(t))}{t}$$

This tells it to you loud and clear.

We now discuss the 7 fundamental properties of the Poisson Process.

Property (1) - Law of Rare Events

I take a half infinite time interval.



At every interval $\frac{1}{j}$ of time, I flip a coin.

I'm doing exactly a Bernoulli Process, except I'm flipping exactly at the interval of time.

Flip at times $\frac{1}{j}, \frac{2}{j}, \dots, \frac{n}{j}, \dots$ a coin which is

biased w/ bias $p_j = \frac{\alpha}{j}$.

What is the probability of getting k heads at time $\frac{n+\epsilon}{j}$?

$$\left\{ \begin{array}{l} \text{immediately after } \frac{n}{j} \text{ and before } \frac{n+1}{j} \\ \left. \begin{array}{l} \epsilon \rightarrow 0 \\ \epsilon > 0 \end{array} \right\} \frac{n}{j} < \frac{n+\epsilon}{j} < \frac{n+1}{j} \end{array} \right\}$$

This probability is just the binomial distribution - one of the first things you learned in this class.

4/24/98.7

You have n tosses and you want k ones.
The bias is $\left(\frac{\alpha}{j}\right)$, so you have:

$$P(k \text{ 1's at time } \frac{n+E}{j}) = \binom{n}{k} \left(\frac{\alpha}{j}\right)^k \left(1 - \frac{\alpha}{j}\right)^{n-k}$$

Now, we let $j \rightarrow \infty$.

Which means, we toss the coin faster and faster.

In the limit, you toss the coin continuously.

$$\lim_{j \rightarrow \infty} \frac{1}{j}, \frac{2}{j}, \dots, \frac{n}{j}, \dots = \text{continuous}$$

$$\text{Set } t = \frac{n}{j} \Rightarrow j = \frac{n}{t}$$

Let $n \rightarrow \infty$,

t remain fixed.

$$\lim_{n \rightarrow \infty} \frac{n+E}{j} = \lim_{n \rightarrow \infty} \left(1 + \frac{E}{n}\right)t = t$$

$$P(N_\alpha(t) = k) = \lim_{n \rightarrow \infty} P(k \text{ 1's at time } \frac{n+E}{j}) \quad \left. \begin{array}{l} \text{in the limit,} \\ P(k \text{ 1's at time } t) \end{array} \right\}$$

$$= \lim_{n \rightarrow \infty} \binom{n}{k} \left(\frac{\alpha}{j}\right)^k \left(1 - \frac{\alpha}{j}\right)^{n-k}$$

Substitute $j = \frac{n}{t}$ and rewriting expressions gives:

$$= \lim_{n \rightarrow \infty} \frac{(n)_k}{k!} \frac{\alpha^k}{\left(\frac{n}{t}\right)^k} \left(1 - \frac{\alpha}{\left(\frac{n}{t}\right)}\right)^n \left(1 - \frac{\alpha}{\left(\frac{n}{t}\right)}\right)^{-k}$$

$$= \lim_{n \rightarrow \infty} \frac{(n)_k}{n^k} \cdot \frac{\alpha^k t^k}{k!} \left(1 - \frac{\alpha t}{n}\right)^n \left(1 - \frac{\alpha t}{n}\right)^{-k}$$

(*)

The numerator and denominator each have k factors. In the limit, the ratios each term is 1.

$$\lim_{n \rightarrow \infty} \frac{\overbrace{n(n-1)(n-2) \dots (n-k+1)}^k}{\underbrace{n \cdot n \cdot n \dots n}_k}$$

$$= \lim_{n \rightarrow \infty} \left(\frac{n}{n}\right) \left(\frac{n-1}{n}\right) \left(\frac{n-2}{n}\right) \dots \left(\frac{n-k+1}{n}\right)$$

$$= 1$$

By definition,

$$\lim_{n \rightarrow \infty} \left(1 + \frac{r}{n}\right)^n = e^r$$

$$\text{So, } \lim_{n \rightarrow \infty} \left(1 - \frac{\alpha t}{n}\right)^n = e^{-\alpha t}$$

4/24/98.8

And we have:

$$P(N_{\alpha}(t)=k) = \frac{\alpha^k t^k}{k!} e^{-\alpha t}$$

And we've just shown that the Poisson Process can be viewed as continuous coin tossing.
This limiting process is called the Law of Rare Events.

We'll talk about this more later.

But for the moment, observe that the Law of Rare Events has notable practical applications.

Example - Mother bakes cookies w/ raisins.

She wants 1 raisin per cookie.

She mixes the dough, throws in the raisins at random, and then bakes the cookies.

What's the probability that a cookie has no raisins?

$$\frac{1}{e}$$

Can't you read?

Repeat:

We have a Poisson Process with intensity $\alpha = 1$

There is, on the average, 1 raisin per cookie.

The probability of getting n raisins in a cookie is given by the Poisson Distribution:

$$\left\{ \begin{array}{l} \text{Average number of blips per} \\ \text{unit area} \\ = \frac{1 \text{ raisin}}{\text{cookie}} \end{array} \right\}$$

$$P(N_{\alpha}(1)=n) = \frac{\alpha^n 1^n}{n!} e^{-\alpha(1)}$$

and $\alpha = 1$, so this is:

$$= \frac{1}{n!} \cdot \frac{1}{e}$$

So the probability that a cookie has no raisins is:

$$P(N_{\alpha=1}(1)=0) = \frac{1}{e}$$

Example - You have to set up my lectures in LaTeX, which I would greatly appreciate 😊 for next years students.

And you make 1 mistake per page.

What's the probability that I read your page and find 3 mistakes?

$\alpha=1$, so we again have:

$$P(N_{\alpha=1}(1)=n) = \frac{1}{n!} \cdot \frac{1}{e} \Rightarrow P(N_{\alpha=1}(1)=3) = \frac{1}{3!} \cdot \frac{1}{e}$$



Why?

Because you approximate the Bernoulli Process of making a mistake for each letter by the Law of Rare Events of the Poisson Process.

Therefore, you can consider this as a Poisson Process with intensity 1.
And, therefore, you can tell immediately what the probability is of having 1000 mistakes per page.

↖ w/ tremendous accuracy.

Now you begin to see what it's about.

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



Property (2) — Limit of Uniform Process

Let's get the Poisson Process as a limit of the Uniform Process

In other words, let's see that the Poisson Process is really dropping α blips per unit length.

We take a uniform process w/ random function $U_a(t)$.

We set a so that:

$$\alpha = \frac{n}{a}$$

We let a vary, together w/ the number of points n that you drop.

But the ratio $\frac{n}{a}$ is constant.

We drop n points in the interval $[0, a]$, where $\frac{n}{a} = \alpha$.

$$P(U_a(t)=k) = \binom{n}{k} \left(\frac{t}{a}\right)^k \left(1 - \frac{t}{a}\right)^{n-k}$$

Now, let's work in $a = \frac{n}{\alpha}$:

$$= \frac{\binom{n}{k}}{k!} \frac{t^k}{\left(\frac{n}{\alpha}\right)^k} \left(1 - \frac{t}{\frac{n}{\alpha}}\right)^n \left(1 - \frac{t}{\frac{n}{\alpha}}\right)^{-k}$$

$$= \frac{\binom{n}{k} \xrightarrow{1}}{n^k} \cdot \frac{\alpha^k t^k}{k!} \left(1 - \frac{\alpha t}{n}\right)^n \left(1 - \frac{\alpha t}{n}\right)^{-k}$$

in the limit, as $n \rightarrow \infty$, this is the same as equation (x).

$$\lim_{n \rightarrow \infty} P(U_a(t)=k) = \frac{\alpha^k t^k}{k!} e^{-\alpha t}$$

$$\boxed{\lim_{n \rightarrow \infty} P(U_a(t)=k) = P(N_\alpha(t)=k)}$$

So, we see that the Poisson Process really is dropping α blips per unit length.

Property (3) - Memorylessness

Connection w/ memorylessness. See [4/22/98.5-8].

Recall that the waiting time for the first blip in the Poisson Process is exponentially distributed:

$$P(T > t) = e^{-\alpha t} \Rightarrow P(T \leq t) = 1 - e^{-\alpha t}$$

Remember, if you waited so long for the first blip, that wait doesn't make expected waiting time any shorter (i.e., the expected waiting time does not change).


The expected waiting time $E(T) = \frac{1}{\alpha}$.

In other words, the higher the intensity α , the shorter the expected waiting time. Period.

Recall that the density of the exponential distribution is:

$$\text{dens}(T=t) = \alpha e^{-\alpha t} \quad \text{exponential density}$$

We reason as follows:



Here are my blips.

I wait an exponential waiting time, as they say in the business, for the first blip.

After the first blip has occurred, I wait another exponential waiting time for the second blip. And so on.

That means that the waiting time for the n^{th} blip is the sum of n independent, exponentially distributed, identically distributed random variables. And that should give us back the Poisson Process. Let's see if that checks.

$T_1, T_2, \dots, T_n$ are i.i.d., exponentially distributed random variables. What is the density of their sum?

$$\text{dens}(T_1 + T_2 + \dots + T_n = t) = ?$$

Let's first do the density of 2 random variables.

Then we'll guess at a more general result.

$$\text{dens}(T_1 + T_2 = t) = \text{where } T_1 \text{ and } T_2 \text{ are independent, we've shown that the density of their sums is the convolution of their density functions.}$$

[4/15/98.8]

4/24/98.11

Since T_1 and T_2 are independent and $T_1 \geq 0$ and $T_2 \geq 0$, we have the convolution:

$$\text{dens}(T_1 + T_2 = t) = \int_0^t \text{dens}(T_1 = s) \text{dens}(T_2 = t-s) ds$$

Since $T_1 + T_2$ are identically distributed, they have the same density function. Namely:

$$\text{dens}(T_1 = x) = \text{dens}(T_2 = x) = \alpha e^{-\alpha x}$$

$$= \int_0^t \alpha e^{-\alpha s} \alpha e^{-\alpha(t-s)} ds$$

$$= \int_0^t \alpha^2 e^{-\alpha t} ds$$

$$= \alpha^2 e^{-\alpha t} \int_0^t ds$$

$$\boxed{\text{dens}(T_1 + T_2 = t) = \alpha^2 t e^{-\alpha t}}$$

Proceeding by induction, we find:

$$\text{dens}(T_1 + T_2 + \dots + T_n = t) = \frac{\alpha^n t^{n-1}}{(n-1)!} e^{-\alpha t}$$

Proof by induction

Basis: $n=1$
 $\text{dens}(T_1) = \frac{\alpha^1 t^0}{0!} e^{-\alpha t}$
 $= \alpha e^{-\alpha t} \checkmark$

Inductive Step: assume induction hypothesis holds for n .

$$\text{dens}(\underbrace{T_1 + \dots + T_n}_U + T_{n+1} = t) \stackrel{?}{=} \frac{\alpha^{n+1} t^n}{n!} e^{-\alpha t}$$

$$\text{dens}(U + T_{n+1} = t) = \int_0^t \text{dens}(U=s) \text{dens}(T_{n+1}=t-s) ds \quad \leftarrow \text{convolution of densities since } U, T_{n+1} \text{ independent}$$

$$= \int_0^t \text{dens}(T_1 + \dots + T_n = s) \text{dens}(T_{n+1} = t-s) ds$$

The induction hypothesis gives

$$= \int_0^t \frac{\alpha^n s^{n-1}}{(n-1)!} e^{-\alpha s} \alpha e^{-\alpha(t-s)} ds$$

This is called the Gamma Distribution

↑
 strictly speaking, a misnomer.
 This should be called the
Gamma Density, not
 distribution.

$$\begin{aligned}
&= \int_0^t \frac{\alpha^{n+1} s^{n-1}}{(n-1)!} e^{-\alpha t} ds \\
&= \frac{\alpha^{n+1}}{(n-1)!} e^{-\alpha t} \int_0^t s^{n-1} ds \\
&= \frac{\alpha^{n+1}}{(n-1)!} e^{-\alpha t} \frac{1}{n} t^n \\
&= \frac{\alpha^{n+1} t^n}{n!} e^{-\alpha t} \quad \checkmark
\end{aligned}$$

Which completes the inductive proof of the Gamma Density:

$$\text{dens}(T_1 + \dots + T_n = t) = \frac{\alpha^n t^{n-1}}{(n-1)!} e^{-\alpha t}$$

What is it the density of?

The Gamma Density, also known as the Gamma Distribution, is the waiting time for the n^{th} blip, we guess, in the Poisson Process.
It should be.

In other words, we should have that:

$$\text{event } (N_\alpha(t) = n) = \underbrace{(T_1 + T_2 + \dots + T_n < t) - (T_1 + T_2 + \dots + T_n + T_{n+1} < t)}$$

event that number of blips in the interval $[0, t)$ is n

The event that at least n blips occur in $[0, t)$ minus the event that at least $n+1$ blips occur in the same interval is exactly the same as saying there are exactly n blips in the interval $[0, a)$.

The events are equal.
Now we want to check the probabilities:

$$P(N_\alpha(t) = n) \stackrel{?}{=} P(T_1 + T_2 + \dots + T_n < t) - P(T_1 + T_2 + \dots + T_n + T_{n+1} < t)$$

Using the fact that:

$$P(X \leq x) = \int_{-\infty}^x \text{dens}(X=s) ds$$

and the Gamma Density (Distribution), we get the probabilities of the RHS above.

$$\begin{aligned}
&= \int_0^t \frac{\alpha^n s^{n-1}}{(n-1)!} e^{-\alpha s} ds - \int_0^t \frac{\alpha^{n+1} s^n}{n!} e^{-\alpha s} ds \\
&= \frac{\alpha^n}{(n-1)!} \int_0^t s^{n-1} e^{-\alpha s} ds - \frac{\alpha^{n+1}}{n!} \int_0^t s^n e^{-\alpha s} ds
\end{aligned}$$

4/24/98.13

We evaluate the 2nd integral by parts, in anticipation of performing miracles of cancellation.

$$\int \underbrace{s^n}_u \underbrace{e^{-\alpha s}}_{dv} ds = \underbrace{s^n \left(\frac{-1}{\alpha} e^{-\alpha s} \right)}_u \underbrace{+ \int \left(\frac{-1}{\alpha} e^{-\alpha s} \right) n s^{n-1} ds}_v$$

$$u = s^n$$

$$du = n s^{n-1} ds$$

$$dv = e^{-\alpha s} ds$$

$$v = \int e^{-\alpha s} ds$$

$$= -\frac{1}{\alpha} e^{-\alpha s}$$

$$\int s^n e^{-\alpha s} ds = -\frac{s^n}{\alpha} e^{-\alpha s} + \frac{n}{\alpha} \int s^{n-1} e^{-\alpha s} ds$$

Using this indefinite integral equality above gives:

$$\begin{aligned} &= \frac{\alpha^n}{(n-1)!} \int_0^t s^{n-1} e^{-\alpha s} ds - \frac{\alpha^{n+1}}{n!} \left[-\frac{s^n}{\alpha} e^{-\alpha s} \right]_{s=0}^t - \frac{\alpha^{n+1}}{n!} \left(\frac{n}{\alpha} \int_0^t s^{n-1} e^{-\alpha s} ds \right) \\ &= \frac{\alpha^n}{n!} \left[s^n e^{-\alpha s} \right]_{s=0}^t \end{aligned}$$

$$P(N_\alpha(t) = n) = \frac{\alpha^n t^n}{n!} e^{-\alpha t} \quad \checkmark$$

This is the Poisson Distribution

$$= P(N_\alpha(t) = n)$$

So, it checks. As desired.

This completes the first 3 properties of the Poisson Process.

So far, we've done only theory.

Now we need applications - various properties of the Poisson Process.

Seven Properties of the Poisson Process (cont'd)

We continue our in-depth study of the Poisson Process, which is one of the fundamental processes of probability.


Intuitively, again, what is the Poisson Process?

You go over a continuum and you sprinkle blips.

The sprinkling is completely random, subject only to one condition

↑
the average number of blips
per unit area equals α
(the intensity).

Nothing else is given.



Under these conditions, there is only one possible formula that gives the formula for the probability and that is the formula for the Poisson Process.

That's extraordinary.
Don't you ever forget it.

In symbols:

Our sample space Ω consists of rare events, which are sequences of blips.

$$\Omega = \{ \omega = (\omega_1, \omega_2, \dots) \}$$

↑
blip

And we define a probability on the set of all such rare sequences.

More generally, we define a random function:

$$N_\alpha(t) = \text{number of blips in } [0, t)$$

subject to the following condition:

$$P(\underbrace{N_\alpha(b) - N_\alpha(a) = n}_{\substack{\text{event that number of blips} \\ \text{in interval } [a, b) \text{ is } n}}) = \frac{\alpha^n (b-a)^n}{n!} e^{-\alpha(b-a)}$$

and furthermore

if $[a, b) \cap [c, d) = \emptyset$, then the random variables

$N_\alpha(b) - N_\alpha(a)$ and $N_\alpha(d) - N_\alpha(c)$ are independent.
 number of blips in $[a, b)$ number of blips in $[c, d)$

This summarizes our description of the sample space, which we have given, here, in terms of events. Equivalent to the description of the event.

We're now trying to get a good feeling for this fundamental process.

As this process is not something we're used to, even though it occurs in nature more frequently than the others.

What have we seen so far?

We have seen 3 ways in which we can obtain the Poisson Process as a limiting process of known processes

- (1) We started tossing coins on intervals that became thicker and thicker, with bias $p_j = \frac{\alpha}{j}$ that varied w/ the thickness of the interval. $(\frac{1}{j}, \frac{2}{j}, \dots, \frac{n}{j}, \dots)$

And we saw that the limiting process that came from this coin tossing process was the Poisson Process.

That was the Law of Rare Events.

Limits of Bernoulli coin tossing

So you can view Poisson Process as some sort of infinitesimal coin tossing.

- (2) We take the uniform process on the interval $[0, a)$.




We let a tend to infinity, subject to the condition that the average number of points per unit length remains equal to α .

Under these conditions, we saw that the limiting process is, again, the Poisson Process.

(3) We considered the relationship between the Poisson Process and the exponential distributions.

The exponential distribution, as we have seen, is the only memoryless continuous distribution.

Taking this as our starting point, we view a rare sequence as follows:



We wait a memoryless time of intensity α .

The expectation is $\frac{1}{\alpha}$.

As soon as there is a blip, you start waiting again, with intensity α for the next blip.

And so forth.

In this way, you create a rare sequence.

In other words, we took the sum:

$$T_1 + T_2 + \dots + T_n$$

n independent random variables, which are identically distributed with an exponential distribution.

We envision the Poisson Process by using the set theoretic identity:

$$\begin{aligned} \text{event } (T_1 + T_2 + \dots + T_n < t) - (T_1 + T_2 + \dots + T_n + T_{n+1} < t) \\ = (N_\alpha(t) = n) \end{aligned}$$

Then by convolution, we obtain the density of the sum of n independent and identically distributed random variables as the Gamma Distribution.

↑ more correctly,
Gamma Density.

By fooling around w/ the Gamma Distribution and integration by parts, we saw that the Poisson Process comes out.


- So we have, to date, 3 heavy duty justifications for the Poisson Process. You should be fairly convinced. Now, let's go to number (4).

Property (4) $P(N_\alpha(t) = k \mid N_\alpha(a) = n) = P(U_\alpha(t) = k)$ } uniform process obtained by conditioning Poisson process

Let's compute an interesting probability.
What does this mean?

$$P(N_\alpha(t) = k \mid N_\alpha(a) = n) = ? \quad t < a$$

assuming we have n blips in the interval $[0, a)$, what then is the probability that k of those blips are in the interval $[0, t)$?




By the definition of conditional probability :

$$P(N_\alpha(t) = k \mid N_\alpha(a) = n) = \frac{P((N_\alpha(t) = k) \cap (N_\alpha(a) = n))}{P(N_\alpha(a) = n)}$$

The events in the numerator are not independent

We can do a slight computation that is done very often, which is the following.
We want to rewrite the intersection in terms of independent events.

All you have to do is realize that:



$$(N_\alpha(t) = k) \cap (N_\alpha(a) = n) = (N_\alpha(t) = k) \cap (N_\alpha(a) - N_\alpha(t) = n - k)$$

these are disjoint events

4/29/98.5

$$= \frac{P((N_\alpha(t)=k) \cap (N_\alpha(a) - N_\alpha(t) = n-k))}{P(N_\alpha(a)=n)}$$

And since the numerator is the probability of the intersection of two disjoint events, that probability is the product of the probabilities.

$$= \frac{P(N_\alpha(t)=k) P(N_\alpha(a) - N_\alpha(t) = n-k)}{P(N_\alpha(a)=n)}$$

$$\text{where } P_\alpha\left(\begin{bmatrix} [c,d] \\ m \end{bmatrix}\right) = \frac{\alpha^m (d-c)^m}{m!} e^{-\alpha(d-c)}$$

$$= \frac{\frac{\alpha^k t^k}{k!} e^{-\alpha t} \frac{\alpha^{n-k} (a-t)^{n-k}}{(n-k)!} e^{-\alpha(a-t)}}{\frac{\alpha^n a^n}{n!} e^{-\alpha a}}$$

$$= \frac{n!}{k!(n-k)!} \frac{t^k (a-t)^{n-k}}{a^n}$$

$\binom{n}{k}$

$$= \binom{n}{k} \frac{t^k (a-t)^{n-k}}{a^n}$$

What's this?

This is the Uniform Process

$$\boxed{P(N_\alpha(t)=k | N_\alpha(a)=n) = P(U_a(t)=k)}$$

$U_a(t)$ is the random function of the uniform process

Ha. Ha. Ha.

Conclusion: When you condition the number of blips to be n in the interval $[0, a)$, then, if you look inside the interval $[0, t)$, it's same as if it were a uniform process.
As you would expect.



4/29/98.6

Therefore, I, in a sense, cheated you.

We can derive all the properties of the uniform process cheapo.

We didn't have to derive them the way we did.

We could have derived them from the fact that:

$$P(N_a(t)=k | N_a(a)=n) = P(U_a(t)=k)$$

But it would have been overwhelming to do it this way, from this stuff.


For example, let's work out an ambitious property of the uniform process.

(We'll do an ambitious property, rather than a Mickey Mouse property, so we're going against the given Law - Always do Mickey Mouse problems first.)

Example - Needles on a stick

Remember how we worked this out elaborately? [4/10/98.1-6]

Now this problem comes out in a couple of lines.




We have \$n\$ needles of length \$h\$. We drop them on a stick of length \$a\$ such that they don't stick out.

What is the probability that no two needles overlap?

$$\frac{(a-nh)^n}{a^n}$$

Let's see if we can derive this from the newly acquired knowledge that the uniform process is obtained by conditioning the Poisson process.



We want the event:

$$((T_1 > h) \cap (T_2 > h) \cap \dots \cap (T_n > h))$$

This is the event that the first \$n\$ blips are at least a distance \$h\$ from each other.

4/29/98.7

What's the probability of this event? (It's not the probability we are after).

$$P((T_1 > h) \cap (T_2 > h) \cap \dots \cap (T_n > h))$$

this is the intersection of independent events, because they are independent random variables that are identically distributed w/ an exponential distribution.

$$= P(T_1 > h) P(T_2 > h) \dots P(T_n > h)$$

$$\left\{ \begin{array}{l} \text{And since [4/24/98.10]:} \\ P(T > h) = e^{-\alpha h} \end{array} \right\}$$

$$= (e^{-\alpha h})^n$$

$$P((T_1 > h) \cap \dots \cap (T_n > h)) = e^{-n\alpha h}$$

But again, that's not the probability we are after.
What we want is:

$$P((T_1 > h) \cap (T_2 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = n)$$

↪ This will give us the probability that the needles will fall on the stick.
Unfortunately, when we expand this, by the definition of conditional probability, we have:

$$P((T_1 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = n) = \frac{P(((T_1 > h) \cap \dots \cap (T_n > h)) \cap (N_\alpha(a) = n))}{P(N_\alpha(a) = n)}$$

I don't see an easy way to obtain this numerator

So, instead, we compute the desired probability indirectly.
Recall, from above, that:

$$P((T_1 > h) \cap \dots \cap (T_n > h)) = e^{-n\alpha h}$$

Observe that the events $(N_\alpha(a) = 0), (N_\alpha(a) = 1), \dots$ are a partition of the sample space.

Therefore, we can compute the above LHS by the Law of Alternatives

{which we use more than
any other formula}



$$\underbrace{P((T_1 > h) \cap \dots \cap (T_n > h))}_{e^{-n\alpha h}} = \sum_{k \geq 0} P((T_1 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = k) \underbrace{P(N_\alpha(a) = k)}_{\frac{\alpha^k a^k}{k!} e^{-\alpha a}}$$

$$e^{-n\alpha h} = e^{-\alpha a} \sum_{k \geq 0} P((T_1 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = k) \frac{\alpha^k a^k}{k!}$$

$$\underbrace{e^{\alpha(a-nh)}}_{e^{\alpha(a-nh)}} = \sum_{k \geq 0} P((T_1 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = k) \frac{\alpha^k a^k}{k!}$$

We expand this exponential into a power series,
 $f(x) = e^x$

$$= \sum_{k \geq 0} \frac{f^{(k)}(0)}{k!} x^k \quad \leftarrow \text{Maclaurin series}$$

$$= \sum_{k \geq 0} \frac{x^k}{k!}$$

Let $x = \alpha(a-nh)$, which gives:

$$e^{\alpha(a-nh)} = \sum_{k \geq 0} \frac{(\alpha(a-nh))^k}{k!}$$

$$\sum_{k \geq 0} \frac{\alpha^k (a-nh)^k}{k!} = \sum_{k \geq 0} P((T_1 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = k) \frac{\alpha^k a^k}{k!}$$

We now have 2 power series in α , which are equal.

Since the Taylor expansion of a function is unique, the coefficients must be individually equal.

Therefore, for $k=n$, we obtain:

$$\frac{\alpha^n (a-nh)^n}{n!} = P((T_1 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = n) \frac{\alpha^n a^n}{n!}$$

which gives us the desired probability:

$$P((T_1 > h) \cap \dots \cap (T_n > h) | N_\alpha(a) = n) = \frac{(a-nh)^n}{a^n} \quad \checkmark$$

Q/A: Where did we use the fact that $a-nh$ is positive?

No where. There is some place where it must be implicitly used.
 I don't know where.

4/29/98.9


An interesting example for you to work out.
The density of the k^{th} order statistic.
Use the same method. } exercise

Property (5) Schrödinger Randomization

(an application of the preceding)

I take the interval $[0, 1)$,
And I take the Poisson Process $N_\alpha(t)$

Now we divide the interval $[0, 1)$ into n equal sub-intervals.



Then, we take:

$$N_\alpha\left(\frac{1}{n}\right) = \Theta_1$$

This is a random variable for given α, n . ↑ we'll call it Θ_1

Similarly, we have the following random variables:

$$N_\alpha\left(\frac{1}{n}\right) = \Theta_1$$

$$N_\alpha\left(\frac{2}{n}\right) - N_\alpha\left(\frac{1}{n}\right) = \Theta_2$$

$\vdots$

$$N_\alpha(1) - N_\alpha\left(\frac{n-1}{n}\right) = \Theta_n$$

These are occupation numbers
when you drop a random
number of blips.

(balls/blips into boxes)
except here, you are dropping
a random Poisson number of
balls into boxes.

Why is this good?

It's terrific.

Because dropping a random Poisson number of balls into the boxes has the property
that the number of balls in different boxes are independent.

Which was not the case when we dropped k balls into n boxes
(from our discussion of Maxwell-Boltzmann statistics [2/10/98.5]).

Here, $P(\Theta_1=32)$, $P(\Theta_2=17)$ and so forth are independent

$$P_{MB}(\theta_i = j) = P(N_\alpha(\frac{1}{n}) = j \mid N_\alpha(1) = k)$$

4/29/98.10

- If you want Maxwell-Boltzmann statistics, we have to condition.
And we can get any problem of Maxwell-Boltzmann statistics, relative to occupation numbers, by conditioning. It's automatic now.
So even that I could have skipped.
We could have started this course w/ the Poisson Process.

Maxwell-Boltzmann statistics is obtained by conditioning the Poisson Process.

Example

$$P((\theta_1 = \{0\}) \cap (\theta_2 = \{0\}) \cap \dots \cap (\theta_n = \{0\}) \mid N_\alpha(1) = k)$$

So we have k balls into n boxes
where each box
has 0 or 1 ball.

The boxes even look like boxes now.

And it better come out the way we computed it when we did the first computation we ever did in this course.
It better come out right.

Again, as in the previous example, I don't know how to compute this probability directly.

However, I do know how to compute:

$$P((\theta_1 = \{0\}) \cap (\theta_2 = \{0\}) \cap \dots \cap (\theta_n = \{0\}))$$

these are all independent events

Let's consider $P(\theta_i = \{0\})$:

$$P(\theta_i = \{0\}) = P((\theta_i = 0) \cup (\theta_i = 1))$$

$$= P(\theta_i = 0) + P(\theta_i = 1)$$

$$\begin{aligned} P(\theta_i = 0) &= P(N_\alpha(\frac{1}{n}) = 0) \\ &= P_\alpha\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}\right) \\ &= \frac{\alpha^0 (\frac{1}{n})^0 e^{-\alpha(\frac{1}{n})}}{0!} \\ &= e^{-\frac{\alpha}{n}} \end{aligned}$$

$$\begin{aligned} P(\theta_i = 1) &= P(N_\alpha(\frac{1}{n}) = 1) \\ &= P_\alpha\left(\begin{bmatrix} 0 \\ 1 \end{bmatrix}\right) \\ &= \frac{\alpha^1 (\frac{1}{n})^1 e^{-\alpha(\frac{1}{n})}}{1!} \\ &= \frac{\alpha}{n} e^{-\frac{\alpha}{n}} \end{aligned}$$

$$P(\theta_1 = \{0\}) = e^{-\frac{\alpha}{n}} + \frac{\alpha}{n} e^{-\frac{\alpha}{n}}$$

Similarly, for $\theta_2, \theta_3, \dots, \theta_n$.
These independent random variables have the same probability:

$$\begin{aligned} P(\theta_i = \{0\}) &= P((\theta_i = 0) \cup (\theta_i = 1)) \\ &= P(\theta_i = 0) + P(\theta_i = 1) \\ &= P_{\lambda}\left(\left[\begin{smallmatrix} \frac{i-1}{n}, \frac{i}{n} \\ 0 \end{smallmatrix}\right]\right) + P_{\lambda}\left(\left[\begin{smallmatrix} \frac{i-1}{n}, \frac{i}{n} \\ 1 \end{smallmatrix}\right]\right) \\ &= \frac{\alpha^0 \left(\frac{i}{n} - \left(\frac{i-1}{n}\right)\right)^0}{0!} e^{-\alpha \left(\frac{i}{n} - \left(\frac{i-1}{n}\right)\right)} + \frac{\alpha^1 \left(\frac{i}{n} - \left(\frac{i-1}{n}\right)\right)^1}{1!} e^{-\alpha \left(\frac{i}{n} - \left(\frac{i-1}{n}\right)\right)} \\ &= e^{-\frac{\alpha}{n}} + \frac{\alpha}{n} e^{-\frac{\alpha}{n}} \end{aligned}$$

Therefore, we have:

because the events are independent:

$$\begin{aligned} P((\theta_1 = \{0\}) \cap \dots \cap (\theta_n = \{0\})) &= P(\theta_1 = \{0\}) P(\theta_2 = \{0\}) \dots P(\theta_n = \{0\}) \\ &= \left(e^{-\frac{\alpha}{n}} + \frac{\alpha}{n} e^{-\frac{\alpha}{n}}\right)^n \\ &= \left(e^{-\frac{\alpha}{n}} \left(1 + \frac{\alpha}{n}\right)\right)^n \\ &= e^{-\alpha} \left(1 + \frac{\alpha}{n}\right)^n \end{aligned}$$

Very nice.

But we want the conditioning.

How do we get the conditioning?

We do it this way.

We use the Law of Alternatives

$$\underbrace{P((\theta_1 = \{0\}) \cap \dots \cap (\theta_n = \{0\}))}_{e^{-\alpha} \left(1 + \frac{\alpha}{n}\right)^n} = \sum_{k=0}^n P((\theta_1 = \{0\}) \cap \dots \cap (\theta_n = \{0\}) | N_{\alpha}(1) = k) \underbrace{P(N_{\alpha}(1) = k)}_{\frac{\alpha^k}{k!} e^{-\alpha}}$$

$$\cancel{e^{-\alpha} \left(1 + \frac{\alpha}{n}\right)^n} = \sum_{k=0}^n P((\theta_1 = \{0\}) \cap \dots \cap (\theta_n = \{0\}) | N_{\alpha}(1) = k) \frac{\alpha^k}{k!} \cancel{e^{-\alpha}}$$

$$\sum_{k=0}^n P(\theta_1=\{0\} \cap \dots \cap \theta_n=\{0\} | N_\alpha(1)=k) \frac{\alpha^k}{k!} = \left(1 + \frac{\alpha}{n}\right)^n$$

we use the binomial theorem
to expand this:

$$\left(\frac{\alpha}{n} + 1\right)^n = \sum_{k=0}^n \binom{n}{k} \left(\frac{\alpha}{n}\right)^k (1)^{n-k}$$

$$= \sum_{k=0}^n \frac{(n)_k}{k!} \frac{\alpha^k}{n^k}$$

And we have that:

$$P(\theta_1=\{0\} \cap \dots \cap \theta_n=\{0\} | N_\alpha(1)=k) = \frac{(n)_k}{n^k} \quad \checkmark$$

$(n)_k \leftarrow \{\text{total \# of ways of selecting } k \text{ unique boxes (equivalent to boxes w/ occupation numbers } 0, 1)\}$

$n^k \leftarrow \{\text{total \# of ways of dropping } k \text{ balls into } n \text{ boxes}\}$

This is one of the first computations we did in this course.

The occupation numbers with 0, 1.

Q: Why is it that we can say: if the sums are equal, the terms are equal?

A: Because if two functions have the same power series for the Taylor expansions, then they are equal. Equality of polynomial functions.

We have 2 functions of the variable α , which are equal.

α doesn't know that it is a variable.

Therefore, the Taylor series expansions are equal, therefore the coefficients are equal.

You can view the Schrödinger method as a probabilistic interpretation of Taylor expansion.

Another problem w/ Schrödinger Randomization:

4/29/98, 13

Example - Any problem related to occupation numbers can be done using the Poisson Process

Trivially. Always using the afore mentioned trick.

What is the probability that all the boxes are occupied? $P((\theta_1 > 0) \cap \dots \cap (\theta_n > 0) | N_\alpha(1) = k) = ?$

$$P((\theta_1 > 0) \cap \dots \cap (\theta_n > 0)) = ?$$

again, the occupation numbers of different intervals are independent random variables.

Hence, this is the intersection of independent events.

$$P((\theta_1 > 0) \cap \dots \cap (\theta_n > 0)) = P(\theta_1 > 0) P(\theta_2 > 0) \dots P(\theta_n > 0)$$

$$\text{for } i=1, \dots, n: P(\theta_i > 0) = 1 - P(\theta_i \leq 0) \leftarrow \text{by definition of probability.}$$

$$= 1 - P(\theta_i = 0) \leftarrow \text{negative occupation numbers are meaningless here.}$$

$$= 1 - P_\alpha\left[\begin{matrix} \lfloor \frac{1}{n} \rfloor \\ 0 \end{matrix}\right]$$

$$= 1 - \frac{\alpha^0 (\frac{1}{n})^0}{0!} e^{-\alpha (\frac{1}{n})}$$

$$P(\theta_i > 0) = 1 - e^{-\frac{\alpha}{n}}$$

$$= (1 - e^{-\frac{\alpha}{n}})^n$$

$$= \left(\frac{e^{-\frac{\alpha}{n}} - e^{-\frac{\alpha}{n}}}{e^{-\frac{\alpha}{n}}} \right)^n$$

$$= \left[e^{-\frac{\alpha}{n}} (e^{\frac{\alpha}{n}} - 1) \right]^n$$

$$P((\theta_1 > 0) \cap \dots \cap (\theta_n > 0)) = e^{-\alpha} (e^{\frac{\alpha}{n}} - 1)^n$$

By the Law of Alternatives:

$$P((\theta_1 > 0) \cap \dots \cap (\theta_n > 0)) = \sum_{k=0}^{\infty} P((\theta_1 > 0) \cap \dots \cap (\theta_n > 0) | N_\alpha(1) = k) P(N_\alpha(1) = k)$$

$$e^{-\alpha} (e^{\frac{\alpha}{n}} - 1)^n \quad \frac{\alpha^k}{k!} e^{-\alpha(1)}$$

4/29/98.14

For convenience, as in the text, let's denote events as follows:

$$A = ((\theta_1 > 0) \cap \dots \cap (\theta_n > 0))$$

and we've already showed that:

$$P(A) = e^{-\alpha} (e^{\frac{\alpha}{n}} - 1)^n$$

$$B_k = ((\theta_1 > 0) \cap \dots \cap (\theta_n > 0) | N_{\alpha}(1) = k)$$

↑ {these are the events whose probability we want to determine.}

Using this notation, we write the expression obtained from the Law of Alternatives as:

$$P(A) = \sum_{k \geq 0} P(B_k) \frac{\alpha^k}{k!} e^{-\alpha}$$

multiplying both sides by e^{α} :

$$e^{\alpha} P(A) = \sum_{k \geq 0} P(B_k) \frac{\alpha^k}{k!}$$

using $P(A)$, as obtained above, we can rewrite the LHS:

$$(*) \quad \underbrace{\left(e^{\frac{\alpha}{n}} - 1\right)^n}_{\text{expand this, using the binomial theorem:}} = \sum_{k \geq 0} P(B_k) \frac{\alpha^k}{k!}$$

expand this, using the binomial theorem:

$$\left((-1) + e^{\frac{\alpha}{n}}\right)^n = \sum_{j=0}^n \binom{n}{j} (-1)^j \left(e^{\frac{\alpha}{n}}\right)^{n-j}$$

$$= \sum_{j=0}^n \binom{n}{j} (-1)^j \underbrace{e^{(n-j)\frac{\alpha}{n}}}_{\text{expand using Maclaurin Series}}$$

$$e^x = \sum_{k \geq 0} \frac{x^k}{k!}, \quad \text{with } x = \frac{(n-j)\alpha}{n}$$

$$= \sum_{j=0}^n \binom{n}{j} (-1)^j \left[\sum_{k \geq 0} \frac{\left(\frac{(n-j)\alpha}{n}\right)^k}{k!} \right]$$

$$= \sum_{j=0}^n \binom{n}{j} (-1)^j \left[\sum_{k \geq 0} \left(1 - \frac{j}{n}\right)^k \frac{\alpha^k}{k!} \right]$$

$$= \sum_{j=0}^n \sum_{k \geq 0} \binom{n}{j} (-1)^j \left(1 - \frac{j}{n}\right)^k \frac{\alpha^k}{k!}$$

$$\left(e^{\frac{\alpha}{n}} - 1\right)^n = \sum_{k \geq 0} \underbrace{\sum_{j=0}^n \binom{n}{j} (-1)^j \left(1 - \frac{j}{n}\right)^k}_{P(B_k)} \frac{\alpha^k}{k!}$$

So now we have:

$$\sum_{k \geq 0} \underbrace{\sum_{j=0}^n \binom{n}{j} (-1)^j \left(1 - \frac{j}{n}\right)^k}_{P(B_k)} \frac{\alpha^k}{k!} = \sum_{k \geq 0} \underbrace{P(B_k)}_{\uparrow} \frac{\alpha^k}{k!}$$

The LHS and RHS are both polynomials in α .
Since they are equal, for each term α^k , the corresponding coefficients must be equal.

$$P(B_k) = \sum_{j=0}^n \binom{n}{j} (-1)^j \left(1 - \frac{j}{n}\right)^k$$

And this is precisely the probability we desire:

$$P(B_k) = P((\theta_1 > 0) \wedge \dots \wedge (\theta_n > 0) / N_{\alpha}(1) = k) = \sum_{j=0}^n \binom{n}{j} (-1)^j \left(1 - \frac{j}{n}\right)^k$$

This is the same as the formula we obtained by using the inclusion-exclusion principle early in the course [2/18/98.4-5].

- This Schrödinger Randomization is one of the most powerful probability techniques known.
- Another solution of equation (*):

$$\underbrace{\left(e^{\frac{\alpha}{n}} - 1\right)^n}_{\text{This is a function of } \alpha.} = \sum_{k \geq 0} \underbrace{P(B_k) \frac{\alpha^k}{k!}}_{\text{Note that we sum over all terms } \alpha^k.}$$

This is a function of α .
 $f(\alpha)$

Note that we sum over all terms α^k .
So this is just the Taylor expansion of the RHS.

4/29/98.16

From course 18.01, we have that:

$$f(x) = \sum_{k \geq 0} \frac{c_k}{k!} x^k,$$

$$\text{where } c_k = \left[D^k f(x) \right]_{x=0}$$

↑ {the k^{th} derivative of $f(x)$,
evaluated at $x=0$.}

Combining these two equations for $f(x)$, we have:

$$\sum_{k \geq 0} \underbrace{P(B_k)}_{\uparrow} \frac{x^k}{k!} = \sum_{k \geq 0} \underbrace{\left[D^k (e^{\frac{x}{n}} - 1)^n \right]_{x=0}}_{\uparrow} \frac{x^k}{k!}$$

This means that:

$$P(B_k) = P((\theta_1 > 0) \wedge \dots \wedge (\theta_n > 0) | N_{\alpha}(1) = k) = \left[D^k (e^{\frac{x}{n}} - 1)^n \right]_{x=0}$$

And this is a nice, closed form expression →

Not bad.

It's better than the formula that we obtained on the previous page, as well as w/ inclusion-exclusion:

It's an interesting exercise to show that this formula is the same as the formula obtained before.

Probability that all n boxes are occupied when k balls dropped:


$$P(B_k) = \left[D^k (e^{\frac{x}{n}} - 1)^n \right]_{x=0} = \sum_{j=0}^n \binom{n}{j} (-1)^j \left(1 - \frac{j}{n}\right)^k$$

4/29/98.17

Property (6) Poisson blips in 2 colors

We have Poisson blips, but now the blips come in 2 colors - red and black.
 We have a machine that only registers red blips
 it misses black blips.

How do we work this in?



Q: What remains to be said?

A: We have to tell what the probability that a given blip will be red or black.
 otherwise, you haven't specified it.

Therefore, we have a Poisson process together with the probability that a blip will be red or black.

How do we say this correctly? Once we say it correctly, of course we can compute it.
 Here's how we say it correctly.

$X_1, X_2, \dots$ = sequence of independent Bernoulli random variables
 where $P(X_i = 1) = p$
 $\begin{cases} 1 = \text{red} \\ 0 = \text{black} \end{cases}$

Then, how do we write that the machine counts only the red blips?
 This is an important realization.

$X_1 + X_2 + \dots + X_{N_\lambda(t)}$ = number of red blips in the interval $[0, t]$

Next time, we will compute probability:

$$P(X_1 + X_2 + \dots + X_{N_\lambda(t)} = n)$$

by the Law of Alternatives.

Note, of course, that $N_\lambda(t)$ is independent of the X_i .

This is an exact way of expressing that we have a machine that measures red blips.



The Poisson Process (Cont'd) / Random Walk and the Normal Distribution

The last lecture we did, on Schrödinger Randomization, was probably the hardest lecture in the whole course.

Basically we did the lecture in one hour.
It took me several years to understand it.

There are 2 properties of the Poisson Process left to describe.
Then we start on the last big topic of this course, which is the normal distribution and the stochastic process associated w/ the normal distribution - namely, Brownian motion.

50 SHEETS
22-141
100 SHEETS
22-142
200 SHEETS
22-144



Property (6) Poisson blips in 2 colors

Suppose we have a Poisson Process where blips are red or black
red w/ probability p
black w/ probability $q = 1-p$

We are only interested in registering the red blips.
How many red blips do we get in a given time interval?

$R_\alpha(t)$ = number of red blips in interval $[0, t]$

↑ this is a random function

Our objective is to compute the probability distribution of this random function:

$$P(R_\alpha(t) = k) = ?$$

$$P(X_i = 1) = p$$

$$P(X_i = 0) = q$$

this is a joint distribution

As we observed last time, for each X_i , you get count of $\begin{cases} 1 \text{ w/ probability } p & \text{(red)} \\ 0 \text{ w/ } & q & \text{(black)} \end{cases}$

$$R_\alpha(t) = X_1 + X_2 + \dots + X_{N_\alpha(t)}$$

All the X_i are independent.

↑ there are, altogether,
 $N_\alpha(t)$ blips in the interval $[0, t]$

So, we want to compute the probability:

$$P(R_\alpha(t) = k) = P(X_1 + X_2 + \dots + X_{N_\alpha(t)} = k) = ?$$

How do you think we compute this?

The only way known to man, The Law of Alternatives.

by the Law of Alternatives:

5/1/98.2

$$P(X_1 + X_2 + \dots + X_{N_d(t)} = k) = \sum_{n \geq 0} P(X_1 + X_2 + \dots + X_{N_d(t)} = k \mid N_d(t) = n) P(N_d(t) = n)$$

Since it is given that $N_d(t) = n$,
we have that event:

$$(X_1 + X_2 + \dots + X_{N_d(t)} = k \mid N_d(t) = n) = (X_1 + X_2 + \dots + X_n = k)$$

$$= \sum_{n \geq 0} P(X_1 + X_2 + \dots + X_n = k) P(N_d(t) = n)$$

$$\frac{\alpha^n t^n}{n!} e^{-\alpha t}$$

the formula for
the Poisson distribution

$$= \sum_{n \geq 0} P(X_1 + X_2 + \dots + X_n = k) \frac{(\alpha t)^n}{n!} e^{-\alpha t}$$

Next, we have to remind ourselves what probability distribution
of the sum of n independent random variables is.

This is called the binomial distribution.

Remember, The good old days, [2/20/98.7-9]

Let's perform an interchange of notation to remind
ourselves of the formula.

$$S_n = X_1 + X_2 + \dots + X_n$$

and we have:

$$P(S_n = k) = \begin{cases} \binom{n}{k} p^k q^{n-k} & , n \geq k \\ 0 & \text{otherwise} \end{cases}$$

$$= \sum_{n \geq 0} P(S_n = k) \frac{(\alpha t)^n}{n!} e^{-\alpha t}$$

$$= \sum_{n \geq k} \binom{n}{k} p^k q^{n-k} \frac{(\alpha t)^n}{n!} e^{-\alpha t}$$

$$= \sum_{n \geq k} \frac{(n)_k}{k!} p^k q^{n-k} \frac{(\alpha t)^n}{n!} e^{-\alpha t}$$

$$= \frac{1}{k!} p^k q^{-k} e^{-\alpha t} \sum_{n \geq k} \frac{(n)_k}{n!} (\alpha t q)^n$$

$$\text{let } j = n - k$$

$$= \frac{1}{k!} p^k q^{-k} e^{-\alpha t} \sum_{j \geq 0} \frac{(j+k)_k}{(j+k)!} (\alpha t q)^{j+k}$$

$$\frac{(j+k)_k}{(j+k)!} = \frac{(j+k) \cdots (j+1)}{(j+k) \cdots (j+1) j!} = \frac{1}{j!}$$

$$= \frac{1}{k!} p^k q^{-k} e^{-\alpha t} (\alpha t q)^k \sum_{j \geq 0} \frac{(\alpha t q)^j}{j!}$$

This is nice.
This is the Taylor expansion
for $e^{\alpha t q}$

$$= \frac{1}{k!} (p\alpha)^k t^k \underbrace{e^{-\alpha t} e^{\alpha t q}}_{e^{-\alpha t} e^{\alpha t q} = e^{-\alpha t(1-q)} = e^{-\alpha t p}}$$

$$\boxed{P(X_1 + X_2 + \dots + X_{N_\alpha(t)} = k) = \frac{(p\alpha)^k t^k}{k!} e^{-(p\alpha)t} \\ P(R_\alpha(t) = k)}$$

What is this trying to tell us?

It's trying to tell us that the random function $R_\alpha(t)$, the count of the red blips, is a Poisson process, where the intensity has changed from α to $p\alpha$.

$$\alpha \rightarrow p\alpha \quad \frac{(p\alpha)^k t^k}{k!} e^{-(p\alpha)t}$$

This is intuitively obvious.
Especially after $\frac{1}{2}$ hour of computation.

5/1/98.4

Conclusion

If you sample only the real blips, you still have a Poisson process, but the intensity is changed by multiplying by the probability of a real blip.

$$\alpha \rightarrow (p\alpha)$$

See, the Poisson Process has all these nice properties.
Last one (but we could go on forever).

Property (7) Laplace Transform

If you are given only the average number of blips in a unit interval, then the only stochastic process that fits these data is the Poisson process, Complete randomness, except for the intensity α .

We've said this several times already. Loudly.

Now you say: Excuse me, that's not completely random.
If it were completely random, then even α would be random.

Let's see what happens if we take the Poisson process and you make the intensity α into a random variable, continuous, with a given density.
We've got the technique.

Suppose α is random with density $\text{dens}(\alpha=s) = f(s)$

{I want to give $\text{dens}(\alpha=s)$ a name.}

Next, we compute the probability that there are no blips in the interval $[0, t)$:

$$P(N_\alpha(t)=0) = \int_{-\infty}^{\infty} P(N_\alpha(t)=0 | \alpha=s) \overbrace{\text{dens}(\alpha=s)}^{f(s)} ds$$

{The only way known to man to do this.
The continuous law of Alternatives.}

$$= \int_0^{\infty} P(N_\alpha(t)=0 | \alpha=s) f(s) ds$$

since it is given that $\alpha=s$,
we have that events:

$$(N_\alpha(t)=0 | \alpha=s) = (N_s(t)=0)$$

$$= \int_0^{\infty} P(N_s(t)=0) f(s) ds$$

$$\frac{s^0 t^0}{0!} e^{-st} = e^{-st}$$

$$P(N_\alpha(t)=0) = \int_0^{\infty} e^{-st} f(s) ds$$

The Laplace Transform

5/1/98.5

This is where the Laplace transform comes from.
This is a probabilistic interpretation of the Laplace transform.

Every fact of the Laplace transform can be interpreted probabilistically from this.
We already saw convolution.
Now we see the Laplace Transform.

- This stochastic process, namely the Poisson process w/ random intensity is, in the literature, called the Cox process.
We don't have enough time in class, but there are wonderful things you can say about this.

That's it kids.
That's the end of the Poisson process.

Now, we start on the most important topic of this course.
The Theory of the Normal Distribution
You could say the rest of this course was preparation for this topic.

Let's start w/ something we have strangely neglected:

- Random Walk

It's really something we have done implicitly.
Let's work it out.

Take the Bernoulli process.
We toss a coin.

If we get heads, write +1
 tails, write -1

} That's the only change.
But this change makes a tremendous psychological difference, because you can graph it.

So we take:

$Y_1, Y_2, \dots$ i.i.d. random variables

Let assume, here, that the coin is fair:

$$P(Y_i = +1) = P(Y_i = -1) = \frac{1}{2}$$

And we define the random variable:


$$\begin{aligned} G_n &= \text{position at time } n \\ &= Y_1 + Y_2 + \dots + Y_n \end{aligned}$$

22-141 50 SHEETS
22-142 100 SHEETS
22-146 200 SHEETS



5/1/98.6

We are pretty much tossing a fair coin, but the language changes.
And we graph it as follows:



every sample $\omega = (\pm 1, \pm 1, \dots)$

The sample space Ω is the set of all sequences of ± 1 .

To every sample point ω , you associate a continuous curve, as illustrated above.
As soon as you do that, you realize there are questions to be asked.
For example:

Q: How long (i.e., number of tosses) does it take to get back to 0?

Q: How long does it take to reach some particular value of x ?

All sorts of questions from this "different" viewpoint.
You'll find all sorts of these answered in the book.

Let's consider:

$$P(G_n = x) = ?$$

probability that at time n , you are at level x

every toss has probability $\frac{1}{2}$ of being $\begin{cases} \text{heads} \rightarrow +1 \\ \text{tails} \rightarrow -1 \end{cases}$

Suppose we have:

$u = \# \text{ of } +1\text{'s}$

$v = \# \text{ of } -1\text{'s}$

$$\left. \begin{array}{l} u = \# \text{ of } +1\text{'s} \\ v = \# \text{ of } -1\text{'s} \end{array} \right\} \begin{array}{l} u + v = n \leftarrow \text{number of tosses} \\ u - v = x \leftarrow \text{level} \end{array}$$

$$\left. \begin{array}{l} v = u - x \\ u + (u - x) = n \end{array} \right\} \Rightarrow u = \frac{n+x}{2}$$

$$\left. \begin{array}{l} u = x + v \\ (x+v) + v = n \end{array} \right\} \Rightarrow v = \frac{n-x}{2}$$



5/1/98.7

Consider the event $(G_n = x)$ in terms of the standard random variable S_n in the Bernoulli process [2/20/98.7-9].
Recall S_n is the number of heads in n tosses.

$$\underbrace{(G_n = x)}_{\substack{\text{after } n \text{ tosses, you} \\ \text{are at level } x}} = \underbrace{(S_n = u)}_{\substack{\text{after } n \text{ tosses, you} \\ \text{have } u \text{ heads}}}$$

$$\begin{aligned} P(G_n = x) &= P(S_n = u) \quad \leftarrow \text{binomial distribution} \\ &= \binom{n}{u} \left(\frac{1}{2}\right)^u \left(\frac{1}{2}\right)^{n-u} \quad \leftarrow n-u=v \\ &= \binom{n}{u} \frac{1}{2^n} \quad \text{and since } u = \frac{n+x}{2}, \end{aligned}$$

$$\boxed{P(G_n = x) = \binom{n}{\frac{n+x}{2}} \frac{1}{2^n}}$$

• What is :

$$P(G_{2n} = 0) = ?$$

why $2n$? Because we can not go to 0 in an odd number of steps. There must be an even number of steps. Figure it out for yourselves.

From above, we have :

$$\begin{aligned} P(G_{2n} = 0) &= \binom{2n}{\frac{2n+0}{2}} \frac{1}{2^{2n}} \\ &= \binom{2n}{n} \frac{1}{2^{2n}} \end{aligned}$$


(*)

Now let's define the random variable :

$F =$ time of first return to origin

5/1/98.8

If you look just at the x axis, and you project the curve onto the x axis



Then you can visualize a random walk as a random walk. You go left, right, right, ...
Hence the name.
So it makes sense to ask when you return to the origin.

Now, let me tell you informally something that we are not going to establish, but that you ought to know.
It's part of your cultural background.

Probability that a random walk that starts at the origin ever comes back to the origin is 1, in 1 dimension and 2 dimensions.
In 3-dimensions, the probability is .397.
This is an incredible fact of life. You can make money on this.
In 1+2 dimensions, all roads go to 1.
In 3 dimensions, not.

dimension	$P(\text{random walk starting at origin returns to origin})$
1	1
2	1
3	.397
4+	weirdo. Less + less weirdo than exponential

This makes it all the more decisive to compute the probability distribution of this random variable F , the first time you come back to the origin.

Let's define:

$$f_{2n} \triangleq P(F=2n)$$

What is the probability distribution?
 $f_{2n} = P(F=2n) = ?$

5/1/98.9

You know, we've covered 3 terms of probability in this course.
 Not that I want to brag.
 But this is 3 terms of probability.
 If you ever take another probability course later 18.313 will be vindicated.
 Because, you will say: why is this guy doing all this trivial stuff?
 Just watch. I promise this is going to happen.

So we next consider the event $(G_{2n} = 0)$.

Now we use a very important trick.

If you returned to 0 at time equal $2n$, that means there must be a first time when you return to 0.

If you don't see this, I stop.

Let's write that down.

$$(G_{2n} = 0) = ((F=2) \cap (G_{2n-2} = 0)) \cup ((F=4) \cap (G_{2n-4} = 0)) \cup \dots \cup (F=2n)$$

at time $2n$ you were at origin first return to origin was at time 2, followed by time $2n-2$, at which time you were at origin. Thus, at origin at time $2+(2n-2)=2n$

$2+(2n-2)=2n$ $4+(2n-4)=2n$

Now we take probabilities.

On the RHS, we have the union of disjoint events.

And the intersectands are independent random variables (F 's and G 's are independent)

$$P(G_{2n} = 0) = P(F=2)P(G_{2n-2} = 0) + P(F=4)P(G_{2n-4} = 0) + \dots + P(F=2n)$$

As is tradition, set $P(G_{2n} = 0) = U_{2n}$. Rewriting, we obtain:

$$(**) \quad U_{2n} = f_2 U_{2n-2} + f_4 U_{2n-4} + \dots + f_{2n}$$

But, from equation (*), we have that:

$$U_{2n} = P(G_{2n} = 0) = \binom{2n}{n} \frac{1}{2^{2n}}$$

Therefore, we can solve recursively for the f 's

$$\begin{cases} U_{2n} = \binom{2n}{n} \frac{1}{2^{2n}} \\ U_{2n} = f_2 U_{2n-2} + f_4 U_{2n-4} + \dots + f_{2n} \end{cases}$$

giving the desired probabilities:

$$f_{2i} = P(F=2i)$$

Alternatively, we can use the z-transform.

$$\text{Set: } U(z) = 1 + U_2 z^2 + U_4 z^4 + U_6 z^6 + \dots$$

$$F(z) = f_2 z^2 + f_4 z^4 + f_6 z^6 + \dots$$

Then we have:

$$U(z) - 1 = U(z) F(z)$$

Proof sketch

First, multiply the two power series on the RHS.

$$\begin{aligned} U(z)F(z) &= \boxed{f_2} z^2 + \boxed{f_4} z^4 + \boxed{f_6} z^6 + \boxed{f_8} z^8 + \dots \\ &\quad + \boxed{U_2 f_2} z^4 + \boxed{U_2 f_4} z^6 + \boxed{U_2 f_6} z^8 + \dots \\ &\quad + \boxed{U_4 f_2} z^6 + \boxed{U_4 f_4} z^8 + \dots \\ &\quad + \boxed{U_6 f_2} z^8 + \dots \end{aligned}$$

$$= \sum_{n \geq 1} \left((U_{2n-2} f_2 + U_{2n-4} f_4 + \dots + U_{2n-2n} f_{2n}) z^{2n} \right)$$

$$= \sum_{n \geq 1} \left(\left(\sum_{j=1}^n U_{2n-2j} f_{2j} \right) z^{2n} \right)$$

$$\underbrace{U(z) - 1} = \sum_{n \geq 1} \left(\left(\sum_{j=1}^n U_{2n-2j} f_{2j} \right) z^{2n} \right)$$

we defined $U(z)$ above. It can be written as:

$$U(z) = 1 + \sum_{n \geq 1} (U_{2n} z^{2n})$$

Subtracting 1 gives the LHS above:

$$\sum_{n \geq 1} (U_{2n} z^{2n}) = \sum_{n \geq 1} \left(\left(\sum_{j=1}^n U_{2n-2j} f_{2j} \right) z^{2n} \right)$$

These two polynomials in z are equal only if the coefficients of each term z^{2n} are equal. Thus this requires that:

$$\text{for all } n \geq 1, \quad U_{2n} = \sum_{j=1}^n (U_{2n-2j} f_{2j}) \quad \left. \begin{array}{l} \text{This is a tautology.} \\ \text{This is precisely equation (**).} \\ \text{Thus we have shown that:} \\ U(z) - 1 = U(z) F(z) \quad \checkmark \end{array} \right\}$$

Since $U(z) - 1 = U(z)F(z)$, we have that:

$$F(z) = \frac{U(z) - 1}{U(z)}$$

$$F(z) = 1 - \frac{1}{U(z)}$$

Another way of writing $U(z)$ is:

$$U(z) = \sum_{n \geq 0} (U_{2n} z^{2n})$$

↑ recall that we proved earlier [5/1/98.7] that:

$$U_{2n} = P(G_{2n} = 0) = \binom{2n}{n} \frac{1}{2^{2n}}$$

$$= \sum_{n \geq 0} \left(\binom{2n}{n} \frac{1}{2^{2n}} z^{2n} \right)$$

"An easy computation shows"

Well, a straight forward computation, anyway, shows that this is the Maclaurin Series expansion of the function: $(1 - z^2)^{-\frac{1}{2}}$

$$U(z) = (1 - z^2)^{-\frac{1}{2}}$$

And we can rewrite $F(z)$ above as:

$$F(z) = 1 - (1 - z^2)^{\frac{1}{2}}$$

from our definition above,

$$F(z) = f_2 z^2 + f_4 z^4 + \dots$$

$$\sum_{n \geq 1} (f_{2n} z^{2n}) = 1 - (1 - z^2)^{\frac{1}{2}}$$

expand by the binomial theorem w/ fractional exponents

$$= 1 - (1 + (-z^2))^{\frac{1}{2}}$$

$$= 1 - \sum_{n \geq 0} \left(\frac{1}{2}\right)^n (-z^2)^n$$

$$= \cancel{1} - \cancel{1} - \sum_{n \geq 1} \left(\frac{1}{2}\right)^n (-z^2)^n$$

$$\sum_{n \geq 1} (f_{2n} z^{2n}) = - \sum_{n \geq 1} \left(\left(\frac{1}{2}\right)^n (-1)^n z^{2n}\right)$$

These z polynomials are equal and, thus, the coefficients for each term z^{2n} are equal.

$$f_{2n} = \pm \left(\frac{1}{2}\right)^n, \text{ for all } n \geq 1$$

$$\left(\frac{1}{2}\right)^n \stackrel{n \text{ terms}}{=} \frac{\frac{1}{2} \left(\frac{1}{2}-1\right) \left(\frac{1}{2}-2\right) \dots \left(\frac{1}{2}-(n-1)\right)}{n!}$$

$$= \frac{-\frac{1}{2} \left(-\frac{1}{2}\right) \left(-\frac{3}{2}\right) \dots \left(-\frac{(2n-3)}{2}\right)}{n!}$$

$$= (-1)^{n-1} \frac{\frac{1}{2} \left(\frac{1}{2}\right) \left(\frac{3}{2}\right) \dots \left(\frac{(2n-3)}{2}\right)}{n!} \times \left(\frac{2^n}{2^n}\right)$$

$$= \frac{(-1)^{n-1}}{2^n} \frac{\overbrace{1 \times 1 \times 3 \times \dots \times (2n-3)}^{n-1 \text{ terms}}}{n!} \times \left(\frac{\overbrace{2 \times 4 \times \dots \times (2n-2)}^{n-1 \text{ terms}}}{2 \times 4 \times \dots \times (2n-2)}\right)$$

$$= \frac{(-1)^{n-1}}{2^n} \frac{(2n-2)!}{n! (2 \times 1)(2 \times 2) \dots (2 \times n-1)}$$

$$\uparrow \frac{2^{n-1} (n-1)!}{2^{n-1} (n-1)!}$$

$$= \frac{(-1)^{n-1}}{2^n} \frac{(2n-2)!}{n! \cdot 2^{n-1} (n-1)!}$$

$$= (-1)^{n-1} \frac{(2n-2)!}{(n-1)! n!} \frac{1}{2^{2n-1}}$$

$$= (-1)^{n-1} \frac{1}{2^{n-1}} \frac{(2n-1)!}{(n-1)! n!} \frac{1}{2^{2n-1}}$$

$$= (-1)^{n-1} \frac{1}{2^{n-1}} \binom{2n-1}{n-1} \frac{1}{2^{2n-1}}$$

$$\left|\left(\frac{1}{2}\right)^n\right| = \left|\frac{1}{2^{n-1}} \binom{2n-1}{n-1} \frac{1}{2^{2n-1}}\right|, \text{ for all } n \geq 1$$

5/1/98.13

And now we have the desired probability distribution we have been after - namely, that of the random variable F , the first time you come back to the origin on a random walk.

$$f_{2n} \triangleq P(F=2n) = \left| \frac{1}{2^{n-1}} \binom{2n-1}{n-1} \frac{1}{2^{2n-1}} \right|, \quad \forall n \geq 1$$

This alternative computation, where we use the z -transform to solve for the probability distribution f_{2n} , is far from a trivial computation.

Q/A: This is related to the Catalan numbers.

Prof. Stanley has just finished the 2nd volume of his book on combinatorics, which includes a chapter called "62 permutations of the z -transform."

Now we come to the real point of this lecture.
The real point is this.

Is there a continuous analogue of random walk?

In other words, from the Bernoulli process, we obtained the Poisson process by a limiting process.


Is there another limiting process where, from the random walk interpretation of the Bernoulli process the steps become shorter and shorter, and (if you don't get too much oscillation, so that the limit is well behaved), you can obtain this other limiting process?

The amazing thing is, one of the best chapters of probability, that there is such a limiting process.

How do you get it?

Didn't you all go to the Museum of Science when you were kids?
When you went there, what did you see?
You saw this.

This is why I drew the random walk this way.



The marbles are dropped and displaced by the pins.
This is a random walk

And the fact that the marbles pile up according to a certain curve, means that there is a limiting process.

This limiting process gives you the Gaussian curve as the limiting curve.

Next: We will study this limiting process.

The Bell Shaped Curve

Let's continue the discourse we initiated last time, when we briefly surveyed the notion of random walk.

Recall that a symmetric random walk is just:

sequence $Y_1, Y_2, \dots$ i.i.d. random variables with the property that:

$$P(Y_i = +1) = \frac{1}{2}, \quad P(Y_i = -1) = \frac{1}{2}$$


just like coin tossing, where heads = +1
tails = -1

Under these circumstances, we define:

$$G_n = Y_1 + Y_2 + \dots + Y_n$$

position at time n

The intuitive interpretation, as we saw last time, is a random walk.



Namely, you look at the x axis w/ integer points.

Then you have a drunkard starting at the origin.

If the coin comes up heads, the drunkard moves to the right +1
tails, " " " " " left -1

This process has infinitely many properties that are still being studied.


We call the random walk symmetric because we use a fair coin.

Actually, if the coin is not fair, the properties of random walk are the same as they would be for a fair coin, except for a drift.

The drift corresponds to the unfairness of the coin.

So the symmetric random walk gives you all the fundamental properties.

We can visualize this by graphing a random walk as a polygonal curve.



5/4/98.2

You can view a stochastic process, i.e., random walk, whose sample space Ω consists of all these paths.

The sample space is the set of all sequences of ± 1 . A sample point $\omega = (\pm 1, \pm 1, \dots)$.


But you can visualize it by a path.

∞ infinite

So you are putting a probability on the set of all paths (i.e., this is equivalent to putting a probability on the set of all sequences).

The paths are infinite. For finite paths, you just condition.

As we saw last time, from our Museum of Science contraption:



When the balls are dropped, they form a curve, which is a 'bell shaped' curve.

The equation of this curve is:

$$\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$$

Standard Normal Distribution ← the explanation of this bell shaped curve.

This is one of the important chapters in probability — in fact, of mathematics.

Let's first verify that the Standard Normal Distribution is indeed a probability distribution.

$$\int_{-\infty}^{\infty} P(X=x) dx = 1$$

So, if $\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}$ is a probability distribution, integral must be 1. $\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \geq 0$.
Checks:

$$\underbrace{\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx}_{\text{call this integral } A} \stackrel{?}{=} 1$$

{ We did this in 18.03 }

use a different variable of integration
for the 2nd term

5/4/98.3

$$A^2 = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy$$

$$= \frac{1}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-\frac{(x^2+y^2)}{2}} dx dy$$

Switch to polar coordinates

$$\int_D f(x,y) dA = \int_{r_0}^{r_1} \int_{\theta_0}^{\theta_1} f(r \cos \theta, r \sin \theta) r d\theta dr$$

$$= \frac{1}{2\pi} \int_{r=0}^{\infty} \int_{\theta=0}^{2\pi} e^{-\frac{r^2}{2}} r d\theta dr$$

$$\int_0^{2\pi} e^{-\frac{r^2}{2}} r d\theta = 2\pi e^{-\frac{r^2}{2}} r$$

$$= \frac{1}{2\pi} \int_0^{\infty} 2\pi e^{-\frac{r^2}{2}} r dr$$

$$= \int_0^{\infty} e^{-\frac{r^2}{2}} r dr$$

$$\text{Let } u = \frac{r^2}{2}$$

$$du = r dr$$

$$= \int_{r=0}^{\infty} e^{-u} du$$

$$= -e^{-u} \Big|_{r=0}^{\infty}$$

$$= -e^{-\frac{r^2}{2}} \Big|_{r=0}^{\infty}$$

$$= 0 - (-1)$$

$$A^2 = 1 \checkmark$$

$$\text{Thus: } A = \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx = 1 \checkmark$$

The bell shaped curve does indeed define a probability distribution.

5/4/98.4

On which sample space?Does there exist a sample space and a random variable on that sample space which has this probability distribution?These questions are easily justified using a "Grant's Tomb" argument.
As follows:Sample space $\Omega = (-\infty, \infty)$
the linesample point $\omega =$ point on the line

$\left\{ \begin{array}{l} \text{this is the only case we've covered where} \\ \text{the sample points are actual points.} \end{array} \right\}$

events are the events generated by the intervals

$\leftarrow$ events like we did for the uniform process
 $(a, b]$, etc.

Define probability by setting:

$$P((a, b]) = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-\frac{x^2}{2}} dx$$

$\leftarrow$ that's a probability.


We just verified that the integral over the whole line is 1:

$$P((-\infty, \infty)) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{x^2}{2}} dx = 1$$

So this is a sample space.

It may not be what you expect for a sample space, but you have to swallow it.

Now let's find a random variable.

The random variable is that we pick a point on the line at random.Choose a point at random, what's the probability distribution? $\leftarrow$ "Who's buried in Grant's Tomb?"random variable $Z =$ choose a point at random

$$P(a < Z \leq b) = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-\frac{z^2}{2}} dz$$

Because we built it that way. It's circular.

5/4/98, 5

Recall the definition of density:

$$P(a < Z \leq b) = \int_a^b \text{dens}(Z=t) dt$$

So, fact that:

$$P(a < Z \leq b) = \int_a^b \underbrace{\frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}}}_{\downarrow} dx$$

means that:

$$\text{dens}(Z=t) = \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}}$$

- This trick works - whenever you have a distribution and you want to construct a sample space.

A random variable, having the given probability distribution, can be constructed this way.

This shows that it is always possible to realize a probability distribution by a random variable.

- Z is called a standard normal distributed random variable

↑
this notation is universal

- Now, we want a justification for this phenomenon.

Here, we get to the frontiers of science.

The problem is not that there is no justification.

The problem is there are too many.

When there are several completely different justifications for the same phenomenon, something is funny.

You are stunned if there is just one.

But, when there are several, from completely different points of view, something may be missing.

I will point out the gaps in everybody's knowledge, which people don't like to talk about.

We'll discuss 3 justifications:

- 1) Gauss
- 2) Maxwell - Einstein
- 3) Central Limit Theorem of Probability

} 3 different justifications

- First, let's see what is it we want to justify.
Let's be clear as to what is it we want to justify.
To do that, we have to introduce a new notion:

Variance ←

- Variance of a random variable

First, we'll give the definition of variance, then we'll make comments on what it is really about.

Character Assassination,
aka Course Evaluations,
done.

1) letter for his eyes
only.

2) MIT evaluation.

Let X = any random variable with expectation

$$E(X) = m$$

The variance, by definition, of X is:

$$\text{Var}(X) = E((X - m)^2)$$

rewriting this using $E(X) = m$:

$$= E((X - E(X))^2)$$

expanding this

$$= E(X^2 - X E(X) + E(X)^2)$$

$$\text{Var}(X) = E(X^2 - E(X)^2)$$

or, if you prefer:

$$\text{Var}(X) = E(X^2) - m^2$$

Before we ask what variance is about,
let's back up and look again at random variables.

- What is a real random variable, really?

It can be one of many things, depending on what you are interested in.

View (1) A random variable describes a random phenomenon

By giving the phenomenon a probability distribution, instead of a specific value.

It's a natural, logical thing to describe ^{such} phenomena, not by a single value, but by a range of values together with their probabilities.

View (2) Random variable is the result of a search

↑ we touched briefly on this in SyarClass 4 [4/24/98]

Every integer random variable defines a partition of the sample space. When you do a search, the random variable is the measurement that tells you which block in this partition the unknown sample point is in.

That's the beginning of search theory, information theory, and all that.

But, now, we are seeing a 3rd concept of a random variable.

View (3) Random Variable as measurement of an imperfectly given quantity

For example, you want to measure people's height.

You measure 1000 people and get their average.

You want to know how close is this average to the average you would get by measuring everyone on the planet's height.

You can view the process of measuring someone's height as:

The process of measuring an average and the deviation from that average.

Or: You have a measurement process, with a measuring device that is imperfect. So, it varies a lot.

The variance of a random variable is related to the interpretation of the random variable as an imperfect measurement in view of the expectations, i.e., View (3) above.

The variance tells you how accurate the random variable is.

You are squaring the deviation of the random variable from its average value.

$$\text{Var}(X) = E((X - E(X))^2)$$

$E(X) = m$

$\text{Var}(X)$ will be big if the random variable is spread out.

small " " " " is concentrated around its average

Variance is a measure of spread

It is the best measure of spread

In the limit, suppose that X is a random variable that takes the value 3 with probability 1.

$$\left. \begin{array}{l} P(X=3) = 1 \\ E(X) = 3 \end{array} \right\} \Rightarrow \text{Var}(X) = E\left(\left(\overset{\substack{\uparrow \\ X=3, \text{ in} \\ \text{the limit}}}{X} - \underbrace{E(X)}_3\right)^2\right)$$

$$= 0$$

So, if the random variable X takes only one value, its variance $\text{Var}(X)$ is 0.

If the variables are all completely spread out, the variance will be big.
We'll see examples w/ the normal distribution.

Properties of Variance

Whereas $E(X+Y) = E(X) + E(Y)$ always,
it is not necessarily true that

$$\text{Var}(X+Y) \stackrel{?}{=} \text{Var}(X) + \text{Var}(Y)$$

However, if X and Y are independent random variables:

$$\text{Var}(X+Y) = \text{Var}(X) + \text{Var}(Y)$$

$\uparrow$ X, Y independent,
then we have equality.

Let's verify this fact that variance is additive for independent random variables.

Proof X, Y are independent random variablesSince $\text{Var}(Z) = E((Z - E(Z))^2)$, the variance of the sum $X+Y$ is, by definition:

$$\text{Var}(X+Y) = E\left(\underbrace{((X+Y) - E(X+Y))^2}_{\text{expand this out}}\right)$$

$$= E\left[(X+Y)^2 - 2(X+Y)\underbrace{E(X+Y)}_{(E(X)+E(Y))} + \underbrace{E(X+Y)^2}_{(E(X)+E(Y))^2}\right]$$

$$= E\left[X^2 + 2XY + Y^2 - 2XE(X) - 2YE(X) - 2XE(Y) - 2YE(Y) + E(X)^2 + 2E(X)E(Y) + E(Y)^2\right]$$

since expectation of a sum is the sum of the expectations:

$$= E(X^2) + E(2XY) + E(Y^2) - E(2XE(X)) - E(2YE(X)) - E(2XE(Y)) - E(2YE(Y)) + E(E(X)^2) + E(2E(X)E(Y)) + E(E(Y)^2)$$

remember that, for c some constant:

$$E(cZ) = E(\underbrace{Z + Z + \dots + Z}_c) = cE(Z)$$

also, the expectation is a constant.

$$E(E(Z)) = E(Z)$$

 $E(Z)$ is a constant, so it goes out of the expectation

For example:

$$E(2YE(X)) = E(\underbrace{2E(X)}_{\text{constant}} Y) = \underbrace{2E(X)}_{\text{constant}} E(Y)$$

$$\begin{aligned}
 &= E(X^2) + 2E(XY) + E(Y^2) \\
 &\quad - \cancel{2E(X)^2} - \cancel{2E(X)E(Y)} - \cancel{2E(Y)E(X)} - \cancel{2E(Y)^2} \\
 &\quad + \cancel{E(X)^2} + \cancel{2E(X)E(Y)} + \cancel{E(Y)^2}
 \end{aligned}$$

Since $X+Y$ are independent, the product rule for expectation states that:

$$E(XY) = E(X)E(Y)$$

For proof, see text p. 145

$$\begin{aligned}
 &= E(X^2) + 2E(X)E(Y) + E(Y^2) \\
 &\quad - E(X)^2 - \cancel{2E(Y)E(X)} - E(Y)^2
 \end{aligned}$$

$$\begin{aligned}
 &= \underbrace{E(X^2) - E(X)^2}_{\text{Var}(X)} + \underbrace{E(Y^2) - E(Y)^2}_{\text{Var}(Y)}
 \end{aligned}$$

$$= \text{Var}(X) + \text{Var}(Y) \quad \checkmark$$

which proves our contention that, if X and Y are independent random variables:

$$\text{Var}(X+Y) = \text{Var}(X) + \text{Var}(Y)$$

Variance: The Normal Distribution

X = a random variable

$E(X) = m$ ← expectation or mean. We'll assume finite expectation

↑ { Note that it is possible, by the way, for a random variable to have infinite expectation.
You can construct examples of this. }

The variance we define as the measure of the deviation from the expectation.
We square so that deviations on either side of the mean count the same.

$$\text{Var}(X) = E((X-m)^2)$$




$$= E(X^2 - 2mX + m^2) = E(X^2) - E(2mX) + E(m^2)$$

$$= E(X^2) - \underbrace{2mE(X)}_{2m^2} + \underbrace{E(m^2)}_{m^2}$$

$$= E(X^2) - m^2$$

$$= E(X^2) - E(X)^2$$

Suppose we have a continuous random variable, and we have density plots as follows:



Variance is important when you study random variables from the point of view of measurement.

From this perspective:

A random variable is an inaccurate measure of the mean

90% of statistics relates to this situation - inaccurate measure of the mean.

As we saw last time:

If X and Y are independent random variables,

then:

$$\text{Var}(X+Y) = \text{Var}(X) + \text{Var}(Y)$$

dependent variables
could have
equality by
accident

This is not necessarily true if X and Y are not independent

(Unlike where w/ expectations (means), where mean of the sum equals sum of the means - no questions asked.)

$$\begin{aligned} \text{Var}(cX) &= E((cX)^2) - E(cX)^2 \\ &\quad \uparrow c \text{ a constant} \\ &= c^2 E(X^2) - (c E(X))^2 \\ &= c^2 (E(X^2) - E(X)^2) \\ &\quad \underbrace{\hspace{1cm}}_{\text{Var}(X)} \end{aligned}$$

$$\boxed{\text{Var}(cX) = c^2 \text{Var}(X)}$$

That means that the variance is not dimensionally the same as the random variable. So we are lead to introduce another quantity that is dimensionally the same as the random variable.

$$\sigma(X) = \sqrt{\text{Var}(X)}$$

Standard Deviation

$$\begin{aligned} \sigma(cX) &= \sqrt{\text{Var}(cX)} \\ &= \sqrt{c^2 \text{Var}(X)} \\ &= |c| \sqrt{\text{Var}(X)} \end{aligned}$$

$$\boxed{\sigma(cX) = |c| \sigma(X)}$$

$$\text{Var}(cX) = c^2 \overset{\text{quadratic}}{\text{Var}(X)}$$


different dimensionality

$$\text{--vs--} \quad \sigma(cX) = |c| \sigma(X)$$

same dimensionality

$$\begin{aligned}
 \text{Var}(X+c) &= E((X+c)^2) - E(X+c)^2 \\
 &= E(X^2 + 2cX + c^2) - (E(X) + E(c))^2 \\
 &\quad \underbrace{(E(X) + E(c))^2}_{(E(X)+c)^2} \\
 &= E(X^2 + 2cX + c^2) - E(X)^2 - 2cE(X) + c^2 \\
 &= E(X^2) + 2cE(X) - E(X)^2 - 2cE(X) \\
 &= E(X^2) - E(X)^2
 \end{aligned}$$

$$\boxed{\text{Var}(X+c) = \text{Var}(X)}$$



Intuitively, this should be clear.

The variance is the deviation from the mean.

When you add a constant to X , you shift the mean, $\leftarrow E(X+c) = E(X)+c$
the deviation around the shifted mean does not change. $\leftarrow \text{Var}(X+c) = \text{Var}(X)$

This is a very important fact.

The variance of a random variable does not depend on a position.

It is an intrinsic number associated with the random variable that measures something about the density distribution, something about its shape.

Now let's compute a couple of variances, just to get a feeling.
Example - X is Bernoulli w/ probability distribution:

$$P(X=1) = p$$

$$P(X=0) = q = 1-p$$

$$\text{Var}(X) = E(X^2) - E(X)^2$$

$$\sum_n n P(X^2=n)$$

$$(X^2=n) = (X=n)$$

$$\sum_n n P(X=n)$$

$$P(X=0) \text{ when } n=0^2, (n=0).$$

$$P(X=1) \text{ when } n=1^2, (n=1).$$

$$0^2 \cdot q + 1^2 \cdot p$$

$$p$$

$$\left[\sum_n n P(X=n) \right]^2$$

$$(0 \cdot q + 1 \cdot p)^2$$


$$p^2$$

$$= p - p^2$$

$$= p(1-p)$$

$$\boxed{\text{Var}(X) = pq}$$

Observe that $pq = p - p^2$ is a quadratic equation w/ maximum at $p = \frac{1}{2}$, $1-p = \frac{1}{2} = q$:



$$\left. \begin{aligned} \frac{d}{dp}(p-p^2) &= 1-2p = 0 \\ \frac{d^2}{dp^2}(p-p^2) &= -2 < 0 \end{aligned} \right\} p-p^2 \text{ maximum at } p = \frac{1}{2}$$

Therefore $pq = p - p^2$ has a maximum value of $(\frac{1}{2})(\frac{1}{2}) = \frac{1}{4}$

$$\text{Var}(X) = pq \leq \frac{1}{4}$$

So the maximum variance of the Bernoulli random variable occurs when the bias of the coin $p = \frac{1}{2}$.

This upper bound is important when simplifying computations, as you will see shortly.

The standard deviation of the Bernoulli random variable is:

$$\sigma(X) = \sqrt{\text{Var}(X)}$$

$$= \sqrt{pq}$$

$$\boxed{\sigma(X) = \sqrt{pq} \leq \frac{1}{2}}$$

$$P(N_\alpha(t) = k) = \frac{\alpha^k t^k}{k!} e^{-\alpha t}$$

5/6/98.5.0

Example - Poisson Process

We have the random function $N_\alpha(t)$ of the Poisson process w/ intensity α .
Fix $\alpha + t$ so we have the random variable $N_\alpha(t)$.

We computed the expectation as [4/24/98.5]:

$$E(N_\alpha(t)) = \alpha t$$

Rewrite this as:

$$\begin{array}{c} \text{average}^* \\ \text{blips} \end{array} \rightarrow E(N_\alpha(t)) = \alpha \leftarrow \text{intensity}$$

$\xrightarrow{\text{per unit length}} t$

This is how we build the Poisson process.
 We want the average number of blips per unit length to be α .

That's how it comes out

A remarkable fact about the Poisson process is that the variance has the same value as the mean.

$$\text{Var}(N_\alpha(t)) = \alpha t$$

It's a gruesome computation,
 I want you to suffer.

Fact:

The Poisson Process is the only process where the mean equals the variance.

This is very remarkable.

It characterizes the Poisson Process in a very deep way.

$$\text{Var}(N_\alpha(t)) = \alpha t$$

5/6/98.5.1

- Proof Following outline by Feller, An Introduction to Probability Theory + Applications, Vol 1, 3rd Ed., Wiley, 1967 p. 228.

For convenience, set $X = N_\alpha(t)$
 $\swarrow$ $\alpha + t$ fixed.

$$\text{Var}(X) = \underbrace{E(X^2)}_{\substack{\text{we already know } E(N_\alpha(t)) = \alpha t. \\ E(X)^2 = (\alpha t)^2}} - \underbrace{E(X)^2}_{(\alpha t)^2}$$

We now focus on this:

$$E(X^2) = \sum_{n \geq 0} n^2 P(X^2 = n)$$

$$= 0 + \sum_{n \geq 1} n^2 P(X^2 = n)$$

$$= \sum_{n \geq 1} n^2 P(X = \sqrt{n})$$

$$\text{Let } k = \sqrt{n} \Rightarrow k^2 = n$$

$$= \sum_{k \geq 1} k^2 P(X = k)$$

$$P(X = k) = P(N_\alpha(t) = k) = \frac{\alpha^k t^k}{k!} e^{-\alpha t}$$

$$= \sum_{k \geq 1} k^2 \frac{\alpha^k t^k}{k!} e^{-\alpha t}$$

$$= e^{-\alpha t} \sum_{k \geq 1} k \frac{\alpha^k t^k}{(k-1)!}$$

$$= e^{-\alpha t} \sum_{k \geq 1} ((k-1) + 1) \frac{\alpha^k t^k}{(k-1)!}$$

$$= e^{-\alpha t} \sum_{k \geq 1} (k-1) \frac{\alpha^k t^k}{(k-1)!} + e^{-\alpha t} \sum_{k \geq 1} \frac{\alpha^k t^k}{(k-1)!}$$

$$= \cancel{e^{-\alpha t}} (\alpha t)^2 \underbrace{\sum_{k \geq 2} \frac{(\alpha t)^{k-2}}{(k-2)!}}_{\text{Maclaurin series for } e^{\alpha t}} + \cancel{e^{-\alpha t}} (\alpha t) \underbrace{\sum_{k \geq 1} \frac{(\alpha t)^{k-1}}{(k-1)!}}_{\text{Maclaurin series for } e^{\alpha t}}$$

So we have that:

$$\text{Var}(N_\alpha(t)) = (\alpha t)^2 + (\alpha t) - (\alpha t)^2$$

$$\text{Var}(N_\alpha(t)) = \alpha t \quad \checkmark$$

Q.E.D. Again, it is a remarkable fact that $E(N_\alpha(t)) = \text{Var}(N_\alpha(t)) = \alpha t$

Example - Exponential Random Variable

T exponential, i.e.,

$$P(T > t) = e^{-\alpha t}$$

T is the memoryless waiting time for the first blip in the Poisson Process.

We already computed the expectation [4/22/98.7-8]:

$$E(T) = \frac{1}{\alpha} \quad \left\{ \begin{array}{l} \text{The greater the intensity of the Poisson Process,} \\ \text{the smaller the waiting time for the first blip.} \end{array} \right.$$

Now let's compute the variance:

$$\text{Var}(T) = \underbrace{E(T^2)} - \underbrace{E(T)^2}_{\frac{1}{\alpha^2}}$$

$$\left\{ \begin{array}{l} E(T^2) = \int_{-\infty}^{\infty} s \, \text{dens}(T^2=s) \, ds \\ (T^2=s) = (T=\sqrt{s}) \\ \text{Let } t=\sqrt{s} \Rightarrow t^2=s \\ = \int_0^{\infty} t^2 \, \text{dens}(T=t) \, dt \end{array} \right.$$

$$= \int_0^{\infty} t^2 \, \text{dens}(T=t) \, dt - \frac{1}{\alpha^2}$$

$$\left\{ \begin{array}{l} \text{We've already shown [4/22/98.6]:} \\ \text{dens}(T=t) = \alpha e^{-\alpha t} \end{array} \right.$$



$$= \underbrace{\int_0^{\infty} t^2 \alpha e^{-\alpha t} dt}_{\text{we need to evaluate this integral. We can obtain this in a clever way. Recall the integral for the expectation of } T:}$$

We need to evaluate this integral. We can obtain this in a clever way. Recall the integral for the expectation of T :

$$\begin{aligned} E(T) &= \int_{-\infty}^{\infty} t \underbrace{\text{dens}(T=t)}_{\downarrow} dt \\ &= \int_0^{\infty} t \alpha e^{-\alpha t} dt = \frac{1}{\alpha} \end{aligned}$$

From this, multiplying both sides by $\frac{1}{\alpha}$, we obtain:

$$\int_0^{\infty} t e^{-\alpha t} dt = \frac{1}{\alpha^2}$$

Differentiate both sides with respect to α :

$$\frac{\partial}{\partial \alpha} \int_0^{\infty} t e^{-\alpha t} dt = \frac{\partial}{\partial \alpha} \left(\frac{1}{\alpha^2} \right)$$

$$\int_0^{\infty} t \frac{\partial}{\partial \alpha} e^{-\alpha t} dt = -2 \left(\frac{1}{\alpha^3} \right)$$

$$-\int_0^{\infty} t^2 e^{-\alpha t} dt = -\frac{2}{\alpha^3}$$

Multiply both sides by $(-\alpha)$:

$$\int_0^{\infty} t^2 \alpha e^{-\alpha t} dt = \frac{2}{\alpha^2} \quad \left. \vphantom{\int_0^{\infty} t^2 \alpha e^{-\alpha t} dt} \right\} \begin{array}{l} \text{This is precisely the integral we} \\ \text{were after.} \\ \text{Pretty Neat.} \end{array}$$

$$= \frac{2}{\alpha^2} - \frac{1}{\alpha^2}$$

$$\boxed{\text{Var}(T) = \frac{1}{\alpha^2}}$$

Suppose we have random variable X with a given expectation:

$$E(X) = m$$

Let's define a new random variable:

$$Y = aX + b$$

↖ ↗
 a, b are constants

What are the mean and variance of Y ?

$$\begin{aligned} E(Y) &= E(aX + b) \\ &= E(aX) + E(b) \end{aligned}$$

$$E(Y) = aE(X) + b$$

$$\begin{aligned} \text{Var}(Y) &= \text{Var}(aX + b) \\ &= \text{Var}(aX) \quad \leftarrow \text{since } \text{Var}(Z + c) = \text{Var}(Z) \\ &= a^2 \text{Var}(X) \quad \leftarrow \text{since } \text{Var}(cZ) = c^2 \text{Var}(Z) \end{aligned}$$

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

Therefore, if $\begin{cases} E(X) = m \\ \text{Var}(X) = \sigma^2 \end{cases}$ then the random variable

$$\frac{X - m}{\sigma} \text{ has mean 0 and variance 1.}$$

Such a random variable is said to be normalized.

{the term standardized
is also used.}

$$\begin{aligned} \text{Proof } E\left(\frac{X - m}{\sigma}\right) &= E\left(\frac{1}{\sigma}X - \frac{m}{\sigma}\right) \\ &= E\left(\frac{1}{\sigma}X\right) - E\left(\frac{m}{\sigma}\right) \\ &= \frac{1}{\sigma}E(X) - \frac{m}{\sigma} \\ &= 0 \quad \checkmark \end{aligned}$$

$$\begin{aligned}
 \text{Var}\left(\frac{X-m}{\sigma}\right) &= \text{Var}\left(\frac{1}{\sigma}X - \frac{m}{\sigma}\right) \\
 &= \text{Var}\left(\frac{1}{\sigma}X\right) \\
 &= \frac{1}{\sigma^2} \text{Var}(X) \xrightarrow{\sigma^2} \sigma^2 \\
 &= 1 \quad \checkmark
 \end{aligned}$$

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



• Another application of the same Mickey Mouse principle is the following:

(*) If Y is a normalized random variable, i.e., $E(Y)=0$ and $\text{Var}(Y)=1$, then the random variable $X = \sigma Y + m$ has:

$$\begin{aligned}
 E(X) &= m \quad \text{and} \\
 \text{Var}(X) &= \sigma^2
 \end{aligned}$$

Proof

$$\begin{aligned}
 E(X) &= E(\sigma Y + m) \\
 &= E(\sigma Y) + E(m) \\
 &= \sigma E(Y) \xrightarrow{0} 0 + m \\
 &= m \quad \checkmark
 \end{aligned}$$

$$\begin{aligned}
 \text{Var}(X) &= \text{Var}(\sigma Y + m) \\
 &= \text{Var}(\sigma Y) \\
 &= \sigma^2 \text{Var}(Y) \xrightarrow{1} \sigma^2 \\
 &= \sigma^2 \quad \checkmark
 \end{aligned}$$

i.e., Z has a standard normal density distribution

• Recall a random variable Z is standard normal whenever the density is:

$$\text{dens}(Z=t) = \frac{1}{\sqrt{2\pi}} e^{-\frac{t^2}{2}} \quad \left\{ \text{See [5/4/98.4-5]} \right\}$$

- Let's perform the preceding computations with standard normal random variables.

By the immediately discussion (*), it follows that:

Let Z = standard normal variable

$$\left. \begin{array}{l} E(Z) = 0 \\ \text{Var}(Z) = 1 \end{array} \right\} \text{Properties of normalized random variables.}$$

Let's check that Z satisfies these properties:

$$E(Z) = \int_{-\infty}^{\infty} z \text{dens}(Z=z) dz$$

$$= \int_{-\infty}^{\infty} z \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} dz$$

$$\text{Let } u = -\frac{z^2}{2}$$

$$du = -z dz$$

$$= \frac{-1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^u du$$

$$= \frac{-1}{\sqrt{2\pi}} e^u \Big|_{-\infty}^{\infty}$$

$$= 0 - 0$$

$$E(Z) = 0 \quad \checkmark$$

Standard normal variable has mean 0.

$$\text{Var}(Z) = E(Z^2) - E(Z)^2$$

$$= \int_{-\infty}^{\infty} s \text{dens}(Z^2=s) ds$$

$$\text{Let } z = \sqrt{s} \Rightarrow z^2 = s$$

$$= \int_{-\infty}^{\infty} z^2 \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} dz$$

$$= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} z^2 e^{-\frac{z^2}{2}} dz$$

← {we obtain this integral by differentiating a known integral.}

Digression: What is $\int_{-\infty}^{\infty} z^2 e^{-\frac{z^2}{2}} dz$?

5/6/98.11

To start with, we have shown that:

$$E(Z) = \int_{-\infty}^{\infty} z \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} dz = 0$$

Thus, we have the following equality:

$$\int_{-\infty}^{\infty} z e^{-\frac{z^2}{2}} dz = 0$$

Differentiating both sides, we get:

$$\frac{d}{dz} \int_{-\infty}^{\infty} z e^{-\frac{z^2}{2}} dz = \frac{d}{dz} 0$$

$$\int_{-\infty}^{\infty} \left(\frac{d}{dz} z e^{-\frac{z^2}{2}} \right) dz = 0$$

$$\int_{-\infty}^{\infty} \left[z \left(-z e^{-\frac{z^2}{2}} \right) + e^{-\frac{z^2}{2}} \right] dz = 0$$

$$\underbrace{\int_{-\infty}^{\infty} z^2 e^{-\frac{z^2}{2}} dz}_{\text{this is the integral we are trying to obtain.}} = \underbrace{\int_{-\infty}^{\infty} e^{-\frac{z^2}{2}} dz}$$

this is the integral we are trying to obtain.

Recall that $\text{dens}(Z=z) = \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}}$
And since we know that, by the definition of density:

$$\int_{-\infty}^{\infty} \text{dens}(Z=z) dz = 1$$

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} dz = 1$$

Which readily provides that:

$$\int_{-\infty}^{\infty} e^{-\frac{z^2}{2}} dz = \sqrt{2\pi}$$

And thus, we have that:

$$\int_{-\infty}^{\infty} z^2 e^{-\frac{z^2}{2}} dz = \sqrt{2\pi}$$

Pretty slick!

$$\begin{aligned} \text{So, } \text{Var}(Z) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} z^2 e^{-\frac{z^2}{2}} dz \\ &= \frac{1}{\sqrt{2\pi}} \sqrt{2\pi} \end{aligned}$$

$$\text{Var}(Z) = 1 \quad \checkmark$$

So it is indeed the fact that:

$$Z = \text{a standard normal random variable} \Rightarrow \begin{cases} E(Z) = 0 \\ \text{Var}(Z) = 1 \end{cases}$$

properties of
normalised random
variables.

Further, from (*), we also know that:

random variable $X = \sigma Z + m$ has

$$E(X) = m$$

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

We want to go a bit farther than that.

We want the density distribution of $\sigma Z + m$.

$$\text{dens}(X=t) = ?$$

$$\uparrow X = \sigma Z + m$$

We spent a whole week on the Algebra of Probability Distributions.

Let's apply it now.

Our chickens are coming to roost.

You say: Well, we know how to do that.

Let's do it.

$$P(X \leq t) = P(\sigma Z + m \leq t)$$

$$\text{event } (\sigma Z + m \leq t) = (Z \leq \frac{t-m}{\sigma})$$

observe that σ is always positive
 $\sigma > 0$

$$= P(Z \leq \frac{t-m}{\sigma})$$

$$= \int_{-\infty}^{\frac{t-m}{\sigma}} \underbrace{\text{dens}(Z=s)}_{\frac{1}{\sqrt{2\pi}} e^{-\frac{s^2}{2}}} ds$$

The cumulative probability distribution is
the integral of the density
distribution.

$$P(X \leq t) = \int_{-\infty}^{\frac{t-m}{\sigma}} \frac{1}{\sqrt{2\pi}} e^{-s^2/2} ds$$

Hence:

$$\begin{aligned} \text{dens}(X=t) &= \frac{d}{dt} P(X \leq t) \\ &= \frac{d}{dt} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{t-m}{\sigma}} e^{-s^2/2} ds \\ &= \frac{1}{\sqrt{2\pi}} \frac{d}{dt} \int_{-\infty}^{\frac{t-m}{\sigma}} e^{-s^2/2} ds \end{aligned}$$

$$\left\{ \begin{array}{l} \text{Recall the 2nd Fundamental Theorem of Calculus:} \\ \frac{d}{du} \left(\int_a^u f(s) ds \right) = f(u) \end{array} \right\}$$

We evaluate the above by the chain rule.

$$\text{Let } u = \frac{t-m}{\sigma}$$

$$g(u) = \int_{-\infty}^u e^{-s^2/2} ds$$

Then:

$$\frac{d}{dt} g\left(\frac{t-m}{\sigma}\right) = \frac{d}{dt} g(u) = \frac{dg(u)}{du} \frac{du}{dt}$$

$$\begin{aligned} \frac{d}{dt} \int_{-\infty}^{\frac{t-m}{\sigma}} e^{-s^2/2} ds &= \left(\frac{d}{du} \int_{-\infty}^u e^{-s^2/2} ds \right) \cdot \frac{1}{\sigma} \\ &= e^{-u^2/2} \leftarrow \\ &= e^{-\left(\frac{t-m}{\sigma}\right)^2/2} \cdot \frac{1}{\sigma} \end{aligned}$$

$$\boxed{\text{dens}(X=t) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(t-m)^2/2\sigma^2}}$$

As this is a density, automatically it integrates to 1. We don't have to check that.

$$\int_{-\infty}^{\infty} \text{dens}(X=t) dt = \int_{-\infty}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} e^{-(t-m)^2/2\sigma^2} dt = 1$$

Because it's the density of a random variable, the integral has to equal 1. It's stupid to check that.

5/6/98.14

$\text{dens}(X=t) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(t-m)^2/2\sigma^2}$ is called the Normal Distribution of mean m and variance σ^2

Conclusion

If random variable Z is standard normal, then $\sigma \overbrace{Z}^{\text{random variable}} + m$ has the density:

$$\text{dens}(\sigma Z + m = t) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(t-m)^2/2\sigma^2}$$

In statistics books, this is abbreviated as saying:

$$\sigma Z + m \text{ is } \mathcal{N}(m, \sigma^2)$$

← normally distributed w/ mean m and variance σ^2

And, of course, if you have a random variable X , where:

$$X \text{ is } \mathcal{N}(m, \sigma^2)$$

then:

$$\frac{X-m}{\sigma} \text{ is standard normal}$$

$$X = \sigma Z + m$$

↓

$$E(X) = m$$

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

$$\mathcal{N}(m, \sigma^2)$$

↓

$$E(Z) = 0$$

$$\text{Var}(Z) = 1$$

$$\mathcal{N}(0, 1)$$

If $X_1, X_2, \dots, X_n$ are i.i.d. of finite variance with

$$E(X_i) = m$$

$$\text{Var}(X_i) = \sigma^2 < \infty$$

since X 's are independent:

$$\text{then } \text{Var}\left(\underbrace{\frac{X_1 + X_2 + \dots + X_n}{n}}_{\text{average}}\right) = \text{Var}\left(\frac{1}{n}X_1\right) + \text{Var}\left(\frac{1}{n}X_2\right) + \dots + \text{Var}\left(\frac{1}{n}X_n\right)$$

$$= \frac{1}{n^2} \text{Var}(X_1) + \frac{1}{n^2} \text{Var}(X_2) + \dots + \frac{1}{n^2} \text{Var}(X_n)$$

$$\text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \frac{\sigma^2}{n}$$

$$\text{Var}(cX) = c^2 \text{Var}(X)$$

What does that mean?
It's a very important result.

It means that if you take n measurements of the same quantity and average them, then the variance of the average:

$$\text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \frac{\sigma^2}{n}$$

↑ goes down as n increases.

and the standard deviation of the average:

$$\sigma\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \frac{\sigma}{\sqrt{n}}$$

↑ goes down as n increases.

In other words, it confirms our intuitive feeling that taking several measurements of the same quantity makes the measurement of the expectation more accurate.

Namely, the variance of the expectation goes down.

$$\hookrightarrow \text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \frac{\sigma^2}{n} \leftarrow$$

Provided that the variance is finite.

Statistics in One Easy Lesson

90% of statistics is studying independent random variables, with finite variance.

The main assumption of statistics is this:

When you make repeated measurements, you might as well assume that the random variables you are measuring are normally distributed.

When you make several measurements of the same quantity, close your eyes and assume that the random variables are normally distributed.

In other words, suppose you have random variables:

$$X_1, X_2, \dots, X_n \quad \leftarrow \text{independent random variables}$$

You take the mean m , and the variance σ^2
and you forget everything else.

Then, you replace the X 's by:

$$X_i = \sigma Z + m$$

↑ standard normal random variable

That's the main assumption of statistics.

In making measurements of a quantity X with $E(X) = m$
 $\text{Var}(X) = \sigma^2 < \infty$

assume that X is $\mathcal{N}(m, \sigma^2)$

Assume that X is normally distributed
 w/ mean m and variance σ^2 .

NB: This assumption has never been fully justified.

Engineers, etc. say who cares? It works.

There are partial justifications, as we will see, by the:

1) Central Limit Theorem

2) Various derivations of the Normal Law,
 axiomatic principles.

None of them provides a complete justification.

To this day, there is a question how to completely, rigorously justify this assumption.

We will assume it.

Why do we make this assumption?

How does it work?

Suppose we make n measurements of a quantity
 $X_1, X_2, \dots, X_n \leftarrow \left\{ \begin{array}{l} \text{We assume } X_1, X_2, \dots, X_n \text{ are independent} \\ \text{random variables.} \\ \text{And that each is } \mathcal{N}(m, \sigma^2) \end{array} \right\}$

Suppose the variance is known and is finite:

$\text{Var}(X_i) = \sigma^2 < \infty \leftarrow \left\{ \begin{array}{l} \text{if it's not known, we can limit it by an} \\ \text{upper bound, like we did w/ the Bernoulli} \\ \text{random variable [5/6/98.4]} \end{array} \right\}$

You want to know what the mean of these random variables is.

Want to estimate $m = E(X_1) = E(X_2) = \dots = E(X_n)$

You don't know the mean.

What do we do?

Use Bayes' Law.

$$P(B|A) = \frac{P(A|B)P(B)}{P(A)}$$

5/6/98.17

We have also a version of Bayes' Law for densities [4/10/98.10]:

$$\underset{\text{posterior}}{\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n)} = \underset{\text{likelihood}}{\text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n | m=s)} \underset{\text{prior}}{\text{dens}(m=s)}$$

How does Bayes' Law for densities help us?

Taking n measurements of the same quantity means:

event $X_1 = t_1 \leftarrow$ measurement 1 gives t_1
 $X_2 = t_2 \leftarrow$ measurement 2 gives t_2
 $\vdots$
 $X_n = t_n \leftarrow$ measurement n gives t_n

↑ the X 's are assumed normal, with mean m and variance σ^2 .

What does Bayes' Law do for you?

Bayes' Law says:

You don't know the mean m , so the mean is a random variable.The density of the mean, given these n measurements

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n)$$

will be equal to the density of these n measurements, given the mean, times the prior times a constant.

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) \propto \text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n | m=s) \text{ (Prior)}$$

{ So, if we establish a prior - namely, some preliminary guess of what the mean we are measuring is going to be like, then the likelihood is computed very simply. }

Likelihood

$$\text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n | m=s)$$

these random variables are independent. So their densities multiply.

{ means that the mean m is fixed to s from the start. }

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS

$$\underset{\substack{\uparrow \\ \text{likelihood}}}{\text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n | m=s)} = \text{dens}(X_1=t_1 | m=s) \text{dens}(X_2=t_2 | m=s) \dots \text{dens}(X_n=t_n | m=s)$$

Consider $\text{dens}(X_i=t_i | m=s)$.

It is given that mean $m=s$.

As discussed, each of the X_i is assumed to be normal with variance σ^2 and we are given mean $m=s$.

So each X_i is $\mathcal{N}(s, \sigma^2)$.

Thus:

$$\begin{aligned} \text{dens}(X_i=t_i | m=s) &= \frac{1}{\sigma \sqrt{2\pi}} e^{-(t_i-s)^2/2\sigma^2} \\ &= \left(\frac{1}{\sigma \sqrt{2\pi}} e^{-(t_1-s)^2/2\sigma^2} \right) \left(\frac{1}{\sigma \sqrt{2\pi}} e^{-(t_2-s)^2/2\sigma^2} \right) \dots \left(\frac{1}{\sigma \sqrt{2\pi}} e^{-(t_n-s)^2/2\sigma^2} \right) \end{aligned}$$

So the likelihood comes out real easy, because they are independent random variables.
So you plug that in to Bayes' Law for densities.

- Now comes the 2nd big assumption of statistics. $\leftarrow \left\{ \begin{array}{l} \text{We won't finish this today, so we'll} \\ \text{just set it up.} \end{array} \right\}$
Unjustified, except that it works.
I wish one of you would work this out. Make me happy.
Justify it.
Write me a letter. Send me a reprint.

With the above as likelihood, Bayes' Law for densities becomes:

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) \propto \frac{1}{\sigma^n (2\pi)^n} e^{-((t_1-s)^2 + (t_2-s)^2 + \dots + (t_n-s)^2)/2\sigma^2} \overbrace{\text{dens}(m=s)}^{\text{the prior}}$$

This looks like a mess.

But it isn't.

We've said 100 times that every density has to integrate to 1.

Now, we're going to see the real culmination of this fact.

Let's focus on the exponent of e :

$$\begin{aligned} (t_1-s)^2 + (t_2-s)^2 + \dots + (t_n-s)^2 &= ns^2 - 2s(t_1+t_2+\dots+t_n) + (t_1^2+t_2^2+\dots+t_n^2) \\ \left\{ (t_i-s)^2 = t_i^2 - 2st_i + s^2 \right\} &\quad \text{multiply by } 1 = \frac{n}{n} \end{aligned}$$

$$= ns^2 - n2s \left(\frac{t_1+t_2+\dots+t_n}{n} \right) + (t_1^2+t_2^2+\dots+t_n^2)$$

$$\uparrow \frac{t_1+t_2+\dots+t_n}{n} = \bar{m} \quad \text{The sample mean}$$

5/6/98.19

We take n measurements, and we average those n measurements.

$$\frac{t_1 + t_2 + \dots + t_n}{n} = \text{the sample mean } \bar{m}$$

The average of the actual numbers we observed.
They are just numbers.
We averaged them.

$$= ns^2 - n2s\bar{m} + (t_1^2 + t_2^2 + \dots + t_n^2)$$

How do we get rid of these t_i ?

$$= \underbrace{n(s - \bar{m})^2}_{n s^2 - n 2 s \bar{m} + n \bar{m}^2} + \text{crud}$$

$$n s^2 - n 2 s \bar{m} + n \bar{m}^2$$

how do we get rid of the crud?
I'll tell you next time.

If you look carefully, you'll realize that it can be removed wonderfully, from the fact that densities must integrate to 1.
We'll see this in detail.

After we do this, we have the statistics we want.
And you can go to town w/ statistics.




The Normal Distribution (cont'd)

Let's continue our study of the Normal Distribution.
Last time, we noted that [5/4/98.4-5]:

if Z = a standard normal random variable,
then:

$$\text{dens}(Z=t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$$



this density is the bell-shaped curve

Let me recite a few facts about this density.
In the old days, people carried around tables.
Today, we do these computations on computer.

$$P(-a < Z \leq b) = P(Z \leq b) - \underbrace{P(Z \leq -a)}$$

since the density is symmetric about $x=0$,

$$P(Z \leq -a) = P(Z \leq \infty) - P(Z \leq a) \\ = 1 - P(Z \leq a)$$

$$= \int_{-\infty}^b \text{dens}(Z=t) dt - 1 + \int_{-\infty}^a \text{dens}(Z=t) dt$$

$$P(-a < Z \leq b) = \int_{-\infty}^b \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt + \int_{-\infty}^a \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt - 1$$

Example:

$$P(-1 < Z \leq 1) = \underbrace{\int_{-\infty}^1 \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt}_{0.8413} + \underbrace{\int_{-\infty}^{-1} \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt}_{0.8413} - 1 \\ = 0.6826$$

These integrals are readily available.
(for example, text p. 177)

$$P(-1 < Z \leq 1) = 0.6826$$

$$P(-1.96 < Z \leq 1.96) = 0.95$$

$$P(-2 < Z \leq 2) = 0.9545$$

$$P(-2.58 < Z \leq 2.58) = 0.99$$

← significant

← very significant

These numbers have names
in statistics.
These are technical terms.

When you say standard normal random variable Z is significant, it means:

$$P(-1.96 < Z \leq 1.96) = 0.95$$

By common agreement.

- We also saw last time:

The random variable $\sigma Z + m$ has the density:

$$\text{dens}(\sigma Z + m = t) = \frac{1}{\sigma \sqrt{2\pi}} e^{-(t-m)^2/2\sigma^2}$$

Random variable $\sigma Z + m$ has mean m and standard deviation σ (variance = σ^2)

22-141 50 SHEETS
22-142 100 SHEETS
22-144 200 SHEETS



- Basic Rule of Statistics

70% of all statistical computations use the following:

Given a "tough" random variable X with $E(X) = m$ and $\text{Var}(X) = \sigma^2$,
replace it by $\sigma Z + m$

In other words, replace X by the simplest possible random variable
that has the same expectation and the same variance.
Close your eyes and hope for the best.
And it works.

I told you last time, nobody really knows, ultimately, why this works.
Rivers of words have been written - by philosophers, physicists, mathematicians, crackpots.
The great French mathematician Henri Poincaré once wrote that:

"The explanation of the normal distribution is something that mathematicians
consider to be the physicists' task. And physicists consider to be
the mathematicians' task."

So the ball bounces back and forth.

- Last time, we began to apply this basic rule, in a problem of estimation of measurement.

We started w/ Bayes' Law for densities:

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) = \frac{\text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n | m=s) \text{dens}(m=s)}{\text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n)}$$

↑ posterior
↑ likelihood
↑ prior

dens($m=s$ given n independent measurements)

This is what a lot of people make their living on.
In very sophisticated ways, to be sure.

Last time, we said "Let's take Bayes' Law and apply the Basic Rule of Statistics"
Let's take n measurements and apply the Basic Rule of Statistics.

When we apply the Basic Rule of Statistics, we will assume that:

$X_1, X_2, \dots, X_n$ are independent normally distributed random variables, whose mean and variance we do not know.

We assume the variance is finite ^{$\sigma^2 < \infty$} and, therefore, we take an upper bound for the variance that is safe.

And we plug that in.

The problem is to estimate the mean of the random variable on the basis of the data given by n measurements.

$$E\left(\frac{X_1 + \dots + X_n}{n}\right) = ?$$

- We saw last time [5/6/98.14-15] that if we took the average of n independent, identically distributed random variables, then the variance goes down by $\frac{1}{n}$.

$$\text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \frac{\sigma^2}{n}$$

So, you get a much more accurate estimate of the average with larger n , provided that the variance is finite.

- Now, let's continue, from last time, our evaluation of Bayes' Law for densities [5/6/98.18]:

If $X_1, X_2, \dots, X_n$ are i.i.d., normally distributed, then:

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) \propto$$

$$\frac{1}{\sigma^n (2\pi)^n} e^{-((t_1-s)^2 + (t_2-s)^2 + \dots + (t_n-s)^2)/2\sigma^2} \text{dens}(m=s)$$

for fixed n , this is
a constant K_1 .

$$\propto e^{-((t_1-s)^2 + (t_2-s)^2 + \dots + (t_n-s)^2)/2\sigma^2} \text{dens}(m=s)$$

Something cute comes up when simplifying this expression.

Call $\bar{m} = \frac{t_1 + t_2 + \dots + t_n}{n}$
 $\uparrow$ the sample mean

In other words, $t_1, t_2, \dots, t_n$ are the actual measurements we obtained in our experiment. Numbers. No nonsense.

We take the average of these numbers,

This should give an approximation to what we want to compute. Let's see how.

Let's focus on the exponent:

$$(t_1 - s)^2 + (t_2 - s)^2 + \dots + (t_n - s)^2 = ns^2 - \underbrace{2s(t_1 + t_2 + \dots + t_n)}_{\text{multiply by } 1 = \frac{n}{n}} + (t_1^2 + t_2^2 + \dots + t_n^2)$$

$$= ns^2 - n2s\left(\frac{t_1 + t_2 + \dots + t_n}{n}\right) + (t_1^2 + t_2^2 + \dots + t_n^2)$$

$$= ns^2 - n2s\bar{m} + (t_1^2 + t_2^2 + \dots + t_n^2)$$

$$\text{Add } 0 = n\bar{m} - n\bar{m}$$

$$= \underbrace{ns^2 - n2s\bar{m} + n\bar{m}}_{n(\bar{m} - s)^2} - \underbrace{n\bar{m} + (t_1^2 + t_2^2 + \dots + t_n^2)}_{\text{cancel Contains no } s. \text{ Call this constant } K_2.}$$

$$(t_1 - s)^2 + (t_2 - s)^2 + \dots + (t_n - s)^2 = n(\bar{m} - s)^2 + K_2$$

So:

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) \propto \underbrace{e^{-(n(\bar{m}-s)^2 + K_2)/2\sigma^2}}_{\substack{\text{this is just} \\ \text{a constant } K_3}} \text{dens}(m=s)$$

$$\underbrace{e^{-K_2/2\sigma^2}}_{\substack{\text{this is just} \\ \text{a constant } K_3}} e^{-n(\bar{m}-s)^2/2\sigma^2}$$

$$\propto K_3 e^{-n(\bar{m}-s)^2/2\sigma^2} \text{dens}(m=s)$$

$$\propto e^{-n(\bar{m}-s)^2/2\sigma^2} \text{dens}(m=s)$$

Let k = constant that incorporates constants K_1, K_2, K_3 .

5/8/98.5

Then Bayes' Law for densities is:

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) = \frac{\text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n | m=s) \text{dens}(m=s)}{\text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n)}$$

Rewrite the denominator using the continuous Law of Alternatives consequence [4/8/98.7]:

$$= \frac{k e^{-(\bar{m}-s)^2 / 2\sigma_n^2} \text{dens}(m=s)}{\int_{-\infty}^{\infty} \text{dens}(X_1=t_1, X_2=t_2, \dots, X_n=t_n | m=s) \text{dens}(m=s) ds}$$

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) = \frac{k e^{-(\bar{m}-s)^2 / 2\sigma_n^2} \cdot \text{prior} \text{dens}(m=s)}{\int_{-\infty}^{\infty} k e^{-(\bar{m}-s)^2 / 2\sigma_n^2} \text{dens}(m=s) ds}$$

Now you say "What shall we choose as prior?"
There is an honest choice and a dishonest choice.

1) The honest choice


Choose a normal distribution w/ a certain variance, which you guess.

That can be done.

It's a hairy computation, but it can be done.

2) The dishonest choice

Choose as prior $\text{dens}(m=s) = 1$ for all s .



Note: This is not a probability density.

$$\int_{-\infty}^{\infty} \text{dens}(m=s) ds \neq 1$$

Who cares? It works anyway.

That's what they do

This is the 2nd big assumption in statistics.

No one has ever justified this.

I tell you, because maybe you will.

It's good for you to know the truth.

Even though this irritates statisticians,

Don't ever tell them you choose as prior a measure that is not a density.

They choose the uniform measure of the line.

Which means you pick a point on the line at random, but it's not a probability.

However, they say, if you choose $\text{dens}(m=s)=1$, then equation (*) is still convergent, because the RHS has a negative exponential.

So it still makes sense.


Let's see what happens.

Close your eyes and make the 2nd big assumption:

Choose as prior, not a probability density, but a measure:

$$\text{dens}(m=s)=1, \text{ for all } s$$

If really pushed against the wall, they will say:



Replace this by a normal distribution w/ very high variance and that is almost uniform.

So why not go to the limit (even though the limit doesn't make sense)?

Then, equation (*) becomes:

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) = K e^{-(\bar{m}-s)^2/2\sigma_n^2}$$

↑ incorporates constant K and integral in the denominator w/ $\text{dens}(m=s)=1$.

Now comes the beauty of this.

We can immediately solve for constant K' , even though it comes from a mess.

Why?

Because the above is a density distribution, therefore the whole thing has to integrate to 1.

$$\int_{-\infty}^{\infty} \text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) ds = \int_{-\infty}^{\infty} K' e^{-(\bar{m}-s)^2/2\sigma_n^2} ds = 1$$

$$\Rightarrow K' = \frac{1}{\int_{-\infty}^{\infty} e^{-(\bar{m}-s)^2/2\sigma_n^2} ds}$$

So we know K' .

5/8/98.7

What's this telling us?

It's telling us that, with a uniform prior (i.e., $\text{dens}(m=s) = 1$, for all s), then the posterior is normally distributed with mean $\bar{m}$ (the sample mean) and variance $\frac{\sigma^2}{n}$

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) = k' e^{-\frac{(\bar{m}-s)^2}{2\frac{\sigma^2}{n}}}$$

posterior

Posterior normally distributed
with mean $\bar{m}$ and variance $\frac{\sigma^2}{n}$

And that's an interesting conclusion, even though obtained not rigorously.

Conclusion:

If you take n measurements,
and we take a prior, which in some intuitive sense means no knowledge whatsoever
then the posterior tells us:

the mean is a random variable that is normally distributed w/ variance $\frac{\sigma^2}{n}$

This really confirms our previous computation of the variance of the mean of n independent random variables [5/6/98.14-15]:

$$\text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \frac{\sigma^2}{n}$$

The formula:

$$\text{dens}(m=s | X_1=t_1, X_2=t_2, \dots, X_n=t_n) = k' e^{-\frac{(\bar{m}-s)^2}{2\frac{\sigma^2}{n}}}$$

is the basic formula of statistics.

In other words, if you have n measurements, and you want to make a prediction about their average, this is the formula you use.

i.e., m is $\mathcal{N}(\bar{m}, \frac{\sigma^2}{n})$

How do you use the formula to make a prediction about the average?
Like this.

Since m is normally distributed with $E(m) = \bar{m}$
 $\text{Var}(m) = \frac{\sigma^2}{n}$

it follows that:

$$Z = \frac{m - \bar{m}}{\frac{\sigma}{\sqrt{n}}} \text{ is standard normal.}$$

see [5/6/98.14]

$$E(Z) = E\left(\frac{m - \bar{m}}{\frac{\sigma}{\sqrt{n}}}\right)$$

$$= E\left(\frac{m}{\frac{\sigma}{\sqrt{n}}}\right) - E\left(\frac{\bar{m}}{\frac{\sigma}{\sqrt{n}}}\right)$$

$$= \frac{1}{\frac{\sigma}{\sqrt{n}}} E(m) - \frac{\bar{m}}{\frac{\sigma}{\sqrt{n}}}$$

given that $E(m) = \bar{m}$

$$= \frac{1}{\frac{\sigma}{\sqrt{n}}} \bar{m} - \frac{\bar{m}}{\frac{\sigma}{\sqrt{n}}}$$

$$E(Z) = 0 \checkmark$$

$$\text{Var}(Z) = \text{Var}\left(\frac{m - \bar{m}}{\frac{\sigma}{\sqrt{n}}}\right)$$

$$= \text{Var}\left(\frac{m}{\frac{\sigma}{\sqrt{n}}} - \frac{\bar{m}}{\frac{\sigma}{\sqrt{n}}}\right)$$

$$= \text{Var}\left(\frac{m}{\frac{\sigma}{\sqrt{n}}}\right)$$

$$= \frac{1}{\frac{\sigma^2}{n}} \text{Var}(m)$$

given that $\text{Var}(m) = \frac{\sigma^2}{n}$

$$= \frac{1}{\frac{\sigma^2}{n}} \cdot \frac{\sigma^2}{n}$$

$$\text{Var}(Z) = 1 \checkmark$$

$\therefore Z = \frac{m - \bar{m}}{\frac{\sigma}{\sqrt{n}}}$ is standard normal.

Hence:

$$P(a < m \leq b) = \int_a^b \underbrace{\text{dens}(m=s)}_{\text{where } m \text{ is } \mathcal{N}(\bar{m}, \frac{\sigma^2}{n})} ds$$

where m is $\mathcal{N}(\bar{m}, \frac{\sigma^2}{n})$, we have (see [5/6/98.13]):

$$= \int_a^b \frac{1}{\frac{\sigma}{\sqrt{n}} \sqrt{2\pi}} e^{-(s - \bar{m})^2 / 2 \frac{\sigma^2}{n}} ds$$

Better is to write the probability in terms of the standard normal random variable. It's easier to use this, because the tables available provide the integral of the standard normal density distribution.

$$P(a < Z \leq b)$$

$$\underbrace{P(a < \frac{m - \bar{m}}{\frac{\sigma}{\sqrt{n}}} \leq b)}_{\text{where } Z \text{ is } \mathcal{N}(0, 1)} = \int_a^b \underbrace{\text{dens}(Z=s)}_{\text{where } Z \text{ is } \mathcal{N}(0, 1)} ds$$

$$= \int_a^b \frac{1}{\sqrt{2\pi}} e^{-s^2/2} ds$$

$$(*) \quad \boxed{P(a < \frac{m - \bar{m}}{\frac{\sigma}{\sqrt{n}}} \leq b) = \frac{1}{\sqrt{2\pi}} \int_a^b e^{-s^2/2} ds}$$

← this is the integral you find in tables. So it's always better to normalize.

Example

I toss a coin 100 times and obtain 41 heads.
Is the coin fair?

We assume $m = \mathcal{N}(\bar{m}, \frac{\sigma^2}{n})$

$$n = 100$$

The sample mean $\bar{m} = \frac{41}{100}$

We don't know σ .

But we do know that if X is a Bernoulli random variable,

$$\text{Var}(X) = \sigma^2 = pg \leq \frac{1}{4} \quad [5/6/98.3-4].$$

So we assume $\sigma = \frac{1}{2}$, to be safe

↑ maximum possible variance.

Assume distribution as widespread as possible, to be safe.

Therefore, we have, at the significant level:

$$P(-1.96 < \frac{m - \bar{m}}{\frac{\sigma}{\sqrt{n}}} \leq 1.96) = \underbrace{\frac{1}{\sqrt{2\pi}} \int_{-1.96}^{1.96} e^{-s^2/2} ds}_{0.95}$$

Hence:

$$P(-1.96 < \frac{m - \frac{41}{100}}{\frac{(\frac{1}{2})}{\sqrt{100}}} \leq 1.96) = 0.95$$

We are assessing whether the coin is fair.
So we want to see if the above holds for $m = \frac{1}{2}$

$$P(-1.96 < \frac{\frac{1}{2} - \frac{41}{100}}{\frac{(\frac{1}{2})}{\sqrt{100}}} \leq 1.96) = 0.95$$

1.8

And, indeed, $-1.96 < 1.8 \leq 1.96$.

A more informative way of setting this up is:

$$P(-1.96 \frac{\sigma}{\sqrt{n}} + m \leq \frac{h}{n} < (1.96) \frac{\sigma}{\sqrt{n}} + m) = 0.95$$

where h = # of heads in n tosses

$$P((-1.96 \frac{\sigma}{\sqrt{n}} + m)n \leq h < (1.96 \frac{\sigma}{\sqrt{n}} + m)n) = 0.95$$

↑ all h that satisfy this inequality $(-1.96 \frac{\sigma}{\sqrt{n}} + m)n \leq h < (1.96 \frac{\sigma}{\sqrt{n}} + m)n$ indicate the coin is fair (i.e., $m = \frac{1}{2}$) to a significant level (0.95).



5/8/98.10

In this particular example :

$$P(40.2 \leq h < 59.8) = 0.95$$



The fact that we got 41 heads out of our $n=100$ tosses means that the inequality: $40.2 \leq 41 < 59.8$ is satisfied.

The fact that we get between 40.2 and 59.8 heads happens 95% of the time.

Therefore, it is not surprising that I got 41 heads.

Therefore, the result 41 heads out of 100 is not significant at the 0.95 confidence level.

In other words, we have no right to conclude that the coin is biased. Because 95% of the time, the number of heads in 100 tosses is between $[40.2, 59.8)$.

This is how statistical reasoning works.

This scheme of reasoning is 75% of statistics.

Don't kid yourselves.

The scheme of reasoning we followed is the scheme of reasoning for statistics.

- We tossed a coin and observed the number of heads.
- We conjectured that the coin was fair (i.e., $m=2$).
- We chose a significance level (in this case, 0.95).
- Using the standard normal random variable, we determined the lower and upper bound on the observations (i.e., # of heads) for which the integral of the standard normal density equals the significance level.
- If the # of observations is w/in the lower + upper bounds, the conjecture is supported.
 " " " " " " not w/in " " " " " , the conjecture is not supported.

If we had tossed the coin 100 times and had obtained 39 heads, then the inequality does not hold.

$$P(40.2 \leq (h=39) < 59.8) = 0.95$$

And this would be significant at the 0.95 level.

And we would conclude that the coin was biased and that $m \neq \frac{1}{2}$.

It's like going to the astrologer.

50 SHEETS
22-141 100 SHEETS
22-142 200 SHEETS
22-144



5/8/98.11

Is there a way of justifying all these steps?
 How far have people gotten?
 Let's consider a number of justifications people have put forward

Justification (1) The Central Limit Theorem

It's not really a justification.
 It's a psychological justification.

This is a major result of probability.

Suppose you have:

$X_1, X_2, \dots$ i.i.d. random variables with:

an infinite sequence

$$E(X_1) = m$$

$$\text{Var}(X_1) = \sigma^2 < \infty \leftarrow \text{finite variance}$$

the main assumption
 ✓

The random variable $X_1 + X_2 + \dots + X_n$ has mean nm :

$$\begin{aligned} E(X_1 + X_2 + \dots + X_n) &= E(X_1) + E(X_2) + \dots + E(X_n) \\ &= nm \end{aligned}$$

So we obtain a standard normal random variable by taking the sum of n random variables, subtracting nm , and normalizing the result by $\sigma\sqrt{n}$:

$$Z = \frac{X_1 + X_2 + \dots + X_n - nm}{\sigma\sqrt{n}}$$

$$E(Z) = E\left(\frac{X_1 + X_2 + \dots + X_n - nm}{\sigma\sqrt{n}}\right)$$

$$= \frac{1}{\sigma\sqrt{n}} \left[E(X_1 + X_2 + \dots + X_n - nm) \right]$$

$$= \frac{1}{\sigma\sqrt{n}} \left[E(X_1) + E(X_2) + \dots + E(X_n) - E(nm) \right]$$

$$= \frac{1}{\sigma\sqrt{n}} \left[nm - nm \right]$$

$$E(Z) = 0 \quad \checkmark$$



5/8/98.12

$$\text{Var}(\bar{Z}) = \text{Var}\left(\frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}}\right)$$

$$= \text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{\sigma \sqrt{n}} - \frac{nm}{\sigma \sqrt{n}}\right) \quad \leftarrow \text{Var}(Y+c) = \text{Var}(Y)$$

$$= \text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{\sigma \sqrt{n}}\right) \quad \leftarrow \text{Var}(cY) = c^2 \text{Var}(Y)$$

$$= \frac{1}{\sigma^2 n} \text{Var}(X_1 + X_2 + \dots + X_n)$$

Since $X_1, X_2, \dots, X_n$ are independent [5/4/98.8-10]:

$$= \frac{1}{\sigma^2 n} \underbrace{[\text{Var}(X_1) + \text{Var}(X_2) + \dots + \text{Var}(X_n)]}_{n\sigma^2}$$

$$= \frac{1}{\sigma^2 n} \cdot n\sigma^2$$

$$\text{Var}(\bar{Z}) = 1 \quad \checkmark$$

So, random variable $\bar{Z} = \frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}}$ has $E(\bar{Z}) = 0$
 $\text{Var}(\bar{Z}) = 1$ ↖

The random variable $\frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}}$ is standard ↖

We want to consider what happens when $n \rightarrow \infty$.

If there's going to be a limit as $n \rightarrow \infty$, it's going to happen as you keep the random variables normalized, for each n .

And, indeed:

$$\lim_{n \rightarrow \infty} P\left(a < \underbrace{\frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}}}_{\text{the normal distribution}} \leq b\right) \rightarrow \frac{1}{\sqrt{2\pi}} \int_a^b e^{-s^2/2} ds \quad \leftarrow$$

In other words, if you add more and more measurements, and keep normalizing them so they are standard normal, then in the limit, they take the bell shaped curve for their density.

That's the Central Limit Theorem.

This is part of the justification why, given several independent measurements of the same quantity, you assume that the X_i are normally distributed.

5/8/98.13

What is tough is that you don't know how far in the limit you have to go until you converge.

In practice, say statisticians, this convergence is extremely fast.

For about 50 years, there was a school of Russian statisticians, whose sole purpose was to investigate the speed of convergence of the Central Limit Theorem.
There are many results on this.

Note: We have not proved The Central Limit Theorem.

There is a bad proof in my book.
There is a good proof in the Italian version of my book.

There are many proofs of this theorem.
We don't have time.
It takes a couple of hours to prove it.

Let's see how it works.

Example

TWA
United

There are 2 airlines flying between Boston and LA.


Every day, 1000 passengers fly between Boston and LA.
Each passenger chooses between TWA and United by flipping a fair coin.
How many seats must each airline have to be safe?

What do you mean by safe?

You assign confidence levels.

You can be safe 95% of the time
or very safe 99% of the time.

You look up in your computer, or p.177 of the text, and find that a standard normally distributed random variable ≤ 1.645 has area under the density distribution equal to 0.95.



$$\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{1.645} e^{-s^2/2} ds = 0.95$$

$$P(Z \leq 1.645) = 0.95$$

Now, you turn the crank.

This is just a Bernoulli Process.

Let X_i = random variable that the i th passenger chooses TWA

$$P(X_i = 1) = p = \frac{1}{2}$$

$$P(X_i = 0) = 1 - p = \frac{1}{2}$$

The X_i are i.i.d

Let $S_n = X_1 + X_2 + \dots + X_n$ here, $n = 1000$
 = total number of passengers that choose TWA

The Central Limit Theorem states:

$$P\left(a < \frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}} \leq b\right) \rightarrow \frac{1}{\sqrt{2\pi}} \int_a^b e^{-s^2/2} ds \quad \text{as } n \rightarrow \infty$$

So, we have:

$$P\left(\frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}} \leq 1.645\right) \approx \underbrace{\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{1.645} e^{-s^2/2} ds}_{0.95}$$

Since this is a Bernoulli process:

a) the mean of each X_i is:

$$E(X_i) = \sum_n n P(X_i = n)$$

$$= 0 \cdot (1-p) + 1 \cdot p \quad \leftarrow p = \frac{1}{2} \text{ (coin is fair)}$$

$$\boxed{m = E(X_i) = \frac{1}{2}}$$

b) as we've shown earlier [5/6/98.3-4], the upper bound on the variance

$$\sigma^2 = \text{Var}(X_i) \leq \frac{1}{4}$$

So, to be safe, we take $\boxed{\sigma = \frac{1}{2}}$

Now we simply need to find S_n that satisfies the inequality in equation (*).
 Namely, we want S_n such that:

$$\frac{S_n - nm}{\sigma \sqrt{n}} \leq 1.645$$

$$S_n \leq 1.645 \sigma \sqrt{n} + nm$$

$$S_n \leq 1.645 \cdot \frac{1}{2} \sqrt{1000} + 1000 \cdot \frac{1}{2}$$

$$\boxed{S_n \leq 526.01 \text{ w/ } 95\% \text{ confidence}}$$

So, w/ 526 seats, you can accommodate passengers, from the 1000 that choose at random to fly TWA, 95% of the time.

5/8/98.15

That's important kids.

No nonsense.

If you provide 1000 seats, you'll go broke.

Suppose you want to be safe 99% of the time.
Plug it all in.

- 1) First, we determine (from the tables) that a standard normally distributed variable Z has 99% of the area under the density distribution when:

$$P(Z \leq 2.33) = 0.99$$

- 2) Then, we use the Central Limit Theorem:

$$P\left(\frac{S_n - nm}{\sigma \sqrt{n}} \leq 2.33\right) \approx \underbrace{\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{2.33} e^{-s^2/2} ds}_{0.99}$$

- 3) Finally, we solve for S_n such that the above inequality holds:

$$\frac{S_n - nm}{\sigma \sqrt{n}} \leq 2.33$$

$$S_n \leq 2.33 \sigma \sqrt{n} + nm$$

$$S_n \leq 2.33 \frac{1}{2} \sqrt{1000} + 1000 \cdot \frac{1}{2}$$

$$S_n \leq 536.84 \text{ w/ } 99\% \text{ confidence}$$

In summary:

min seats	confidence
526	95%
537	99%

With only 11 more seats, you go from 95% to being 99% sure you can handle passengers who randomly choose TWA.

All of this is non trivial.

Because we do not know the rate of convergence of the Central Limit Theorem. We need to address how good an approximation we have for small n (i.e., $n \neq \infty$).

Q/A after class:

There is a nice proof of the following: $D = \frac{d}{dx}$

$$[D^k x^n]_{x=0} = [D^n x^k]_{x=0}$$

$$[p(D)g(x)]_{x=0} = [g(D)p(x)]_{x=0}$$

For example: $[D^4(x^5 + 2x^3)]_{x=0} = [(D^5 + 2D^3)x^4]_{x=0}$



The Normal Distribution (cont'd)

Recall that:

$$E(Z) = 0$$

$$\text{Var}(Z) = 1$$

a random variable Z w/ standard normal distribution means that the density is:

$$\text{dens}(Z=t) = \frac{1}{\sqrt{2\pi}} e^{-t^2/2}$$

In other words:

$$P(Z \leq t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-s^2/2} ds$$

These statements are equivalent, as we know.

The normal distribution is also known as the Gaussian Distribution

• More generally [5/6/98.9-12]:

The random variable $\sigma Z + m$ has a normal distribution w/:

$$E(\sigma Z + m) = m$$

$$\text{Var}(\sigma Z + m) = \sigma^2$$

• And we computed [5/6/98.12-14]:

$$P(\sigma Z + m \leq t) = \frac{1}{\sigma \sqrt{2\pi}} \int_{-\infty}^t e^{-(s-m)^2/2\sigma^2} ds$$

• We also observed, when we were studying variance, that [5/6/98.14]:

if $X_1, X_2, \dots, X_n$ are i.i.d. random variables w/ $\text{Var}(X_i) = \sigma^2 < \infty$

then:

i.e., $\text{Var}(X_i)$ finite

$$\text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \text{Var}\left(\frac{1}{n}X_1\right) + \text{Var}\left(\frac{1}{n}X_2\right) + \dots + \text{Var}\left(\frac{1}{n}X_n\right)$$

$$\underbrace{\hspace{10em}}_{\text{variance of the average}} = \frac{1}{n^2} \text{Var}(X_1) + \frac{1}{n^2} \text{Var}(X_2) + \dots + \frac{1}{n^2} \text{Var}(X_n)$$

$$= \frac{1}{n^2} (n \text{Var}(X_1))$$

$$= \frac{1}{n} \text{Var}(X_1)$$

$$\text{Var}\left(\frac{X_1 + X_2 + \dots + X_n}{n}\right) = \frac{\sigma^2}{n}$$

- Last time, we concluded with an exploration of this fact. Namely, if you have n independent, identically distributed variables w/ finite variance, suppose that:

$$E(X_1) = m$$

$$\text{Var}(X_1) = \sigma^2$$

Then:

$$\begin{aligned} E(X_1 + X_2 + \dots + X_n) &= E(X_1) + E(X_2) + \dots + E(X_n) \\ &= n E(X_1) \\ &= nm \end{aligned}$$

$$\begin{aligned} \text{Var}(X_1 + X_2 + \dots + X_n) &= \text{Var}(X_1) + \text{Var}(X_2) + \dots + \text{Var}(X_n) \\ &= n \text{Var}(X_1) \\ &= n\sigma^2 \end{aligned}$$

How do we normalize the random variable $X_1 + X_2 + \dots + X_n$?

Normalizing means: 1) subtracting the mean
and 2) dividing by the standard deviation
so that the resulting random variable has:

$$\text{mean} = 0$$

$$\text{variance} = 1$$

$$\frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}} \leftarrow E(X_1 + X_2 + \dots + X_n)$$

$\sigma \sqrt{n}$
 standard deviation $(X_1 + X_2 + \dots + X_n)$

this random variable is standardized

So, for every n , we keep standardizing it, so that it doesn't blow up. This is the only hope we have of tending to a limit for the distribution.

And, as a matter of fact, under these circumstances:

finite variance

independent random variables

You have:

$$P\left(a < \frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}} \leq b\right) \rightarrow \frac{1}{\sqrt{2\pi}} \int_a^b e^{-t^2/2} dt$$

Which means that the limit of the distribution of this standardized random variable tends to the standard normal distribution.

$$P\left(\frac{X_1 + X_2 + \dots + X_n - nm}{\sigma \sqrt{n}} \leq t\right) \rightarrow \frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-s^2/2} ds$$

This is the content of the Central Limit Theorem. $\leftarrow$ {we don't have time to prove this}

The Central Limit Theorem provides yet another justification for statistics.
Namely, when you take several "repeated measures" of "same quantity", you assume that the quantity in question is given by the a normal distribution, whose mean = sample mean
variance = sample variance

You plug it in and do computations.

For example, we discussed last time the problem of passengers choosing, at random, one of two airlines [5/8/98.13-15].

What is the tolerance level, as they say, for an airline to have full seats?

$\uparrow$ Say 95% of the time, we want to accommodate all passengers.
For the rest of the time, the hostess will come in and say "Will someone wish to give up a seat?"

As we showed, we have:

$$P\left(\frac{S_n - nm}{\sigma \sqrt{n}} \leq 1.645\right) \approx 0.95 \quad \Rightarrow S_n = 526$$

$\uparrow$ we look this up in a table

So this is another justification for statistics.


We saw the original justification from using Bayes' Theorem and choosing $\text{dens}(m=s) = 1$ [5/8/98.6].
Now let's see two additional justifications.

Justification (2) Maxwell-Einstein Derivation of the Normal Distribution

→ some of his best papers were on the normal distribution.

They say this:

Let's take a plane and suppose we pick a point at random on the plane.
Let's try to give meaning to picking a point at random on the plane.



Now we know we can not use the measure of the plane (area).
Because that gives you infinite area and that can not be normalized.
That's a sticky point that people have never been able to fix.

(Remember, people use measures as priors and they work, but they have no idea why it works.)

We want a real probability density in 2 variables (X, Y) which are random.
What are we to mean for random?
We mean the following things:

- (1) X and Y (coordinates of the point you pick) are independent random variables
↖ you pick an independent X and an independent Y


Then, the joint density equals the product of the densities [4/15/98.7]:

$$\text{dens}(X=t, Y=s) = \text{dens}(X=t) \text{dens}(Y=s)$$

and for reasons of symmetry, of course:

$$\text{dens}(X=t) = \text{dens}(Y=t)$$

- (2) switch to polar coordinates



The radius R and angle θ are random variables.

$$\text{dens}(X=t, Y=s) dt ds = \text{dens}(R=r, \Theta=\theta) r dr d\theta$$

⇒ The assumption is that:

$$\text{dens}(R=r, \Theta=\theta) = \text{dens}(R=r)$$

↑ assumption is that the density is independent of Θ .

With just this assumption, that's it.
Out comes the normal distribution


The only density that satisfies the above is the normal distribution
 $\text{dens}(R=r) \quad \frac{1}{r\sqrt{2\pi}} e^{-r^2/2\sigma^2}$

Lets see why.
We've actually done most of the work.

Say $\text{dens}(R=r) = \underbrace{g(r)}_{\text{some function } g \text{ of } r}$
 and $\text{dens}(X=t) = f(t)$

From items (1) + (2) above, we have:

$$\begin{aligned} \text{dens}(X=t, Y=s) dt ds &= \text{dens}(R=r, \Theta=\theta) r dr d\theta \\ \text{dens}(X=t) \text{dens}(Y=s) dt ds &= \text{dens}(R=r) r dr d\theta \\ \underbrace{\text{dens}(X=t) \text{dens}(Y=s) dt ds}_{f(t)f(s)} &= \underbrace{\text{dens}(R=r) r dr d\theta}_{g(\sqrt{t^2+s^2})} \end{aligned}$$



This is a functional equation to solve.

(*) $f(t)f(s) = g(\sqrt{t^2+s^2})$

You have to find all functions $f + g$ that satisfy this functional equation.
That's not hard.

Set $s=0$:

$$\left. \begin{aligned} f(t)f(0) &= g(\sqrt{t^2+0}) \\ &= g(t), \quad t \geq 0 \end{aligned} \right\} \Rightarrow f(t) = \frac{g(t)}{f(0)}, \quad t \geq 0$$

Similarly, set $t=0$:

$$\left. \begin{aligned} f(0)f(s) &= g(\sqrt{t^2+s^2}) \\ &= g(s), \quad s \geq 0 \end{aligned} \right\} \Rightarrow f(s) = \frac{g(s)}{f(0)}, \quad s \geq 0$$

So, we can rewrite functional equation (*) as:

$$\frac{g(t)}{f(0)} \frac{g(s)}{f(0)} = g(\sqrt{t^2+s^2})$$

$$\frac{g(t)g(s)}{f(0)^2} = g(\sqrt{t^2+s^2})$$

This is the functional equation we already had, when we were doing exponential distributions [4/15/98.14-16], [4/22/98.5].

$$\text{Set } h(t) = \frac{g(t)}{f(0)}$$

Then:

$$h(t)h(s) = h(\sqrt{t^2+s^2})$$

This is the functional equation for the exponential

$$h(t) = e^{ct}$$

$$\text{Since } h(t) = \frac{g(t)}{f(0)},$$

$$g(t) = e^{c't^2}$$

Recall that $g(t) = \text{dens}(R=t)$ is just a density distribution. So we know that:

$$\int_{\theta=0}^{2\pi} \int_{r=0}^{\infty} e^{c'r^2} r dr d\theta = 1$$

So we can determine c' .

If you work this out, you end up with:

$$g(t) = \frac{1}{\sigma\sqrt{2\pi}} e^{-t^2/2\sigma^2}$$

The normal distribution.

$$\Rightarrow \left\{ \begin{aligned} h(t^2)h(s^2) &= h(\sqrt{t^2+s^2}) \\ \frac{g(t^2)}{f(0)} \frac{g(s^2)}{f(0)} &= \frac{g(\sqrt{t^2+s^2})}{f(0)} \\ \frac{g(t)g(s)}{f(0)^2} &= \frac{g(\sqrt{t^2+s^2})}{f(0)} \\ \frac{g(t)g(s)}{f(0)} &= g(\sqrt{t^2+s^2}) \\ 2f(0) &= f(0)^2? \end{aligned} \right.$$

So the normal distribution comes out just from this consideration of items (1) + (2) above.

Really wonderful stuff.

The last justification we shall consider:
Justification (3) Wiener's Characterization

5/11/98.7

This is the best one - not just because he taught this course,

Wiener says:

- When we dealt w/ tossing a coin, we didn't characterize just coin tossing. We studied the Bernoulli Process - sequences of coin tossing.
- Similarly, when we studied the waiting time for a phone call to end, namely the exponential distribution, we didn't just study the exponential distribution, we studied sums of independently distributed exponential random variables, and we were thereby lead to the Poisson Process!
- When we studied picking a point at random in the interval $[0, a]$, we immediately went over to picking n points at random in the interval $[0, a]$ and then we started doing hairy computations

In each case, when we were confronted w/ the problem of studying a probability distribution, we tried to go over to a stochastic process associated w/ that distribution.

probability distribution $\rightarrow$ stochastic process

That's what Wiener does w/ the Normal Distribution.

normal distribution $\rightarrow$ Wiener Process

The Wiener Process or Brownian motion

It's based on the following remark.
A very important fact.

Suppose X_1 is $\mathcal{N}(m_1, \sigma_1^2)$
 X_2 is $\mathcal{N}(m_2, \sigma_2^2)$ } $X_1 + X_2$ independent

X_1 is normal, with mean m_1 ,
variance σ_1^2

So we have 2 independent random variables, which are normally distributed.

Then the basic fact is that their sum is normal, with mean the sum of their means and their variance, the sum of their variances.

$X_1 + X_2$ is $\mathcal{N}(\underbrace{m_1 + m_2}, \underbrace{\sigma_1^2 + \sigma_2^2})$

To be expected, since $\nearrow$
expectation of sum is sum
of the expectations.

$\nwarrow$ But the interesting thing is that the
variances behave like this.

As a matter of fact, one can show that the normal distribution is the only distribution where:

$$X_1 \text{ with } E(X_1) = m_1 \\ \text{Var}(X_1) = \sigma_1^2$$

$$X_2 \text{ with } E(X_2) = m_2 \\ \text{Var}(X_2) = \sigma_2^2$$

$$X_1 + X_2 \text{ has } E(X_1 + X_2) = m_1 + m_2 \\ \text{Var}(X_1 + X_2) = \sigma_1^2 + \sigma_2^2$$

The means add and the variances add

We could check this by the convolution of 2 hairy integrals

Alternatively:

Proof

$$\text{We may assume that } m_1 = 0 \\ m_2 = 0$$

Since if we shift the random variable:

$$X_1 \leftarrow X_1 - m_1$$

$$\begin{aligned} E(X_1) &= E(X_1 - m_1) \\ &= E(X_1) - E(m_1) \\ &= m_1 - m_1 \\ &= 0 \end{aligned}$$

So we only need to check:

X_1, X_2 independent, normal with:

$$E(X_1) = 0 \quad E(X_2) = 0$$

$$\text{Var}(X_1) = \sigma_1^2 \quad \text{Var}(X_2) = \sigma_2^2$$

By the continuous law of alternatives consequence [4/8/98.7], the joint density is:

$$\text{dens}(X_1 + X_2 = t) = \int_{-\infty}^{\infty} \text{dens}(X_1 + X_2 = t / X_2 = s) \text{dens}(X_2 = s) ds$$

The event $X_1 + X_2 = t$, given that $X_2 = s$ is the same as: $X_1 + s = t$
 $X_1 = t - s$

$$= \int_{-\infty}^{\infty} \text{dens}(X_1 = t - s) \text{dens}(X_2 = s) ds$$

$$= \int_{-\infty}^{\infty} \frac{1}{\sigma_1 \sqrt{2\pi}} e^{-(t-s)^2 / 2\sigma_1^2} \frac{1}{\sigma_2 \sqrt{2\pi}} e^{-s^2 / 2\sigma_2^2} ds$$

$$= \frac{1}{\sigma_1 \sqrt{2\pi}} \frac{1}{\sigma_2 \sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-(t-s)^2/2\sigma_1^2} e^{-s^2/2\sigma_2^2} ds$$

Let's work out the exponents:

$$\begin{aligned} \frac{-(t-s)^2}{2\sigma_1^2} - \frac{s^2}{2\sigma_2^2} &= \frac{-t^2 + 2ts - s^2}{2\sigma_1^2} - \frac{s^2}{2\sigma_2^2} \\ &= \frac{-t^2\sigma_2^2 + 2ts\sigma_2^2 - s^2\sigma_2^2 - s^2\sigma_1^2}{2\sigma_1^2\sigma_2^2} \end{aligned}$$

The claim is that this is a perfect square of the form:

$$-\frac{(s\sqrt{\sigma_1^2 + \sigma_2^2} + \beta)^2}{2\sigma_1^2\sigma_2^2}$$

Note how tantalizingly close:

$$-\frac{(s\sqrt{\sigma_1^2 + \sigma_2^2} - t\sigma_2)^2}{2\sigma_1^2\sigma_2^2}$$

gets.

The claim is, that if you work this all out, you eventually obtain:

(*)

$$\text{dens}(X_1 + X_2 = t) = \frac{1}{\sqrt{\sigma_1^2 + \sigma_2^2} \sqrt{2\pi}} e^{-t^2/2(\sigma_1^2 + \sigma_2^2)}$$


Normal distribution w/ variance $\sigma_1^2 + \sigma_2^2$

additive.

So, if you take 2 random variables that are normally distributed, if you take their sum, then their sum is normally distributed with mean the sum of their means and with variance the sum of their variances.

The main idea from Wiener.


Let's construct a continuous random walk



Can you give a consistent probability density for this situation?
The answer is very simple.

$\Omega =$ all continuous functions $f(t)$, where $f(0)=0$

events are defined as follows:



at each time t , the balls have moved around, like the Museum of Science display [5/1/98.13]. So at each time t , we have a normal distribution.

Let $W(t)$ be a random function.

The density is:

$$\text{dens}(W(t)=x) = \frac{1}{\sqrt{2\pi t}} e^{-x^2/2t}$$

$$\frac{1}{\sigma\sqrt{2\pi}} e^{-x^2/2\sigma^2}$$


Variance is t
 $\sigma^2 = t$

There is one major thing to check.

Is this definition consistent?

Consistency

Check:




Let $t = t_1 + t_2$

By the Law of Alternatives consequence:

$$\text{dens}(W(t)=x) = \int_{-\infty}^{\infty} \underbrace{\text{dens}(W(t)=x | W(t_1)=y)}_{\text{It is given that } W(t_1)=y.} \underbrace{\text{dens}(W(t_1)=y) dy}_{\frac{1}{\sqrt{2\pi t_1}} e^{-y^2/2t_1}}$$

It is given that $W(t_1)=y$.

Thus, for each $y \in (-\infty, \infty)$, we have:



The resulting curve will be shifted. And this will be normally distributed, with expectation y . It will be centered at y .

$$\text{dens}(W(t)=x | W(t_1)=y) = \frac{1}{\sqrt{t_2} \sqrt{2\pi}} e^{-(x-y)^2/2t_2}$$

$$= \int_{-\infty}^{\infty} \underbrace{\frac{1}{\sqrt{t_2} \sqrt{2\pi}} e^{-(x-y)^2/2t_2}}_{g(x-y) = \text{dens}(Z_2=x-y)} \underbrace{\frac{1}{\sqrt{t_1} \sqrt{2\pi}} e^{-y^2/2t_1}}_{f(y) = \text{dens}(Z_1=y)} dy$$

$$g(x-y) = \text{dens}(Z_2=x-y) \quad f(y) = \text{dens}(Z_1=y)$$

This is exactly the convolution of two normal distributions.

[4/15/98.9]

Thus:

$$h(x) = \int_{-\infty}^{\infty} f(y) g(x-y) dy$$

Where:

$$h(x) = \text{dens}(Z_1 + Z_2 = x)$$

$$= \text{dens}(Z_1 + Z_2 = x)$$

And this is exactly the density distribution we obtained earlier in equation (x) [5/11/98.8-9]

$$= \frac{1}{\sqrt{(t_1) \sqrt{2\pi}} \sqrt{(t_2) \sqrt{2\pi}}} e^{-x^2/2((t_1)^2 + (t_2)^2)}$$

And, since $t = t_1 + t_2$

$$\boxed{\text{dens}(W(t)=x) = \frac{1}{\sqrt{2\pi t}} e^{-x^2/2t}} \quad \checkmark$$

So it's consistent.

Therefore, because it is consistent, it therefore exists.

There are more technical things to verify, but this is a basic fact.

If you stop a random walk in between, and look at it in between, and measure densities, it's consistent w/ the density measurement of the entire random walk.

(Just the same idea as we used to show the memoryless property for the Poisson Process [4/29/98.10-12])

Therefore, we have a new stochastic process that obtains probabilities by a random function.
Instead of a random walk [5/11/98.5-13],

$$P(F=2n)$$

you have a random function.

$$P(W(t) \leq x)$$

This is how a random walk tends to this Wiener Process in the limit.

The Central Limit Theorem tells you how random walk approaches Wiener Process.

All computations are done in the Wiener Process, once you discover the Wiener Process exists, and you forget about random walks.

You just do everything in the Wiener Process.

This is the basic process of the theory of stochastic processes, which you study in 18.445.

Law of Large Numbers $\leftarrow$ {A strengthening of the fact that the variance goes down as n increases.}

$X_1, X_2, \dots$ of i.i.d. random variables
with finite variance

$$\uparrow \text{Var}(X_1) = \sigma^2 < \infty$$



$$P\left(\left|\frac{X_1 + X_2 + \dots + X_n}{n} - m\right| > \varepsilon\right) \rightarrow 0, \text{ as } n \rightarrow \infty, \\ \text{for any } \varepsilon > 0.$$

This is the probability that the average deviates from the mean by more than ϵ .

Note that this law is true for any ϵ .

And this means, more or less, that w/ probability 1, the sample average $\frac{X_1 + X_2 + \dots + X_n}{n}$ tends to the mean m .

A strengthening of this statement would be:

$$P\left(\lim_{n \rightarrow \infty} \frac{X_1 + X_2 + \dots + X_n}{n} = m\right) = 1$$

The Strong Law of Large Numbers

So, now you know all the probability that you need.
I hope you take another probability course and you find it real easy.

Last Quiz on Wednesday.

I hope you all take UROPs w/ me and help me w/ my problems.

— The End —



{ Problem #9 is from a research paper. }

Q: Can you use the law of successive conditioning to prove the inclusion-exclusion principle?

A: Yes. You can prove the inclusion-exclusion principle using conditional probability. I want to give that to you as a problem.

There are dozens of proofs of the inclusion-exclusion principle. Most proofs use induction on the number of events.

First you prove that:

$$P(A_1 \cup A_2) = P(A_1) + P(A_2) - P(A_1 \cap A_2)$$

which is an axiom of probability.

Then you assume that it is true for n events and prove:

$$P(A_1 \cup A_2 \cup \dots \cup A_n \cup A_{n+1}) = \dots$$

- 1) condition relative to the last event
- 2) juggle the properties of conditional probability
- 3) reduce it to the case of an event with conditional probability

There's another way of proving the inclusion-exclusion principle, which is given in the book, which is the simplest.

Let's discuss it here, since I'm not going to discuss it in class.

His favorite proof of the inclusion-exclusion principle

This is based on the idea of an indicator random variable

Indicator Random variable

You have a sample space Ω and an event A .

You define a random variable I_A , which is called the indicator random variable of the event A .

Let me define it first probabilistically, then set theoretically. Equivalent - two languages.



$$I_A = \begin{cases} 1 & \text{if } A \text{ happens} \\ 0 & \text{if } A \text{ does not happen} \end{cases}$$

read: Indicator A

Probabilistic

What does this mean in set theoretic language?
i.e.,

$$I_A(\omega) = \begin{cases} 1 & \text{if } \omega \in A \\ 0 & \text{if } \omega \notin A \end{cases}$$

Set Theoretic

I_A at a sample point ω .

We see a striking contrast between using probabilistic language and set theoretic language.

I_A is a random variable.
What is its probability distribution?

$$P(I_A = 1) = P(A)$$

$$P(I_A = 0) = 1 - P(A)$$

What is the expectation of this random variable?

$$\begin{aligned} E(I_A) &= \sum_n n P(I_A = n) \\ &= 1 \cdot P(I_A = 1) + 0 \cdot P(I_A = 0) \\ &= P(A) \end{aligned}$$

from above,
since $P(I_A = 1) = P(A)$

$E(I_A) = P(A)$

This is a fundamental identity.

What does it do for you?

It reduces computations of probabilities to computations of expectations.

Why are computations w/ expectations nicer?

Because expectation of a sum is the sum of the expectations.

No questions asked.

So now, let's use the technique of indicator random variables to prove the inclusion-exclusion principle.

To do that, we have to do the following.
We have to set up a correspondence between the operations of boolean algebra of events and the operations of indicator random variables, which are the ordinary sum + times.

This is very nice. We get rid of $\cap$, $\cup$ and they get replaced by $+$, $\times$.

Let's see how that works.

First, what is the indicator of A^c ?

That's easy.

$$I_{A^c} = \underset{\substack{\uparrow \\ \text{the function } 1}}{1} - I_A$$

Q/A: $1 = I_\Omega$
 1 is a function that takes the value 1 at every sample point.

Next, The indicator of A and B .

$$I_{A \cap B} = I_A \cdot I_B$$

$\uparrow$
ordinary times

$I_A \cdot I_B = 1$ iff both A and B happen

Things get a little kinky w/ the union.

$$I_{A \cup B} = I_A + I_B - I_A \cdot I_B$$

Let's check.

If $I_{A \cup B} = 1$ then either I_A or I_B or both must be 1

	I_A	$+$	I_B	$-$	$I_A \cdot I_B$	
$1 =$	1	+	0	$-$	0	
$1 =$	0	+	1	$-$	0	
$1 =$	1	+	1	$-$	1	✓

But this can be written in a much more striking way.
And this is the key to the whole story.

$$I_{A \cup B} = 1 - (1 - I_A)(1 - I_B) \Leftrightarrow \begin{cases} 1 - [1 - I_A - I_B + I_A I_B] \\ I_A + I_B - I_A I_B \end{cases}$$

By the way, this is DeMorgan's Law:

$$\begin{aligned} I_{A \cup B} &= 1 - I_{(A \cup B)^c} \\ &= 1 - I_{A^c \cap B^c} \\ &= 1 - I_{A^c} I_{B^c} = 1 - (1 - I_A)(1 - I_B) \end{aligned}$$

This extends to several unions.

More generally:

$$I_{A_1 \cup A_2 \cup \dots \cup A_n} = 1 - (1 - I_{A_1})(1 - I_{A_2}) \dots (1 - I_{A_n})$$

Indicator of $A_1 \cup A_2 \cup \dots \cup A_n$

This is the inclusion-exclusion principle.

How does the inclusion-exclusion principle come out?

By taking expectations of both sides.

Let's expand RHS by rules of algebra.

$$I_{A_1 \cup A_2 \cup \dots \cup A_n} = \sum_i I_{A_i} - \sum_{i < j} I_{A_i} I_{A_j} + \sum_{i < j < k} I_{A_i} I_{A_j} I_{A_k} - \dots + \dots$$

Taking expectations:

$$E(I_{A_1 \cup A_2 \cup \dots \cup A_n}) = E\left(\sum_i I_{A_i} - \sum_{i < j} I_{A_i} I_{A_j} + \sum_{i < j < k} I_{A_i} I_{A_j} I_{A_k} - \dots + \dots\right)$$

expectation is additive

~~the expectation of the indicator of an event is the probability of that event.~~

$$= \sum_i E(I_{A_i}) - \sum_{i < j} E(\overbrace{I_{A_i} I_{A_j}}^{I_{A_i \cap A_j}}) + \sum_{i < j < k} E(\overbrace{I_{A_i} I_{A_j} I_{A_k}}^{I_{A_i \cap A_j \cap A_k}}) - \dots + \dots$$

Recall the fundamental identity [sc 2/27/98.2].

The expectation of the indicator of an event is the probability of that event.

$$E(I_A) = P(A)$$

So we have:

$$P(A_1 \cup A_2 \cup \dots \cup A_n) = \sum_i P(A_i) - \sum_{i < j} P(A_i \cap A_j) + \sum_{i < j < k} P(A_i \cap A_j \cap A_k) - \dots + \dots$$

That's it. We've proved the inclusion-exclusion principle.

So now you say: "What's the point of going through all this stuff w/ Rényi's Principle if there is a simpler way?"

• More on Rényi's Principle (from a different point of view)

Let's wax theoretical.

Suppose you are given:

$A_1, A_2, \dots, A_n = \text{events (finite)}$

And suppose they are allowed $\cup, \cap, ^c$, but not infinite unions (the 3 boolean operations, but not countable additivity)

How do we describe the most general event that is expressible in terms of $A_1, A_2, \dots, A_n$?

In other words, you're given these events.
What do they determine?

There are 2 ways of answering this question.

Cheapo answer

A: anything that can be expressed in terms of the events with $\cup, \cap, ^c$.
That's called a boolean function.

Answer — any boolean function of $A_1, A_2, \dots, A_n$.

A boolean function is any expression that can be made up of $\cup, \cap, ^c$.


Ex:

$$((A_1 \cup A_2) \cap (A_3 \cup A_1^c) \cap (A_4 \cup A_2^c)) \cup (A_3 \cup A_4 \cup A_5^c)$$

So, in a certain sense, this answers the question.
But it's not a satisfactory answer.

The real answer

Suppose we have 3 events in our sample space Ω



- $A_1 \cap A_2 \cap A_3$
- $A_1 \cap A_2 \cap A_3^c$
- $A_1 \cap A_2^c \cap A_3$
- //// $A_1 \cap A_2^c \cap A_3^c$
- etc

} 2^n smallest possible sets

You see that there is a better answer to the question when you look at this picture.

Consider the smallest possible sets determined by these events, including the one outside all the events.

Any event determined by ~~the~~ the A_i will be some union of the smallest possible events.

How do we determine precisely what it means to be the smallest possible set?

Take all the events and intersect the event or the complement.

Every event expressible in terms of $A_1, A_2, \dots, A_n$ is a disjoint union of events of the form:

$$\rightarrow A_1^{\pm} \cap A_2^{\pm} \cap \dots \cap A_n^{\pm} \leftarrow \left\{ \begin{array}{l} \pm \text{ is Reyni's notation for a} \\ \text{set or its complement.} \end{array} \right\}$$

$$\text{where } A_i^+ = A_i$$

$$A_i^- = A_i^c$$

These are the smallest possible sets.

These are sometimes called the atoms.

So any event expressible in terms of n events is the union of disjoint atoms. This is called the disjunctive normal form.

Every boolean function of $A_1, A_2, \dots, A_n$ is a disjoint union of atoms. This is called disjunctive normal form.

How many distinct events can be expressed in terms of n events?

$$\underbrace{2^{2^n}}_{\text{number of distinct atoms for } n \text{ events}}$$

$$\text{atom}_1 \cup \text{atom}_2 \cup \dots \cup \text{atom}_{2^n}$$

For every atom, you can either take it or not take it.

$$\text{Total \# of disjoint unions of atoms} = 2^{\# \text{ distinct atoms}}$$


$$= 2^{2^n}$$

In logic, this result is called the construction of the free boolean algebra w/ n generators.



One thing I did not mention in class that is covered in the book is the circuit theoretic interpretation of the boolean operations. This was worked out by Shannon when he took this course for his S.M. thesis.

Circuit-theoretic interpretation

every event is considered as a unit.



You can piece together units in series or in parallel




The interpretation is this.
This is a channel through which a message goes.
These are units that can break down.

If you put two units in series, the message will go through iff both of them work.
If you put two units in parallel, the message will go through iff at least one works.

Complementation - they invented something

Here is the main application of this idea:



The circuit is made up of units that may break down independently.
The probability that a unit breaks down is p .
" " " " " does not break down is $q = 1 - p$.

All units have the same probability of breaking down.
Under these conditions, you want the probability that a message sent from s will arrive at t .

This computation is an important theory, which people work full time on. It's called reliability theory.

There's a journal called the IEEE Transactions on Reliability. That is dedicated exclusively to computing stuff like that.



Reliability of a series, parallel circuit is the probability that the circuit works.

So, in general, there are two independent problems:

- 1) You are given a circuit and you have to compute its reliability
- 2) You are given so much money and you need a given reliability (e.g., probability = .99), and you have to design the circuit. Of course, if you put everything in parallel, it works. But that's too expensive. That's a very complicated chore. There's a whole art to it. There are an incredible number of techniques to simplify things.

What's the basic computation of reliability?

Inclusion - exclusion.

Let's see, using the above circuit,

random variable

$$X_i = \begin{cases} 1 & \text{with probability } p \\ 0 & \text{with probability } q \end{cases}$$

This is the most important application of the inclusion - exclusion principle there is. Reliability of the circuit is the probability that at least one path goes through.

Define events:

$$A_1 = ((X_1=1) \cap (X_4=1))$$

$$A_2 = ((X_2=1) \cap (X_4=1))$$

$$A_3 = ((X_3=1) \cap (X_4=1))$$

$$\text{reliability} = P(A_1 \cup A_2 \cup A_3)$$

The inclusion-exclusion principle gives:

$$\begin{aligned} P(A_1 \cup A_2 \cup A_3) &= P(A_1) + P(A_2) + P(A_3) \\ &\quad - P(A_1 \cap A_2) - P(A_1 \cap A_3) - P(A_2 \cap A_3) \\ &\quad + P(A_1 \cap A_2 \cap A_3) \end{aligned}$$



$$\begin{aligned}
&= P((X_1=1) \cap (X_4=1)) + P((X_2=1) \cap (X_4=1)) + P((X_3=1) \cap (X_4=1)) \\
&\quad - P((X_1=1) \cap (X_2=1) \cap (X_4=1)) \\
&\quad - P((X_1=1) \cap (X_3=1) \cap (X_4=1)) \\
&\quad - P((X_2=1) \cap (X_3=1) \cap (X_4=1)) \\
&\quad + P((X_1=1) \cap (X_2=1) \cap (X_3=1) \cap (X_4=1))
\end{aligned}$$

And since all the X_i are independent:

$$\begin{aligned}
&= p^2 + p^2 + p^2 - p^3 - p^3 - p^3 + p^4 \\
P(A_1 \cup A_2 \cup A_3) &= 3p^2 - 3p^3 + p^4
\end{aligned}$$

that's the reliability of that circuit

Q: Can you describe the circuit as the sum or product of their indicators and then take the expectation?

A: No, you can't.

You can't expand this. There is no way.

These are like the atoms.

You really have to use inclusion-exclusion.

In general, we can have a complicated circuit and computing the reliability by inclusion-exclusion can be extremely laborious.

It's a big business to simplify inclusion-exclusion computations.

Even to get approximate expressions by breaking off the inclusion-exclusion sums after the first terms. You have to see what the error is when you break it off.

How do you represent the distributive law?

You can, but it's a little contrived.

Algebra of Partitions

a partition $\pi = \{B\}$ with the properties:
family of events B . B is called a block.

(1) if $B \in \pi$ then $P(B) > 0$

(2) if $B, B' \in \pi$ and $B \neq B'$ then
 $P(B \cap B') = 0$

(3) $P(\bigcup_{B \in \pi} B) = 1$




We want to define operations on partitions that are analogues of boolean $\cup$ and $\cap$ of sets.

If π and σ are two partitions of Ω, P , the meet $(\pi \wedge \sigma)$ is the partition whose blocks are all intersections


$B \cap C$, for $B \in \pi, C \in \sigma$, whenever $P(B \cap C) > 0$

We will use this later when we do entropy.


Let's see a picture



partition π





partition σ



partition $\pi \wedge \sigma$

the meet is all the rectangles

$\sigma = \sim$ partition





There is also an analogue of set union for partitions, but it's harder to visualize.

If π and σ are partitions of Ω , P , then we say π is a finer partition than σ :

$$\pi \leq \sigma$$

when every block of π is contained in some block of σ .

Example



$$\pi \leq \sigma$$

π is finer than σ

How do I get a smaller partition?
I just cut up the blocks of σ and then I get π .

Start w/ 0.

Cut into smaller blocks to get π .

Next time we'll discuss the analogue of set union, the join ($\vee$) of partitions. There is a beautiful algebra of partitions, which parallels the algebra of sets. But partitions do not satisfy the distributive law.

$\pi \vee \sigma =$ smallest partition of which both π and σ are finer.
 $\uparrow$ join

Many different priors occur in practice.

Sometimes you really have information that says that:

- number of good guys is even
- or information that there is a bump in the middle
- or convex
- or monotonic, whatever.

There's a whole array of techniques for computing posterior probabilities.

You may notice that the notation we use for Bayesian statistics is not totally rigorous.

Because you use the random variable U for the prior and the random variable U for the posterior.

$$P(U=j|A=i) = \frac{P(A=i|U=j)P(U=j)}{P(A=i)}$$

These two random variables are different, because they have different distributions. I followed standard statistical treatment, but it's not completely rigorous.


To be rigorous, you would have to give them different names, like U_0 and U

$$P(U=j|A=i) = \frac{P(A=i|U_0=j)P(U_0=j)}{P(A=i)}$$

prior $\nearrow$ posterior $\nearrow$


Trees

In practice, a lot of sample spaces are defined in terms of trees.



$$P(\omega_1) = P(A_1)P(B_1|A_1)P(D_1|B_1)$$

Note that you could have an infinite product for the probability of a sample point, $P(\omega)$. For example:



If you sum up the finite number of branches, then you have the probability of the event at a given vertex.

These probabilities are consistent with the Law of Successive Conditioning [3/2/98.2-3]
(ex: $D_1 \supseteq B_1 \supseteq A_1$, etc.).

The most notable example, which we will cover partly in class, but not completely. So we'll cover it here.

Urn Models

This is a model that is used to cover all sorts of situations.

Traffic accidents, spread of a disease, etc.

You start w/ an urn w/ r red balls and b black balls.

Then you extract a ball:


$\begin{matrix} r \text{ red} \\ b \text{ black} \end{matrix}$

$$\text{Let } X_1 = \begin{cases} 1 & \text{if 1st ball is red} \\ 0 & \text{if " " " black} \end{cases}$$

You set up a replacement discipline whereby:

$\left\{ \begin{array}{l} \text{if you extract a red ball,} \\ \quad \text{you put in } r_1 + 1 \text{ red and } b_1 \text{ black balls} \\ \text{if you extract a black ball,} \\ \quad \text{you put in } r_1' \text{ red and } b_1' + 1 \text{ black balls} \end{array} \right.$

Note: r_1, b_1, r_1', b_1' may be negative.



Then you perform successive extractions, following the above discipline. For each extraction, you add so many red balls and so many black balls.

So, for example, if the first 3 extractions are red, the number of red balls is:

$$\begin{aligned} & r - \underbrace{1 + r_1 + 1}_{1^{\text{st}}} - \underbrace{1 + r_1 + 1}_{2^{\text{nd}}} - \underbrace{1 + r_1 + 1}_{3^{\text{rd}}} \text{ extraction and discipline} \\ & = r + 3r_1 \text{ red balls after } 3^{\text{rd}} \text{ extraction and discipline} \end{aligned}$$

You can vary the discipline.
 You can indicate the spread of a disease by saying that if the disease occurs (i.e., a black ball is extracted), you can make the disease more probable by adding more black balls.
 And you can study how the disease progresses.

The case we will study in class (I don't want to spoil it now) is:

You extract one ball
 You replace $c+1$ balls of the color extracted

That's the Pólya Urn Model.

So this is the sample space.


A succession of extractions, with a very complex way of computing the probabilities.

Let $X_2 = \begin{cases} 1 & \text{if the 2nd ball extracted is red} \\ 0 & \text{if " " " " " " black} \end{cases}$

By the Law of Alternatives:

$$P(X_2=1) = P(X_2=1|X_1=0)P(X_1=0) + P(X_2=1|X_1=1)P(X_1=1)$$

or from the probability tree:



$P(X_2=1) = \text{sum of these disjoint probabilities.}$

$$P(X_1=0) = \frac{b}{r+b}$$

$$P(X_2=1|X_1=0) = \frac{r+r_1'}{r+r_1'+b+b_1'}$$

$$P(X_1=1) = \frac{r}{r+b}$$

$$P(X_2=1|X_1=1) = \frac{r+r_1}{r+r_1+b+b_1}$$

In general computations of probabilities of these types do not simplify. But we will see that in the Polya Urn case, where we replace the extracted ball w/ $C+1$ of the same color, there is a simplification.

$$P(X_2=1) = \left(\frac{r+r_1}{r+r_1+b+b_1} \right) \left(\frac{b}{r+b} \right) + \left(\frac{r+r_1}{r+r_1+b+b_1} \right) \left(\frac{r}{r+b} \right)$$

This can be rewritten as:

$$P(X_2=1) = \frac{br(r+r_1)(r+r_1')}{(r+b)^2(r+r_1'+b+b_1')(r+r_1+b+b_1)}$$

In the case of urn models, where you extract and replace balls w/ various disciplines, the conditional probabilities allow you to answer basic questions.

An example where you study the sample space by conditional probability.

A few words about balls into boxes.

$\circ \circ \circ \dots \circ$ k balls
 $\_ \_ \_ \dots \_$ n boxes

We've covered:

	Maxwell-Boltzmann	Bose-Einstein	Fermi-Dirac
balls	distinguishable	indistinguishable	indistinguishable, occupation number 0 or 1
boxes	distinguishable	distinguishable	distinguishable

Let's consider the case where:

Balls - distinguishable
 Boxes - indistinguishable



A bunch of balls goes into a box.
 Another bunch of balls goes into a box.

Clearly unoccupied boxes make no difference - the boxes are indistinguishable. It makes no difference.

The count is which balls are in the same box.
 You might as well assume that all the boxes are occupied.

You don't count the number of configurations.

You count the number of configurations in which all the boxes are occupied.

What is this?

a partition

A partition is a set of non-empty sets of balls belonging to the same class.

Placing distinguishable balls into indistinguishable boxes is the same as partitioning the set of balls.

Because all you have for the data is which balls go together.

balls distinguishable
boxes indistinguishable
(all boxes occupied) } partitions of the set of balls

Let's consider another case:

Balls - indistinguishable
Boxes - indistinguishable } both indistinguishable



You might as well assume all the boxes are occupied.

That means, to every box, there is an integer associated with it (i.e., the number of indistinguishable balls that occupy the box).

But the boxes are indistinguishable.

That means you have the set of integers, which is unordered and which add up to K .

This corresponds to:

partition of the integer K into $\leq n$ summands

Ex: $K=5 \Rightarrow$ partitions (the order of the summands doesn't matter)

5
 $4+1$
 $3+2$
 $3+1+1$
 $2+2+1$
 $2+1+1+1$
 $1+1+1+1+1$

} $\Sigma = 5$

indistinguishable balls
indistinguishable boxes, all boxes covered } partition of an integer

What about the situation where the order of the summands matters?

Let's take partitions of k (say $k=5$) that are ordered into (say 3) parts.

unordered partitions
of size 3

$$3 + 1 + 1$$

$$2 + 2 + 1$$

Q/A: remember, all the boxes are occupied. You are not allowed 0.

if order mattered
then:

$$3 + 1 + 1$$

$$1 + 3 + 1$$

$$1 + 1 + 3$$

$$2 + 2 + 1$$

$$2 + 1 + 2$$

$$1 + 2 + 2$$

6

what does this correspond to "balls into box"-wise?

Bose-Einstein statistics requiring at least one ball in each box.

All occupation numbers > 0 .

$\left\{ \begin{array}{l} \text{balls indistinguishable} \\ \text{boxes indistinguishable,} \\ \text{all boxes occupied} \end{array} \right\} \xrightarrow{\text{followed by}} \left\{ \begin{array}{l} \text{balls indistinguishable} \\ \text{boxes distinguishable,} \\ \text{all boxes occupied} \end{array} \right\} = \text{partition of integer } k \text{ into ordered summands.}$

• Most general case - Balls into Boxes



given: partition π of the set B $\leftarrow \pi =$ blocks of B , which are disjoint, whose union is B .
 partition σ of the set U . $\leftarrow \sigma =$ likewise.

Informally:

Place the balls into the boxes.

Two balls belonging to the same block of the π partition are considered identical.

Two boxes belonging to the same block of the σ partition are considered identical.

How many ways are there of putting the balls into the boxes?

This is the most general case.

I know how to solve it, but I'm just too lazy to write it up.

There is a very simple formula for generating this.

If you want to work on this, you are welcome.

It's not hard.

It should be written up somewhere, but it isn't.

Balls in the same block are indistinguishable.
 Balls in different blocks are partially distinguishable.

This is the most general balls into boxes problem, if placing the balls into boxes is a function.

We've considered one case where placing the balls into boxes is not a function.
 Disposition is not a function

↑ where we place balls into boxes and the balls in the box are configured a certain way.

Then all hell breaks loose.

You have a tremendous variety of placements that have to adhere to certain configurations.
 You want to see an example?

One of the most famous examples is this:

22-141 50 SHEETS
 22-142 100 SHEETS
 22-144 200 SHEETS



Reluctant Functions (This is pretty sophisticated)



Informally:

Take the 1st ball. It has a choice. Either it goes into a box, or it goes to another ball.

If it goes to a ball, wherever that ball goes, the 1st ball goes too.

This 2nd ball carries on its back the 1st ball.

Then this ball now decides - either it goes to another ball, or it goes into a box.

And so on.


Once a ball goes to another ball or into a box, it is no longer available.

Eventually, the balls have to drop into a box.

In other words, if you follow the history of any ball, it eventually drops (together w/ any balls it has gone to) into a box.

Q: if ball 1 goes to ball 3 and then ball 2 goes to ball 3, does ball 2 have an option where it goes when it piggybacks on ball 3?

A: No. The balls are not ordered when they go to other balls.




The above illustrates the occupancy model.
What would be the distribution model?

You view the balls in a box.

This corresponds to the balls going into a box and being configured in a certain way in the box.

The balls in the box form a forest of rooted trees



forest from our example

↑ together, you can have several trees going into the same box.

There are a number of different ways of doing this that are very complicated.
But it's an extremely simple problem (if you can number it).
This is in a paper of mine.

The number of trees, we know.

I want the number of consequences of a reluctant function. There is a very simple formula.
I leave it to you as an exercise.

The total number is not that trivial, if you restrict occupied boxes to a single tree, rather than a rooted forest, then it becomes tough.

Remarkably enough, this does occur.

Put the balls in the boxes and configure them in a forest of rooted trees.

Disposition is a special case - the forest is restricted to one tree, one stump



Another special case that occurs very frequently is where the balls in each box are configured into several linearly ordered trees.

disjoint linear orders



← {Laguerre polynomials count the number of configurations for this special case of several linearly ordered trees.

We can go on forever with special cases of reluctant functions (i.e., restricting the configurations of forests to have certain properties).

The theory of Abel polynomials gives the formula for counting the number of configurations in the general case.

There's more to this than meets the eye.

You can make your living on this balls into boxes business.

Partitions

$\{ \dots \}$ a set of n elements

There are 2^n subsets

$\binom{n}{k}$ subsets with k elements

We can ask the same questions about partitions:

What is the total number of partitions of a set with n elements?

- and -

What is the total number of partitions with k blocks of a set with n elements?

The formula for the total number of partitions is very remarkable.

total number of partitions:

B_{n+1} = Bell number
= number of partitions of a set w/ n elements

$$B_{n+1} = \frac{1}{e} \left(\frac{1^n}{0!} + \frac{2^n}{1!} + \frac{3^n}{2!} + \dots \right)$$

← {This is in a brown book, which is a collection of his papers.}

this is an integer!

How do you see that this is an integer intuitively?

Dobinsky's formula

This is a very remarkable formula. It's not clear that this is an integer. And yet it gives the number of partitions of a set. (Graham, Knuth, Patashnik have a closed form (asymptotically) for this - Ex 9.46)

When I was young, one of my first papers was to give a very simple proof of this.


There are many other formulas for Bell numbers, but this is the only closed form formula.

We will see this later when we do the theory and the reasoning of the Poisson Process.

{ It's very important to be almost right. It's more important to be almost right than right. Why? Because, if you are almost right, then you know it's stable under perturbation. But if a small deviation makes a big difference, then your statement is unstable. }

total number of partitions with k blocks $\leftarrow$ (partition n things into k non-empty subsets)

This can be computed by inclusion-exclusion




You take all functions of Maxwell-Boltzmann where every box is occupied.
Then you make the boxes indistinguishable by dividing by $k!$.
So the number of ways of placing these balls into boxes is:

- 1) all the boxes are occupied
- 2) divide by $k!$ to make the boxes indistinguishable


This is equal to:

$S(n, k)$ - Stirling Numbers of the second kind

$S(n, k)$ - Stirling Numbers of the second kind



1) There are $S(n, k)$ ways to partition n balls into k blocks



} balls into blocks

2) Now the blocks become balls.
Now we place balls into boxes.



} blocks into boxes

$$(x)_k = x(x-1)(x-2)\dots(x-k+1)$$

Lower factorial gives number of ways of placing blocks into boxes

So we have:

$$x^n = \sum_{k=0}^n S(n, k) (x)_k$$

Recall the difference of a function Δ :

$$\Delta f(x) = f(x+1) - f(x)$$

$$\Delta(x)_k = (x+1)_k - (x)_k$$

$$= \underbrace{(x+1)(x)\dots(x-k+2)}_{(x)_{k-1}} - \underbrace{(x)(x-1)\dots(x-k+2)(x-k+1)}_{(x)_{k-1}}$$

$$= [(x+1) - (x-k+1)] (x)_{k-1}$$

$$\Delta(x)_k = k(x)_{k-1}$$

From $x^n = \sum_{k=0}^n S(n,k) (x)_k$, we have:

$$\Delta x^n = \sum_{k=0}^n S(n,k) k (x)_{k-1}$$

Multiple applications of Δ gives:

$$\Delta^j x^n = \sum_{k=0}^n S(n,k) (k)_j (x)_{k-j}$$

From this, we can show that:

$$[\Delta^k x^n]_{x=0} = k! S(n,k) \implies S(n,k) = \frac{[\Delta^k x^n]_{x=0}}{k!}$$

Where do Stirling numbers of the second kind come up?

All over the place. For example:

Recall the differential operator D , which is $\frac{d}{dx}$. Then we can show that:

$$(xD)^n = \sum_{j=0}^n S(n,j) x^j D^j$$

Stirling number of the second kind

Combinatorial Identities
John Riordan
Wiley + Sons 1968 p. 218

Proof by induction.

Basis: $(xD)^0 = 1$

Inductive step:

Assume $(xD)^n = \sum_{j=0}^n S(n,j) x^j D^j$

Then:

$$\begin{aligned} (xD)^{n+1} &= (xD)(xD)^n \\ &= (xD) \sum_{j=0}^n S(n,j) x^j D^j \\ &= \sum_{j=0}^n S(n,j) (xD) x^j D^j \end{aligned}$$

Note that:

$$\begin{aligned} (xD)(xDf(x)) &= x(x'Df(x) + xD^2f(x)) \\ &= xDf(x) + x^2D^2f(x) \\ (xD)^2 &= xD + x^2D^2 \end{aligned}$$

applying the differential operator gives:

SC 3/13/18, 13

$$= \sum_{j=0}^n S(n, j) x \left[j x^{j-1} D^j + x^j D^{j+1} \right]$$

$$= \sum_{j=0}^n S(n, j) \left[j x^j D^j + x^{j+1} D^{j+1} \right]$$

$$= \sum_{j=0}^n \left(j S(n, j) x^j D^j + \underbrace{S(n, j) x^{j+1} D^{j+1}}_{S(n, j-1) x^j D^j} \right)$$

$S(n, j-1) x^j D^j \leftarrow$ change of variable
 $j \rightarrow j-1$
 $j+1 \rightarrow j$

$$= \sum_{j=0}^n \underbrace{\left[j S(n, j) + S(n, j-1) \right]}_{S(n+1, j)} x^j D^j$$

$S(n+1, j) \leftarrow$ identity for Stirling numbers of the second kind.

$$(x D)^{n+1} = \sum_{j=0}^n S(n+1, j) x^j D^j \quad \checkmark$$

Note that:

$$\begin{aligned} (x D) x^k &= x k x^{k-1} \\ &= k x^k \end{aligned}$$

$$\begin{aligned} (x D)^2 x^k &= (x D) (x D) x^k \\ &= (x D) k x^k \\ &= x k k x^{k-1} \\ &= k^2 x^k \end{aligned}$$

And we have:

$$(*) \quad (x D)^n x^k = k^n x^k$$

as well as:

$$(**) \quad x^j D^j x^k = (k)_j x^{k-j}$$

$$\begin{cases} x^j (D^j x^k) = x^j (k)_j x^{k-j} \\ = (k)_j x^k \end{cases}$$



Taking the original identity:

$$(xD)^n = \sum_{j=0}^n S(n,j) x^j D^j$$

multiply both sides by x^k :

$$\underbrace{(xD)^n x^k}_{\text{from } (*), \quad k^n x^k} = \sum_{j=0}^n S(n,j) \underbrace{x^j D^j x^k}_{\text{from } (**), \quad (k)_j x^k}$$

Which gives the combinatorial identity involving the Stirling numbers of the second kind:

$$k^n x^k = \sum_{j=0}^n S(n,j) (k)_j x^k$$


This is an analytic derivation.

All the formulas of { random variables,
expectation,
probability }

can be obtained as purely probabilistic formulas, in addition to the analytic derivations.


• Needles on a stick (again)

A one liner.




We have the interval $[0, a]$.

I take another interval of length $a-nh$.



On this other interval, I pick n points at random.

Then, for each point, I take the interval $[0, a]$ and select the same points.



What is the number of ways this can be done?

$$\frac{(a-nh)^n}{a^n}$$

$\leftarrow \left\{ \begin{array}{l} \text{number of ways of dropping } n \text{ points} \\ \text{on the interval } [0, a-nh] \end{array} \right\}$
 $\leftarrow \left\{ \begin{array}{l} \text{number of ways of dropping } n \text{ points} \\ \text{on the interval } [0, a] \end{array} \right\}$

Note: I said number of ways when, in fact, it's a measure.

I wish we had time to work this in 2- dimensions.

• Geometric Probability


From the text:

"The theory of geometric probability is concerned w/ random points on the plane, spheres, and other geometric objects, as well as with random lines, circles, and other figures."

Let's start w/ the most famous problem of Geometric Probability, which is the Buffon Needle Problem.

Buffon Needle Problem

We draw an infinite number of straight lines on the plane, parallel and distance d apart.



You have a needle of length l .
You take the needle and drop it on the plane at random.

What is the probability that the needle will intersect at least one of the lines?

This can be solved many ways.

It can be solved by the continuous Law of Alternatives trigonometrically:

{ You drop the center of the needle,
then you rotate the needle.
And you compute the probability of intersection by the continuous Law of Alternatives.

But now, let's give a completely probabilistic solution, using no computations.
This is a test of your understanding of probability.
It's almost pure handwaving.

Let $A_l =$ event that needle meets at least one of the straight lines.

$X_l =$ number of straight lines the needle meets when dropped at random.

↑ random variable

Now, we consider the expectation of X_l :

$$E(X_l) = \sum_n n P(X_l = n) \\ = 0 \cdot P(X_l = 0) + 1 \cdot P(X_l = 1) + 2 \cdot P(X_l = 2) + \dots$$

if the needle is short ($l < d$):

In this case, the needle can meet at most 1 line.

$$E(X_l) = 0 \cdot P(X_l = 0) + 1 \cdot P(X_l = 1) \\ = P(X_l = 1) \leftarrow (X_l = 1) \text{ is just the event } A_l \\ = P(A_l)$$

Therefore, by this reasoning, we have reduced the computation of the probability of an intersection of a straight line to the computation of the expectation.

This is a very frequently used technique in probability.

Expectations are easier to compute than probabilities.

Why?


Because the expectation of the sum of two random variables is the sum of the expectations. No questions asked. [2/25/98.8-9]

How do we exploit this fact?

Like this.

Take the original needle of length l .

Then tack on an additional parallel needle of length l'



Now, I claim that the number of times that the needle of length $l+l'$ meets a straight line is the number of times the first segment meets a straight line, plus the number of times the second segment meets a straight line. If you don't see that, I can not explain it.

$$X_{l+l'} = X_l + X_{l'}$$

Therefore:

$$E(X_{l+l'}) = E(X_l) + E(X_{l'})$$

Now, just for psychological reasons, let's set:

$$f(l) = E(X_l)$$

which gives:

$$f(l+l') = f(l) + f(l')$$

Now we ask, which kinds of functions satisfy this condition?

Constant times l .

Therefore,

$$f(l) = cl$$

$$f(l+l') = f(l) + f(l')$$

$$c(l+l') = cl + cl'$$

universal constant c ,
independent of l ,
but depending on d (the distance between straight lines).

$$E(X_L) = cL$$

Now we use the Charles Lamb principle.

He wrote a short story - The Discovery of Roast Pork.

How was roast pork discovered?

Very simply.

One day, in a hamlet in England, a house accidentally burned down.

There were pigs in it. They were roasted.

Someone accidentally tasted the meat and liked it.

It was fantastic. It was terrific.

So what did they do?

They started stuffing pigs into houses and burning the houses down.
They liked the meat.

Until one day, a great genius came along and said:

You don't have to put them into the houses.


This is how roast pork was invented.

This is the
method of
scientific
discovery

We're going to go through this crisis right now.

Now I say, who says the needles being combined need to be parallel to each other.

They could be crooked and the same conclusion holds



↑ that's the roast pork part.

(provided the needles are rigid, of course)

The above argument works just the same for this situation.


Exactly the same conclusion.

Because if you took a crooked needle of length $l + l'$, the number of times this crooked needle meets any of the straight lines on the plane is the same as the number of times it meets on the stretch of length l + the number of times it meets on the stretch of length l' .

We still have:

$$X_{l+l'} = X_l + X_{l'}$$

- Proceeding along this line, we can do something like this:



$$X_{l_1 + l_2 + \dots + l_n} = X_{l_1} + X_{l_2} + \dots + X_{l_n}$$

needle

And it's still true.

$$E(X_l) = cl$$

- You can pass to the limit:



Any rigid thing like this, you drop it.
The expected number of straight lines it meets is proportional to its length.
For any curve, whatsoever



$$E(X_l) \propto \overbrace{\text{Length of curve}}^l \\ = cl$$

The universal constant c in all the above, is the same.
It's not the probability, but the expectation.

$$E(X_l) = cl \quad \text{for any curve, whatsoever.}$$

- What we are missing is the value of the constant c .

We need to drop the most favorable shape that will reveal the constant c .
So we drop the needle (i.e., curve) of that shape.
So I take this needle.



No matter how you drop each rigid needle (circle), it always intersects the straight lines twice.

$$\text{So } E(X_{\pi d}) = 2$$

$$E(X_{\pi d}) = 2$$

$$= c(\pi d) \quad \leftarrow \text{since } E(X_l) = cl$$

$$\Rightarrow \boxed{c = \frac{2}{\pi d}}$$

And we have the universal constant.

So, for any curve whatsoever,

$$(*) \quad \boxed{E(X_l) = \frac{2}{\pi d} l}$$

Now, we go back to the straight curve, where we have a short needle ($l < d$). We had shown [SC 4/10/98.2]:

short
straight
needle

$$E(X_l) = P(X_l = 1) = P(A_l)$$

And we have formula (*) above, compliments of the universal constant:

$$E(X_l) = \frac{2}{\pi d} l$$

Combining the above, we have that:

$$\underbrace{P(A_l) = \frac{2}{\pi d} l}$$


This is an experimental method for computing π .
You drop needles.

[Q: What happens w/ longer straight needles that can intersect more than one straight line in the plane?
A: I've never worked it out.]



This argument can be adapted to give a much more general result.

A convex set in the plane is a set that has no bumps.



Convex



NOT convex

(points on this piece of segment $\overline{xy}$ $\notin B$)

Convex is an elegant way of saying no bumps.

A nice way of saying:


given any two points x and y in the set,
if you join these two points by a straight line,
the entire segment is included in the set.

} Convex set

This is just an elegant way of excluding bumps.

Sylvester's Theorem

Example 1: You take 2 convex sets A and B in the plane, one inside the other.
Big set B , small set A .



$$A \subseteq B$$

Notes for this example, sets $A+B$ need not be convex

Suppose I pick a point arbitrarily in B .
Then I ask, given the condition that I've picked a point in B ,
what is the probability that the point also belongs to A ?

The answer is obvious, It's the ratio of the area of A divided by the area of B .

$$P(x \in A | x \in B) = \frac{P((x \in A) \cap (x \in B))}{P(x \in B)}$$

$$= \frac{P(x \in A)}{P(x \in B)}$$

$$P(x \in A | x \in B) = \frac{\text{area}(A)}{\text{area}(B)} \quad \left. \begin{array}{l} \text{see commentary on} \\ \text{[SC 4/10/98.9-10]} \end{array} \right\}$$

We implicitly use the concept of a measure here. Measures can provide conditional probabilities but not probability of basic events. e.g., not $P(x \in A)$



Now, what's the next best thing in the world after a point?

A line.


Now, instead of picking a point at random, I'm picking a line at random.
I ask the same question.

Note - for this example, the sets $A+B$ must be convex.

For 2 convex sets A, B , where $A \subseteq B$,

I pick a line at random.


Assuming that the line intersects B , what is the probability that it also intersects A ?



This can also be done by pure handwaving.


- Let's go back and analyze what we really did in the Button Needle Problem.
What is really going on:

- ① We draw straight lines
- ② Then dropped a needle




But this is statistically
the same as:

- ① drawing the needle
- ② Then dropping the straight lines



From this analysis,
let's consider the case ($l < d$), i.e., where the needle is of sufficiently small size.



event $A_2 =$ needle meets at least one straight line
this is the same as saying:
event that a straight line intersects the needle

So,

$P(\text{a straight line intersects the needle}) \propto l \leftarrow \text{length of the needle}$

$$P(\text{a straight line intersects the needle}) = P(A_1) = \frac{2}{\pi d} l \propto l$$

- How do we get back to Sylvester's problem?
Let's ask the following questions:

What does it mean to pick a random line?

Let's go back and see what we did in Example 1, where we picked a random point in the plane.

How did we interpret the process of picking a point at random, in terms of picking points?

We said that we pick a point in the plane.

That's not a uniform probability.

It's a measure, called the area, which is infinite.

So we implicitly use the concept of a measure when picking the points.

Then we took the ratio of two measures, and that gave us the probability.
So we need the notion of a measure.

Measure

What's a measure?

A measure is something that has all the properties of probability, except that unlike probabilities which lie between $[0, 1]$, measures lie between $[0, \infty)$.

measure μ - Properties

- $\mu(A) \geq 0$
- $\mu(A \cup B) = \mu(A) + \mu(B) - \mu(A \cap B)$
- $\mu(\emptyset) = 0$

But $\mu(\Omega) \neq 1 \leftarrow \text{unlike probability, where } P(\Omega) = 1$

- Countable additivity $A_1, A_2, \dots$ is a finite or infinite sequence of pairwise disjoint sets
 $\mu(\bigcup_n A_n) = \sum_n \mu(A_n)$

example: area(A) is a measure in the plane

example: length(A) is a measure on the line

You get probabilities from measures by taking conditional probabilities.

If you have a measure, instead of probability, you can only talk about conditional probability.

You can not talk about the probability of a basic event (i.e., $P(A)$).

This is a little tricky, but that's the way it is.

Let me tell you a little secret.

One of the great unsolved problems of statistics is the fact that in Bayes' Law, you can use, as the prior, measures instead of probabilities.

$$P(A|B) = \frac{P(B|A) \cancel{P(A)}}{P(B)} \xleftarrow{\text{prior}} \mu(A) \quad \triangleq \quad \frac{\mu(A)}{\mu(B)}$$

Conditional Probability in terms of measures

For example, instead of tossing a coin (i.e., $P(A) = p$), we have a random variable on the whole straight line and we estimate Bayes by taking measurements.

If you can measure the whole straight line, you have to take, as prior, the measure on the whole line, because it can be anywhere.
Without taking probabilities.

People have realized this since 1939, when a book by Sir Harold Jeffries appeared, trying to justify why this is.

Why can you take as prior a measure instead of a probability?

No one has come up w/ a suitable justification.

Do you want to become famous? Justify this.

So measures occur as priors in Bayes' Law.

Q: What utility is there in taking the measure as a prior?

A: You want a uniform prior.

You are estimating a random variable (e.g., $P(X=n)$). Suppose you don't know the whole distribution.

You can use measures.

Note that the posterior is a probability: $P(A|B) = \frac{\mu(A)}{\mu(B)}$
even if the prior is a measure.

Back to our re-examination of Example 1.

In answering the question about picking points in A and B, we implicitly used the fact that the right probability to use is the ratio of two measures.

$$P(x \in A | x \in B) = \frac{\text{area}(A)}{\text{area}(B)} \xleftarrow{\text{prior}} \frac{\mu(A)}{\mu(B)}$$



Now, we have the same problem w/ lines.

Now, we have to visualize the sample space Ω as the set of all straight lines in the plane.

$\Omega =$ all straight lines in the plane

To answer the question posed in Example 2, we need a measure.

But this measure is not any old measure.

This measure has to be independent of position, or invariant of position, as they say.

And it's remarkable that there is one, and only one, measure on a set of straight lines, which is just like the analogue of area. It's easier than you think.

Is there an invariant measure on the set of all straight lines in the plane?

Yes. There is.

And it is unique, except for normalization factors.

Why?

it depends on what you call measure = 1
With area, for example, it's up to you
what you call area = 1.

We use exactly the same reasoning as we did in the case of the Buffon Needle Problem.

Now our sample space Ω , remember, is the set of all straight lines.

A sample point $\omega \in \Omega$ is a straight line.



We've already picked our needle, so we picked a segment of length l .



event straight line meets l



event straight line does not meet l

Now we consider the event that a straight line meets l .

What is the measure of that event?

First, we note that if we tack on l' , $l' \parallel l$, then the measure is the sum.



Let $B_l =$ event that a straight line meets a segment of length l

Then we take as the measure:

$$\mu(B_l) = c l \leftarrow \text{length}$$

$$\uparrow \left\{ \begin{array}{l} \text{the universal constant } \frac{2}{\pi d} \\ [4/10/98.5-6] \end{array} \right\}$$

Every straight line either meets l , or it doesn't meet.
And note that this measure is independent of the position of l .

So this gives you the measure of several events, namely, the event that a straight line intersects l .

That's enough. (There is another unsolved problem here, which I'll tell you about later).^(*)

All you need is that. You don't want the formula for the measure of a set of straight lines.
All you need is the measure of a set of straight lines that meet a given segment l .
this is the important part

So we have the measure $\mu(B_l) = c l$.

Now, we use exactly the same argument as before, w/ the Buffon Needle Problem.



If I take any shape whatsoever, with length $(s) = l$,
then define random variables:

$$X_l(\omega) = \begin{cases} 1 & \text{if straight line } \omega \text{ meets shape } S \\ 0 & \text{otherwise} \end{cases}$$


Then $X_l(\omega)$ is the indicator random variable where a straight line meets the shape.

And using the same argument presented in [SC 4/10/98.2-6], we have the expectation:

$$E(X_l) = c \cdot l$$


average of times the shape S , of length l , is intersected.

Back to example 2 [SC 4/10/98.8], where I have two convex sets A and B , with $A \subseteq B$.



Since I have a convex set, a line that intersects B intersects A either twice (case 2) or not at all (case 1).

We ignore the extreme case of Probability = 0 that the line intersects B at a single point




Therefore, the probability that the straight line that meets B also meets A equals:

$$P(\text{line meets A} \mid \text{line meets B}) = \frac{\mu(B_A)}{\mu(B_B)}$$

$$= \frac{\chi \cdot \text{perimeter}(A)}{\chi \cdot \text{perimeter}(B)}$$

We take, as measure of a convex set, twice its perimeter. Note that we require convex sets to ensure that any line intersects twice or not at all (ignoring, again, the extreme case w/ probability = 0, of a single intersection).

For example, we want to avoid the pathological case:



$$P(\text{line meets A} \mid \text{line meets B}) = \frac{\text{perimeter}(A)}{\text{perimeter}(B)}$$

This gives you a beautiful way of computing perimeters.

This gives you a practical way of computing perimeters of weird convex sets. You take the average number of lines that meet it, within a known convex set B.

$$\text{perimeter}(A) = P(\text{line meets A} \mid \text{line meets B}) \text{perimeter}(B)$$

This is the beginning of tomography (tomography is done in 3-dimensions). An object has a smaller object inside, the probability that a straight line that meets the outer object will also meet the inner object is the ratio of their surface areas. By the same argument really.

You can read my recently published book on the subject, published in 1950.

- The famous unsolved problem I referred to (*):
It's been around quite a while.

Ω = all straight lines in the plane

Let's define:


$$I_A(\omega) = \begin{cases} 1 & \text{if line } \omega \text{ meets } A \\ 0 & \text{otherwise} \end{cases}$$

For indicator random variables of events A and B , we've shown that (Text p. 90-91):

$$I_A + I_B = I_{A \cup B} + I_{A \cap B} \quad \left. \begin{array}{l} \text{ } \\ \text{events} \end{array} \right\} \text{ due to Boolean algebra}$$

For the indicator random variables associated with lines, you have other identities, besides this one.

When you are dealing with lines, you have additional identities that come in because of geometry.




No one has been able to describe all these fundamental geometric identities.

✓ This problem is formally described in my book.

The algebra of indicator random variables for lines is something more than Boolean algebra, due to the geometry.

Entropy and Information

One of the most beautiful chapters in probability.
It's too bad I can't cover it in the lectures, there's only so much one can do.



Suppose we have sample space Ω and an integer random variable X .

As we briefly mentioned [SC 2/27/98, 10-11], you can have the situation where the sample space is partitioned into blocks.

The events $(X=n)$, n an integer with $P(X=n) > 0$, are a partition of the sample space Ω if they disjointly cover Ω $\uparrow$ strictly

This gives us another way of looking at random variables.

We view a random variable as a question.

A bad guy has chosen a sample point ω .



We try to determine the unknown sample point by asking questions.

The answers to our questions will tell us which block the chosen sample point belongs in.

The more questions we ask, the more we narrow down where the element, chosen by the bad guys, is.

Thus, a random variable is a question that delivers a certain amount of information on the search problems of finding particular sample points.

Entropy is a numerical measurement on the amount of information that a random variable has.

Let's see what the properties of such information ought to be.

Entropy is a numerical measure of the amount of information in a partition π of Ω

{ At this point, we will drop reference to random variable and replace it w/ partition.
Because all that matters is the partition. }

First, I give you the formula:

For $\pi = \{B : B \in \pi\}$,
 $\uparrow$ Block

$$H[\pi] \triangleq \sum_{B \in \pi} P(B) \log_2 \frac{1}{P(B)}$$

entropy of a partition $\uparrow$ by definition

{ any block has positive probability ($P(B) > 0, B \in \pi$);
intersection of any 2 blocks has probability 0
 $P(B \cap B') = 0, B, B' \in \pi, B \neq B'$;
probability of union of all the blocks is 1
 $P(\bigcup_{B \in \pi} B) = 1$ }

Q: Why $H[\pi]$?

A: Because Boltzmann called it H.


This formula is sculpted on Boltzmann's grave.
I'm not kidding.

Now you say: This is weirdissimo. Where did you get this?

The purpose of this lecture is to show you this is not weirdissimo. It's absolutely natural.

The units of entropy are called bits $\leftarrow$ after Prof. Shannon
bits = binary digits

Example



Take the partition of two blocks with:

$$\pi_0 = \{B, C\}, \quad P(B) = \frac{1}{2} \Rightarrow P(C) = \frac{1}{2}$$

The entropy of this partition is:


$$\begin{aligned} H[\pi_0] &= P(B) \log_2 \frac{1}{P(B)} + P(C) \log_2 \frac{1}{P(C)} \\ &= \frac{1}{2} \log_2 \frac{1}{(\frac{1}{2})} + \frac{1}{2} \log_2 \frac{1}{(\frac{1}{2})} \end{aligned}$$

$H[\pi_0] = 1 \text{ bit}$ } Therefore, a partition of 2 blocks w/ equal probability has 1 bit of information.

The simplest question is answered Yes or No w/ probability $\frac{1}{2}$.

Now you say, what if we had two blocks w/ different probabilities?

Example



$$\pi_1 = \{B, C\} \text{ with } P(B) = \epsilon \Rightarrow P(C) = 1 - \epsilon$$

$$H[\pi_1] = \epsilon \log_2 \frac{1}{\epsilon} + (1 - \epsilon) \log_2 \frac{1}{(1 - \epsilon)}$$

What happens as $\epsilon \rightarrow 0$?

$$\lim_{\epsilon \rightarrow 0} H[\pi_1] = \cancel{\epsilon^0 \log_2 \frac{1}{\epsilon}} + (1 - \epsilon) \log_2 \frac{1}{(1 - \epsilon)} \rightarrow 0$$

$$\lim_{\epsilon \rightarrow 0} H[\pi_1] = 0$$

In other words, a partition w/ one teeny weeny block and one big block has very small entropy.

In particular, if you take the partition w/ a single block of probability 1,

$$\pi = B, P(B) = 1$$

$$H[\pi] = P(B) \log_2 \frac{1}{P(B)}$$

$$= 1 \log_2 1$$

$$= 0 \text{ bits}$$

the entropy is 0.

It doesn't give you any information as to where the unknown sample point is.

Now we come to the fundamental properties of entropy:

Property (1) if π and π_2 are independent partitions then

{independent = every block of π independent of π_2 }


$$H[\pi \wedge \pi_2] = H[\pi] + H[\pi_2]$$

↑ meet

Meet for partition is the analogue of intersection for sets.

What is $\pi \wedge \pi_2$?


It's the partition obtained by intersecting every block of π w/ every block of π_2 .



π



π_2



$\pi \wedge \pi_2$

Meet gives you a partition that is finer.

- Let's verify this.
This is the crucial computation of this theorem.

Suppose $\pi = \{B_i\}$, $\pi_2 = \{C_j\}$

Then

$$H[\pi \wedge \pi_2] = \sum_{B, C} P(B \cap C) \log_2 \frac{1}{P(B \cap C)} \quad \left. \vphantom{\sum_{B, C}} \right\} \text{This is for the meet of any two partitions.}$$

sum over all
B's and C's

- The blocks of the meet $\pi \wedge \pi_2$ are the intersections of the blocks.
When considering all pairs of blocks in the cartesian product $B \times C$, blocks that do not intersect have $B \cap C = \emptyset$. In which case:

$$P(B \cap C) \log_2 \frac{1}{P(B \cap C)} = 0 \log_2 \frac{1}{0} = 0, \text{ by convention.}$$

- Since the partitions are independent, then blocks B and C are independent.

$$P(B \cap C) = P(B) P(C)$$

$$= \sum_{B, C} P(B) P(C) \log_2 \frac{1}{P(B) P(C)}$$

$$= \sum_{B, C} P(B) P(C) \left(\log_2 \frac{1}{P(B)} + \log_2 \frac{1}{P(C)} \right)$$

split into 2 sums:

$$= \sum_{C \in \pi_2} P(C) \underbrace{\sum_{B \in \pi} P(B) \log_2 \frac{1}{P(B)}}_{H[\pi]} \left. \vphantom{\sum_{C \in \pi_2}} \right\} \leftarrow \sum_{B, C} P(B) P(C) \log_2 \frac{1}{P(B)}$$

$$+ \sum_{B \in \pi} P(B) \underbrace{\sum_{C \in \pi_2} P(C) \log_2 \frac{1}{P(C)}}_{H[\pi_2]} \left. \vphantom{\sum_{B \in \pi}} \right\} \leftarrow \sum_{B, C} P(B) P(C) \log_2 \frac{1}{P(C)}$$

$$= \sum_{C \in \pi_2} \overbrace{P(C)}^1 H[\pi] + \sum_{B \in \pi} \overbrace{P(B)}^1 H[\pi_2]$$



$$H[\pi \wedge \pi_2] = H[\pi] + H[\pi_2] \quad \checkmark$$

Q/A: This has to be a partition w/ a countable number of blocks.

It can not be a continuous partition.


This is one case where the integral is not the limit of sums.
Very weird.

Recall there is a partial order on partitions.
The partial order of finer:

$\pi \leq \pi_3$ means π is finer than π_3

(i.e., every block of π is contained in some block of π_3)

[SC 2/27/98.11] (Take the blocks of π_3 and cut them up to get π .)



~~~~~ =  $\pi$   
—— =  $\pi_3$

Property (2) if  $\pi \leq \pi_3$

$$H[\pi] \geq H[\pi_3]$$

The finer the partition is, the bigger the entropy.

If you take a finite sample space, the maximum entropy is for the partition where each sample point is a block (i.e., the finest partition).

Property (3)  $H[\pi \wedge \pi_2] \leq H[\pi] + H[\pi_2]$

This is true in general, even if  $\pi$  and  $\pi_2$  are not independent.  
You have equality only if  $\pi$  and  $\pi_2$  are independent.

• These 3 properties characterize entropy completely.

Q: Is there a property like inclusion-exclusion, something like that for entropy?

A: I wish we did.

People have tried. Very hard.

That's a very good question.

# Maximum Entropy Principle

There is a discrete version, which we are about to see.

There is a continuous version, which is much deeper.

The discrete version is this: (i.e., fix the number of blocks)

Among all partitions with  $n$  blocks, the ones with maximum entropy are the ones for which all blocks have probability  $\frac{1}{n}$ .

This can be proved and if we have time I'll prove it later.

$$P(B) = \frac{1}{n}, \forall B \in \pi$$

all blocks have equal probability

There is a deep interpretation of the maximum entropy principle.

If you don't know anything about the distribution of the random variable that takes  $n$  values, then you assume that all values are equally likely.

You assume maximum entropy.

$$B_i = (X=i), \text{ for } n \text{ values of } i \quad \pi = \{B_i\} = \{(X=i)\}$$

block integer random variable

$$P(B_i) = \frac{1}{n} \Rightarrow P(X=i) = \frac{1}{n}$$


maximum entropy

↑  
for all non-null events  $(X=i)$

Let  $X$  = integer random variable

$\pi$  is the partition defined by  $X$

$$(\text{i.e., } \pi = \{(X=i), \forall i \text{ s.t. } (X=i) \neq \emptyset\})$$



Define  $H(X) = H[\pi]$

People like to talk about entropy of random variables, even though entropy is a property of partitions.

Similarly, we have the joint entropy of 2 integer random variables  $X$  and  $Y$ :

$$H(X+Y) = H[\pi \wedge \pi_Y]$$

partition of  $X$  partition of  $Y$



So, the random variable maximum entropy is the one of the uniform distribution. Provided the random variable takes  $n$  values.

- Now we play 20 questions.  
I think of a number between 0 and  $2^n - 1$ .  
And you have to guess it.  
The question discipline is that you can only ask Yes or No questions.  
(Usually the questions are restricted).

What is the minimum number of questions?

- First, we'll do it abstractly, even though we know what the answer is.

We can write the number in binary digits.

0110 ... 110  
n binary digits

The number has at most  $n$  binary digits (number  $\in [0, 2^n - 1]$ ).

So you can guess the number in  $n$  questions.


1) Is the 1<sup>st</sup> digit 1? A: No  $\rightarrow 0$   
Yes  $\rightarrow 1$

2) Is the 2<sup>nd</sup> digit 1?

n) Is the  $n$ <sup>th</sup> digit 1?

$\Rightarrow$  n digits  
0110 ... 110  
 $\uparrow$   $\uparrow$   $\uparrow$   
 $A_n = \text{No}$  ...  $A_2 = \text{Yes}$   $A_1 = \text{No}$

- Now let's do it over, jazzed up.  
We repeat this in purely probabilistic terms.



We have a sample space of size  $2^n$ , with sample points being the numbers 0 through  $2^n - 1$ . (i.e.,  $n$  information)

Since we have no prior, each sample point has probability:

$$P(\omega) = \frac{1}{2^n}$$

A question is an arbitrary random variable over the number of entries.  
The answer, with the unknown point  $\omega$ , tells you which block the unknown point is in.

What does it mean to find the point w/ questions?

Well, you have questions  $\pi_1, \pi_2, \dots, \pi_k$  (i.e., partitions)

To find a point means that the meet of these partitions is a trivial partition with one block with one point.

That's what it means to find a point.

So the problem of search is to find the minimum number of partitions whose meet is the trivial partition.

Let's call the trivial partition  $\hat{O}$ .

find min  $k$  such that:

$$\pi_1 \wedge \pi_2 \wedge \dots \wedge \pi_k = \hat{O} \quad \text{read "zero hat"}$$

But, we have that:

$$H[\hat{O}] = H[\pi_1 \wedge \pi_2 \wedge \dots \wedge \pi_k]$$

and from the inequality from property 3 above:

$$\leq H[\pi_1] + H[\pi_2] + \dots + H[\pi_k]$$

Let's calculate the entropy of the trivial partition.

$$\begin{aligned} H[\hat{O}] &= \sum_{B \in \hat{O}} P(B) \log_2 \frac{1}{P(B)} \\ &= \sum_{\omega} P(\omega) \log_2 \frac{1}{P(\omega)} \\ &= 2^n \left( \frac{1}{2^n} \right) \log_2 \frac{1}{\left( \frac{1}{2^n} \right)} \end{aligned}$$

$$\boxed{H[\hat{O}] = n}$$

And, therefore:

$$n \leq H[\pi_1] + H[\pi_2] + \dots + H[\pi_k]$$

So whatever question search scheme you devise, it has to consist of partitions, or random variables, whose entropies add up to at least  $n$ .

This is a constraint.

You can not guess the point unless the entropy adds up to  $n$ .

Preferably, you want to make the questions (partitions) independent, because you save entropy:

$$n = H[\hat{\sigma}] = H[\pi_1 \wedge \pi_2 \wedge \dots \wedge \pi_k] = H[\pi_1] + H[\pi_2] + \dots + H[\pi_k]$$

↑  
equality only  
when  $\pi_1, \dots, \pi_k$   
all independent  
(by property 1)

RHS is min when  $\pi_1, \dots, \pi_k$  all independent

- Viewed in our first abstraction, each digit of the number is an independent random variable.

$$\begin{array}{ccccccc} 0 & 1 & 1 & \dots & 1 & 1 & 0 \\ \uparrow & & & & \uparrow & \uparrow & \\ A_n & & & & A_2 & A_1 & \end{array}$$

Q/A: There is a concept of conditional entropy, which we will not cover.

- Philosophically, entropy : partitions :: probability : sets  
Partitions are much tougher than sets.


This stuff takes a whole term in Course 6.

I had some notes written, but I've lost them 😞

- Famous Problem of the Scales

You have  $3^n$  coins, all of them gold, but one of them is fake

↑ it's lighter



scale

You are allowed the following experiment:

That is a question { You take 2 subsets of k coins and you put them on the scale.  
The scale either doesn't tilt, or it tilts to the right, or it tilts to the left.

What is the minimum number of weighings needed to determine the fake coin?



*(we could change this to any number  
and modify things accordingly)*

SC 4/24/98.10

The sample space  $\Omega$  has  $3^n$  sample points with:

$$P(\omega) = \frac{1}{3^n}$$

The entropy of the trivial partition  $\hat{\sigma}$  is:

$$H[\hat{\sigma}] = \sum_{B \in \hat{\sigma}} P(B) \log_2 \frac{1}{P(B)}$$

$$= \sum_{B \in \hat{\sigma}} \left( \frac{1}{3^n} \right) \log_2 \left( \frac{1}{\left( \frac{1}{3^n} \right)} \right)$$

$$= \log_2 3^n$$

$$= (\log_2 3) n$$

universal constant  $c_u$

$$H[\hat{\sigma}] = c_u \cdot n$$


That means the number of weighings you need, where each weighing gives you some entropy, must have a total entropy adding up to  $c_u \cdot n$ , otherwise you can not determine the fake coin.

(Identifying the fake coin is equivalent to realizing the trivial partition)

So what's a weighing?

Weighing


A weighing is a partition w/ 3 blocks:



$\Omega$

2 blocks have the same number of elements

The coins in these 2 blocks are placed on the plates of the scale:



In each weighing, you know in which of the 3 blocks the fake coin is.

$\Rightarrow$  { If the scale tilts, then you know which block of  $B_1, B_2$  the fake coin is in.  
If the scale doesn't tilt, the fake coin is in block  $B_3$

Every weighing gives you some information as to where the fake coin is.

- By the Maximum Entropy Principle [sc 4/24/98.6], we are best off if we split the coins into 3 equal piles. {abuse of notation:  $B_i$  = event that fake coin is in block  $B_i$ }

$$P(B_i) = \frac{1}{3}, \quad i = 1, 2, 3$$

and the entropy of this partition  $\pi = \{B_1, B_2, B_3\}$  is:

$$H[\pi] = \sum_{B \in \pi} P(B) \log_2 \frac{1}{P(B)}$$

$$= \sum_{i=1}^3 \left( \frac{1}{3} \right) \log_2 \frac{1}{\left( \frac{1}{3} \right)}$$

$$= (\log_2 3) \cdot 1$$

universal constant  $c_u$

$$H[\pi] = c_u \cdot 1$$

This is the entropy of a single weighing.

- Since we have that the entropy of the trivial partition is:

$$H[\delta] = c_u \cdot n$$

then we need at least  $n$  weighings, each w/ entropy  $H[\pi] = c_u \cdot 1$ , to determine which is the fake coin.

$$H[\delta] = c_u \cdot n \iff n H[\pi] = n (c_u \cdot 1)$$

You can not determine the fake coin in less than  $n$  weighings. That's what information theory tells you.

- Now that we know the min number of weighings, let's devise a way of doing it. How do you do this in practice? Easy.

Label each coin in the ternary system:

| Coin   | Label    |
|--------|----------|
| 0      | 0        |
| 1      | 1        |
| 2      | 2        |
| 3      | 10       |
| 4      | 11       |
| 5      | 12       |
| 6      | 20       |
| 7      | 21       |
| 8      | 22       |
| 9 etc. | 100 etc. |

there are  $3^n$  coins

← Coding (i.e., labelling)  
sample points

Thus each coin will have a label consisting of a sequence of 0's, 1's, and 2's.  
For example:

0 1 2 1 0 0 1 0 1 ← coin

0 1 2

three blocks


For the 1<sup>st</sup> weighing, take the first digit and split the 3<sup>n</sup> coins into one of 3 piles, depending on this digit.

For each successive weighing, proceed similarly by splitting the 3<sup>n</sup> coins based on the successive digit.

### • Shannon Coding Theorem

In general, the Shannon Coding Theorem, which is the main theorem in information theory, goes something like this.

You have a discrete sample space  $\Omega$ , with entropy of the trivial partition some number, say  $n$ .



$\Omega$

$$H[\hat{\sigma}] = n$$

You want to code the sample space, by assigning to each sample point, a sequence of 0, 1 digits, in such a way that the sequence is the shortest possible.

The Shannon Coding Theorem tells you that, if you are allowed shorter and longer digits and you count only the average number of digits, you can always code the sample space (regardless of its probabilities) with an average of  $n$  digits.

That's just the beginning of information theory.

The fun comes when you have noise.

In other words: Back to the "20 questions" example of guessing a number [SC 4/24/98.7]

But, now I'm allowed 10 lies.

What's the minimum number of questions?

} This is a channel w/ noise

SC 4/24/98.13

You have to work out the entropy you need and the main theorem of Shannon's Theorem. How you work out the entropy, taking into account the noise.

Even w/ 1 lie, it is tough.

This gives you an idea about information theory.

Next week, a specialist on information theory will give the Super-Class. Professor Viola of Course 6.

He works full-time using entropy to solve interesting problems

Image Reconstruction } all done  
Pattern Recognition } by entropy

There is a fundamental question. People don't really know why it works.  
Why does the Maximum Entropy Principle [SC 4/24/98.6] work?

Q: Is there a relationship between the entropy of the conjugate prior and the entropy of the posterior in Bayes' Law?

A: Lot's of words.  
Rivers of words.

The relation between Bayes' and entropy is unsolved.

Work is still going on. Research is going on.

People have not understood the relationship.

It's one of the major unsolved problems of science.

Perhaps you will solve it, while I'm still alive.

Please.

50 SHEETS  
100 SHEETS  
200 SHEETS

22-141  
22-142  
22-144



Within the text, pagination is of the form *lecture\_date.page*. Note that a different convention is used in both the table of contents and the index, where pagination is of the form *lecture\_number.page*. The table of contents provides the mapping between *lecture\_number* and *lecture\_date*. For example, page 1.1 in the table of contents and index corresponds to page 2/4/98.1 in the text, where the topic anomaly detection is discussed.

## Index

- Abel polynomials, SC2.8
- Alternatives
  - law of, 11.8, 13.11, 14.1
  - law of (continuous analogue), 23.5
- Anomaly detection, 1.1
- Average measurements, 33.15
- Average number of blips, 29.6
- Axioms of probability, 1.3
- Balls into boxes, 3.11, 4.10, SC2.3
  - forest of rooted trees, SC2.8
- Bayes' estimate, 25.4
  - using conjugate prior, 25.7
- Bayes' law, 12.5, 13.11, 14.8, 15.9
  - (continuous analogue), 24.10, 25.1
  - densities, 33.18
- Bayes' theorem
  - sampling without Replacement, 16.5
- Bell numbers, SC2.9
- Bernoulli process, 1.7, 2.6, 28.1, 28.2
  - conditional probability, 11.5, 13.4, 14.2
  - density plots, 24.13
  - random variables, 8.7, 8.8, 8.10, 9.3
  - sample space, 1.7
- Bernoulli run, infinitely often, 2.11, 3.2
- Beta function, 25.1
  - discrete, AN16.7, AN16.8
  - integral, 18.11, 24.11, 24.12, 25.2
- Binomial
  - coefficients, 4.2, 6.5
  - distribution, 8.9
  - identity, 6.5, 6.6, 6.10, 7.7, 10.5, 10.11, 10.13, 10.15, 15.13, 16.1, 18.14
  - theorem, 4.4, 4.8, 8.9, 9.8
- Bits, SC4.2
- Blips, 28.12, 29.1
  - colored, 30.17, 31.1
- Blocks, 11.8
- Borel sets, 17.5
- Borel-Cantelli lemma, 3.2
- Bose-Einstein statistics, 5.7, 16.4
- Buffon needle problem, SC3.2, SC3.8
- Cauchy's functional equation, 26.14, 28.5
- Central limit theorem, 34.11
- Circuit theoretic interpretation
  - Boolean algebra, SC1.7
- Cluster analysis, 1.1
- Coin
  - fair, 34.9



- tossing, 1.7
  - continuous, 29.8
  - gaps, 9.3
- Computing
  - $\pi$ , SC3.6
  - perimeters, SC3.12
- Conditional probability, 11.1, 11.7
  - continuous, 22.1
  - density, 22.11
  - law of alternatives, 14.1
- Conditioning event of probability
  - zero, 23.1
- Confidence intervals, 24.15, 25.3
- Continuity property, 2.3
- Convex set, SC3.7
- Convolution, 26.8
- Countable additivity, 2.2
- Cumulative distribution, 17.2, 18.2
- DeFinetti's theorem, AN16.10
- Density, 17.3, 18.1
  - algebra, 26.1
  - function
    - using inverse function, 25.13
  - joint, independent random variables, 26.7
  - plots, 19.4, 24.13
- Difference operator, 4.3, SC2.11
  - backwards, 4.4
- Dirichlet distribution, 18.5
- Dirichlet process, 17.5, 22.2
  - random function, 19.3
- Disjoint intervals
  - Poisson process, 29.4
- Dispositions, 4.11, 5.3, SC2.8
- Distribution interpretation, 15.6
  - Maxwell-Boltzmann, 4.6, 6.1, 15.4
- Dropping points
  - on a circle, 18.12
  - on a triangle, 20.1
  - on an interval, 23.1
- Entropy
  - independent partitions, SC4.3
  - partitions, SC4.2, SC4.5
- Event, 1.3, 2.1
  - not certain, 3.10
  - not impossible, 3.10
- Expectation, 10.8, 15.1
  - random variable plus constant, 27.4
  - sum of random variables, 10.8
- Exponential density, 29.10
- Exponential distribution, 28.8, 28.15
- Factorial
  - lower, 4.3
  - rising, 4.4, 5.4
- Fair coin, 24.7
- Fermi-Dirac statistics, 6.4, 9.10, 15.12, 16.1
- Flags on poles, 4.11
- Gamma density, 29.12
- Gamma distribution, 29.11
- Gaps
  - coin tossing, 9.3
  - order statistics, 18.12
  - uniform process, 23.8
- Gaussian distribution, 35.1
- Geometric distribution, 8.12
- Hard spheres, 1.1
- Hypergeometric distribution, 9.11, 14.9
  - multivariate, 10.5
- Identically distributed, 8.3
- Inclusion-exclusion principle, 6.7, 7.1, 30.15, SC1.1

- Independence of
  - disjoint intervals, 29.5
  - disjoint intervals (Poisson process), 28.14
  - events, 1.5, 2.4
  - random variables, 8.4, 17.2
- Independent and identically distributed, 9.2
- Indicator random variables, 26.9
- Information Search, SC4.7
- Inspector's paradox, 27.11
- Intensity, 29.2, 31.3
- Joint
  - density, 19.6, 26.3
  - sum, 26.10
  - distribution, 9.13, 10.1
- Joint cumulative distribution
  - continuous random variables, 19.5
- Kolmogorov zero-one law, 3.8
- Laplace law of succession, 25.5
- Laplace transform, 31.4
- Large numbers
  - law of, 35.12
- Likelihood, 14.11, 33.17
- Marginal
  - density, 19.6, 20.4
  - distribution, 10.1
- Matching, 7.1
- Maximum entropy principle, SC4.6
- Maxwell-Boltzmann statistics, 4.1, 5.1, 8.7
  - conditional probability, 11.4
- Maxwell-Einstein
  - normal distribution derivation, 35.4
- Measure, 34.6, SC3.9
  - lines in the plane, SC3.11
- Measurements
  - assumptions, 33.16
  - average, 33.15
- Median, 18.3
- Memoryless property, 28.4
- Memorylessness
  - waiting time, 29.10
- Multinomial coefficient, 4.4, 4.10
- Multiset, 5.10
  - coefficients, 4.4, 5.5, 5.8, 6.5
  - identity, 16.4
- Needles on a stick, 24.1, 30.6, SC3.1, SC3.10
- Negative binomial distribution, 8.13
- Normal distribution, 33.14, 35.1
  - theory of, 28.1
- Normalize, 35.2
- Normalized random variable, 33.8
- Null probability, 9.1
- Occupancy interpretation, 15.6
  - Maxwell-Boltzmann, 4.6, 6.1, 15.5
- Occupation numbers, 30.10, 30.13
  - Bose-Einstein, 6.5, 7.6
  - Dirichlet process, 17.7
  - expectation, 10.10
  - Maxwell-Boltzmann, 4.7, 7.4
- Order statistics, 17.9, 18.7, 19.1, 20.6
  - conditional probability, 22.7
  - density, 18.8
  - expectation, 18.14, 20.11
  - joint density, 22.5
  - joint distribution, 22.4
- Origin
  - time to first return, 31.13
  - time to return, 31.7

- Pólya urn model, SC2.3
- Partitions, 11.8, SC2.9, SC4.1
  - algebra of, SC1.10
  - entropy, SC4.2, SC4.5
  - integer, SC2.5
- Perimeters, SC3.13
- Permutation, 4.2
- Pointless probability, 8.2, 26.9
- Poisson distribution, 29.4
- Poisson process, 28.1
  - events, 28.11
  - fundamental properties, 29.6
  - probability, 28.13, 29.2
  - sample space, 28.9
- Posterior, 14.11, AN16.4, 25.6
  - in statistics, 34.7
- Prior, 14.11
  - conjugate, 16.6, AN16.2, AN16.8, 24.16, 25.6
  - in statistics, 34.5
  - uniform, 15.11, 16.6, 16.10, AN16.2
  - uniform(continuous), 24.11
- Probability, 2.2
  - axioms, 1.3
  - distribution, 8.1
  - space, 1.4
- Probability distributions
  - algebra of, 25.9
- Problem of scales, SC4.9
- Quantum probability, 9.14
- Rényi's principle, 6.9, 7.8, SC1.5
- Rademacher's theorem, 3.9
- Random function, 17.9, 29.4, 35.10
- Random variable
  - continuous, 17.1
  - indicator, 26.9, SC1.1
  - integer, 8.1
  - interpretation of, 32.7
  - normalized, 33.8
  - standard normal, 32.5
  - standardized, 33.8
- Random walk, 31.5
  - as Bernoulli Process, 31.7
  - continuous analogue, 31.13
  - self avoiding, 1.2
  - symmetric, 32.1
- Rare events, 28.10
  - law of, 29.6
- Record values, 27.1
- Reluctant functions, SC2.7
- Run, 2.7
- Runners on a track, 27.8
- Sample
  - point, 1.7, 14.11
  - space, 1.3, 2.1
    - Poisson process, 28.9
  - space (infinite), 15.3
- Sampling
  - with replacement, 5.10, 9.6, 9.8, AN16.6
  - with replacement (infinite), 9.12
  - without replacement, 9.7, 9.9, AN16.1
- Schrödinger randomization, 30.9
- Scientific induction, 14.8
- Shannon coding theorem, SC4.12
- Standard deviation, 33.2
  - Bernoulli process, 33.3
- Standard normal
  - distribution, 32.2
  - random variable, 32.5, 33.10
  - random variable density, 34.1
  - table, 34.8
- Standardize, 35.2
- Standardized random variable, 33.8
- Statistics
  - assumptions, 33.15, 33.18, 34.3

- basic rule, 34.2
- Stirling number's of the second kind, SC2.11
- Stochastic process, 10.2
- Successive conditioning, 11.11, 12.5
- Successive probability, 11.7
- Sylvester's theorem, SC3.7
- Tail event, 3.8
- Test Problem, 12.6
- Tree
  - probability, 12.3, 13.5, SC2.1
- Uniform distribution, 18.5
- Uniform process, 17.5, 28.1
  - by conditioning Poisson process, 30.1
  - limit, 29.9
- Urn models, SC2.2
- Urn sampling
  - conditional probability, 13.1, 14.8
  - with replacement, 12.9
- Vandermonde's identity, 10.7
- Variance
  - addition formula, 32.8
  - Bernoulli process, 33.3
  - exponential random variable, 33.6
  - Poisson process, 33.5
  - properties of, 33.2
  - random variable, 32.6
- Waiting time, 14.3
  - Bernoulli process, 8.10, 10.14
  - memorylessness, 28.4
- Whitworth problems, 18.6
- Wiener process, 35.7
- z transform, 31.10
- Zero-one law, 3.8