2018 Edition

NOTES ON STOCHASTIC OPERATIONS RESEARCH

Mark Huber

©2018 Mark Huber Version 2018.000

Contents

Contents		iii
Pro	Preface	
1	What is Stochastic Operations Research?	1
2	Queues	6
3	Little's Law	11
4	Idle time for $G/G/s$ queue	15
5	Balance and detailed balance	19
6	Bottlenecks and Systems	24
7	Introduction to statistics	29
8	Simulation	33
9	Flowchart for DES	38
10	Implementing DES in R	44
11	Stochastic Petri Nets	48
12	How to make decisions	54
13	Statistical decision making	59
14	Decision Trees	64
15	More VN-M Utility Theorem	69
16	Psychology of decision making	73

17	Expected value of perfect information	77
18	Framing and a two-envelope problem	81
19	Creating utility functions to test beliefs	86
20	Shannon Entropy	90
21	Game Theory	95
22	Two person Zero-sum games	98
23	Nash Equilibria	103
24	Randomized algorithms	107
25	Randomized verification	111
26	Randomized QuickSelect	116
27	Randomization for IP's and LP's	122
28	Forecasting	127
29	Simple Exponential Smoothing	131
30	Seasonality	135
31	Marginal Analysis	139
32	News vendor problem with continuous demand	143
33	Inventory	146
34	Solving Inventory Problems	151
35	Probabilistic Dynamic Programming	155
36	Markov Decision Processes	159
37	Finding average reward	163
38	Variance Reduction	167
A	Probability review A.1 Elementary facts	171 171 173

	A.3	A short guide to counting	-74
	A.4	How to find $\mathbb{E}[X]$	76
	A.5	How to find $\mathbb{V}(X)$	77
	A.6	Distributions	77
	A.7	Discrete distributions	78
	A.8	Continuous Distributions	80
	A.9	How to use the Central Limit Theorem (CLT)	81
B	Proofs of Theorems 1		
	B.1	The Mixture Space Theorem	83

Preface

These lecture notes cover a one semester undergraduate course in Stochastic Operations Research. The goal here is to introduce students to a broad range of ideas used in Stochastic OR. The notes start with queuing theory, then move on to simulation techniques. This is followed by decision theory and game theory. After that is a brief foray into randomized algorithms, before settling on simple techniques for forecasting from time series.

The preparation needed for this course is a typical undergraduate course in probability. In particular, knowledge of expected value together with exponential and uniform distributions is a must. The text *Probability: Theory and Exploration* is OpenAccess and free to use, and covers all the ideas from Probability needed for the course. For those needing a less comprehensive review, the first Appendix provides a rapid introduction to probability.

Chapter 1 What is Stochastic Operations Research?

Question of the Day What is stochastic Operations Research?

Today

- Overview of the course
- Some probability review

OR is part of applied mathematics

- What decisions have you made lately?
- What queues have you been in?
- Have you ever tried to predict the future based on probabilities?

What is Operations Research

- Mathematics behind running an operation efficiently
- Perfect information (each crate of apples costs \$3.45) uses deterministic OR
- Partial information (apples cost random) uses stochastic OR

Probability encodes partial information

- Variable $x \in \mathbb{R}$, could be anything
- Random variable $X \in \mathbb{R}$ we have more information
- For example, if X is a roll of a fair six sided die:

$$X \sim \text{Unif}(\{1, 2, 3, 4, 5, 6\}).$$

• We know more about X than $x \in \{1, 2, 3, 4, 5, 6\}$

$$\mathbb{P}(X \le 2) = \mathbb{P}(X = 1) + \mathbb{P}(X = 2) = \frac{1}{6} + \frac{1}{6} = \frac{2}{6}$$

Partial information arises because

- Based on future events (ex: what will next customer arrive)
- It costs too much to obtain information (ex: how many units will the PS4 sell?)
- Probability measures/distributions give a simple way to model exactly how much info we do have.

Example The bookstore models the number of books needed for a class as a binomial random variable w/ parameters n = 20, p = 0.6

How many books should be ordered to have at least a 95% chance of meeting demand?

$$X \sim \text{Bin}(20, 0.6)$$

[Each student independently has a 0.6 chance of buying the book at the bookstore.] Q: What is smallest $a: \mathbb{P}(X < a) > 0.95$?

$$\mathbb{P}(X=i) = \binom{\binom{n}{i}}{i} (1-p)^{i} = \binom{20}{i} (0.6)^{i} (0.4)^{20-i}$$

We will use the statistical software R in this course. Open source (and free) from www.r-project.org

pbinom(seq(0,20),20,0.6)

tells us $\mathbb{P}(X \leq a)$ for all $a \in \{0, 1, 2, \dots, 20\}$

```
[1] 1.099512e-08 3.408486e-07 5.041261e-06 4.734497e-05 3.170311e-04
 [6] 1.611525e-03 6.465875e-03 2.102893e-02 5.652637e-02 1.275212e-01
[11] 2.446628e-01 4.044013e-01 5.841071e-01 7.499893e-01 8.744010e-01
[16] 9.490480e-01 9.840388e-01 9.963885e-01 9.994760e-01 9.999634e-01
[21] 1.000000e+00
```

So the 17th entry is the first where $\mathbb{P}(X \le a) > 0.95$. The 17th entry of seq(0, 20) is 16. Test with:

```
pbinom(15,20,0.6)
pbinom (16, 20, 0.6)
```

Types of random variables

Definition 1

```
A random variable X is discrete if there is a set \{x_1, x_2, x_3, \ldots\} such that \mathbb{P}(X \in \{x_1, x_2, \ldots\}) = 1.
```

Examples of discrete distributions

- $X \sim \text{Unif}(\Omega)$, Ω is finite
- $X \sim \text{Bern}(p)$ Number of successes in 1 trial.
- $X \sim Bin(n, p)$ Number of successes in n independent trials.
- $X \sim \text{Geo}(p)$ Number of trials until first success.
- $X \sim \text{NegBin}(r, p)$ Number of trials until r successes.

Definition 2

A random variable X is **continuous** if for all $a \in \mathbb{R}$, $\mathbb{P}(X = a) = 0$.

Fact 1

Continuous random variables X all have a density f_X that satisfies for all a < b,

$$\mathbb{P}(a \le X \le b) = \int_{a}^{b} f_X(s) \, ds$$

Definition 3

The **indicator function** 1(expression) is 1 if the expression is true, and 0 if the expression is false.

Example



Examples of continuous distributions

- $X \sim N(\mu, \sigma^2), f_X(s) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-(s-\mu)^2/[2\sigma^2]\right)$
- $X \sim \operatorname{Exp}(\lambda), f_X(s) = \lambda \exp(-\lambda s) \mathbb{1}(s \ge 0)$

•
$$X \sim \text{Unif}([a,b]), f_X(s) = \frac{1}{b-a} \mathbb{1}(s \in [a,b]).$$

Example Our bookstore models the time until the arrival of the next customer as Exp(2/hr). What is the chance that the next customer does not arrive in the first hour?



A: Let T be time of first arrival.

Then $T \sim \text{Exp}(2)$ (do everything in hours)

$$\mathbb{P}(T \ge 1) = \int_{1}^{\infty} 2 \exp(-2s) \, ds$$

= $-\exp(-2s)|_{1}^{\infty} = \lim_{b \to \infty} -e^{-2b} + e^{-2(1)} = e^{-2}$
 $\approx \boxed{13.53\%}$

Another example Back to the bookstore, what is

$$\mathbb{P}(T \ge 1 | T \ge 1/3) = ?.$$

Well,

$$\mathbb{P}(T \ge 1|T \ge 1/3) = \frac{\mathbb{P}(T \ge 1, T \ge 1/3)}{\mathbb{P}(T \ge 1/3)} = \frac{\mathbb{P}(T \ge 1)}{\mathbb{P}(T \ge 1/3)} = \frac{e^{-2(1)}}{e^{-2(1/3)}} = e^{-4/3}.$$

In general...

Mark Huber Notes on Stochastic Operations Research

Fact 2 For $X \sim \text{Exp}(\lambda), t \ge 0, s \ge 0$:

 $\mathbb{P}(X \ge t + s | X \ge t) = \mathbb{P}(X \ge s).$

Once $X \ge t$, "forgets" about past. Call it memoryless

Definition 4

The continuous **memoryless** distribution is the exponential distribution.

Sequences Let X_1, X_2, \ldots be a sequence of r.v.'s. Often they have the same distribution... and often any subset of the r.v.'s are independent.

iid = independent and identically distributed

Certain random variables X have a finite expected value $\mathbb{E}[X]$.

Fact 3 (Strong Law of Large Numbers) Suppose *X* has expected value $\mathbb{E}[X]$. Then for X_1, X_2, \ldots iid with $X_i \sim X$:

$$\lim_{n \to \infty} \frac{X_1 + X_2 + \dots + X_n}{n} = \mathbb{E}[X].$$

[Note: it is not necessary for the variance of *X* to be finite for this to hold.] Calculating $\mathbb{E}[X]$:

If X is discrete:
$$\mathbb{E}[X] = \sum_{s:\mathbb{P}(X=s)>0} s \cdot \mathbb{P}(X=s)$$

If X is continuous: $\mathbb{E}[X] = \int_{-\infty}^{\infty} sf_X(s) \, ds$

Calculating $\mathbb{E}[g(X)]$:

If X is discrete:
$$\mathbb{E}[X] = \sum_{s:\mathbb{P}(X=s)>0} g(s) \cdot \mathbb{P}(X=s)$$

If X is continuous: $\mathbb{E}[X] = \int_{-\infty}^{\infty} g(s) f_X(s) \, ds$

Chapter 2

Queues

Question of the Day Customers arrive at a help desk at rate 10/hour randomly w/ interarrival distributions exponential. On average, how many customers arrive in an eight hour workday?

Today

- Queues
- Queue notation
- Stopping time

Queues (aka waiting lines)

- Inevitable part of life with limited resources
- Service providers must decide who to serve first
- Time sensitive
 - Medical care
 - Other emergencies
 - Google response (too long, user leaves
- More servers (emergency personel, computer servers, etcetera) cost money
- Basic queuing models can be analyzed mathematically
- Complex queuing models studied using simulations

Definition 5

A **queue** consists of a set of customers waiting for a service.

Queueing terms

Definition 6

The **queue discipline** is the method by which the next customer in the queueing system is selected for service.

Definition 7

The **service time** is the time needed to serve a single customer.

Definition 8

The **interarrival time** is the time between two arrivals of customers to the queue.

channels	the # of servers
FIFO	first in, first out
LIFO	last in, first out
priority	highest priority served first
capacity	Maximum # of customers in queue

Queuing networks Often a queue feeds into other queues



Definition 9

A **Queueing network** is a collection of queues set up as a graph where each node is a server with a queue. If server i has any outgoing arcs, then when a customer completes service i, it must join one of the queues at the end of an outgoing arc.

Examples:

- Traffic systems
- Internet
- Telephone system

Qotd

interarrival time $\sim \text{Exp}(10/\text{hr})$



Assume interarrival times independent unless specified otherwise

 $A_1, A_2, A_3, \ldots \stackrel{\text{iid}}{\sim} \mathsf{Exp}(10).$

[Recall: iid means independent, identically, distributed.]

Time of 1st arrival A_1 2nd arrival $A_1 + A_2$ 3rd arrival $A_1 + A_2 + A_3$

What is the expected number of arrivals in [0, 8]? To answer, need Wald's Equation.

Wald's Equation This is a result covered in Stochastic Processes. In this course, will use results like these Will rarely prove them.

Fact 4 (Wald's Equation) Let $X_1, X_2, \ldots \stackrel{\text{iid}}{\sim} X$ where $\mathbb{E}[X]$ is finite. Let T be a stopping time with respect to X_1, X_2, \ldots If

- 1. $\mathbb{P}(X \ge 0) = 1$ or
- 2. $\mathbb{E}[T] < \infty$, then

$$\mathbb{E}\left(\sum_{n=1}^{T} X_n\right) = \mathbb{E}[T]\mathbb{E}[X_i].$$

Wald's equation uses notion of stopping times

Definition 10

T is a **stopping time** for a sequence X_1, X_2, \ldots if for all *n*, it is possible to determine if $\{T \le n\}$ using only the values X_1, \ldots, X_n .

• Example:

 $T = \inf\{n : A_1 + A_2 + \dots + A_n > 8\}.$

• The event $\{T \leq 5\}$ consists of

$$\cup_{i=1}^{5} \{T = i\} = \{A_1 > 8\} \cup \{A_1 < 8 < A_2\} \cup \{A_2 < 8 < A_3\} \cup \{A_3 < 8 < A_4\} \cup \{A_4 < 8 < A_5\} = \{A_5 > 8\}.$$

- To determine if $T \le 5$ requires knowing if $\{A_5 > 8\}$.
- Note that # of arrivals in [0, 8] is T 1.

Since $A_i \ge 0$, Wald's Equation applies:

$$\sum_{i=1}^{T} A_i = \mathbb{E}[T]\mathbb{E}[A_i].$$

Since $A_i \sim \text{Exp}(10)$, $\mathbb{E}[A_i] = 1/10$. Also

$$\sum_{i=1}^{T} A_i = \text{ time of first arrival after time 8.}$$

For example, in this picture T = 4, $\sum_{i=1}^{T} A_i = 9.1$:

9.1

8

In general,

$$\sum_{i=1}^{T} A_i = 8 + \text{time from 8 until next arrival}$$

Use the memoryless properties of exponential random variables

Fact 5 If $X \sim \operatorname{Exp}(\lambda)$, then

$$[X - r | X > r] \sim \mathsf{Exp}(\lambda).$$

Fact 6 (Markov property for exponential interarrival times) Let $T_i = A_1 + \cdots + A_i$ be the *i*th arrival time where the A_i are iid $Exp(\lambda)$. Then

$$[A_i - a | A_{i-1} < a, A_i > a] \sim \mathsf{Exp}(\lambda).$$

This also holds when working with stopping times:

Fact 7 (Strong Markov property for exponential interarrival times) Let $T_i = A_1 + \cdots + A_i$ be the *i*th arrival time where the A_i are iid $\text{Exp}(\lambda)$, and T be a stopping time. Then

$$[A_T - a | A_{T-1} < a, A_T > a] \sim \mathsf{Exp}(\lambda).$$

So, $\sum_{i=1}^{T} A_i = 8 + X$, where $X \sim \text{Exp}(10)$. Putting this all together:

$$\mathbb{E}\left[8+X\right] = \mathbb{E}[T](1/10).$$

The left hand side is 8 + 1/10, so $\mathbb{E}[T] = 8(10) + 1$, and

$$\mathbb{E}[T-1] = 8(10) = \boxed{80}.$$

Generally:

Fact 8 (Expected arrivals for exponential interarrival times) With iid interarrival times with $A_i \sim \text{Exp}(\lambda)$, the average number of arrivals in time [a, b] is $(b - a)\lambda$.

Chapter 3 Little's Law

Question of the Day Customers arrive at a store at rate 5 per hour, and stay an average for 20 minutes. What is the long term average number of customers in the store at any point in time?

Today

- Queue notation
- Little's Law

Last time

- · Arrivals came exponentially
- Can analyze exactly in this case
- What if interarrivals arbitrary?

Notation for queues General queue notation has six slots:



For example:

$M/M/1/\infty/\infty/FIFO$

is a queue with 1 server, an infinite length queue and set of customer arrivals, and uses the first in-first out system of queue discipline.

The code for the first two entries is as follows:

- M Markovian=Memoryless=Exponential distribution
- D Deterministic (constant) times
- E_k Erlang w/ parameter k
- G_I General independent
- *G* General (possibly dependent)

Reminder, if A_1, \ldots, A_k are iid $Exp(\lambda)$, then

 $A_1 + \cdots + A_k$

are Erlang with parameters k and $\lambda.$

Often only the first three slots used:

- Consider an M/D/2 queue
- The arrivals are exponential
- The service times are determinstic (constant)
- There are 2 servers

Some queue disciplines

- FIFO (first in, first out) (aka FCFS=first come, first served)
- SIRO service in random order

The simplest queue (and easiest to analyze) is

 $M/M/1/\infty$

Commonly used variables with queues are:

 $\lambda = \text{ arrival rate (expected arrivals per unit time)}$ $\mu = \text{departure rate (expected services per unit time by each server)}$ s = # of servers $\rho = \frac{\lambda}{\mu \cdot s} \text{capacity utilization}$ $p_n = \text{steady state chance of queue length being } n$ L = expected long run length of line $W_q = \text{expected long run wait for a customer}$ $W = W_q + \text{expected service time for customer}$

Use the term "steady-state" to describe asymptotic behavior. Queues take time to reach equilibrium if start empty.

Mark Huber | Notes on Stochastic Operations Research

In the 1950's, practitioners started to notice something. No matter what the arrival and service distributions were,

average line length = arrival rate times average time spent in system

Conjectured by Morse, proved by John Little in 1961.

Theorem 1 (Little's law)

The average # of customers in a queue equals the long term arrival rate times the average time a customer spends in the system. That is,

 $L = \lambda W.$

- This is an amazing theorem!
- Works no matter what arrival and service distributions are!
- Also independent of queue discipline!

Qotd Arrival rate 5 per hour. Each averages 20 minutes in store.

$$\lambda = 5/\text{hr}, W = 1/3\text{hr} \Rightarrow L = 5/3 \approx 1.666$$

is the average # of customers in the store.

Example: Archytas consulting measures a queue at various times during the day and finds the average queue length is 11 individuals. By sampling random customers they estimate W to be 5 minutes. What is the customer arrival rate?

$$L = \lambda W \Rightarrow \lambda = L/W = 11/(5 \text{ min}) = 2.200/\text{min}$$

Proof idea

Do idea rather than formal proof.

Graph the length of the queue versus time.

Each rectangle of area corresponds to a person arrived in queue.

Let W_i denote the time person i spends in system.

Then area under queue in [0, t] is

$$\sum_{i=1}^{T} W_i - \sum_{i=1}^{T} \min(W_i - t, 0).$$

where T is the number of arrivals in [0, t].

Example picture



So the average length of the queue is

$$\frac{1}{t}\sum_{i=1}^{T} W_i - \frac{1}{t}\sum_{i=1}^{T} \max(T_i + W_i - t, 0).$$

If $\mathbb{E}[W_i] < \infty$, the last term goes to o as $t \to \infty$ Now $\mathbb{E}[T/t] \to \lambda$, that's what the steady state arrival rate is. If W_i independent,

$$\mathbb{E}\left[\sum_{i=1}^{T} W_i\right] = \mathbb{E}[T]\mathbb{E}[W_i],$$

by Wald's Equation, so

$$\lim_{t \to \infty} \mathbb{E}\left[\frac{1}{t} \sum_{i=1}^{T} W_i\right] = \lim_{t \to \infty} \mathbb{E}[T/t]W = \lambda W.$$

When W_i not independent (and depending on queue discipline, they might not be!) the proof is much harder.

Chapter 4 Idle time for G/G/s queue

Question of the Day A Google server receives on average 200 requests a second, each of which takes on average 0.001 seconds to resolve. What percentage of time is the server idle?

Today

- Idle time for G/G/s queue
- Steady state behavior for M/M/1 queue

Qotd

- Consider an interval of time [0, t].
- During this time there are on average 200t requests
- Each request takes on average 0.001 time to resolve.
- Let S_i denote time needed to service request i.
- By Wald: $\mathbb{E}\left[\sum_{i=1}^{T} S_i\right] = \mathbb{E}[T]\mathbb{E}[S_i]$
- So the average time to resolve requests received in $\left[0,t\right]$ is

$$(200t)(0.001) = t/5.$$

- Not all of those requests are served inside interval $\left[0,t\right]$



- When t large, extra on end negligible

- Percentage time server busy $\approx (t/5)/(t-0) = 1/5$
- Server idle roughly 80% of time.

Capacity utilization

Definition 11

For a G/G/s queue, the **capacity utilization** is

$$o = \frac{\lambda}{\mu \cdot s}$$

Fact 9

For a G/G/s queue with finite expected service and interarrival times, and $\rho < 1$, the long run (steady state) busy time is ρ , and the long run idle time is $1 - \rho$.

Proof idea

- The average number of customers arriving in [0, t] is $t \cdot \lambda$.
- If $\rho < 0,$ the queue returns to empty after some fixed expected time R (proof uses renewal theory)
- Expected length of one service $1/[s\cdot\mu]$
- Expected length of service needed for customers arriving in [0, t] is $t\lambda/[s \cdot \mu]$
- Expected length of service needed for customers arriving in [0, t] that happens in [0, t] is between $(t R)\lambda/[s \cdot \mu]$ and $t\lambda/[s \cdot \mu]$.
- As t goes to infinity, percentage time busy is $\lambda/[s\cdot\mu]$

Stochastic Processes

Definition 12 A collection of random variables $\{X_t\}$ is called a **stochastic process**.

- Example: Let L_t be the length of a queue at time t.
- Note L_1 and $L_{1.1}$ are *not* independent
- Often interested in steady state behavior

Definition 13

For a stochastic process $\{X_t\}_{t\geq 0}$, the steady state time spent in state *i* is

$$\lim_{t \to \infty} \frac{1}{t} \int_0^t \mathbb{1}(X_t = i) \, dt.$$

Example

- Suppose $X_0 = 0$. Then X_t spends 1.5 seconds in state 0 and jumps to 1.
- From 1, X_t spends 1 second in state 1, then jumps to 0.
- The steady state time X_t spends in state o is 1.5/(1.5+1) = 60%
- Note: Replace 1.5 seconds with an exponential random variable with rate 1.5 and 1 with an exponential random variable with rate 1.
- Then still X_t spends 60% of time in steady state in state o.

Definition 14

A stochastic process $\{X_t\}_{t\geq 0}$ is a **jump process** if the path $t \mapsto X_t$ is constant with a number of discontiuities that is countably infinite with probability 1.

Definition 15

A stochastic process over state space Ω is a **continuous time Markov chain** if it is a memoryless jump process. This means there is a function $\lambda : \Omega \to [0, \infty)$ where at time t, given $\{X_{t'}\}_{t' \leq t}$, the time τ of the next jump satisfies $\tau - t \sim \text{Exp}(\lambda(X_t))$, and the distribution of X_{τ} only depends on X_t .

Fact 10

Suppose $X \sim \text{Exp}(\lambda_X)$ and $Y \sim \text{Exp}(\lambda_Y)$ are independent. Then $\min\{X, Y\} \sim \text{Exp}(\lambda_X + \lambda_Y)$.

Proof. Recall that $A \sim \text{Exp}(\lambda)$ if and only if $\mathbb{P}(A > a) = \exp(-\lambda a)$. So

$$\mathbb{P}(\min\{X,Y\} > a) = \mathbb{P}(X > a, Y > a)$$

= $\mathbb{P}(X > a)\mathbb{P}(Y > a)$ [independent]
= $\exp(-\lambda_X a)\exp(-\lambda_Y a)$
= $\exp(-[\lambda_X + \lambda_Y]a)$

which means $\min\{X, Y\} \sim \mathsf{Exp}(\lambda_X + \lambda_Y).$

17 184

Fact 11

Let X_t be the number of customers in an M/M/1 queue at time t. Then X_t is a continuous time Markov chain.

Proof. Fix t. Consider the time until the next jump. If $X_t = 0$, this is just the time of the next arrival, which is $\text{Exp}(\lambda)$. If $X_t > 0$, this is the time of the next arrival or service, which is the minimum of two independent exponential random variables, hence also an exponential random variable.

When Ω discrete, continuous time Markov chains can be represented graphically

- Nodes = states
- Arcs labeled with rate of jump



Chapter 5 Balance and detailed balance

Question of the Day In an M/M/2 queue w/ average wait between arrivals of 1 min and average service time for each server of 1.5 min, what is the percentage of time exactly 1 server is idle?

Today

- · Balance and detailed balance equations
- M/M/s equations

Let $\lambda(i, j)$ be the rate at which state *i* jumps to state *j*.

Theorem 2 (Ergodic Theorem)

Suppose there exists a probability distribution π over countable Ω such that for all $i \in \Omega$

$$\sum_{j\in\Omega}\pi(i)\lambda(i,j)=\sum_{i\in\Omega}\pi(i)\lambda(j,i).$$

(These are called the *balance equations*.) Then the steady state amount of time spent in state i is $\pi(i)$.

Fact 12 For the M/M/1 queue,

$$\pi(i) = (1-\rho)\rho^i$$

satisfies the balance equations.

Proof. Start with state o:

 $\pi(0)\lambda(0,1) = \lambda(1-\rho) = (\lambda/\mu)\mu(1-\rho) = \pi(1)\lambda(1,0).$

Now suppose i > 0:

$$\pi(i)\lambda(i,i+1) + \pi(i)\lambda(i,i-1) = (1-\rho)[(\lambda/\mu)^i\lambda + (\lambda/\mu)^i\mu] = (1-\rho)[\lambda^{i+1}/\mu^i + \lambda^i/\mu^{i-1}]$$

and

$$\pi(i+1)\lambda(i+1,i) + \pi(i-1)\lambda(i-1,i) = (1-\rho)[(\lambda/\mu)^{i+1}\mu + (\lambda/\mu)^{i-1}\lambda]$$
$$= (1-\rho)[\lambda^{i+1}/\mu^{i} + \lambda^{i}/\mu^{i-1}]$$

 \square

So the balance equations hold!

This should look familiar!

Fact 13

Let π be the steady state distribution of an M/M/1 queue. Then for $L_{\infty} \sim \pi$, the steady state line length satisfies

$$L_{\infty} + 1 \sim \text{Geo}(1-\rho).$$

In particular $L = \mathbb{E}[L_{\infty}] = \rho/(1-\rho)$.

Example Suppose our Google server $\lambda = 200$, $s\mu = 1000$ is an M/M/1 queue. What is the average steady state queue length?

$$L_{\infty} + 1 \sim \text{Geo}(1 - 200/1000),$$

so

$$\mathbb{E}[L+1] = 1/.8 = 1.25 \Rightarrow \mathbb{E}[L] = \boxed{0.2500}$$

What is the average time a request stays in the system? Little's law:

$$L = \lambda \cdot W \Rightarrow 0.25 = 200W \Rightarrow W = 0.001250$$
 seconds.

What changes with multiple servers

- Let N = # of customers in system
 - When $N \ge 2$ both servers are working
 - When N = 1 one server is working
 - When N = 0 no servers are working
- So what is the rate with 2 servers.

- First server time $T_1 \sim \text{Exp}(\mu), T_2 \sim \text{Exp}(\mu)...$
- ...so first service complete $\min\{T_1, T_2\} \sim \mathsf{Exp}(2\mu)$

CTMC graph looks like:



Balance

- Last time introduced notion of balance
- Probability flow out of a node = probability flow in

Definition 16
For a CTMC, the **balance equations** are
$$(\forall i \in \Omega) \left(\sum_{j \in \Omega} \pi(i)\lambda(i, j) = \sum_{j \in \Omega} \pi(j)\lambda(i, j) \right).$$

- Detailed balance is a stronger condition
- Says that prob flow from i to j must equal prob flow from j to i.

```
Definition 17
For a CTMC, the detailed balance equations are
```

$$(\forall i, j \in \Omega)(\pi(i)\lambda(i, j) = \pi(j)\lambda(j, i)).$$

Fact 14 A CTMC that obeys detailed balance for π also obeys the balance equations.

Proof. Suppose π obeys detailed balance. Then

$$\sum_{j\in\Omega} \pi(i)\lambda(i,j) = \sum_{j\in\Omega} \pi(j)\lambda(j,i)$$

so balance is also satisfied.

For queues, easier to use detailed balance to solve for π . M/M/2 queue:

$$\pi(0)\lambda = \pi(1)\mu & 2\rho\pi(0) = \pi(1) \\ \pi(1)\lambda = \pi(2)2\mu & \rho\pi(1) = \pi(2) \\ \pi(2)\lambda = \pi(3)2\mu & \rho\pi(2) = \pi(3) \\ \vdots & \vdots & \vdots & \vdots \\ \end{pmatrix}$$

which makes:

$$\pi(i) = \pi(0)2\rho^i.$$

Since the probabilities have to add up to 1:

$$\pi(0) \left[1 + 2\rho + 2\rho^2 + \cdots \right] = 1$$

$$\pi(0) \left[1 + \frac{2\rho}{1 - \rho} \right] = 1$$

$$\pi(0) = \left[1 + \frac{2\rho}{1 - \rho} \right]^{-1} = \frac{1 - \rho}{1 + \rho}$$

Fact 15

For the M/M/2 queue with $\rho=\lambda(2\mu)^{-1}<1,$ the steady state distribution is

$$\pi(i) = 2\rho^i \frac{1-\rho}{1+\rho}$$

and the expected steady state line length is

$$\frac{2\rho}{1-\rho^2}.$$

Proof. The expected steady state line length is

$$\sum_{i=0}^{\infty} i\pi(i) = \sum_{i=1}^{\infty} 2\frac{1-\rho}{1+\rho}i\rho^i.$$

This converges absolutely for $\rho < 1$ by the ratio test.

Now

$$\begin{split} \sum_{i=1}^{\infty} i\rho^i &= \sum_{i=1}^{\infty} \sum_{j=1}^{i} \rho^i \\ &= \sum_{j=1}^{\infty} \sum_{i=j}^{\infty} \rho^i \text{ [by Tonelli's Theorem]} \\ &= \sum_{j=1}^{\infty} \rho^j / (1-\rho) \\ &= \rho / (1-\rho)^2. \end{split}$$

Plugging into the earlier equation gives:

$$\mathbb{E}_{\pi}[ext{line length}] = rac{2
ho}{(1+
ho)(1-
ho)}.$$

Note as $\rho \rightarrow 1,$ expected line length goes to infinity.

Chapter 6 Bottlenecks and Systems

Question of the Day Suppose that cars obey a 2 second rule, leaving 2 seconds between them and the car in front of them. A freeway has on the eastbound side 4 lanes that narrow to 2 lanes, leading to a backup of 1400 vehicles. If traffice on the 2 lane part is moving at 70 mph, how long will a vehicle that enters the backup wait before reaching the 2 lane part of the road.



Today

• Bottlenecks in queuing systems



Bottlenecks

• Where does the queue form?

- At paperwork–much slower to get through than Front Desk
- Say that Front Desk and Paperwork are in series
- Capacity for Paperwork much lower than capacity for Front Desk
- (Wait in line will usually be higher for paperwork)
- Overall capacity is $\min\{1/7, 1/2\} = 1/7$ per minute.

Definition 18

If customers served by one queue become customers for another queue, the queues are in **series**.

Fact 16

For servers connected in series, the overall capacity to serve is equal to the minimum of the capacities of the servers.

Parallel servers

Definition 19

When a queue is being serviced by more than one server simultaneously, the servers are in **parallel**.

• Example: Parallel computing



average service time 7 min

• To determine service rate, ask: if queue at Enter has many customers, on average how many are served in *t* minutes?

$$\underbrace{\frac{t}{5}}_{\text{from Server 1}} + \underbrace{\frac{t}{7}}_{\text{from Server 2}} = \text{service rate} \cdot t$$

Mark Huber | Notes on Stochastic Operations Research

So

service rate
$$=$$
 $\frac{1}{5} + \frac{1}{7} = \frac{12}{35} \approx 0.3428$ per min

Fact 17

Servers in parallel add their service rates.

Relation to other mathematics

- Related to tropical geometry named in honor of Brazilian mathematician Imre Simon.
- Two operations:

$$x \oplus y = \min\{x, y\}$$
 Series
 $x \otimes y = x + y$ Parallel

Definition 20

Consider servers $\{1, \ldots, n\}$ each with service rate μ_i . If there is a subset of servers with equal service rate where raising that rate would increase the overall service rate of the system, call that subset a **bottleneck**.



- Pouring out of a bottle like a series of servers
- Overall flow=min{flows in series}

Q of Day

- The bottleneck is the entry to the two lane highway
- Two second rule



• How fast are cars "being served"

$$\left(\frac{2 \text{ cars}}{2 \text{ sec}}\right) \left(\frac{3600 \text{ sec}}{1 \text{ hour}}\right) = \frac{3600 \text{ cars}}{\text{hour}}$$

- Note: the speed of the cars did not come into it!
- Flow capacity on 4 lane highway:

$$\left(\frac{4 \text{ cars}}{2 \text{ sec}}\right) \left(\frac{3600 \text{ sec}}{1 \text{ hour}}\right) = \frac{7200 \text{ cars}}{\text{hour}}$$

- So if 4 lane highway operating at more than 50% capacity, get backup at merge
- On average every 2 sec, 2 cars are served by two lane highway = 1 car/s.
- So if there are 1400 cars in queue, 1400 s
- Again, doesn't matter what speed the 2 lanes are moving at!
- 70mpg or 30 mph, what matters is the 2 s rule.

Computer controlled cars

- One way to increase capacity of freeways
- Lower time between cars
- Computer driven cars could (hopefully!) do this safely
- · Would save billions in infrastructure costs

What about accidents

- Suppose people slow down
- But maintain same physical distance between cars
- If they slow 10%, and road was at 90% capacity before
- Causes backup

Chapter 7 Introduction to statistics

Question of the Day Douglas Consulting records queue length at 5 times during the day of 17, 21, 8, 7, 14. What is the average queue length? What is the standard eviation?

Today

- Basic statistics
- Introduction to R

Statistics

- The science of data analysis
- Interacts with probability

Fact 18 (Strong Law of Large Numbers) For X_1, X_2, \ldots iid with finite mean,

$$\lim_{n \to \infty} \frac{X_1 + \dots + X_n}{n} = \mathbb{E}[X_i] \text{ w/ prob } 1.$$

[The sample average converges to the true average with prob 1.]

Putting data in R

- Download R: www.r-project.org
- x <- c(17,21,8,7,14)
- <- is assignment operator
- c standard for concatenate or combine

Once data is in R

- mean(x)
- ?mean gives related topics, open help window

13.40 is estimate of mean queue length .

Can use R to generate random variates

- runif(n = 10, min = 0, max = 1)
- Generates $X_1, ..., X_{10} \sim Unif([0, 1])$ iid
- mean(runif(n = 10, min = 0, max = 1))
- Can sometimes be quite far away from 1/2

Variance

- $\mathbb{V}(X) = \mathbb{E}((X \mathbb{E}(X))^2)$
- On average, square of distance from average value.
- (Square so always nonnegative)

•
$$\mathrm{SD}(X) = \sqrt{\mathbb{V}(X)}$$

- In R: sd(x)
- Gives 5.941 as estimate of standard deviation

$$\hat{\mu} = \frac{X_1 + \dots + X_n}{n}, \ \hat{\sigma} = \sqrt{\frac{(X_1 - \hat{\mu})^2 + \dots + (X_n - \hat{\mu})^2}{n - 1}}.$$

Fact 19

 $\mathbb{E}(\hat{\mu}) = \mathbb{E}(X_i)$ and $\mathbb{E}(\hat{\sigma}^2) = \mathbb{V}(X_i)$. These are *unbiased* estimates for the mean and variance. (Note, in general $\mathbb{E}(\hat{\sigma}) \neq \text{SD}(X_i)$.)

Subsets of data

- Suppose the consultant decides only middle three data points are accurate
- Use x [2:4] (here [2:4] is the sequence 2,3,4)
- mean(x[2:4]) gives 12.00
- sd(x[2:4]) gives 7.810
Scripts in R

- A *script* is a set of commands for software to execute
- Put them in a text file (end with .R for R)
- Be sure to put them in directory that R is accessing
- getwd() gets the working directory from R
- setwd() sets the working directory for R
- Note \ is an escape character in R, always use / for directories (even in windows)
- A line that begins with the number symbol # will be ignored
- Used for commentating your code

Example script

• Put in file test.R

```
# This is a script
x <- c(17,21,8,7,14)
print(mean(x))
print(sd(x))</pre>
```

- print prints the argument to the screen
- · Fancier ways exist
- Typing source("test.R") is the same as if I'd typed in those three commands directly.

Functions

- Take inputs and run some commands, then return output
- Example, making $f(x, y) = 2xy^2$

```
f <- function(x,y) {
   return(2*x*y^2)
}</pre>
```

- Try f(3,1) and f(-2,3) and f(y = 3, x = 2)
- I can make default values for the inputs as follows:

```
31 184
```

```
f <- function(x = 1, y = 3) {
    return(2*x*y^2)
}</pre>
```

• If you don't enter values for x or y, these take over

```
• Try f(), f(x = 3)
```

There are five types of variables in R

- 1. Vectors
 - x < -c(4, 3, 2) or y < -4 (one dimensional vector)
- 2. Character strings

```
• y <- "abc"
```

- 3. Matrices
 - Create by binding vectors together.
 - Ex: A <- rbind(c(1,3),c(2,-2))
 - det(A), eigen(A)

4. Lists

- · Like vectors, but components can be different data types
- x<- list(u = 2,v = "abc")
- Try x and x\$u
- (Often used by functions to return different pieces of info)
- Many lists built into R
- Try Nile and h <- hist (Nile) and h and str(h)

5. Data frames

- Like matrices, but components can be mixed type
- d <-data.frame(list(kids=c("Jack", "Jill"), ages=c(12,10)))
- d
- · Usually data frames are read in from files

Chapter 8 Simulation

Question of the Day Say arrivals at a queue are Unif([0,2]) while service is Unif([0,1]). What is the long term average waiting time of a customer?

Today

- Simulation
- Discrete event simulation

Complex models

- Most models are too complicated to solve exactly.
- Qotd: G/G/1 queue–tough to get exact answers.
- Simulation provides an easier approach

Simulation

- Generate a random instance of the queue on a computer
- Use that as a "random draw" from the world
- Then use our basic statistical tools to estimate true average

Types of simulation

- Some track entities moving through system
 - Particle simulations
 - Traffic simulations
 - Airplane boarding
- Some track the state of the system
 - Queueing networks

Example: G/G/1 queue

- State of the system entirely described by...
 - Length of the queue
 - Current time
 - Next service time (if any)
 - Next arrival time
 - Is a server free?

Definition 21

An event in a DES is anything that changes the state of a system.

- Two events in G/G/1:
 - Service (reduces queue by 1)
 - Arrival (increases queue by 1)

Discrete Event Simulation (DES)

Definition 22

A **Discrete Event Simulation (DES)** is a method of simulation which focuses on events that change the state of the system. The key component is an eventlist which keeps track of currently scheduled events and the times they change.

- Keeps track of clock (current time) and system value
- Keeps list of events that are scheduled to occur
- Each event does two things
 - Changes the state
 - Possibly schedules other events

Graphical representation

- Nodes represent events
 - Below node, write the changes in state
- · Arcs represent new schedule events
 - Labeled with the time delay before event goes off



Event A schedules event B after t time

- Note that *t* might be a random variable.
- Can also have a condition attached (if crossed with $\sim)$



Event A schedules event B after t time only if condition holds

Qotd DES

- Three events associated with a queue
 - 1. Someone joins the queue
 - 2. Someone starts service
 - 3. Someone leaves service
- One extra event
 - Run event to initialize the system
- State of the system: # in queue = Q, # of servers free = S

Run event

- Initializes $Q \leftarrow 0, \, S \leftarrow 1$
- Schedules the first arrival.
- In graph form:



Enter event

- Increases the queue by 1
- If a server is free, schedules a start of service
- Then schedules next arrival
- That way never run out of arrivals!





- Note: no time on Enter \rightarrow Start means instantaneous

Start event

- Decreases the queue by 1
- Decreases the available servers by 1
- Schedules the end of service = leave event



Leave event

- · Frees up a server
- If queue nonempty, should schedule a service





Putting everything together





Chapter 9 Flowchart for DES

Question of the Day How should ties be resolved in a DES?

Today

- Master flowchart for DES
- · Ties and deadlocks
- Running through an example

Last time



Working through example

• Event list at begining:

tstatetype of event0.00(0,1)Run

• Execute Run event at t = 0:

Poll t get $t = 0.12$	t	state	type of ev
Koli ι_a , get $\iota_a = 0.12$	0.12	(0, 1)	Enter
Schedule Enter t_a time later			
Remove Run from event list			
• Execute Enter event at $t = 0.12$:			
Increase Q by 1	+	state	type of e
Roll $t_a \leftarrow 1.20$	0.12	(1, 1)	Start
Schedule Enter at $t + t_a$	1.32	(1,1)	Enter
Remove Enter from event list			
• Execute Start event at $t = 0.12$:			
$S \leftarrow S - 1, Q \leftarrow Q - 1$	4	stata	turna of or
Roll $t_s \leftarrow 0.93$		(0, 0)	Leave
Schedule Leave at $0.12 + 0.93$	1.32	(0,0)	Enter
Remove Start from event list			
• Execute Leave event at $t = 1.05$:			
$S \leftarrow S + 1$			
Since $Q = 0$, don't schedule Start		t	state typ
Remove Leave from event list		1.32	

Notes

- Always have an Enter in the list
- Usually quite after t greater than some fixed $t_{\rm end}.$

Programming a DES

- The central object is event list.
- Each row of the event list is a time for an event and an event name.
- To run a DES, event list must be updated properly.
- The following flowchart outlines this procedure.

Event Scheduling Flowchart



Variation: Queue with limited capacity



Two nonidentical servers in parallel



Possible to extends DES further

· Can turn variables into arrays

- Ex: $S[1],\ldots,S[k]$ for servers of k different types
- Can pass parameters along on arcs

Chapter 10 Implementing DES in R

Question of the Day Use simulation to determine for G/G/1 queue with interarrivals iid Unif([0, 2]), service times Unif([0, 1]), what is the average waiting time for a customer.

Today

• Putting DES into R

DES Framework

- 1. Pull earliest event from event list
- 2. Execute the event
 - a) Update state
 - b) Schedue new events
- 3. Remove event from event list
- 4. If any events remain, go to step 1

[Go through DES.R code]

What do we want to know?

- What is the average queue length?
- What is idle/busy percentage?

Example: Queue length



Recall Average of f(t) over $t \in [a, b]$ is:

average =
$$\frac{\int_{a}^{b} f(t) dt}{b-a}$$

In example:

$$\frac{1}{4-0}[3(1.8-0)+4(2.6-1.8)+3(4-2.6)=3.2.$$

Summary statistics

- · Add variables to keep track of these things
- [Show extra statistics in DESstat.R]

DESstat.r: R routines for discrete-event simulation (DE # Written by: Mark Huber

Data frame sim holds the event list

MAIN DES LOOP

```
# main loop of the simulation
mainloop <- function(maxsimtime) {
  time <- 0  # record original time
  while((nrow(sim) > 0) & (sim[1,1] < maxsimtime)) {
    # Now take the first event off of the event list and e
    event <- sim[1,]  # take first event
    sim <<- sim[-1,] # drop the first event</pre>
```

```
# execute the event
    executeEvent(event)
    # update statistics and clock time
    time <- min(sim[1,1], maxsimtime)</pre>
    statistics["Average queue length"] <<- statistics["Ave</pre>
       (time - event$time)*((1-state["s"]) + state["q"])
    statistics["Idle percentage"] <<- statistics["Idle per
       (time - event$time) * (state["s"] > 0)
  }
}
# EXECUTE CODE
# The general code above needs an event processing functic
executeEvent <- function(event) {</pre>
  if (event$type == "RUN") {
    sim <<- rbind(sim,data.frame(time = ta(),type = "ENTEF</pre>
  }
  if (event$type == "ENTER") {
    state["q"] <<- state["q"] + 1;</pre>
    sim <<- rbind(sim, data.frame(time = event$time + ta(),</pre>
    if (state["s"] > 0)
      sim <<- rbind(sim,data.frame(time = event$time,type)</pre>
  }
  if (event$type == "START") {
    state["q"] <<- state["q"] - 1;</pre>
    state["s"] <<- state["s"] - 1;</pre>
    sim <<- rbind(sim, data.frame(time = event$time + ts(),</pre>
  }
  if (event$type == "LEAVE") {
    state["s"] <<- state["s"] + 1;</pre>
    if (state["q"] > 0)
      sim <<- rbind(sim,data.frame(time = event$time,type</pre>
  }
  sim <<- sim[order(sim$time),] # sort event list by time</pre>
}
# RANDOM VARIALBES
# Interarrival times Unif([0,2])
```

Mark Huber Notes on Stochastic Operations Research

```
#ta <- function() return(runif(n=1,min=0,max=2)) # intera</pre>
# Service times Unif([0,1])
#ts <- function() return(runif(n=1,min=0,max=1)) # servic</pre>
ta <- function() return(rexp(n=1, rate=5/6))</pre>
ts <- function() return(rexp(n=1, rate=1))</pre>
# TO RUN SIMULATION
runsim <- function(maxsimtime) {</pre>
  # Initialize event list
  sim <<- data.frame(time = 0.0, type = "RUN")</pre>
  # Initialize state of the system
  state <<- c(1,0)
  names(state) <<- c("s", "q") # 1 server, 0 in queue</pre>
  # Initialize statistics for the system
  statistics <<- c(0,0)
  names(statistics) <<- c("Average queue length","Idle per</pre>
  # Run simulation
  mainloop(maxsimtime)
  statistics <<- statistics / maxsimtime
  print(statistics, digits = 4)
}
```

Chapter 11 Stochastic Petri Nets

Question of the Day Can we build a simulation that can be analyzed completely?

Today

- Petri Nets
- Stochastic Petri Nets

Petri nets

- Invented by Carl Petri 1962
- Bipartite graph with two types of nodes: places and transititons
- Arcs run from place to transition, and transition to place.
- When a transition "fires", it puts a token on the place it points to
- A transition needs all of its incoming places occupied to fire.

Example



- The black rectangle is a transition
- The circle is a place
- A token allows the transition to fire, putting a token on outgoing place.



Uses

- Easier to write down mathematical description than DES
- State just number of tokens at each place
- Can ask questions like: What states are reachable from initial state?
- Great for inventory systems

Example

- A manufacturing station is combining the inside and case for a thumbdrive into the finished product
- Looks like this:



• Next step



• Next step



- Reachable states from $\left(P1,P2,P3\right)=\left(0,0,0\right)$ are

 $\bigcup_{i\{0,1,2,3,\ldots\}}\{(1,1,i)\}$

Timed transitions

- Extension: Can add a time to each transition
- Now the new token does not appear until a later time
- Transition cannot be reactivated until time expires
- Example: D/D/1 queue



Stochastic Petri Nets

- Allow the timing for transitions to be random variables
- Example: G/G/1 queue



50 184

Ties

- Doesn't solve problem of ties
- Have to specify priority of transitions

Labeled arcs

- Suppose to build product requires two of one component and one of another
- Label arcs with number of tokens consumed (or produced)



• How the Petri Net progresss:

 $\begin{array}{c} (P1,P2,P3)\\\hline \\(0,0,0)\\(1,1,0)\\(2,2,0)\\(1,2,1)\\(2,3,1)\\(1,3,2)\\(2,3,2)\\\vdots\end{array}$

- Call a Petri Net k-bounded if the number of tokens in any place never exceeds k
- A 1-bounded Petri Net is called *safe*.
- Above example is not *k*-bounded: both P₂ and P₃ grow forever.
- Change: store product from P3

• Change: If more than 4 in P2, store excess elsewhere



Petri Nets Advantages

- Easy to visualize.
- Built in animation procedure.
- Can often prove that a Petri net is safe or at least k-bounded.

Petri Nets Disadvantages

- Have to keep track of which transitions active
- Have to check every transition to see which fire
- Large models can be slow to execute

Inhibitor arc (Optional)

- Works in the opposite fashion from regular arc
- If a token there, prevents transition from firing
- Draw with a circle at arc tip



• In this example, tokens in P2 have priority for service over tokens in P1

Chapter 12 How to make decisions

Question of the Day How should decisions be made?

Today

• Decision analysis

What do decide/herbicide/homocide have in common?

- cide comes from Latin caedere = "to cut"
- de comes from Latin for off
- decide = to cut off (possibilities)

Terminology

- *Decision analysis* is the problem of how best to select from a set of possible course of action.
- State variables (state of nature) are things you cannot control
- Courses of action (*decision variables*) are things you can control
- The result of making a decision once state variables known is outcome or payoff



Example The final in the course might be easy or hard. Can choose to study or not study:

	Study	Don't Study
Hard	Satisfaction	Humiliation
Easy	Wasted Time	Relief

- Usually state variables are random
- [Only have partial information]

Utility

- First step in decision analysis is to assign *utility* to each outcome
- High utility means favorable outcome
- · Low utility means despised outcome
- Note: monetary payoff can be rough substitute for utility
- But in general *not* the same!
- Utility difference in going from \$0 to \$100...
- ...much more than in going from \$999 900 to \$1 000 000
- This effect (first dollar has more utility than last dollar) called *dimishing marginal returns*



• Actions also called strategies

Definition 23

The **strategic form** of a decision problem is a triple (X, Ω, A) where

- 1. X is the nonempty set of strategies for the decider.
- 2. Ω is the nonempty set of states of nature.
- 3. $A: (X \times \Omega) \to \mathbb{R}$ is the payoff function.

Dominance

- Some decisions you never want to make
- Consider payoff matrix with $X = \{a_1, a_2, a_3, a_4\}$ and $\Omega = \{\theta_1, \theta_2, \theta_3\}$.

	a_1	a_2	a_3	a_4
θ_1	5	0	2	1
θ_2	5	3	3	2
θ_2	0	4	2	1

- Note that a_3 always beats a_4 , no matter what the state of nature is!
- There is never a need to use a dominated strategy.

Definition 24

For $a_i, a_j \in X$, say that a_i **dominates** a_j if

 $(\forall \theta \in \Omega)(A(a_i, \theta) \ge A(a_j, \theta)).$

Ways to decide

- Pessimist: *Maximin* Find the minimum payoff over all states for each choice. Use strategy than has the maximum minimum.
- For earlier example: worst outcome gives:

a_1	a_2	a_3
0	0	2

• So the Maximin choice is a_3 , since that maximizes the minimum value

Definition 25	
The Maximin strategy is	
	\cdot $\Lambda(-0)$

```
\arg\max_{a\in X}\min_{\theta\in\Omega}A(a,\theta).
```

Next strategy: Maximax

- The optimist takes a different view
- First, find the largest payoff from each possible strategy
- Take the choice that maximizes the largest payoff
- For example:

a_1	a_2	a_3
5	4	3

• Choice that maximizes is: *a*₁

Definition 26		
The Maximax strategy is		
	$\arg\max_{a \in X} \max_{\theta \in \Omega} A(a, \theta).$	

Laplace: Principle of insufficient reason

- Suppose you know nothing about a situation
- It this case, reasonable to use a uniform distribution
- Assume each state equally likely: θ now a random variable
- More recently: call this a Bayesian approach with noninformative prior
- · Choose strategy that maximizes expected utility
- Example problem:

$$\mathbb{E}(A(a_1,\theta)) = (1/3)(5) + (1/3)(5) + (1/3)(0) = 10/3.$$

In general:

$$a_1 \quad a_2 \quad a_3$$

10/3 6/3 6/3

• Go with strategy a_1

Definition 27 The principle of insufficient reason chooses

$$\underset{a \in X}{\operatorname{arg\,max}} \sum_{\theta \in \Omega} A(a, \theta).$$

Hurwicz Principle (Degree of optimism) Linear convex combination of optimist and pessimist

Definition 28

For $\alpha \in (0, 1)$, the **Hurwicz principle** is

a

$$\operatorname{rg\,max}_{a \in X} \left[\alpha \max_{\theta \in \Omega} A(a, \theta) + (1 - \alpha) \min_{\theta \in \Omega} A(a, \theta) \right]$$

- Note: $\alpha = 0$ is maximin
- $\alpha = 1$ is maximax
- $\alpha \in (0,1)$ something in between
- Decision can change with α
- Continuing example:

 $\begin{array}{ccc} a_1 & a_2 & a_3 \\ 5\alpha & 4\alpha & 3\alpha + 2(1-\alpha) = \alpha + 2 \end{array}$

• $5\alpha \ge 4\alpha$,

$$5\alpha \ge \alpha + 2 \Leftrightarrow \alpha \ge 1/2.$$

- For $\alpha \in [0, 1/2]$, decision a_3
- For $\alpha \in [1/2, 1]$, decision a_1
- Can eliminate α_2

Chapter 13 Statistical decision making

Question of the Day How to make decisions in light of partial information?

Today

- Savage regret
- Bayesian approach

Regret

- The *regret* is the difference between what you get as payoff and best payoff
- Example:

Pay	off N	latrix		Reg	ret N	latrix	2
	a_1	a_2	a_3		a_1	a_2	a_3
θ_1	5	0	2	$ heta_1$	0	4	1
θ_2	5	3	3	θ_2	0	1	0
θ_2	0	4	2	θ_2	5	0	1

• Therefore, maximum regret is:

$$a_1 \quad a_2 \quad a_3 \\ 5 \quad 4 \quad 1$$

• Decision which minimizes regret is: *a*₃

д

Definition 29

The Savage regret decision is

$$\underset{a \in A}{\operatorname{trg\,min}} \left[\max_{\theta \in \Omega} A(a, \theta) - \min_{\theta \in \Omega} A(a, \theta) \right]$$

Using partial information

- All decisions so far do not use any info about state of nature.
- Often know the prob. dist. of θ .

	Study	Don't Study
Final is Hard	8	2
Final is Easy	6	10
Sun explodes	0	12

- Laplace: (1/3, 1/3, 1/3) prob. vector for states
- Maybe (just maybe) should have slightly lower than 1/3 chance for sun exploding.
- Maximin: makes $\mathbb{P}(\text{worst option}) = 1$.
- But $\mathbb{P}(\text{sun exploding}) \neq 1$.
- Decision: must do "Not Study"
- Also fails row linearity:
 - If you add a constant to a row, it should not affect decision
 - Add -3 to first row

	Study	Don't Study
Final is Hard	5	-1
Final is Easy	6	10
Sun explodes	0	12

- New minimum:

S	NS
0	-1

Expected Utility Hypothesis

- Goes back to Daniel Bernoulli 1738
- Satisfies row linearity
- More general than Laplace
- Idea:

[–] New decision: a_1

- Assign a probability to each state indicating partial information individual has about the likelihood [Called a *prior* distribution]
- 2. Calculate the expected payoff from each decision
- 3. Choose the decision that maximized utility
- Example: If 60% chance final is hard, $40\%-\epsilon$ that it is easy, and ϵ chance Sun explodes

$$\mathbb{E}[\text{payoff}|\text{Study}] = (0.6)(8) + (0.4 - \epsilon)(6) + \epsilon(0) = 7.2 - 6\epsilon.$$

$$\begin{array}{cc} \mathrm{S} & \mathrm{NS} \\ 7.2-6\epsilon & 5.2+2\epsilon \end{array}$$

- So for small ϵ , should Study.

Row linearity

- Suppose add C to a row that occurs with probability p_i .
- Adds $p_i C$ to each of the mean payoffs (regardless of decision)
- Does not change max expected payoff decision

Does this work

- 1947: John von Neumann and Oskar Morgenstern
- Gave 4 axioms that if true for an individual...
- ...must be a utility function
- Individual should prefer strategy that maximizes expected utility

Notation

- Call a decision with options a *lottery*
- Another way to state: a *lottery* is just a probability distribution on outcomes
- Example: $\Omega = \{\omega_1, \omega_2, \omega_3\}$ lottery L = (1/3, 1/2, 1/6) gives $\mathbb{P}_L(\omega_1) = 1/3$
- So a lottery is a vector \boldsymbol{p} whose entries add up to 1
- Also called the *probability simplex*.
- For lotteries *L*, *M* say:

- L = M if the actor is indifferent to the playing lottery L or M: this does *not* necessarily mean the probability vectors are equal!
- $L \prec M$ if actor prefers lottery M to L
- For $p \in [0,1]$, let N = pL + (1-p)M be the lottery where the agent has a p chance of getting to play L and and 1-p chance of getting to play M.
- Lotteries form a convex set

Axioms of Utility Theory (vN-M Axioms)

Definition 30

The von Neumann-Morgenstern Axioms are:

1. Completeness: For all lotteries L and M exactly one of the following is true:

$$L \prec M, L = M, \text{ or } M \prec L.$$

- 2. Transitivity: If $L \leq M$, $M \leq N$, then $L \leq N$.
- 3. Continuity: If $L \leq M \leq N$, there exists a $p \in [0, 1]$ such that

$$pL + (1-p)N = M.$$

4. Independence If $L \prec M$, then for any lottery N and $p \in (0, 1]$,

 $pL + (1-p)N \prec pM + (1-p)N.$

Theorem 3 (Von Neumann-Morgenstern Utility Theorem)

A complete and transitive preference relation \leq on a finite set of lotteries satisfies continuity and independence if and only if there is a random variable U such that for all lotteries L and M:

 $L \prec M \Leftrightarrow \mathbb{E}_L[U] < \mathbb{E}_M[U]$ $L = M \Leftrightarrow \mathbb{E}_L[U] = \mathbb{E}_M[U].$

Another way to state it: these four axioms characterize when the expected utility hypothesis holds.

Shape of the utility curve

- Straight line
 - 1st dollar exactly as valuable as millionth dollar.
 - Useful model when money of outcomes small compared to net worth

- Risk Averse (RA)
 - Shape concave (convex down)
 - Avoids big risks
 - Typical function when outcomes comparable to net worth
- Risk Seeker (RS)
 - Shape convex (convex up)
 - Millonth dollar worth more than 1st dollar
 - Trying to be first to reach a milestone
 - Going for the big score
 - Monopoly will often behave this way
 - Venture capitalists
 - "With \$100,000, I could open a restaurant!"
- Kahneran & Tversky 1979 designed experiments to see what participants curve looked like
- Found they could alter the curve by rephrasing the question.

Chapter 14 Decision Trees

Question of the Day Archytas Electronics must decide whether to build a new tablet. The product is expected to have a demand that is either high, low or medium. Since the money involved is small relative to the earnings of the company, utility is just taken to be money here. If the tablet is not built, the company loses nothing. But if it is built, then

		Demand	
	Good	Moderate	Bad
Probability	20%	70%	10%
Payoff	\$1 000 000	\$500 000	\$-500 000

Decision Trees

- · Often, decisions involve seugences of simple decisions
- · A decision tree keeps track of the different possibilities
- Three different type of branches:
 - 1. Decision fork (representd by a square) is a place where decision maker makes a decision
 - 2. Chance fork (represented by a circle) is a place where the state of nature determined

Example

• Archytas Electronics must decide whether to build a new tablet. The product is expected to have a demand that is either high, low or medium. Since the money involved is small relative to the earnings of the company, utility is just taken to be money here. The decision tree looks like:



- Should the tablet be launched?
- Only one decision to make.
- If Launch, expected payoff:

$$(20\%)(10^6) + (70\%)(0.5 \cdot 10^6) + (10\%)(-0.5 \cdot 10^6)$$

= 10⁶(0.2 + 0.35 - 0.05)
= (0.5)10⁶ > 0

• Decision is to launch the product!

To analyze decision trees

- 1. Replace random outcomes at leaves of tree with expected value
- 2. Take decision which maximizes expected value
- 3. Repeat if necessary

Example Suppose the decision tree looks like:



• Replace random outcomes in leaves with expectation:

$$(50\%)(2) + (50\%)(0) = 1$$

 $(60\%)(4) + (40\%)(2) = 3.2$


- Now consider the decision in the upper right: each outcome of the decision is a number, take the decision that maximizes the value
- In this case, "up branch" for 3 utility.



• Again trim by finding expectation:

$$(20\%)(3) + (70\%)(3) + (10\%)(-1) = 2.6$$

- Now the first decision is easy: take the "down branch"
- Any decision tree can be simplified using this approach!

Chapter 15 More VN-M Utility Theorem

Question of the Day How is the VN-M Utility Theorem proved?

Today

- Proving vN-M Utility Theorem
- Examples

Axioms of Utility Theory (vN-M Axioms)

1. Completeness: For any lottery L and M exactly one of the following is true:

$$L \prec M, L = M$$
, or $M \prec L$.

- 2. Transitivity: If $L \preceq M$, $M \preceq N$, then $L \preceq N$.
- 3. Continuity: If $L \leq M \leq N$, there exists a $p \in [0, 1]$ such that

$$pL + (1-p)N = M.$$

4. Independence If $L \prec M$, then for any lottery N and $p \in (0, 1]$,

$$pL + (1-p)N \prec pM + (1-p)N.$$

Theorem 4 (Von Neumann-Morgenstern Utility Theorem)

A complete and transitive preference relation \leq on a finite set of lotteries satisfies continuity and independence if and only if there is a random variable U such that for all lotteries L and M:

$$L \prec M \Leftrightarrow \mathbb{E}_L[U] < \mathbb{E}_M[U]$$
$$L = M \Leftrightarrow \mathbb{E}_L[U] = \mathbb{E}_M[U].$$

Proof: Utility implies axioms. Suppose such a utility function/random variable U exists. Let $L \leq M \leq N$. If L = N then any $p \in [0, 1]$ satisfied the equation. If $L \prec N$ let

$$p = \frac{\mathbb{E}_N[U] - \mathbb{E}_M[U]}{\mathbb{E}_N[U] - \mathbb{E}_L[U]}.$$

 $L\prec N$ means the denominator is not zero, and $L\preceq M\preceq N$ gives $p\in[0,1].$ Note

$$\mathbb{E}_{pL+(1-p)N}[U] = \sum_{\omega \in \Omega} [p\mathbb{P}_L(\omega) + (1-p)\mathbb{P}_N(\omega)]U(\omega)$$
$$= p\mathbb{E}_L[U] + (1-p)\mathbb{E}_N[U]$$
$$= \mathbb{E}_N[U] - p(\mathbb{E}_N[U] - \mathbb{E}_L[U])$$
$$= \mathbb{E}_M[U]$$

Hence M = pL + (1 - p)N and continuity is satisfied. For independence: suppose $L \prec M$, so $\mathbb{E}_L[U] < \mathbb{E}_M[U]$. Then

$$\mathbb{E}_{pL+(1-p)N}[U] = p\mathbb{E}_{L}[U] + (1-p)\mathbb{E}_{N}[U] < p\mathbb{E}_{M}[U] + (1-p)\mathbb{E}_{N}[U]$$

= $\mathbb{E}_{pM+(1-p)N}[U].$

The other direction is more difficult! Need to understand affine functions.

Affine transformations

Fact 20

Definition 31 A function f is affine if for all vectors x and y and $\lambda \in [0, 1]$,

$$f(\lambda x + (1 - \lambda)y) = \lambda f(x) + (1 - \lambda)f(y).$$

Example: f(x) = 3x - 2. Then

$$f(\lambda x_1 + (1 - \lambda)x_2) = 3(\lambda x_1 + (1 - \lambda)x_2) - 2(\lambda + (1 - \lambda)))$$

= $\lambda(3x_1 - 2) + (1 - \lambda)(3x_2 - 2)$
= $\lambda f(x_1) + (1 - \lambda)(x_2)$

Suppose U is a utility function. Then so is U' = aU + b where $a, b \in \mathbb{R}$ and a > 0.

[Any affine transformation of *U* is also a utility function.]

Proof. Since a > 0, U' is an increasing function of U, so maintains the same ordering. Suppose

$$\mathbb{E}_{pL+(1-p)N}[U] = \mathbb{E}_M[U].$$

Then

$$\mathbb{E}_{pL+(1-p)N}[U'] = \mathbb{E}_{pL+(1-p)N}[aU+b]$$
$$= a\mathbb{E}_{pL+(1-p)N}[U] + b$$
$$= a\mathbb{E}_M[U] + b = \mathbb{E}_M[aU+b] = \mathbb{E}_M[U']$$

so continuity is preserved.

Turns out that all utility functions are affine transformations of each other!

Fact 21 (Mixture Space Theorem, Herstein and Milnor) A preference relation \leq on Π is independent and continuous if and only if there exists an affine utility representation $U : \Pi \to \mathbb{R}$ of \leq .

Moreover, if $U : \Pi \to \mathbb{R}$ is an affine representation of \leq , then $U' : X \to \mathbb{R}$ is an affine representation of \leq iff there exist a > 0 and $b \in \mathbb{R}$ such that U' = aU + b.

See the appendix for a proof. One more small fact:

Fact 22

Suppose f is affine and p_1, \ldots, p_n add up to 1. Then for x_1, \ldots, x_n ,

$$f(p_1x_1 + \dots + p_nx_n) = \sum_{i=1}^n p_i f(x_i).$$

Proof. When n = 2, this is just the definition of affine! So use induction on n.

Base Case: n = 1, this gives a tautology.

Induction hypothesis: assume true for n, consider n + 1. Let $\lambda = p_1 + \cdots + p_n$ and $y = (p_1 x_1 + \cdots + p_n x_n)/\lambda$. Then

$$f(\lambda y + (1 - \lambda)x_{n+1} = \lambda f(y) + (1 - \lambda x_{n+1}) \text{ (affine)}$$
$$= \lambda \sum_{i=1}^{n} \frac{p_i}{\lambda} x_i + (1 - \lambda)x_{n+1} \text{ (induction)}$$
$$= \sum_{i=1}^{n+1} p_i x_i,$$

which completes the induction.

Proof: Axioms imply utility (Part II). Collecting our facts:

- Let Π denote the set of lotteries (probability distributions over outcomes)
- By the Mixture Space Theorem, there is an affine function U that represents \preceq .
- For $\omega \in \Omega$, call the *Dirac delta lottery* the lottery $\delta(\omega)$ where $\mathbb{P}(\omega) = 1$.
- Let $V(\omega) = U(\delta(\omega))$.
- Let L be any lottery, then note

$$L = \sum_{\omega} L(\omega)\delta(\omega).$$

where $\sum_{\omega} L(\omega) = \mathbb{P}_L(\Omega) = 1.$

- Then since U is affine,

$$U(L) = U\left(\sum_{\omega} L(\omega)\delta(\omega)\right) = \sum_{\omega} L(\omega)U(\delta(\omega)) = \mathbb{E}_L(V).$$

• Hence V is a utility representation.

Chapter 16 Psychology of decision making

Question of the Day What the common decision making mistakes?

Today

- Zero Illusion
- Allais Paradox
- Gambler's Fallacy
- Anchoring
- Adjustment
- Sunk cost Fallacy

Zero Illusion

- Often people have financial target of breaking even
 - "Must balance the budget"
 - Debt viewed as always bad
 - Savings are not "extra money"
 - leads to a utility curve:



- Counters:
 - Understand that o is just another number, nothing special
 - Rescale or shift problem to move away from o

Allais paradox

Consider two lotteries

Lottery 1A100% win \$1 millionLottery 1B89% win \$1 million, 1% win nothing, 10% win \$5 million

- 1987 study of Machina: people prefer Lottery 1A
- Now change the lotteries:

Lottery 2A89% win nothing, 11% win \$1 millionLottery 2B90% win nothing, 10% win \$5 million

- Same study: people prefer Lottery 2B
- Problem: this is inconsistent

Why inconsistent?

- Consider a function U.
- Then 1A > 1B implies $\mathbb{E}_{1A}[U] > \mathbb{E}_{1B}[U]$:

$$\mathbb{E}_{1A}[U] = U(1) > 0.89U(1) + 0.01U(0) + 0.1U(5) = \mathbb{E}_{1B}[U].$$

• Now let's play with $\mathbb{E}_{2A} < \mathbb{E}_{2B}[U]$:

$$\mathbb{E}_{2A}[U] = 0.89U(0) + 0.11U(1) < 0.9U(0) + 0.1U(5) = \mathbb{E}_{2B}[U]$$

• Add 0.89U(1) - 0.89U(0) to both sides:

$$U(1) < 0.89U(1) + 0.01U(0) + 0.1U(5)$$

What's going on?

- Comes back to regret: 11% \$1 million < 10% \$5 million when likely to get nothing anyway.
- 11% \$1 million > 10% \$5 million when likely to get something anyway.
- Under expected utility, should treat 1% chance $0 \rightarrow 1$ for 10% chance 5 for 1 the same way regardless of what happens in the other 90%.

Gambler's Fallacy

- Flip a fair coin heads 40 times in a row
- $\mathbb{P}(\text{next flip heads}) = 1/2$
- Gambler's Fallacy: we are "due" for a tail.
- By the way, don't pick 1,2,3,4,5,6 for lottery numbers.
- You would hate to split the lottery with 10 other self-satisfied people.
- Pick numbers above 30 to avoid dates.
- Reverse Gambler's Fallacy: bubbles

Anchoring

- Initial assessment/information biases behavior
- Milliken was first to measure charge on an electron
- If other experimenters independent, would be just as likely to be above true answer as below.
- Instead, see more like a limit



• Doesn't take overt planning: just throw out "outliers" that move the mean closer to the old average.

Tvesky & Kahneman 1982 Experiment

- Rolled a number X uniform over $\{1,2,\ldots,100\}$
- Asked participants: Consider percent of African countries are members of the U.N., is it *X* percent?
- Found that people would bias their guess towards the garbage number

Be aware

- Anchoring to expert opinion
- Anchoring to mean

Sunk Cost Fallacy

- Start a project, invest X dollars
- At current time, to complete project, need Y more dollars
- Should the project be completed?
- Answer should only depend on Y, not on X
- Decision should ony depend on future utility, not how much spent
- Easy to think that the amount already put in matters
- Seen in:
 - Military conflicts
 - Large scale real estate projects
 - Majors

Really smart people fall for this

- Newton: South Sea Tulip Bubble
- D'Alembert (Wave Equation) Gambler's Fallacy
- Leibniz thought 12 and 11 were equally likely on sum of two dice

Chapter 17

Expected value of perfect information

Question of the Day Suppose that the economy is expected to be hot, neutral, or cold over the next year. A investor is considering two stocks, or the money market. The payoff matrix is:

	Payoff Matrix			
	Stock 1	Stock 2	MM	
hot	2000	900	600	
neutral	200	300	600	
cold	-600	-200	600	

Given a prob vector for the market of (40%, 30%, 30%), how much should the investor pay in order to see the future?

Today

• Expected value of perfect information

Definition 32

Let U be the utility gained by making an optimal decision. Let W(I) be the optimal utility gained given the knowledge of some information encoded in I. The **expected value of perfect information** is

$$\mathbb{E}[W(I)] - \mathbb{E}[U],$$

the expected amount that having the information in I increases the utility of the optimal decision.

To find EVPI

- 1. Find mean utility with no info.
- 2. Find mean utility with info.

3. Subtract 1 from 2

Question of the day:

- First find mean utility
- Three investments:

$$\mathbb{E}[U_1] = 2000(0.4) + 200(0.3) - 600(0.3) = 680$$
$$\mathbb{E}[U_2] = 900(0.4) + 300(0.3) - 200(0.3) = 390$$
$$\mathbb{E}[U_3] = 600.$$

- So with no information, optimal decision to pick Stock 1.
- What is optimal decision with information?
- If the economy is hot: Stock 1, if neutral or cold, MM

$$\begin{split} \mathbb{E}[W(I)] &= \mathbb{E}[\mathbb{E}[W(I)|I]] \\ &= \mathbb{E}[2000 \cdot \mathbb{1}(I = \text{hot}) + 600 \cdot \mathbb{1}(I = \text{cold})] \\ &= 2000(40\%) + 600(60\%) = 1160. \end{split}$$

• So in this case, EVPI = 1160 - 680 = 480.

What does this mean?

- EVPI measures the expected value of reducing uncertainty
- In this case, if fortune teller could tell you what the market was going to do, you should be willing to pay up to 480 to learn the answer
- Gives a way of measuring the important of various pieces of information.

A continuous example

(Winston 1991, p. 718) Cards Inc. belives the # of a new type of card sold (call it S) to be normal with mean μ (unknown) and variance 100. There is a fixed cost of \$57 to introduce the card, and they make 0.60 per card sold. What is the EVPI for μ_{prior}?

So far

- State of nature is discrete
- Often parameters are continuous
- Cannot draw decision tree

Model

- Need continuous prior for μ , σ
 - How to get such a prior?
 - One method: $\mu \sim N(\mu_{\text{prior}}, \sigma_{\text{prior}}^2)$
 - Overall model: (Use S for sales of the card)

$$\mu \sim \mathsf{N}(\mu_{\text{prior}}, \sigma_{\text{prior}}^2)$$
$$S|\mu \sim \mathsf{N}(\mu, 100).$$

- This is called a *hierarchical model*
- Get μ_{prior} , σ_{prior}^2 by asking experts.

Eliciting $\mu_{\text{prior}}, \sigma_{\text{prior}}$

- Could ask questions like:
 - What are average sales like?
 - What is a value k such that 50% of time sales are within k of average? (Easier for most to understand than standard deviation.)
- Example: Answers are $\mu_{\text{prior}} = 100, k = 10$
- So what's μ_{prior}?
- Consider a standard normal. $\mathbb{P}(Z \leq 0.674898) = 75\%$:



$$\begin{split} \mu &\sim \mathsf{N}(100,\sigma_{\mathrm{prior}}^2)\\ Z &= \frac{\mu - 100}{\sigma} \sim \mathsf{N}(0,1)\\ \mathbb{P}(\mu \leq 110) = 75\% \text{ so}\\ \mathbb{P}(Z \leq (110 - 100)/\sigma) = 75\%. \end{split}$$

Hence

$$\frac{110 - 100}{\sigma} = 0.674898 \Rightarrow \sigma = \frac{10}{0.674898} \approx 14.82.$$

Recall

• Bayesian decision making priciple:

Make the decision that maximizes expected utility

- (Small dollar amounts = utility)
- Let *P* be profit
- Don't make card, P = 0. Do make card:

$$\mathbb{E}[P] = \mathbb{E}[-57 + 0.60S]$$

= $\mathbb{E}[\mathbb{E}[-57 + 0.60S|\mu]]$
= $\mathbb{E}[-57 + 0.60\mu]$
= $-57 + 0.60(100) = 3.$

• So proper (Bayesian) decision: should make the card!

Expected Value of Perfect Information

- What if we know μ_{prior} before hand?
- If $\mu_{\text{prior}} = 95$, $\mathbb{E}_{\mu=95}[P] = -57 + 0.60(95) = 0$
- If $\mu_{\text{prior}} > 95$, $\mathbb{E}_{\mu > 95}[P] = -57 + 0.60 \mu > 0$

- Make the card

- If $\mu_{\text{prior}} < 95$, $\mathbb{E}_{\mu > 95}[P] = -57 + 0.60 \mu < 0$
 - Don't make the card
- Call 95 the break even point for μ
- $W(\mu) = \mathbb{1}(\mu \ge 95)(-57 + 0.60\mu)$

$$\mathbb{E}(W(\mu)) = \int_{-\infty}^{\infty} \mathbbm{1}(u \ge 95)(-57 + 0.60u) \frac{1}{\sqrt{2\pi \cdot 14.82^2}} \exp\left(-\frac{(u-100)^2}{2}\right)$$
$$= \int_{95}^{\infty} (-57 + 0.60u) \frac{1}{\sqrt{2\pi \cdot 14.82^2}} \exp\left(-\frac{(u-100)^2}{2 \cdot 14.82^2}\right)$$
$$\approx 5.24876.$$

• Hence the EVPI is about \$5.24876 - \$3 = \$2.248.

Chapter 18 Framing and a two-envelope problem

Question of the Day Suppose two envelopes contatins 2 different positive amounts of money. You are allowed to look inside one envelope, then either keep it or switch to the other. Is there a way to choose the envelope with more money more than 1/2 of the time?

Today

- Framing
- Two envelopes

Framing

- Extends the Zero Illusion
- People tend to avoid risk when outcomes good
 - "A bird in the hand is worth two in the bush"
- People embrace risk when outcomes bad
 - "He who hesitates is lost"
 - "In for a penny, in for a pound"

Example

- Suppose disease has two treatments
- Treat A: 200 people saved w/ prob 1
- Treat B: 600 people saved w/ prob. 1/3, 0 w/ prob 2/3
- Most doctors go for A: sames expected value, lower risk.

Reframed example:

- Treat A: 400 people die w/ prob. 1
- Treat *B*: 600 people die w/ prob. 2/3, 0 people die w/ prob. 1/3
- Now most doctors go for ${\cal B}$

This effect used a lot in marketing

- 2-week trial w/ money back gaurantee
- Discount for cash: rarely say surcharge for credit

Also can affect happiness with outcome

- Movie Theater promotion
 - June wins \$100 for being the millionth customer
- Movie Theater promotion
 - Betty wins \$10 000 for being the millionth customer
 - Mike wins \$1 000 for being 1,000,001 customer
- Which is happier, June or Mike?

How to neutralize framing effect

- In surveys, ask questions in random order
- Try to give both good and bad outcomes
 - Treat A: 400 die and 200 live
 - Treat B: 600 die and 0 live w/ prob 2/3, 0 die and 600 live w/ prob 1/3

Other approaches to minimizing irrational behavior

- Make target aware of biases
- Feedback from results of model
- Redundant questioning
- Don't let humans decide (linear regression routinely outperforms experts

Two envelope problem

- Money in one is x
- Money in other is y > x
- · Allowed to look in one envelope and see amount of money
- Which envelope should you pick?
- Turns out you can always do better than 50%!

Procedure Start with a probability density f that is positive over $[0, \infty)$, like Exp(1) or |N(0, 1)|.

- 1. Look in random envelope, call result A
- 2. Draw $B \sim f$
- 3. If $B \leq A$, keep the envelope with A, otherwise switch.

What is the chance that you end up with the good envelope?

- Ways to win:
 - Pick y, choose $B \leq y$.
 - Pick x, choose B > x.
- Total chance:

$$(1/2) \int_0^y f(s) \, ds + (1/2) \int_x^\infty f(s) \, ds$$

= (1/2) $\int_0^y f(s) \, ds + (1/2) \int_x^y f(s) \, ds + (1/2) \int_y^\infty f(s) \, ds$
= (1/2) + (1/2) $\int_x^y f(s) \, ds$.

- Since y > x and f(s) > 0, this second term is positive!
- Note: If y = x, then 100% chance of picking the good envelope!

Two envelope paradox

- Now suppose that you know that y = 2x, so the envelopes contain x and 2x
- Say you look inside and see \$1000
- Then you know the other envelope has \$500 or \$2000
- Since each are equally likely, if you pick the other envelope:

 $\mathbb{E}(\text{switch}) = (1/2)500 + (1/2)2000 = 1500.$

- · So you should always switch
- But that argument works no matter what amount you saw in the envelope!
- So we can switch without even seeing what's in the envelope!

Solving the paradox

- No one accepted solution
- One problem: it is not equally likely to have \$500 or \$2000
- The envelops either were (500, 1000) or (1000, 2000), but who said they had to be equally likely?
- Need a model for how the initial x was decided.
- Example: Suppose $X \sim \text{Geo}(1/3)$, for $i \ge 1$, $\mathbb{P}(X = i) = (2/3)^{i-1}(1/3)$.
- Let N be amount seen, M be amount in other envelope
- Suppose see N = 4. What is distribution of M?

$$\mathbb{P}(M = 2|N = 4) = \frac{\mathbb{P}(M = 2, N = 4)}{\mathbb{P}(N = 4)}$$
$$= \frac{\mathbb{P}(X = 2)(1/2)}{\mathbb{P}(X = 2)(1/2) + \mathbb{P}(X = 4)(1/2)}$$
$$= \frac{(1/3)(2/3)(1/2)}{(1/3)(2/3)(1/2) + (1/3)(2/3)^3(1/2)}$$
$$= 2 \cdot 3^2 / [2 \cdot 3^2 + 8] = 18/26 = 9/13 \approx 69.23\%/$$

• So
$$\mathbb{E}[\text{switch}] = (2)(9/13) + (8)(4/13) = 40/13 < 4$$

• Don't switch!

- Suppose see N = 2, then

$$\mathbb{P}(M=1|N=2) = \frac{1}{1+(2/3)} = \frac{3}{5} = 60\%,$$

so

$$\mathbb{E}[\text{switch}] = (4)(0.4) + 1(0.6) = 2.2 > 2.$$

- So with this model switch when N = 2 (or N = 1), but don't switch if N is larger!

Chapter 19

Creating utility functions to test beliefs

Question of the Day Consider the multiple choice question:

What is the capital of Louisiana?

- a) New Orleans
- b) Bon Temps
- c) Baton Rouge
- d) Lousiannaville

How can true beliefs about answer be elicited?

Today

• Building utility functions to discover beliefs

Multiple choice questions

• Common payoff function:

```
payoff(answer) = 1(answer is correct).
```

• Suppose my information/knowledge is:

a) 50%, b) 20%, c) 15%, d) 15%.

Payoff Matrix

	a	b	c	d
a	1	0	0	0
b	0	1	0	0
c	0	0	1	0
d	0	0	0	1

- To maximize expected return, choose a

Sharing knowledge

- Suppose I let you spread 1 unit over a), b), c), d)
- One answer:

a) 30%, b) 40%, c) 10%, d) 20%.

or more compactly:

$$m = (0.3, 0.4, 0.1, 0.2)$$

Expected payoff is:

$$(0.5)(0.3) + (0.2)(0.4) + (0.15)(0.1) + (0.15)(0.2) = 0.275.$$

- Put 1 on a), expected payout 0.5
- So best option still to put one unit on a)

What if I change the payoff?

- Suppose payoff(answer) = $\sqrt{m(\text{correct})}$.
- So (since *a* is correct):

$$f((0,0,1,0)) = 1$$
 point
 $f((1/2,0,1/2,0)) = \sqrt{1/2}$ point

- But what if no true answer, just prior beliefs *p*?
- Expected payoff given beliefs p = (0.5, 0.2, 0.15, 0.15):

$$\mathbb{E}_p[f(0.6, 0.4, 0, 0)] = (0.5)\sqrt{0.6} + (0.2)\sqrt{0.4} = 0.5137\dots$$
$$\mathbb{E}_p[f(1, 0, 0, 0)] = (0.5)\sqrt{1} = 0.5.$$

- So (0.6,0.4,0,0) better choice than (1,0,0,0)

Maximization

• What is the best choice of *m*?

maximize
$$\mathbb{E}_p[f(m)] = \sum_{i=1}^n p_i \sqrt{m_i}$$

subject to $m_1 + m_2 + \dots + m_n = 1$
 $m_i \ge 0$

• Nonlinear optimization

- · Continuous function over closed, compact set, global max must exist
- Easy to see if $p_i = 0$, optimal solution has $m_i = 0$, so assume all $p_i > 0$.
- Lagrange multipliers?
 - Lagrange requires continuous first partial derivatives on open set containing $g_i(m) = 0$
 - But $\sqrt{m(i)}$ not even defined for m(i) < 0
- Because objective function is of form $\sum_i g(m_i)$ can concentrate on two components:

$$\begin{array}{ll} \text{maximize} & p_i \sqrt{m_i} + p_j \sqrt{m_j} \\ \text{subject to} & m_i + m_j = c \ c \in [0,1] \\ & m_i \geq 0, \ m_j \geq 0 \end{array}$$

• Maximum either $m_i = 0, m_i = c$, or somwhere in between

• Let
$$f(a) = p_i \sqrt{a} + p_j \sqrt{c-a}$$

- Then $f'(a) = (1/2)p_i/\sqrt{a} (1/2)p_i/\sqrt{c-a}$
- Let $a^* = c \left[1 + \frac{p_j^2}{p_i^2} \right]^{-1}$.
- Solving $f'(a^*) = 0$ gives:

$$a^* = c \left[1 + \frac{p_j^2}{p_i^2} \right]^{-1} = c p_i^2 (p_i^2 + p_j^2)^{-1}$$

That makes

$$c - a^* = cp_j^2(p_i^2 + p_j^2)^{-1},$$

so

$$f(a^*) = p_i \sqrt{cp_i^2 (p_i^2 + p_j^2)^{-1}} + p_j \sqrt{cp_j^2 (p_i^2 + p_j^2)^{-1}}$$
$$= \sqrt{c} [(p_i^2 + p_j^2)(p_i^2 + p_j^2)^{-1/2}] = \sqrt{c} \sqrt{p_i^2 + p_j^2}$$

The max occurs either at the boundary or at a critical point:

$$\begin{array}{ccc} a & f(a) \\ \hline \mathbf{o} & p_j \sqrt{c} \\ a^* & \sqrt{c} \sqrt{p_i^2 + p_j^2} \\ \mathbf{1} & p_i \sqrt{c} \end{array}$$

Since p_i and p_j are nonegative, the max value is

$$f(a^*) = \sqrt{c(p_i^2 + p_j^2)}.$$

Hence a maximum value is found at

$$m_i = \frac{(m_i + m_j)p_i^2}{p_i^2 + p_j^2}, \ m_j = \frac{(m_i + m_j)p_j^2}{p_i^2 + p_j^2},$$

or equivalently:

$$\frac{m_i}{m_j} = \frac{p_i^2}{p_j^2}$$

• Apply to m_1, m_2 to get:

$$m_2 = m_1(20/50)^2, \ m_3 = m_1(15/50)^2, \ m_4 = m_1(15/50)^2.$$

• Since $m_1 + m_2 + m_3 + m_4 = 1$,

$$m_1^* = \frac{1}{1 + (20/50)^2 + (15/50)^2 + (15/50)^2} = \frac{100}{134}.$$

So

$$m^* = \left(\frac{100}{134}, \frac{16}{134}, \frac{9}{134}, \frac{9}{134}\right).$$

Getting right probabilities

- Note m^* closer to true probabilities
- But not quite right
- Q: is there a utility function where the best answer is to give correct probabilities?
- A: YES! Payoff = $\ln(m(\text{correct answer}))$
- We'll see why next time

Chapter 20 Shannon Entropy

Question of the Day Suppose you assign m_1, \ldots, m_n to choices $\{1, \ldots, n\}$. If c is the correct answer, you receive reward

 $r(m_c)$.

Given beliefs p about the true answer, is there a way to choose r so that

$$\operatorname*{arg\,max}_{m} \mathbb{E}_{p}[r(m_{c})] = p$$

subject to $m_1 + \cdots + m_n = 1$?

Today

- Making probabilities maximize utility
- Shannon Entropy

Last time

$$\begin{array}{c|c} i & \mathbb{P}(c=i) \\ \hline 1 & 50\% \\ 2 & 20\% \\ 3 & 15\% \\ 4 & 15\% \end{array}$$

For $r(\alpha) = \sqrt{\alpha}$,

$$m^* = \frac{1}{p_1^2 + \dots + p_n^2} (p_1^2, p_2^2, \dots, p_n^2).$$

Q of the D

• Is there a way to choose *r* so that the probability vector for *c* maximizes expected utility?

Theorem 5

Suppose c has probability vector p with positive entries. The solution to

$$m^* = \arg \max \mathbb{E}_p(\ln(m_c))$$
 subject to $m_1 + \cdots + m_n = 1$

is $m^* = p$.

Example

- Student gives (0.2, 0.6, 0.1, 0.1) for multiple choice (a, b, c, d).
- If answer a is correct, gets reward $\ln(0.2) = -1.60944$
- Best reward possible is o for (1, 0, 0, 0)
- Treating $\ln(0) = -\infty$, reward for (0, 0.6, 0.2, 0.2) is minus infinity! (Instant Flunk!)
- Theorem says: way to maximized expected utility is to report (0.2, 0.6, 0.1, 0.1)

Proof. Let $f(m) = \mathbb{E}_p(\ln(m_c)) = \sum_{i=1}^n p_i \ln(m_i)$.

Then $\lim_{m_1\to 0} \to -\infty$. In particular, there exists ϵ such that $f(m) < f(1/n, 1/n, \ldots, 1/n)$ for $m_i < \epsilon$. Hence we can restrict our search to $m : m_i \ge \epsilon$ for all i.

Since f(m) is continuous over a closed, compact set, there exists a global maximum and a global minimum. Consider a point m in the region. Suppose $m_i + m_j = k$. Now consider trying to maximize

$$g(m_i) = p_i \ln(m_i) + p_j \ln(m_j) = p_i \ln(m_i) + p_j \ln(k - m_i)$$

over $m_i \in [0, k]$. Differentiating $g(m_i)$ with respect to m_i gives:

$$g'(m_i) = p_i/m_i - p_j/(k - m_i)$$

which is positive for $m_i < kp_i/(p_i + p_j)$, zero for $m_i = kp_i/(p_i + p_j)$, and negative for $m_i > kp_i/(p_i + p_j)$.

Hence there is a unique maximum at $m_i = kp_i/(p_i + p_j)$. Since $k = m_i + m_j$,

$$m_i = \frac{(m_i + m_j)p_i}{p_i + p_j} \Rightarrow m_i = m_j p_i / p_j.$$

Since i and j where arbitrarly, this means any maximum solution must have

$$m_i = m_1 p_i / p_1$$

Since $m_1 + m_2 + \dots + m_n = 1$,

$$m_1\left[\frac{p_1}{p_1} + \frac{p_2}{p_1} + \frac{p_3}{p_1} + \dots + \frac{p_n}{p_1}\right] = 1$$

which means $m_1 = p_1/[p_1 + \dots + p_n] = p_1$, and $m_j = p_1p_j/p_1 = p_j$ for all $j \in \{1, \dots, n\}$.

Notes

- For $\alpha \in (0, 1]$, $\ln(\alpha) \leq 0$, reward always nonpositive
- Can extend to $m_i = 0$ by saying $\ln(0) = -\infty$ (maximum negative reward.

Shannon Entropy

• Measures how spread out a probability distribution is

Definition 33 Suppose that *X* is a discrete random variable. The **Shannon Entropy** of *X* is

$$H(X) = \sum_{i:\mathbb{P}(X=i)>0} -\log_2(\mathbb{P}(X=i))\mathbb{P}(X=i)$$

• High entropy = low information, low entropy = high information

Example: total information

- Suppose X = 5 with probability 1.
- Then $H(X) = -\log_2(1) \cdot 5 = 0.$
- H(X) = 0 for any constant random variable

Example: Spread out over $\{1, \ldots, n\}$

- Suppose $X \sim \text{Unif}(\{1, \dots, n\})$
- Then $H(X) = \sum_{i=1}^{n} -\log_2(1/n)(1/n) = \log_2(n)$
- Reflects fact that a number from $\{1, \ldots, n\}$ requires $\log_2(n)$ bits to encode.

Example: one bit

- Suppose $X \sim \text{Bern}(p)$
- $H(X) = -\log_2(p)p \log_2(1-p)(1-p)$
- $\lim_{p\to 0} H(X) = \lim_{p\to 1} H(X) = 0$
- When p = 0 or p = 1, total information
- H(X) maximized when p = 1/2, least information about bit

iid r.v's

Fact 23 Suppose X_1, X_2, \ldots, X_n are independent random variables. Then $H(X_1, X_2, \ldots, X_n) = H(X_1) + H(X_2) + H(X_3) + \cdots + H(X_n).$

Proof. It is easier to start with two independent random variables:

$$H(X_1, X_2) = \sum_{(x_1, x_2)} -\log(\mathbb{P}(X_1 = x_1, X_2 = x_2))\mathbb{P}(X_1 = x_1, X_2 = x_2)$$

= $\sum_{(x_1, x_2)} -\log(\mathbb{P}(X_1 = x_1)\mathbb{P}(X_2 = x_2))\mathbb{P}(X_1 = x_1)\mathbb{P}(X_2 = x_2)$
= $\sum_{(x_1, x_2)} [-\log(\mathbb{P}(X_1 = x_1) - \log(X_2 = x_2)]\mathbb{P}(X_1 = x_1)\mathbb{P}(X_2 = x_2)$
= $\left[\sum_{x_1} -\log(\mathbb{P}(X_1 = x_1))\mathbb{P}(X_1 = x_1)\right] \left[\sum_{x_2} -\log(\mathbb{P}(X_2 = x_2))\mathbb{P}(X_2 = x_2)]\mathbb{P}(X_2 = x_2)$
= $H(X_1)H(X_2)$

The general case for n > 2 then follows from an induction.

Practical consequences

- Shannon Entropy tells how hard it is to compress information
 - iid coin flips cannot be compressed
 - Suppose know pairs of bits are always 01 or 11

Example: 0111110101

- Compression: map o1 to 0, 11 to 1,

Example:
$$0111110101 \rightarrow 01100$$

- Compressed by a factor of 2!
- Security of passwords
 - Most secure password: uniform among 8 letters (26⁸ combinations)
 - Less secure if half the time you use the word "password"

Alphabet notation

- Superscript * on a set indicates all finite words producible using letters from the set.
- Ex: $A = \{a, b, c\}, cbba \in A^*, abbbccba \in A^*, adb \notin A^*.$

Theorem 6 (Shannon source coding Theorem (1948))

Let A_1, A_2 be two finite alphabets. Suppose that X is a r.v. taking values in A_1 , and f is any uniquely decodable code from A_1 to A_2 . Let S_f be the word length of f(X). Then

$$\frac{H(X)}{\log_2(\#A_2)} \le \mathbb{E}[S_f].$$

Moreover, there exists a code f^\ast such that

$$\mathbb{E}(S_{f^*}) < \frac{H(X)}{\log_2(\#A_2)} + 1.$$

Uniquely decodable

- Consider $f_1: a \mapsto 0, b \mapsto 01, c \mapsto 011$
- Given sequence 0110101010001, can recover unique cbbcaab
- Say f_1 is unique decodable.
- Consider $f_2: a \mapsto 0, b \mapsto 01, c \mapsto 10$
- Given sequence 01001 is it acb or bac?

Chapter 21 Game Theory

Ben Polak, Game Theory (Yale University: Open Yale Courses), http://oyc.yale.edu (Accessed 24 Jan, 2014). License: Creative Commons BY-NC-SA

Question of the Day Suppose you are randomly paired with another student in the class. Each of you secretly writes down either α or β . The payoffs are as follows:

- If you put α , other β , you get an A, other gets a C.
- If both put α , both get B-.
- If you put β , other α , you get C, other gets a C.
- If both put β , both get B+.

What should a rational person do?

Today

• Game Theory (interactive decision theory)

What is game theory

- Decision theory is decider versus nature
- · Game theory is decider versus another decider
- That makes things interesting!
- Earlier: nature fixed
- Now other can *cooperate* or *compete*
- Economics, political science, psychology, ecology

Similarities to decision theory

• Have outcome (me,other) matrix

$$\begin{array}{c|c} & \text{other} \\ \hline \\ \mathbf{me} & \frac{\alpha & \beta}{\alpha & (B-,B-) & (A,C)} \\ \beta & (C,A) & (B+,B+) \end{array}$$

- What should a player do?
- Depends on the environment you find yourself in

Self-interest reigns!

- Suppose everyone out for themselves
- Assuming A > B + > B > C:

$$A = 3, B + = 1, B - = 0, C = -1.$$

$$\begin{array}{c|c} & \text{other} \\ \hline \alpha & \beta \\ \hline \alpha & (0,0) & (3,-1) \\ \beta & (-1,3) & (1,1) \end{array}$$

- What should I choose?
 - If other chooses $\alpha :$ I choose α get 0, choose β get -1
 - If other chooses $\beta:$ I choose α get 3, choose β get 1
 - In either case, α better than β
- Choosing α is a dominating strategy.
- So every one (sensible) chooses $\alpha,$ and suffers for it.
- But (β, β) better for both!
- Econ. term: Pareto inefficient
- This payoff matrix has a name: Prisioners dilemma (Flood & Dresher at RAND in 1950)
- Story: Two criminals are taken prisioner. Each can talk or stay silent.
- Both silent: each serves 1 year

- One talks: talker freed, other serves 3 years
- Both talk: both serve 2 year
- In this scenario, dominant strategy is for both to talk.

Some cooperation

- Psychologists have run this test
- People cooperate much more than in the self-interested case
- How that happens: empathy
- Suppose each player not only cares about their own grade...
- ...but also want their partner to do well.
- This changes the payoff matrix
- Suppose if any player gets A while other gets C, guilt changes reward to -1:

$$\begin{array}{c|c} & \text{other} \\ & \alpha & \beta \\ \hline \alpha & (0,0) & (-1,-1) \\ \beta & (-1,-1) & (1,1) \end{array}$$

- Now there is no dominating strategy
- If other picks $\alpha,$ best choice α
- If other picks β , best choice β

Selflessness

- Take it to extreme: add guilt if get A, other gets C...
- ... add happiness for other if get C, other gets A

$$\begin{array}{c|c} & \text{other} \\ & \alpha & \beta \\ \hline \alpha & (0,0) & (-1,-1) \\ \beta & (3,-1) & (1,1) \end{array}$$

- Now β is a dominating strategy: it's better no matter what other chooses
- If everyone thinks this way, leads to best solution!
- In experiments: 30% choose β

Chapter 22

Two person Zero-sum games

Question of the Day In Rock Paper, Scissors, Rock beats Scissors, Paper beats Rock, and Rock beats Paper. For payoff matrix (entry (a, b) means player I wins a and player 2 wins b):

$$\begin{array}{c|c|c} & & & & \text{Player II} \\ \hline & r & p & s \\ \hline r & (0,0) & (-1,1) & (1,-1) \\ p & (1,-1) & (0,0) & (-1,1) \\ s & (-1,1) & (1,-1) & (0,0) \end{array}$$

what is the best strategy for players to use.

Today

- · Adding uncertainty to the mix
- · Zero sum games
- Solution of all 2 by 2 matrix games

Last time: Prisoner's dilemma

- Arises in practice
 - Joint project (shirk or not to shirk)
 - Setting prices: if both set prices low both suffer
- Not enought to have outcomes
- Need the payoffs
- Payoffs like (-3, 1).

Simpler problem: analyze zero sum games.

Definition 34

In a **zero sum game**, the outcomes (a, b) satisfy a + b = 0.

Example Odd-Even game: I pick a number, my partner picks a number. If the sum is even, I pay partner \$1, otherwise the partner pays me \$1.

- If I always pick odd, pretty soon I'll start losing
- Want to pick odd 50% of time, even 50% of time
- No inherent advantage to either player.
- (Rock, paper, scissors) Han Dynasty (206 BC-220AD)
- (Rock,paper,scissors,lizard,spock) (Kass & Bryla)

Definition 35

The strategic form (aka normal form) of a two-person zero-sum game is a triple (X, Y, A) where

- 1. X is the nonempty set of strategies for Player I
- 2. Y is the nonempty set of strategies for Player II
- 3. $A: (X \times Y) \to \mathbb{R}$ is the payoff function

Game proceeds as follows:

- Simultaneously: Player I chooses $x \in X$ and Player II chooses $y \in Y$
- Player II pays A(x, y) dollars to Player I

Example: Modified Odd-Even

- $X = \{1, 2\}$ (1 = odd, 2 = even)
- $Y = \{1, 2\}$ (1 = odd, 2 = even)
- Suppose payoff matrix is:

Player II
Player I
$$\begin{array}{c|c} & Player II \\ \hline 1 & 2 \\ \hline 1 & -2 & 3 \\ 2 & 3 & -4 \end{array}$$

Does one side have an advantage in this game?

- Suppose Player I plays 1 with probability 3/5
- What is expected payoff:

II plays 1 :
$$(3/5)(-2) + (2/5)(3) = 0$$

II plays 2 : $(3/5)(3) + (2/5)(-4) = 1/5$

- Remember this is 1/5 that II pays I, so at best II can break even with this strategy.
- Can Player I do better?
- Let p be probability Player I plays 1:

II plays 1 :
$$p(-2) + (1-p)(3) = 3 - 5p$$

II plays 2 : $p(3) + (1-p)(-4) = 7p - 4$

- First equation wants p large, second equation wants p small
- Maximum occurs when they are equal:

$$3 - 5p = 7p - 4 \Rightarrow p = 7/12.$$

• Player I optimal strategy: call 1 with prob 7/12

What should Player II do?

- If Player I playing optimally, losing 3 5(7/12) = 7(7/12) 4 = 1/12 each play
- Can Player II prevent Player I from doing any better?
- Now let p be prob Player II plays 1

I plays 1 :
$$p(-2) + (1-p)(3) = 3 - 5p$$

I plays 2 : $p(3) + (1-p)(-4) = 7p - 4$

- So again, minimize loss when p = 7/12.
- By Player II playing 1 w/ prob 7/12, ensures that one average loses at most 1/12 each play no matter what I does
- Turns out this always happens: value of optimal winning strategy for Player I equals the value of the optimal winning strategy for Player II.

Terminology

- Strategies in X or Y called *pure strategies*
- Choosing a strategy in X (or Y) at random is called a *mixed strategy*

Theorem 7 (The Minimax Theorem)

For every finite two person zero sum game

- 1. There is a number V called the **value** of the game
- 2. There is a mixed strategy for Player I such that I wins on average V regardless of what II does.
- 3. There is a mixed strategy for Player II such that II loses average V regardless of what I does.

Fairness

- If V = 0, then the game is fair
- If V > 0, then the game favors Player I
- If V<0, then the game favors Player II

Example Change the Odd-Even payoff matrix again to A(1,2) = 2:

		other		
		1	2	
me	1	-2	2	
	2	3	-4	

What is the value of this game?

Answer

- Can analyze from either Player I or Player II perspective.
- Player I:

II plays 1 :
$$p(-2) + (1-p)(2) = 2 - 4p$$

II plays 2 : $p(3) + (1-p)(-4) = 7p - 4$

- Setting them equal gives: $2-4p=7p-4 \Rightarrow p=6/11$
- This makes the value V = 2 4(6/11) = -2/11

Answer 2: Player II perspective

• From Player II's point of view:

I plays 1 :
$$p(-2) + (1-p)(3) = 3 - 5p$$

I plays 2 : $p(2) + (1-p)(-4) = 6p - 4$

- So $3 5p = 6p 4 \Rightarrow p = 7/11$
- So the value is V = 6(7/11) 4 = -2/11
- Same (as guaranteed by Minimax Theorem)

Relation to linear programming

- For those who have had Math 187
- Minimax Theorem is a special case of LP duality
- Can solve games with #(X) = n, #(Y) = m using linear programming
Chapter 23 Nash Equilibria

Question of the Day Is there an optimal mixed strategy when dealing with 3 or more players? How about for non zero sum games?

Today

• Nash equilibria

Last time

- For two person, zero sum games, unique value of game
- Use minimax strategy
- For nonzero sum, or three or more players, not always an optimal solution
- · Something weaker

Definition 36

The normal form (aka strategic form) of an *n*-player game is any list $G = (S_1, \ldots, S_n; u_1, \ldots, u_n)$ where S_i is the set of strategies for player *i*, and $u_i : (S_1 \times \cdots \times S_n) \to \mathbb{R}$ is player *i*'s payoff function.

Assume each player is trying to maximize their expected utility. Consider the following game:

	Player 2		
		a	b
Player 1	a	(1,2)	(0, 0)
	b	(0,0)	(2, 1)

No pure strategy for Player 1 or 2 is dominant.

Note:

- For strategy (a, a), neither can switch without lowering payoff
- For strategy (b, b), neither can switch without lowering payoff
- Call (a, a) and (b, b) Nash equilibria

Now look at the following game

Player 2
Player 1
$$\begin{array}{c|c}
 & a & b \\
\hline
 & a & (1,2) & (4,0) \\
 & b & (3,0) & (2,1) \\
\end{array}$$

From any pure strategy, some player wants to switch. Means no pure strategies are Nash equilibria.

Definition 37 A probability distribution σ_i on S_i (the strategy set for player *i*) is called a **mixed** strategy.

Notation

- Let σ_i be the mixed strategy played by player i
- Let $\sigma_{-i} = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$ be the mixed strategies played by every player other than i
- Let $u_i(\sigma_1, \ldots, \sigma_n)$ be the expected payoff for Player *i* when each player uses mixed strategy σ_i

Definition 38

 σ is a **Nash equilibrium** if for all i

$$u_i(\sigma) = \max_{\sigma'_i} u_i(\sigma'_i, \sigma_{-i}).$$

In other words, a set of mixed strategies is a Nash equilibrium if changing one player's mixed strategy leads to a lower expected payoff for that person.

Theorem 8 (Nash 1951) Any finite game has at least one Nash equilibrium.

[Proof uses Kakutani's fixed point theorem from analysis.]

Example Another way to say it is that at the Nash equilibrium, every player is insensitive to changes in strategy.

Look again at

Player 2
Player 1
$$\begin{array}{c|c}
 & a & b \\
\hline
 & a & (1,2) & (4,0) \\
 & b & (3,0) & (2,1) \\
\end{array}$$

First, no pure strategy by the players is a Nash equilibrium.

Suppose Player I decides to use a mixed strategy where a is played with probability p. Then the expected payoff for Player II depends on whether or not Player II plays a or b:

> $\mathbb{E}[\text{Payoff for Player II} | \text{Player II plays } a] = 2p + 0(1-p)$ $\mathbb{E}[\text{Payoff for Player II} | \text{Player II plays } b] = 0p + 1(1-p)$

Now suppose Player I chooses p in order to make these expected returns equal!

$$2p = 1 - p \Rightarrow 3p = 1 \Rightarrow 1/3.$$

When Player I plays a with probability 1/3, then no matter what strategy Player II uses (random or deterministic), Player II has the same expected return.

But two can play at that game (so to speak)!

Suppose Player II plays *a* with probability *p* and *b* with probability 1 - p.

 $\mathbb{E}[\text{Payoff for Player I} | \text{Player I plays } a] = 1p + 4(1-p)$ $\mathbb{E}[\text{Payoff for Player I} | \text{Player I plays } b] = 3p + 2(1-p)$

Choosing p to make these expected returns equal:

$$1p + 4 - 4p = 3p + 2 - 2p \Rightarrow 2 = 4p \Rightarrow p = 1/2.$$

So a Nash equibilium is:

Example High dimensional example: Calling the police

- *n* identical players
- Each player can choose to call the police, or not to report a crime
- Benefit to all if someone calls the police is \boldsymbol{x}
- Cost of calling the police is 1 (assume x > 1)

105 184

• Example: n = 3, players 1 and 2 call the police, 3 does not

payoff
$$= (x - 1, x - 1, x)$$
.

- Because of symmetry, look for a Nash equilibrium where each player has mixed strategy of calling the police with probability p.
- p chance call police, get x 1
- (1-p) chance no police, $(1-p)^{n-1}$ chance no one else calls police either.
- When does indifference to choices occur?

$$x - 1 = x(1 - (1 - p)^{n-1}).$$

- Player i in different when $\boxed{p = 1 - (1/x)^{1/(n-1)}}$
- So what is the chance the police are called?

$$1 - (1 - p)^n = 1 - (1/x)^{n/(n-1)}.$$

- Decreasing in *n* [More people means *less* likely to call!]
- Approaches 1 1/x as $n \to \infty$

Chapter 24 Randomized algorithms

Question of the Day Suppose that at least one element of an array of *n* elements has value *a*. What is the fastest way to find such an element?

Today

• Randomized algorithms

Definition 39

A randomized algorithm uses randomness as part of its running procedure.

Many types, all named after famous gambling locations: Monte Carlo, Las Vegas, Atlantic City.

Definition 40 A **Monte Carlo** randomized algorithm returns a random result.

Definition 41

A Las Vegas randomized algorithm always returns the correct result.

Definition 42

An Atlantic City algorithm has a 2/3 chance of returning the correct result.

Later on, computer science got more stodgy in their names:

Definition 43

An algorithm is in **Randomized polynomial time (RP)** if

- It takes a polynomial number of steps in input size [More precisely, if a probabilistic Turing machine always runs in time polynomial in the input size.]
- 2. If correct answer is False, then returns False.
- 3. If correct answer is True, then returns True with probability at least 1/2.
- Note the 1/2 is unimportant. If you want 1/8, just run algorithm 3 times, only report T if TTT.
- Always guaranteed to get F answers right.

Definition 44

An algorithm is in **bounded-error probabilisitic polynomial time (BPP)** if it

- It takes a polynomial number of steps in input size to answer T or F. [More precisely, if a probabilistic Turing machine always runs in time polynomial in the input size.]
- 2. The chance the answer is correct is at least 2/3.
- Can always get answer correct 1/2 of time: just flip a coin.
- As with RP, run BPP multiple times to get 2/3 arbitrarily close to 1.

Definition 45

Say that a problem is in RP (or BPP) if there exists an RP (respectively BPP) algorithm for the problem.

- Hopefully clear from context whether talking about the set of problems...
- ... or the set of algorithms.

Definition 46

A problem is in **Nondeterministic Polynomial time (NP)** if there exists a proof of the answer that can be checked in polynomial time.

• Millenium problem: Does P = NP? (Can every problem whose answer can be checked in polynomial time be solved in polynomial time?)

- Would be happy with RP = NP or BPP = NP.
- Also unknown: can all randomized algorithms be derandomized? Does BPP = P?

Qotd

- How many steps does deterministic algorithm take?
- Basic Alg: look at array elements in order 1 through n, quit when find value a
- Could take 1 step, could take $n \ {\rm steps}$
- Usually worried about worst case behavior, so n steps.

Randomized search

- Suppose I use the following algorithm
 - Choose a random number uniformly from 1 to n
 - If value is *a*, quit, otherwise goto first step.
- Always returns correct answer, so Las Vegas algorithm.
- Number of steps $T, T \sim \text{Geo}(1/n)$, so $\mathbb{E}[T] = n$.
- Unbounded possible number of steps.
- Not any better than original!

Better rand. alg.

- Improvement:
 - Choose a random number i uniformly from 1 to n
 - Look at elements i, i + 1, ..., n, 1, 2, ..., i 1 in order, quit when you find value a.
- Let T be number of steps.
- Let x be location of value i. Then $\mathbb{P}(T = j)$ is probability $i + j = x \mod n = 1/n$.
- So $T \sim \text{Unif}(\{1, ..., n\}, \mathbb{E}[T] = (n+1)/2.$
- Nearly twice as fast as worst case scenario!
- Can't be worse than n steps.
- So simple randomness cuts average running time in half.

Basic randomness principle

- Using randomness "prevents" worst case scenario.
- Identical to using deterministic algorithm on randomly sorted data.

Suppose there are two elements of array with value \boldsymbol{a}

- Original algorithm
 - Could be at positions n 1 and n.
 - Worst case running time n-1.
- First rand. alg.
 - Now 2/n chance of picking a value a element.
 - So $T \sim \text{Geo}(2/n)$, $\mathbb{E}[T] = n/2$.
- Second rand. alg.
 - Could be at positions n 1 and n.
 - So $T \sim \text{Unif}(\{n-1, n-2, n-3, \dots, 1, 1\})$
 - So

$$\mathbb{E}[T] = \frac{1}{n}(1) + \left(1 - \frac{1}{n}\right)\frac{1 + (n-1)}{2} = n/2 + 1/n - 1/2.$$

- Not much improvement.
- Can we do better?

Random order

- Doing even better:
 - Randomly permute the elements
 - Use deterministic algorithm from there.
- Let T be position of first a.
- Then $\mathbb{P}(T \ge i) = \frac{n-i}{n} \frac{n-i-1}{n-1}$.
- Tail sum formula:

$$\mathbb{E}[T] = \sum_{i=1}^{n} \mathbb{P}(T \ge i) = \frac{n-2}{3}.$$

• [Uses $\sum_{i=1}^{n} i = n(n+1)/2$ and $\sum_{i=1}^{n} i^2 = (1/6)n(n+1)(2n+1)$.]

Chapter 25 Randomized verification

Question of the Day Given *n* by *n* matrices *A*, *B*, and *C*, does $A \cdot B = C$?

Today

- Order notation
- Freivalds' Algorithm

Running time

- How long does it take to compute an answer?
- Consider finding

$$\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

• To calculate the upper left element of answer, need to find:

$$1 \cdot 0 + 0 \cdot 1 + 1 \cdot 1,$$

so three multiplications and two additions.

- There are 9 entries to the product, so need 27 mult. and 18 additions
- For any n by n matrix, need

$$n^3$$
 multiplications and $n^2(n-1)$ additions

• Total time: $2n^3 - n^2$ steps.

Order notation

- Idea: the $-n^2$ is unimportant, small compared to $2n^3$.
- The 2 unimportant since we don't know how long a single arithmetic step takes
- The n^3 part tells us how many steps are taken
- For f(n) the number of steps needed for compute n by n matrices using the basic approach

$$f(n) = \Theta(n^3).$$

[Abuse of notation alert: technically should write $f(n) \in \Theta(n^3)$]

- Informally O(g(n)) means $f(n) \leq cg(n)$ for some c...
- ... $\Omega(g(n))$ means $f(n) \ge cg(n)$ for some c...
- ... $\Theta(g(n))$ means $c_1 \leq f_n \leq c_2$ for some c_1, c_2

Definition 47

Say that f(n) is **Big-O** of g(n) if

 $f(n) = O(g(n)) \Leftrightarrow (\exists c > 0) (\exists N) (\forall n \ge N) (f(n) \le cg(n)).$

Definition 48

Say that f(n) is **Big-Omega** of g(n) if

$$f(n) = \Omega(g(n)) \Leftrightarrow (\exists c > 0) (\exists N) (\forall n \ge N) (f(n) \ge cg(n)).$$

Definition 49

Say that f(n) is **Big-Theta** of g(n) if

$$f(n) = \Theta(g(n)) \Leftrightarrow (\exists c_1 > 0)(\exists c_2 > 0)(\exists N)(\forall n \ge N)(c_1g(n) \le f(n) \le c_2g(n)).$$

Example:

- Fact: $2n^3 n^2$ is $\Theta(n^3)$
- Proof: Let $N = 1, c_1 = 1, c_2 = 2$. Let $n \ge 1$. Then

$$n^3 \le 2n^3 - n^2 \le 2n^3,$$

which completes the proof.

Multivariable order notation

- Need to be a little more careful in defining order notation for multiple variables.
- Example: Want $3mn^2 = O(mn^2)$.

Definition 50

Say that $f(n_1, n_2, \ldots, n_d)$ is **Big-** $O/\Omega/Theta$ of $g(n_1, n_2, \ldots, n_d)$ for all $i \in \{1, \ldots, d\} \exists N_{-i}$ such that for all n_{-i} term by term at least N_{-i} , $f(n_1, \ldots, n_d)$ is Big-O/Omega/Theta of $g(n_1, \ldots, n_d)$ as a function of n_i .

Matrix multiplication

- Basic method: $\Theta(n^3)$
- Strassen algorithm (1969): $\Theta(n^{2.807})$
- [small improvements]
- Coppersmith-Winograd (1990) $\Theta(n^{2.376})$
- Stothers (2010) $\Theta(n^{2.3736})$.
- Williams algorithm (2011) $\Theta(n^{2.3727})$

Verifying

- So is there a faster way to verify that $A \cdot B = C$?
- Freivald's algorithm (1977) $O(n^2)$
- More precisely: after $\Theta(kn^2)$ steps, the probability of returning the wrong answer is at most $(1/2)^k$.
- Idea: Basic multiplication $A\vec{x}$ of n by n matrix times n by 1 vector takes $\Theta(n^2)$ steps.
- [Can't improve since matrix has n^2 entries and must look at each entry.]
- Pick random vector \vec{X} .
- Find $A(B\vec{X})$ using two matrix vector multiplications.
- Find $C\vec{X}$ using one matrix vector multiplication.
- If AB = C, then $AB\vec{X}$ always equals $C\vec{X}$.
- But if $AB\vec{X} \neq C\vec{X}$, have proof that $AB \neq C$

Freivald's method. Input: A, B, C, three n by n matrices

- 1) Generate $\vec{X} \sim \text{Unif}(\{0,1\}^n)$
- 2) Calculate $Y \leftarrow B\vec{X}$. Calculate $W \leftarrow A\vec{Y}$. Calculate $Z \leftarrow C\vec{X}$
- 3) If W = Z return true, otherwise return false.

Example: Does

$$\begin{pmatrix} -2 & 1 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & -3 \end{pmatrix} = \begin{pmatrix} -2 & -4 \\ 3 & 1 \end{pmatrix}$$

Pick random vector from $\{0,1\}^2$, perhaps $\vec{X} = (1,0)^T$:

$$\begin{pmatrix} -2 & 1 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & -3 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} -2 & 1 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} -2 \\ 3 \end{pmatrix}$$

whereas:

$$\begin{pmatrix} -2 & -4 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} -2 \\ 3 \end{pmatrix}.$$

So that would return TRUE. But what if $\vec{X} = (1, 1)$:

$$\begin{pmatrix} -2 & 1 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & -3 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} -2 & 1 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 2 \\ -3 \end{pmatrix} = \begin{pmatrix} -7 \\ 3 \end{pmatrix}$$

while

$$\begin{pmatrix} -2 & -4 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} -6 \\ 4 \end{pmatrix}.$$

so result is FALSE.

When returns FALSE, always correct, when returns TRUE, only sometimes correct.

Fact 24

The chance that Freivald's Algorithm is incorrect when it returns TRUE is at most 1/2.

Proof. The only way Freivald can be wrong is if $A \cdot B \neq C$. Hence for some (i, j), $[AB](i, j) \neq C(i, j)$. Now

$$[AB\vec{X}](i,j) = \sum_{k=1}^{n} (AB)_{ik} X_k,$$

and

$$[C\vec{X}](i,j) = \sum_{k=1}^{n} C(i,i)X(k).$$

114	184
-----	-----

When $[ABX](i, j) \neq [CX](i, j)$, that's the same as $[(AB - C)(X)](i, j) \neq 0$. [(AB - C)(X)](i, j) can be written as

$$\sum_{k=1}^{n} (AB)_{ik} X_k - \sum_{k=1}^{n} C(i,k) X(k) = ([AB](i,j) - C(i,j)) X(j) + Y$$

where X(j) and Y are random variables.

Now either Y = 0 or $Y \neq 0$, so

$$\begin{split} \mathbb{P}(([AB](i,j) - C(i,j))X(j) + Y &= 0] \\ &= \mathbb{P}(([AB](i,j) - C(i,j))X(j) = 0, Y = 0] \\ &+ \mathbb{P}(([AB](i,j) - C(i,j))X(j) + Y = 0, Y \neq 0] \\ &= \mathbb{P}(([AB](i,j) - C(i,j))X(j) = 0|Y = 0]\mathbb{P}(Y = 0) \\ &+ \mathbb{P}(([AB](i,j) - C(i,j))X(j) + Y = 0|Y \neq 0]\mathbb{P}(Y \neq 0) \end{split}$$

Since $[AB](i, j) - C(i, j) \neq 0$, the only way for ([AB](i, j) - C(i, j))X(j) = 0 is if X(j) = 0. This has probability 1/2.

Similarly, if $Y \neq 0$, then at most one choice of X(j) can make

$$([AB](i,j) - C(i,j))X(j) = 0,$$

and that choice happens with probability at most 1/2. Hence

$$\mathbb{P}(([AB](i,j) - C(i,j))X(j) + Y = 0]$$

$$\leq (1/2)\mathbb{P}(Y = 0) + (1/2)\mathbb{P}(Y \neq 0)$$

$$= 1/2.$$

 \square

Improving	further
-----------	---------

- Suppose $X \sim \{0, 1, ..., k 1\}$
- Chance of giving wrong answer at most 1/k
- Suppose $X \sim \text{Unif}([0,1])$
- Then chance of giving wrong answer is o
- Does fit Turing machine model

Chapter 26 Randomized QuickSelect

Question of the Day What is the fastest way to find a median of a group of numbers?

Today

· Find medians quickly

Definition 51

A **median** of a finite well-ordered set A is any $a \in A$ such that $\#\{b : a \le b\} \ge (1/2) \# A$ and $\#\{b : a \ge b\} \ge (1/2) \# A$

Comments

- Median, mode, and mean often called measures of central tendancy
- Mode easiest to find, but most worthless estimator
- Median most robust estimator (insensitive to outliers)
- · For well behaved functions, higher variance in median than in the mean
- When #(A) odd, median value is unqiue
- When #(A) even, can be two median values
- Often average the results

The problem

- Finding median(s) of $A \subseteq \mathbb{R}, \#A = n$
- Brute force: sort elements, then pick middle one (or two)
- Sorting n elements takes $O(n\ln(n))$ time

116 184

- Can we do better?
- Use randomness!

Split your set randomly

- Pick an element $a \sim \mathsf{Unif}(A)$
- Compare each remaining element of A to a
- If at most *a*, put to left of *a*, otherwise, put on right
- Example: $A = \{1, 9, 3, 7, 2, 4, 2\}$
- Choose a = 2 randomly
- Sort to get $A_{\text{left}} = \{1, 2, 2\} \le 2 \le \{9, 3, 7, 4\} = A_{\text{right}}$
- Now instead of wanting median of A, want smallest element of A_{right}

To find the kth smallest element of A:

Quio	ckSelect	Input:	A, k
1)	Repeat		
2)	$a \leftarrow Unif(A), r \leftarrow 0, \ell \leftarrow 0, A_L \leftarrow \{$	$A_R \leftarrow$	- {}
3)	For each $b \in A \setminus \{a\}$		
4)	If $a \ge b$		
5)	$A_L \leftarrow A_L \cup \{b\}, \ell \leftarrow \ell + 1$		
6)	else		
7)	$A_R \leftarrow A_R \cup \{b\}, r \leftarrow r+1$		
8)	If $\#(A_L) = k - 1$		
9)	$A \leftarrow \{a\}$		
10)	If $\#(A_L) > k - 1$		
11)	$A \leftarrow A_L$		
12)	If $\#(A_L) < k-1$		
13)	$A \leftarrow A_R, k \leftarrow k - \#(A_L) - 1$		
14)	Until $\#(A) = 1$		
15)	Return the only element of A .		

Fact 25

The total number of comparisions used by QuickSelect is at most n(n-1)/2. The expected number of comparisons used by QuickSelect is at most 2n.

Proof. Suppose k = 1, so we are trying to find the maximum value. The first time we run the repeat loop requires n - 1 comparisions. If we are unlucky and had chosen the largest value of A, then A_L has size n - 1 and k stays the same. The next step would then require n - 2 comparisions, and so on down to the last step that requires only 1 comparision. Hence the total number of comparisions is:

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

Now bound the average number of comparisions. The rough idea is that the size of the set A is being chopped in half at each step, so after t iterations of the repeat loop, the average value of n is $n(3/4)^t$. The total number of comparisons at step t is one fewer than the size of A, making the total expected number of comparisons at most

$$n + n\frac{3}{4} + n\frac{9}{16} + \dots = n(1/[1 - 3/4]) = 4n.$$

Now let's make that precise! Let T denote the amount of comparisons used in a run of the program. Our goal is to find $\mathbb{E}[T|n]$.

Recall the notion of order statistics, which are just the sorted values of *A*:

$$A = a_{(1)} \le a_{(2)} \le a_{(3)} \le \dots \le a_{(n)}.$$

Ties end up helping us, but they make the analysis trickier, so assume for now that there are no ties among the elements of A. Then since the element $a \sim \text{Unif}(A)$, if $a = a_{(i)}$, then $i \sim \text{Unif}(\{1, 2, \ldots, n\})$.

No matter what happens, after we choose a, there are n-1 comparisons. What happens next depends on the value of i.

Case 1: i = k. Then we have found the element, and no more comparisons are necessary. Case 2: i > k. Then $A \leftarrow A_L$. The size of A_L is just i - 1.

Case 3: i < k. Then $A \leftarrow A_R$. The size of A_R is n - i.

Putting this together gives:

$$\begin{array}{c|cccc} i \text{ where } a = a_{(i)} & \text{Size of next } A & \text{Comparisons needed} \\ \hline 1 & n-1 & \mathbb{E}[T|n-1] \\ 2 & n-2 & \mathbb{E}[T|n-2] \\ \vdots & \vdots & \vdots \\ k-1 & n-k+1 & \mathbb{E}[T|n-k+1] \\ k & 1 & \mathbf{0} \\ k+1 & k & \mathbb{E}[T|k] \\ k+2 & k+1 & \mathbb{E}[T|k] \\ k+2 & k+1 & \mathbb{E}[T|k+1] \\ \vdots & \vdots & \vdots \\ n & n-1 & \mathbb{E}[T|n-1]. \end{array}$$

Each of these possibilities occurs with probability 1/n. This gives the following recurrence relation:

$$\mathbb{E}[T|n] = \frac{1}{n} \left[\sum_{i=1}^{k-1} \mathbb{E}[T|n-i] + \sum_{i=1}^{n-k} \mathbb{E}[T|n-i] \right]$$

Let's prove that $\mathbb{E}[T|n] \leq 4n$ using strong induction. When n = 1, the number of comparisons used is 0, so that works.

For our strong induction hypothesis, suppose $\mathbb{E}[T|n'] \leq 4n'$ for all $n' \leq n$. Consider $\mathbb{E}[T|n+1]$. From the recurrence above and using the induction hypothesis:

$$\begin{split} \mathbb{E}[T|n+1] &= n-1 + \frac{1}{n+1} \left[\sum_{i=1}^{k-1} \mathbb{E}[T|n+1-i] + \sum_{i=1}^{n-k} \mathbb{E}[T|n+1-i] \right] \\ &\leq n + \frac{1}{n+1} \left[\sum_{i=1}^{k-1} 4(n+1-i) + \sum_{i=1}^{n-k} 4(n-i) \right] \\ &\leq n + \frac{4}{n+1} \left[\frac{(n+1)(n)}{2} - \frac{(n-k+2)(n-k+1)}{2} \right] \\ &+ \frac{(n)(n-1)}{2} - \frac{(k)(k-1)}{2} \right] \\ &\leq n + \frac{4}{n+1} \left[\frac{(n+1)(n)}{2} - \frac{(n-k+2)(n-k+1)}{2} \right] \\ &+ \frac{(n+1)(n)}{2} - \frac{(k+1)(k)}{2} \right] \\ &\leq n + \frac{4}{n+1} \left[(n+1)(n) - \frac{(n-k)(n-k)}{2} - \frac{k^2}{2} \right] \end{split}$$

It is an easy maximization problem to show that the right hand side is largest for $k \in [1, n]$ when k = n/2. Hence

$$\mathbb{E}[T|n+1] \le n + \frac{4}{n+1} \left[(n+1)(n) - \frac{(n/2)^2}{2} - \frac{(n/2)^2}{2} \right]$$
$$\le n + \frac{4}{n+1} \left[(n+1)(n) - (n+1)(n)(1/4) \right]$$
$$\le n + 4(3/4)n$$
$$= 4n$$

which completes the strong induction.

Trickier question

- What is the variance?
- Easy to resolve through Monte Carlo simulation

Bounding tails easier

- Suppose take 8n comparisons
- By Markov's inequality, only 1/2 chance algorithm requires that many
- n always decreases in algorithm
- So for T number of comparisons...

 $\mathbb{P}(T \ge 8n) \le 1/2, \ \mathbb{P}(T \ge 16n) \le (1/2)^2, \ \mathbb{P}(T \ge 24n) \le (1/2)^3.$

Fact 26
If *T* is the number of samples used by QuickSelect, and
$$\alpha > 0$$
,

 $\mathbb{P}(T \ge \alpha) \le 2e^{-\alpha/11.55}.$

Proof. From the above Markov inequality argument,

$$\mathbb{P}(T \ge \alpha n) \le (1/2)^{\lfloor \alpha n/(8n) \rfloor}$$
$$\le (1/2)^{\alpha/8-1}$$
$$= 2e^{\ln(1/2)\alpha/8}$$
$$\le 2e^{-\alpha/11.55}.$$

	-	-	-
_			

Improvements

- Can slightly improve analysis to get better than 4n by also analyzing how k changes
- Can't do much better than that.
- · Floyd-Rivest Select Algorithm, average comparisons

$$n + \min\{k, n - k\} + o(n).$$

Definition 52

Say that f(n) is **little-o** of g(n) if

$$(\forall \epsilon > 0)(\exists N)(\forall n \ge N)(f(n)/g(n) < \epsilon).$$

Fact 27 It is true that f(n) = o(g(n)) if $\lim_{n\to\infty} f(n)/g(n) = 0$.

Idea of FR-Select

- 1. Let B be $n^{3/4}$ random items from A
- 2. Sort the elements of B
- 3. Find two order statistics u and v of B: $(k/n)n^{3/4} \sqrt{n}$ and $(k/n)n^{3/4} + \sqrt{n}$
- 4. If k > 1/2, compare A against u then v if necessary
- 5. If $k \leq 1/2$, compare A against v then u if necessary
- 6. Now have $A_L < u < A_M < v < A_L$ where A_m roughly $2\sqrt{n}$.
- 7. Sort everything in A_m , then know exactly where $a_{(k)} \in A_m$.

Chapter 27 Randomization for IP's and LP's

Question of the Day Consider the set $\mathcal{U} = \{1, 2, 3, 4, 5, 6, 7\}$ and subsets of \mathcal{U} :

$$S_1 = \{1, 6\}, S_2 = \{3, 4\}, S_3 = \{5\}, S_4 = \{1, 2, 3\}$$

$$S_5 = \{2, 4, 6, 7\}, S_6 = \{4, 5, 6\}, S_7 = \{7\}.$$

Find the smallest collection of S_1 whose union is \mathcal{U} .

Today

- Set cover problem
- Integer programming
- Linear programming
- Randomized rounding

Definition 53

Given \mathcal{U} with *n* elements, and a collection S_1, S_2, \ldots, S_m of subsets of \mathcal{U} , find the smallest subset of $\{S_1, \ldots, S_m\}$ whose union is \mathcal{U} . This is the **Set Cover Problem**

Brute force

- Each S_i is either in or out of the collection.
- So at most $2^6 1$ (-1 since empty set doesn't work) to try
- In general $2^m 1$.
- Total running time $O(n2^m)$.

Greedy

• Grab the S_i next that covers more values than any other

 S_5 then S_1 then S_2 then S_3 .

• Not always optimal! $S_4 \cup S_5 \cup S_7 = \mathcal{U}$.

Problem is in NP

- Decision Set Cover problem is in NP
- DSC: Given $k \leq m$, are there k different S_i whose union is \mathcal{U} ?

Definition 54

A problem is in **NP** if the answer can be checked in time polynomial in the input size of the problem.

- For Decision Set Cover, easy to check if an answer is correct
- Given Set Cover solution, get Decision Set Cover answer
- Given m runs of Decision Set Cover, get Set Cover solution

Integer Program

- A program is just a list (music, computer commands, constraints)
- Integer refers to fact that variables must be integers.
- For set cover, let $x_i = \mathbb{1}(S_i \text{ is part of collection})$
- Example: x = (1, 1, 1, 0, 1, 0, 1) says collection is S_1, S_2, S_3, S_5, S_7 .
- How can we be sure that element 3 is covered?
- Element 3 appears in S_2 and S_4
- So Element 3 is covered iff $x_2+x_4\geq 1$
- Want to use as few S_i as possible.
- So the integer program (IP) is:

$$\max x_{1} + x_{2} + \dots + x_{7}$$

subject to $x_{1} + x_{4} \ge 1$
 $x_{4} + x_{5} \ge 1$
 $x_{2} + x_{4} \ge 1$
 $x_{2} + x_{5} + x_{6} \ge 1$
 $x_{3} + x_{6} \ge 1$
 $x_{1} + x_{5} + x_{6} \ge 1$
 $x_{5} + x_{7} \ge 1$
 $x_{i} \in \{0, 1\} \ \forall i.$

Randomized rounding

- With $x_i \in \{0, 1\} \forall i$, it is an IP
- Without that constraint, it is a Linear Program (LP)
- LP can be solved in polynomial time! (Karmarkar 1979)
- Idea: solve the LP to get x^*
- For each i, draw U_i
- If $U_i \leq x_i^*$, make $y_i = 1$
- Note $\mathbb{E}[y_1 + y_2 + \dots + y_n] = x_1^* + x_2^* + \dots + x_n^*$

Consider one constraint

- $x_1 + x_5 + x_6 \ge 1$
- So what is $\mathbb{P}(y_1 + y_5 + y_6 \ge 1)$? At least:

$$\begin{array}{ll} \min & 1-(1-y_1)(1-y_5)(1-y_6)\\ \text{subject to} \; y_1+y_5+y_6=1\\ & y_1,y_5,y_6\geq 0 \end{array}$$

- · Continuous function over closed, bounded regions, global max exists
- Note, if $y_i \neq y_j$, replacing both with $(y_i + y_j)/2$ increases the objective function.
- So global max has $y_1 = y_5 = y_6 = 1/3$
- Which makes $\mathbb{P}(y_1 + y_5 + y_6 \ge 1) \ge 1 (1 1/3)^3 = 1 8/27.$

124 184

- In general, with at most k variables in constraint, chance of not being covered is at most:

$$(1 - 1/k)^k \le e^{-1}.$$

- If one if covered, remaining are more likely to be covered, so chance all are covered is at least (1 - e⁻¹)ⁿ.
- Too small!

Improving the probability

- Draw $U_1, U_2, ..., U_t \sim Unif([0, 1])$ iid
- Make $y_i = 1(\min\{U_1, \dots, U_t\} \le x_i^*)$
- Now chance of element *i* not being covered is at most

 $1/e^t$

• Recall $\mathbb{P}(\cup A_i) \leq \sum \mathbb{P}(\cup A_i)$. So chance $\exists i$ that is not covered is at most

$$n/e^t$$
.

• Let $Y = \sum_{i=1}^{n} y_i$. Then

$$\mathbb{E}[Y] = \sum_{i=1}^{n} 1 - (1 - x_i)^t \le t \sum_{i=1}^{n} x_i.$$

[Slope of $1 - (1 - x_i)^t$ at most t.]

- By Markov, $\mathbb{P}(Y \ge 2\mathbb{E}[Y]) \le 1/2$.
- If $t = \ln(4n)$, then probability all are not covered is at most 1/4.
- So

 $\mathbb{P}(Y \geq 2\mathbb{E}[Y] \text{ or some elements not covered}) \leq 1/2 + 1/4 = 3/4.$

• So expect to need 4 tries to get both covered and within factor of $\ln(4n)$ of $\sum x_i^*$

Other places Randomized Rounding used

• Facility location problem

Simplex method

- Randomness also helpful in solving LP's quickly
- The simplex method is a way to solve LP's
- The method is fast in practice, but slow in theory
- At some steps, you have choice of where to move next
- By choosing randomly, can make Simplex expected polynomial time

Chapter 28 Forecasting

Question of the Day [Winston 1991, p. 1164] Suppose the following data for T.V. Sales are taken:

Month	Sales
1	30
2	32
3	30
4	39
5	33
6	34

Today

Forecasting

Predicting the future

- 2 main methods
 - Extrapolation
 - * moving averages
 - * smoothing
 - Causal Methods
 - * build statistical model of data
 - * linear regression

Moving averages

- Simplest moving average is just the average of last n observations

Qotd

Month	Sales	Prediction	Error
1	30	-	-
2	32	-	-
3	30	-	-
4	39	$\frac{30+32+30}{3} = 30.66$	8.333
5	33	$\frac{32+30+39}{3} = 33.66$	-0.6666
6	34	$\frac{30+39+33}{3} = 34.33$	-0.3333

This type of analysis perfect for spreadsheets!

How should we decide n?

- Most of time, use $\mathrm{SD}(X)$ to measure "spread" in X
 - Easy to compute

-
$$X_1, X_2$$
 indep., $SD(X_1 + X_2) = \sqrt{SD(X_1)^2 + SD(X_2)^2}$

- Other problem: even if $\mathbb{E}[X]$ exists, $\mathrm{SD}(X)$ might not
- Alternative: mean absolute deviation

Definition 55

The **mean absolute deviation** of X is

 $MAD(X) = \mathbb{E}[|X - \mathbb{E}(X)|].$

Ex:
$$X \sim \text{Exp}(\lambda), \ f_X(s) = \lambda e^{-\lambda s} \mathbb{1}(s \ge 0)$$

$$\mathbb{E}(X) = \int_{-\infty}^{\infty} s f_X(s) \ ds = 1/\lambda.$$

Mark Huber Notes on Stochastic Operations Research

$$\mathbb{E}(|X - \mathbb{E}(X)|) = \int_{-\infty}^{\infty} |s - 1/\lambda| f_X(s) \, ds$$
$$= \int_{-\infty}^{\infty} |s - 1/\lambda| \lambda e^{-\lambda s} \mathbb{1}(s \ge 0) \, ds$$
$$= \int_{0}^{\infty} |s - 1/\lambda| \lambda e^{-\lambda s} \, ds$$
$$= \int_{0}^{1/\lambda} -(s - 1/\lambda) \lambda e^{-\lambda s} \, ds$$
$$+ \int_{1/\lambda}^{\infty} (s - 1/\lambda) \lambda e^{-\lambda s} \, ds$$
$$= \frac{2}{e} \cdot \frac{1}{\lambda}$$

Fact 28 When $\mathbb{E}(X)$ exists, so does MAD(X).

Remarks

- 1. MAD(X), SD(X), $\mathbb{E}(X)$ all have the same units.
- 2. Alternate notation: MD(X) = MAD(X)
- 3. MAD(cX) = c MAD(X)
- 4. $MAD(X_1 + \cdots + X_n)$ difficult to calculate, even for X_i indep: can use Monte Carlo to find

Some MAD values:

Dist	$\mathbb{E}(X)$	$\mathrm{SD}(X)$	MAD(X)
$Exp(\lambda)$	$1/\lambda$	$1/\lambda$	$(2/e)(1/\lambda)$
Unif([a,b])	$\frac{a+b}{2}$	$\frac{b-a}{\sqrt{12}}$	$\frac{b-a}{4}$
$N(\mu,\sigma)$	μ	σ	$\sqrt{2/\pi}\sigma$

Estimating MAD

- Each error in forecast assumed to have mean o.
- Error is data minus prediction

$$E_i = X_i - F_i.$$

• If $\mathbb{E}(E_i) = 0$, MAD $(E_i) = \mathbb{E}(|E_i|)$

129 184

• To estimate MAD, take sample averages of absolute error

$$\widehat{MAD} = \frac{1}{N} \sum_{i=1}^{N} |E_i|$$

Q of Day When N = 3:

$$M\hat{A}D = \frac{|8\ 2/3| + |-2/3| + |-1/3|}{3} = \frac{29}{9} = 3.222$$

- Rule of thumb for choosing ${\cal N}$
 - Pick values of N that minimizes ${\rm M}{\rm \hat{A}}{\rm D}$
- [Show how to estimate MÂD w/ spreadsheet]
- Gives N = 4 for this data.

When to use

• Moving averages work best when

$$x_i = b + \epsilon_i$$

- b is the best level. ϵ_i is random fluctuation
- Work worst when
 - Seasonality
 - Trend

Chapter 29 Simple Exponential Smoothing

Question of the Day How can we smooth out small fluctuations in time series data?

Today

• Exponential Smoothing

The model

- Suppose model is $X_i = b + \epsilon_i$
 - One high value can throw off base for a long time
- · Idea: Let last observation affect prediction less



- This makes A_t a convex linear combination of X_t and A_{t-1}
- Call the error in prediction:

$$e_t = X_t - A_{t-1}$$

131 184

• For exponential smoothing:

$$A_t = A_{t-1} + \alpha (X_t - A_{t-1}) = A_{t-1} + \alpha e_t.$$

• So new prediction is old prediction plus fraction of error from last prediction

Example:

Sales	Forecast	e_t
32		
30	32	-2
32	31.8	0.20
30	31.82	-1.82
39	31.64	7.36
33	32.37	0.63
34	32.44	1.56
	Sales 32 30 32 30 32 30 32 30 32 30 32 30 32 30 32 30 33 34	Sales Forecast 32 32 30 32 32 31.8 30 31.82 39 31.64 33 32.37 34 32.44

Here $\hat{MAD} = 3.04$.

How to choose α

- Before, choose N to minimize ${\rm M}{\rm \hat{A}}{\rm D}$
- Do the same thing here
- Spreadsheet very helpful here
- For α integer multiple of 0.05, best $\alpha = 0.25$

α	MÂD
0.05	3.20
÷	÷
0.20	2.89
0.25	2.88
0.30	2.90
:	÷

Remarks

- Called "smoothing" because variation e_t reduced to αe_t ($\alpha < 1$)
- + $\alpha = 2/(N+1)$ approximately same as moving average w/ N observations

• Called "Exponential smoothing because:

$$A_{t} = \alpha X_{t} + (1 - \alpha)A_{t-1}$$

= $\alpha X_{t} + (1 - \alpha)[\alpha X_{t-1} + (1 - \alpha)A_{t-2}]$
= $\alpha X_{t} + (1 - \alpha)\alpha X_{t-1} + (1 - \alpha)^{2}[\alpha X_{t-2} + (1 - \alpha)A_{t-3}]$
= $\left[\sum_{i=0}^{k} \alpha (1 - \alpha)^{i} X_{t-i}\right] + (1 - \alpha)^{k+1} A_{t-k-1}.$

· So past data effect on present prediction declines exponentially

What is a good α ?

- Could use $\hat{\mathrm{MAD}}$ to choose
- In practice $\alpha=0.1, 0.3, 0.5$ commonly used
- + $\alpha > 0.5$ indications some other trend/seasonality present

When to use

• Moving average works when you have a baseline:

$$X_t = b + \epsilon_t.$$

• Exponential smoothign works when baseline wanders:

$$X_t = m_t + \epsilon_t$$
$$m_t = m_{t-1} + \delta_t$$

[Here $\mathbb{E}[\delta_t] = 0$, m_t is a random walk.]

Holt's method: Exponential smoothing with Trend

• Works well when baseline has a linear trend:

$$X_t = m_t + \epsilon_t$$
$$m_t = m_{t-1} + \gamma + \delta_t$$

• Idea: Use L_t to predict m_t and T_t to predict γ

Holt's Method:

$$L_t = \alpha X_t + (1 - \alpha)(L_{t-1} + T_{t-1})$$
$$T_t = \beta (L_t - L_{t-1}) + (1 - \beta)T_{t-1}$$

133 184

Next months forecast: $f_{t,1} = L_t + T_t$ k months in future forecast: $f_{t,k} = L_t + kT_t$

What's going on?

- Here T_t is a convex linear combination of $L_t L_{t-1}$ and T_{t-1}
- L_t is a convex linear combination of X_t and $L_{t-1} + T_{t-1}$

Example:	Blu-ray sales	(1000's)	a) with $\alpha = 0.3$ and $\beta =$	0.1
L'Aumpre.	Dia iny buieb	(1000 0)	p = 0.0 and $p = 0.0$	U •

Month	Sales	L_t	T_t	Prediction	e_t
		25.3	7.4		
1	32	32.49	7.379	32.7	-0.7000
2	40	39.9083	7.38293	39.869	0.1310
3	38	44.503861	7.1041931	47.29123	-9.291
4	56	52.92563787	7.235951477	51.6080541	4.391
5	67	62.21311254	7.441103797	60.16158935	6.838

Economics data

• Often Econ data is growing exponentially:

$$X_t = ab^t \epsilon_t.$$

• Just take logarithm, and smooth as before:

$$\ln(X_t) = \ln(a) + t\ln(b) + \ln(\epsilon_t).$$

^{Chapter 30} Seasonality

Question of the Day How can we incorported seasonality into forecasting?

Today

· Winter's method for seasonal forecasting

Idea:

 $L_t = \text{baseline}$ $T_t = \text{trend}$ $S_t = \text{multiplier for season}$ c = # of time periods in season

Example: Deck furniture

- Suppose number sold in June is 1.8 times average, $s_6 = 1.8$
- Number sold in January is 0.3 times average, $S_1 = 0.3$
- Use α for L_t , β for T_t , γ for S_t
- For months, c = 12
- For current month t, S_{t-12} was seasonal multiplier last year

Winter's Model

$$L_{t} = \alpha \frac{X_{t}}{S_{t-c}} + (1-\alpha)(L_{t-1} + T_{t-1})$$
$$T_{t} = \beta(L_{t} - L_{t-1}) + (1-\beta)T_{t-1}$$
$$S_{t} = \gamma \frac{X_{t}}{L_{t}} + (1-\gamma)S_{t-c}.$$

135 184

Mark Huber Notes on Stochastic Operations Research

Next month forecast: $f_{t,1} = (L_t + T_t)S_{t+1-c}$ k month forecast: $f_{t,k} = (L_t + kT_t)S_{t+k-c}$

Again: convex combinations of current data and past predictors.

Initializing

- · Need some initial estimates of trends, seasonal factors
- Multiple ways to do this
- Example: suppose have last two years of sales:

Year -2 26 38 4 3 10 14 5 40 28 17 16 13 Year -1 9 6 18 27 48 75 33 50 77 52 31 24

$$\hat{\mu}_{\text{Year-2}} = \frac{450}{12}$$

$$\hat{\mu}_{\text{Year-1}} = \frac{234}{12}$$

$$T_0 = \frac{\hat{\mu}_{\text{Year-1}} - \hat{\mu}_{\text{Year-2}}}{12 \text{ months}} = 1.500/\text{month}$$

$$L_0 = \hat{\mu}_{\text{Year-1}} + 6 \text{ month}(1.5/\text{month})$$

$$= 37.5 + 6(1.5) = 46.5$$

• Now let's get sasonal estimates: Average for Year -2 are 19.5

Year -2 in Jan. =
$$\frac{\text{actual sold}}{\text{average for year}} = \frac{4}{19.5} = 0.205$$

Year -1 in Jan. = $\frac{9}{37.5} = 0.240$.

• Note: s_{1-24} should be relatively close to s_{1-12} or model might be bad.

$$s_{-11} = \frac{0.205 + 0.240}{2} = 0.2225.$$

Spreadsheet time!

- Now that we have model, place in spreadsheet
- · Can calculate accuracy of forecast

More about MAD

- Measure of how good is forecast
- Normality an often used model
- For one reason: easy to calculate with

For normal r.v., MAD =
$$\sqrt{\frac{2}{\pi}}\sigma$$

So
$$\hat{\sigma} = \sqrt{\frac{\pi}{2}} M \hat{A} D$$

- Can use this to construct confidence intervals for forecast
- For normal data, a 95% confidence interval looks like:

[prediction $-1.960\sqrt{\pi/2}$ MÂD, prediction $+1.960\sqrt{\pi/2}$ MÂD]



- This CI includes several assumtions
 - $\mathbb{E}[\text{forecast}] = \text{true value}]$
 - forecast ~ N(true value, $(\sqrt{\pi/2} \text{ MAD})^2)$

Remarks

- You have a three dimensional optimization problem to find α,β,γ that minimizes MAD
- Usually α , β are at most 0.5 (as in Holt)
- However, γ can be larger: seasonal data is rare, so earlier data more important

More sophisticated

- Linear regression
- k different predictor variables
- k predictors each measured n times to get X an m by k matrix
- Results measured n times to get Y a k dimensional vector
- k coefficients β
- Model:

$$Y = X\beta + \epsilon,$$

- Can use least squares or other methods to estimate β
- More detail in Math 152
Chapter 31 Marginal Analysis

Question of the Day [Winston 1991] A gift shop buys plastic St. Louis arches for \$2 and sells them for \$4.50. Unused arches can be sold to smaller shops for \$0.75. Suppose the model is:

$\mathbb{P}(sales=i)$
30%
20%
30%
15%
5%

How many should they buy?

Today

· Optimization through marginal analysis

Notation

- Let d = demand and q = amount ordered
- Let c(d,q) = cost incurred by vendor
- Let D = the random demand (sales)

Goal (Utility = -cost)

$$\min_{q} \mathbb{E}(c(D,q)) = \sum_{d: \mathbb{P}(D=d) > 0} \mathbb{P}(D=d)c(D,q).$$

Usually $f(q) = \mathbb{E}(c(D,q))$ is a convex up function of q.

Definition 56

A function f is **convex** [up] if the line segment connecting any two points on the graph of the function lies at or above the function. That is,

$$(\forall a < b)(\forall \lambda \in [0,1])(\lambda f(a) + (1-\lambda)f(b) \ge f(\lambda a + (1-\lambda)b))$$

Fact 29 (Jensen's inequality) for all random variables X and convex up functions $f: \mathbb{E}[f(X)] \ge f(\mathbb{E}(X))$.

Classic example: $\mathbb{E}[X^2] \ge \mathbb{E}(X)^2$.

Definition 57

When q < d the lost income the lost income is called *understocking cost*. When q > d the expense of buying too many is called *overstocking cost*.

Qotd

- Each calendar not sold (understocked) incurs cost of 4.50 2.50 = 2.00
- Each calendar overstocked incurs cost of 2-1.25



• Each value of d gives a different convex function

Fact 30

If c(d, q) is convex for all d, then for any random variable D,

$$f(q) = \mathbb{E}[c(D,q)]$$

is also convex.

Optimization

• Convex functions in 1D easy to optimize

$$\operatorname*{arg\,min}_{q\in\mathbf{Z}}f(q).$$

- Convexity gives local min = global min
 - Repeat - If f(q+1) < f(q) increase q by 1 - Until f(q+1) > f(q)

Qotd

• $D \in \mathbf{Z}, q \in \mathbf{Z}$

$$f(q+1) - f(q) = \mathbb{E}(c(D,q+1)) - \mathbb{E}(c(D,q))$$
$$= \mathbb{E}(c(D,q+1) - c(D,q))$$

• Note:

$$c(D, q+1) - c(D, q) = \begin{cases} 1.25 & \text{if } q \ge D\\ -2.5 & \text{if } q < D \end{cases}$$
$$= -2.5 + 3.25 \cdot \mathbb{1}(D \le q)$$

$$\mathbb{E}(c(D, q+1) - c(D, q)) = -2.5 + 3.25 \cdot \mathbb{P}(D \le q).$$

So

$$\begin{split} f(q+1) < f(q) \Leftrightarrow -2.5 + 3.75 \cdot \mathbb{P}(D \leq q) < 0 \\ \Leftrightarrow \mathbb{P}(D \leq q) < 2.5/3.75 = 2/3. \end{split}$$

For this problem

$$\mathbb{P}(D \le 100) = 0.30, \ \mathbb{P}(D \le 150) = 0.50, \ \mathbb{P}(D \le 200) = 0.80.$$

Optimal strategy is to order

$$q^* = 200 \text{ calendars}$$

General solution

Let $c_u = \text{cost of 1}$ unit of understock

 $c_o = \text{cost of 1}$ unit of overstock

$$q^* = \inf \left\{ q : \mathbb{P}(D \le q) \ge \frac{c_u}{c_u + c_o} \right\}.$$

News vendor problem

- Children used to buy papers from the publisher
- Would sell what they could
- (After a strike) they could resell unused copies back to printer

Note

• If $\mathbb{E}(c(D,q))$, simple optimization doesn't work



• Gets trapped at first local minima

Chapter **32**

News vendor problem with continuous demand

Question of the Day [Winston 1991 based on Virts & Garrett (1970)] A TV manufacture estimates annual demand as normal with mean 6000, st. dev 2000. Sets cost \$100 to build and sell for \$250. How many should they build?

Today

Continuous Marginal Analysis

Why use \mathbb{R} instead of \mathbb{Z} ?

- Last time treated case where $q\in\mathbb{Z}$
- Normal model not accurate
 - Can't build 1/2 a RV
 - Also 0.1% chance that demand < o!
- Continuous often easier to calculate with than discrete
- [LP's much easier than IP's]
- Can lead to analytical solutions in terms of variables

Convexity

- Both D and q are in $\mathbb R$
- Good news: still have if c(d,q) convex in q for all d, then $\mathbb{E}[c(D,q)]$ is still convex in q for all r.v. D
- Continuous convex functions have unique minimum value over [a, b]

- Let $f(q) = \mathbb{E}(c(D,q))$
- Let h > 0 then

$$\begin{split} f(q+h) - f(q) &= \mathbb{E}(c(D,q+h)) - c(D,q) \\ &= \mathbb{E}([c(D,q+h) - c(D,q)] \mathbb{1}(D \leq q) \\ &+ [c(D,q+h) - c(D,q)] \mathbb{1}(D \in (q,q+h)) \\ &+ [c(D,q+h) - c(D,q)] \mathbb{1}(D \geq q+h)). \\ &= \mathbb{E}(c_oh \mathbb{1}(D \leq q) + \epsilon h \mathbb{1}(D \in (q,q+h)) \\ &- c_u h \mathbb{1}(D \geq q+h)), \end{split}$$

where c_o is overstock cost, c_u is understock cost, $|\epsilon| \le \max\{c_o, c_u\}$. Factoring out an h gives:

$$f(q+h) - f(q) = h[c_o \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}(D \in (q, q+h)) - c_u \mathbb{P}(D \le q) + \epsilon \mathbb{P}$$

- As $h \to 0,$ the middle term goes to 0, and the right term converges to $\mathbb{P}(D > q).$
- So if $c_o \mathbb{P}(D \le q) c_u(1 \mathbb{P}(D \le q)) < 0$...
- ...then f(q) is decreasing, otherwise it is increasing
- So minimum occurs when

$$q^* = \inf\left\{q: \mathbb{P}(D \le q) \ge \frac{c_u}{c_o + c_u}\right\}$$

• For continuous D, $q^* = \{q : \mathbb{P}(D \le q) = c_u/(c_o + c_u)\}.$

For TV

$$c_{0} = \$100$$

$$c_{u} = \$150$$

$$\frac{c_{u}}{c_{o} + c_{u}} = \frac{150}{150 + 100} = \frac{3}{5} = 0.60$$

$$Z \sim N(0, 1)$$

$$6000 + 2000Z \sim N(6000, 2000^{2})$$

$$\mathbb{P}(6000 + 2000Z \leq q) = 0.60$$

$$\mathbb{P}\left(Z \leq \frac{q - 6000}{2000}\right) = 0.60$$

$$\frac{q - 6000}{2000} = \Phi^{-1}(0.6)$$

$$q = 6000 + 2000\Phi^{-1}(0.6)$$

$$q \approx 6506.694...$$

144 184

Should we report $\lfloor q \rfloor$ or $\lceil q \rceil$?

- Answer 1: It doesn't matter-you're kidding yourself if you think your demand model is that accurate
- Answer 2: Calculate $f(\lfloor q \rfloor)$ and $f(\lceil q \rceil)$ and take whichever is smaller.

$$\mathbb{E}(C(D,q)) = \int_0^q c_o \cdot (q-d) \cdot f_D(s) \, ds \\ + \int_q^\infty c_u \cdot (d-q) \cdot f_D(s) \, ds.$$

• For Q of the day:

$$q = 6506 \Rightarrow \cos t \approx 106674 + 57028.5$$

 $q = 6507 \Rightarrow \cos t \approx 106734 + 56988.5$

Beyond Marginal Analysis

- [Winston 1991] Suppose Seuss Construction is bidding on a job which costs \$2 million to complete. Their only competitor bid is believed to be uniform over [2, 4] million dollars. What should they bid?
- Let B = opponent bid, q = Seuss bid

$$\begin{aligned} \text{profit} &= (q-2)\mathbb{1}(q > B) + 0 \cdot \mathbb{1}(q \le B) \\ \mathbb{E}(\text{profit}) &= p(q) = \mathbb{E}((q-2)\mathbb{1}(q > B)) \\ &= (q-2)\mathbb{P}(q > B) \\ &= (q-2)\frac{4-q}{4-2} = \frac{1}{2}\left[-q^2 + 6q - 8\right] \end{aligned}$$

• Maximize profit:

$$p'(q) = \frac{1}{2} [-2q+6]$$
 $p''(q) = -1$
 $p'(q) = 0 \Rightarrow q = 3.$

Since second derivative negative any local minimum is a global minimum.

• End result: Should bid \$3 million

^{Chapter} 33 Inventory

Question of the Day What should the reorder point and ordering quantity be for a reorder point inventory model?

Today

• Inventory models

Inventory

- Revolution over the last 20 years
- Global supply chains: giant network of transport
- Just-in-time ordering
- Amazon driven prices down through technology

Factors to consider

- Lag time between when you order and when you receive inventory
- Opportunity cost of lost sales when your inventory empty
- Cost associated with ordering new product
- Cost associated with storing product

Reorder point model

- Constants you can't control
 - L =lead time for each order
 - K =ordering cost
 - h = holding cost/unit/year
- $c_B = \cos t$ for each unit short
- D =total demand for the year
- 146 184

- Decisions to make
 - r = reorder point (if inventory below this level, reorder)
 - q = Amount of inventory to order when reorder
- Simplifying assumption: never lose a sale (customer always comes back for an item)
- Because order takes time to fill, inventory continues to drop after passing r:



Two effects

- Minimize cost for holding product: low $r \ {\rm low} \ q$
- Minimize cost underserved customers: high r high q

$$\begin{split} TC(q,r) &= \mathbb{E}(\text{total costs}) \\ &= \mathbb{E}(\text{holding costs}) + \mathbb{E}(\text{ordering costs}) + \mathbb{E}(\text{shortage cost}). \end{split}$$

Inventory levels

- Simplifying assumption: Given N customers in cycle, arrival times uniformly distributed

Fact 31

Consider a cycle from t_0 to t_1 . Let I(t) denote the inventory level at time t. Then when customer arrivals are uniformly distributed over a cycle and the demand from each customer is iid, for $T \sim \text{Unif}([t_0, t_1])$,

$$\mathbb{E}(I(T)) = \frac{1}{2} [\mathbb{E}(I(t_0)) + \mathbb{E}(I(t_1)).$$

Proof. Fix t_0 and t_1 the start and end of the cycle. Let N be the demand during the cycle. Then let D_1, D_2, \ldots be the demand (purchase) made by customer *i*. Then

$$I(t_0) - I(t_1) = D_1 + D_2 + \dots + D_N.$$

Hence by Wald's equation

E

$$\mathbb{E}[I(t_1) - I(t_0)] = \mathbb{E}[D_i]\mathbb{E}[N]$$

Let T_i be the *n* uniformly distributed times in $[t_0, t_1]$. Then

$$(I(T)) = \mathbb{E}(\mathbb{E}(I(T)|N))$$

= $\mathbb{E}(\mathbb{E}(I(t_0) - \sum_{i=1}^N \mathbb{1}(T \le T_i)D_i|N))$
= $\mathbb{E}(I(t_0)) - \mathbb{E}(\sum_{i=1}^N \mathbb{E}(D_i)(1/2))$
= $\mathbb{E}(I(t_0)) - \mathbb{E}(N) \left(\sum_{i=1}^N \mathbb{E}(D_i)\right) (1/2)$
= $\mathbb{E}(I(t_0)) - (1/2)(\mathbb{E}(I(t_0)) - \mathbb{E}(I(t_1)))$
= $(1/2)\mathbb{E}(I(t_1) + I(t_0))$

Holding costs

- At beginnig of cycle has $I(t_0)$
- Let X be demand between reordering and getting order of \boldsymbol{q}
- So beginning of next cycle, inventory is r X + q

- At end of cycle, inventory is r-X

$$\mathbb{E}(I(T)) = (1/2)\mathbb{E}[r - X + q + r - X] = \frac{q}{2} + r - \mathbb{E}[X].$$

• This is mean inventory level

$$\mathbb{E}(\text{holding cost}) \approx h(q/2 + r - \mathbb{E}(X)).$$

(Approximate because don't get "negative holding" when inv. below o.)

Ordering costs

- D demand a year, so on average, number of orders is about

$$\frac{\mathbb{E}[D]}{q}$$

(ignores starting/ending inventory)

$$\mathbb{E}(\text{order cost}) = k \frac{\mathbb{E}[D]}{q}.$$

where k is the cost to reorder.

Shortage cost

- Shortage for a cycle: $(X r)^+ = (X r)\mathbb{1}(X r \ge 0)$
- $\mathbb{E}(\text{# of cycles}) = \mathbb{E}[D]/q$
- So

$$\mathbb{E}(\text{shortage cost}) = \frac{\mathbb{E}[D]}{q} \mathbb{E}[(X - r)^+].$$

Total costs are approximately

$$h\left(\frac{q}{2} + r - \mathbb{E}(X)\right) + \frac{c_B \mathbb{E}[D]}{q} \mathbb{E}[(X - r)^+] + \frac{k \mathbb{E}[D]}{q}.$$

- Optimize by ocnsidering what happens when r increases by 1
- Holding costs go up by h
- Shortage costs go down by $c_B(\mathbb{E}[D]/q)\mathbb{1}(X \ge r)$
- So when they are equal, you have the right \boldsymbol{r}
- Next minimize for q

Fact 32

The inventory Rule of Thumb: Choose r^\ast and q^\ast so that

$$q^* = \sqrt{2k\mathbb{E}(D)/h}$$
$$\mathbb{P}(X \ge r^*) = hq^*/[c_B\mathbb{E}(D)].$$

Next time we'll do an example!

Chapter 34 Solving Inventory Problems

Question of the Day [Winston 1991 (updated)] A computer store sells a # of memory cards that is $N(1000, 40.8^2)$ per year. A regional distributor charges \$50 per order which takes 2 weeks to fill. Holding costs for 1 year is \$10, stockout cost (shortage cost) is \$20. What is the proper reorder point and quantity. (Assume demand is normal in any time period.)

Today

• Using the reorder point Rule of Thumb

Economic Order Quantity (EOQ)

• Last time, rule of thumb:

$$q = \sqrt{\frac{2k\mathbb{E}(D)}{h}},$$

where k is reorder cost, h is holding cost , $\mathbb{E}[D]$ is expected demand

QotD

$$h = \$10/\text{card/year}$$
$$K = \$50$$
$$\mathbb{E}(D) = 1000 \text{ cards/year}$$
$$q = \sqrt{\frac{2 \cdot 50 \cdot 1000}{10}} = 100$$

- So now we know the quantity, when do we reorder?
- Lead time 2 weeks

• For our rule of thumb, want

$$r: \mathbb{P}(X \ge r) = \frac{hq}{c_B \mathbb{E}[D]}$$
$$= \frac{10 \cdot 100}{20 \cdot 10000} = 0.05$$

- So choose r so that $\mathbb{P}(X \geq r) = 0.05$
- For normals, subtract mean, dived by sides by standard deviation:

$$\mathbb{P}\left(\underbrace{\frac{X - 1000(2/52)}{40.8\sqrt{2/52}}}_{\mathsf{N}(0,1)} \ge \frac{r - 1000(2/52)}{40.8\sqrt{2/52}}\right) = 0.05$$

So want

$$\frac{r - 1000(2/52)}{40.8\sqrt{2/52}} = \Phi^{-1}(0.95)$$

or equivalently:

$$r = 1000 \cdot \frac{2}{52} + 40.8\sqrt{\frac{2}{52}}\Phi^{-1}(0.95).$$

(so $\Phi^{-1}(0.95)$ is how many standard deviations we are away from the mean) In this case:

$$r = 51.2$$

- Note: since $r \geq \mathbb{E}(X)$, carrying extra stock (on average) to guard against stockouts

Definition 58

The **safety stock** is the reorder point minus the expected demand over the reorder time.

Example

• For the question of the day:

safety stock =
$$\Phi^{-1}(0.95) \cdot 40.8 \sqrt{\frac{2}{52}} = \boxed{13.16}$$
.

Measuring service level

- Previous analysis required that we know c_B the stockout cost
 - Can be hard to quantify
 - What is cost in \$ of upset user?
- Alternate way of measuring service

Definition 59 Service level measure 1 is

 $SLM_1 = \mathbb{E}(\text{demand met})/\mathbb{E}(\text{demand}).$

Service level measure 2 is

 $SLM_2 = \mathbb{E}(\# \text{ of cycles/year where shortage occurs}).$

Example Suppose demand during lead time has following dist:

i	$\mathbb{P}(X=i)$
20	0.2
30	0.2
40	0.2
50	0.2
60	0.2

$$EOQ: q^* = 100, r = 30$$

What are SLM_1 and SLM_2 ?

- Reorder when inventory at most 30
- If demand is 40, 50, 60, shortage!
- Chance demand that high is 60%
- Mean demand unmet in a cycle:

Demand	20	30	40	50	60
Demand unmet	0	0	10	20	30

 $\mathbb{E}[\text{demand unmet in cycle}] = 0.2(10) + 0.2(20) + 0.2(30) = 12.$

$$\mathbb{E}[\text{\# of cycles}] \approx \frac{\text{mean demand}}{\text{order amount}} = \frac{1000}{100} = 10$$

 $\mathbb{E}(\text{unmet demand}) = 12 \cdot 10 = 120.$

$$\mathbb{E}(\text{demand}) = 1000$$

$$SLM_1 = \frac{1000 - 120}{1000} = \boxed{88.00\%}$$

(Note that SLM_1 is unitless.)

- What is expected number of cycles per year where shortage occurs?
- Each of 12 cycles have 60% chance of shortage
- Binomial number of shortages: $SLM_2 = (12)(0.6) = 9.6$ per year.

Using service level to make decisions

- By raising r (reorder level), raise SLM_1 , lower SLM_2 (That is both good!)
- Let B_r be a moount of unmet demand during a cycle
- Let C be # of cycles in a year.

$$\mathbb{E}[C] = \frac{\mathbb{E}[D]}{q}$$

$$\mathbb{E}(\text{unmet demand in a year}) = \mathbb{E}(B_r)\mathbb{E}(C) = \frac{\mathbb{E}(B_r)\mathbb{E}(D)}{q}$$

$$SLM_1 = 1 - [\mathbb{E}(B_r)\mathbb{E}(D)/q]/\mathbb{E}(D) = 1 - \mathbb{E}[B_r]/q.$$

- As r goes up, $\mathbb{E}(B_r)$ goes down, SLM_1 goes up.
- As q goes up, SLM_1 goes up.
- Need to know dist. of X (demand during lead time) in order to say more.

Chapter 35 Probabilistic Dynamic Programming

Question of the Day Abernathy Grocery has 3 stores in an area, w/6 gallons of milk to distribute. Each gallon sells for \$4, or the dairy will buy back for \$1 at the end of the day. Daily demand is:

i	Store 1	Store 2	Store 3
1	60%	50%	40%
2	0	10%	30%
3	40%	40%	30%

What is the optimal distribution of milk to stores?

Today

• Probabilistic Dynamic Programming

Dynamic Programming

- · Solves problem given by recursion from bottom up
- Example: Fibonacci sequence F(n) = F(n-1) + F(n-2)
- Bad way to find F(6), recursively:

Fib	onacci Input: n
1)	$a \leftarrow F(n-1)$
2)	$b \leftarrow F(n-2)$
3)	Return $a + b$



- · Lots of repetition, wasted effort
- Running time is $\Theta(F(n)) = \Theta(\phi^n)$
- A better way: fill in starting from n = 1 and n = 2 upwards

n	1	2	3	4	5	6
F(n)	1	1	2	3	5	8

- Better way is linear in \boldsymbol{n}

Using this principle with milk

- Recursion more complex
 - Store 3 receives either $g_3 \in \{0, 1, 2, 3, 4, 5, 6\}$ gal.
 - Remaining milk is $6 g_3$, Remaining stores to distributed to is 2.
- Use linearity of expectation

 $r_t(g_t) =$ expected revenue from store t when it gets g_t gal. $f_t(x) =$ max expected revenue for x gal to give to stores $1, 2, \ldots, t$

• With this notation, qotd becomes: What is $f_3(6)$?

Bellman equation

- Suppose that I give 2 gal to store 3
- When demand 1, sell a gal to cust, sell one back to dairy

•
$$r_3(2) = 40\%(1 \cdot 4 + 1 \cdot 4) + 60\%(2)(4) = 2 + 4.8 = 6.8$$

- Then $f_3(6) \ge 6.8 + f_2(4)$
- In fact,

$$f_3(6) = \max_{g_3 \in \{0, 1, 2, \dots, 6\}} r_3(g_3) + f_2(6 - g_3).$$

- This is just the same old decision trees from before as an equation
- More generally:

$$f_t(x) = \max_{g_t} r_t(g_t) + f_{t-1}(x - g_t).$$

- Called the Bellman equation or dynamic programming equation
- First step: build the reward table for $r_t(g_t)$:

Reward Table: $r_t(g_t)$									
Gallons									
		0	1	2	3	4	5	6	
	1	0	4	$r_1(2) = 6.2$	8.4	9.4	10.4	11.4	
Stores	2	0	4	6.5	7.5	8.5	9.5	10.5	
	3	0	4	$r_3(2) = 6.8$	8.7	9.7	10.7	11.7	

- Now can find $f_t(x)$
- First row is just give all gallons to store 1
- Second row maximizes over choice of gallons to store 2
- Example:

$$f_2(3) = \max\{r_2(0) + f_1(3), r_2(1) + f_1(2), r_2(2) + f_1(1), r_2(3) + f_1(0)\}\$$

= max{0 + 8.4, 4 + 6.2, 6.5 + 4, 7.5 + 0}

Maximum occurs when give 2 gallons to store 2 and 1 to store 1

	Optimal Table: $f_t(g_t)$									
						Gallons				
			0	1	2	3	4	5	6	
•		1	0	4	6.2	8.4	9.4	10.4	11.4	
	Stores	2	0	4^{0}	8^1	10.5^{2}	12.7^{2}	14.9^{2}	15.9^{3}	
		3	0	4^{0}	8^{0}	12^{1}				
					157	184				

• Now work backward to get optimal solution:

1 gal to store 3 (leaves 2 gal), 1 gal to store 2 (leaves 1 gal), 1 gal to store 1.

Notes

- Any recursive equation can be solved using Dynamic Programming
- When using Dynamic programming to maximum expected value, called Probabilistic Dynamic Programming, or PDP
- Computational complexity
 - t choices to make
 - m possible values for each choice
 - Calculating maximum takes $\theta(m)$ time
 - Size of table is # of states times t
 - Total time $\Theta(tm^2)$

Chapter 36 Markov Decision Processes

Question of the Day At the beginning of each week, a machine is either in Excellent, Good, Average, or Bad shape.

The revenue earned in the week given the state is:

$$E = \$100, \ G = \$80, \ A = \$20, \ B = \$10.$$

At the beginning of each week, it is possible to instantaneously return the state to excellent at a cost of \$200. What should the repair policy be if the machine state evolves as a Markov chain with:



Today

- Markov chains
- Markov decision processes

Markov decision process

Definition 60 A discrete time stochastic process X_0, X_1, X_2, \ldots is a **Markov chain** if $(\forall t)([X_t|X_0, \ldots, X_{t-1}] \sim [X_t|X_{t-1}]).$ • Idea: X_t only depends on X_{t-1} , not on total history

Example

- Say $D_1, D_2, \ldots \sim \mathsf{Unif}(\{-1, 1\})$
- Then let $X_0 = 0$, $X_{t+1} = X_t + D_{t+1}$
- Then X_0, X_1, X_2, \ldots for a Markov chain
- $[X_t|X_0,\ldots,X_{t-1}] \sim [X_t|X_{t-1}] \sim \text{Unif}(\{X_{t-1}-1,X_{t-1}+1\})$

Definition 61

A Markov decision process (MDP) consists of

- 1. State space: Ω which is finite.
- 2. Decision set: For each $i \in \Omega$, a set of decisions D(i).
- 3. Transition probabilities. Use notation:

$$(\forall i, j)(\mathbb{P}(X_{t+1} = j | X_t = i, d) = p(j|i, d))$$

4. Expected Rewards: Start in state i and make decision d, expected reward is r_{id} .

Qotd

- 1. $\Omega = \{E, G, A, B\}$
- 2. $D(i) = \{$ replace,do not replace $\} = \{R, NR\}$
- 3. p(j|i,d):

	NR							R				
		next state j								j		
		Е	G	А	В				E	G	А	В
	Е	0.7	0.3	-	-		-	Е	0.7	0.3	-	-
current state	G	-	0.7	0.3	-			G	0.7	0.3	-	-
i	А	-	-	0.4	0.6		i	А	0.7	0.3	-	-
	В	-	-	-	1			В	0.7	0.3	-	-

Now can calculate the expected rewards from each decision and current state:

- Example: if current state G, no repair, make \$50
- Example: if current state A (or anything), and repair, make \$100-\$200 = -\$100

• Altogether:

		R	NR
	Е	-100	100
i	G	-100	80
	А	-100	50
	В	-100	10

Definition 62

A **policy** is a rule that specifies how each period's decision is made.

Definition 63

A policy is **stationary** if the decision only depends on the current state, and not past history.

For an MDP, no reason not to use a stationary policy!

Infinite time horizon

- Think about model running infinite # of steps
- That makes all expected returns infinite!
- Here are two ways to deal with that:
 - 1. Maximize average expected reward per period

$$\mathbb{E}\left(\lim_{n\to\infty}\frac{\text{rewards from period }1,\ldots,n}{n}\right)$$

2. Discount future reward

reward k in future = $\beta^k \cdot$ reward now

(Here $\beta \in (0, 1)$.)

- Reflects uncertainty in future
- Max reward in one period is 100
- So max average reward also 100
- Max discounted reward:

$$100 + \beta 100 + \beta^2 100 + \dots = \frac{100}{1 - \beta}.$$

161 184

Goal

- Let δ be a policy, and Δ the set of all policies.

 $X_t =$ state of MDP at beginning of period t $X_1 =$ initial state

 d_t = decision made under policy δ at period t given X_1, \ldots, X_t

 $V_{\delta}(i) =$ expected total discounted reward under δ w/ $X_1 = i$

• So

$$V_{\delta}(i) = \mathbb{E}_{\delta} \left(\sum_{t=1}^{\infty} \beta^{t-1} r_{X_t, d_t} | X_1 = i \right).$$

- Let V(i) be $\max_{\delta \in \Delta} V_{\delta}(i)$ (or min as appropriate)
- This is an infinite dimensional optimization problem!

Definition 64

If δ^* satisfies $(\forall i \in \Omega)(V(i) = V_{\delta^*}(i))$, then δ^* is an **optimal policy**.

Fact 33

(Blackwell 1962) If the r_{ij} values are bounded, then an optimal policy exists, moreover a stationary optimal policy exists.

Consequences

- Don't have to check all policies, only stationary ones!
- For QotD, $-100 \le r_{id} \le 100$ so some stationary optimal policy exists.
- Only 2⁴ possible stationary policies!
- Brute force method:
 - 1. Evaluate $V_{\delta}(i)$ for all $i \in S, \delta \in \Delta_{\text{stat}}$
 - 2. Pick the best

Better ways

- 1. Policy iteration
- 2. Linear Programming
- 3. Value iteration

Chapter 37 Finding average reward



Today

Average reward

General decision rule

- Choose decision d when state is $i \le i \le q_i(d)$
- Let π_i denote the fraction of time that the state is i
- Then $\pi_{id} = \pi_i q_i(d)$ is the fraction of time that we are in state *i*, and make decision *d*
- Goal: max expected reward per period. Let $N=\#\Omega.$

$$\sum_{i=1}^{N} \sum_{d \in D(i)} \pi_{id} r_{id}$$

163 184

- Since π_{id} are fractions of time, have

$$\forall i, d, \ \pi_{id} \ge 0 \text{ and } \sum_{i=1}^{N} \sum_{d \in D(i)} \pi_{id} = 1.$$

• Key idea is steady state constraint:

Fraction of time leave
$$j =$$
 Fraction of time enter j

In equation form:

$$\begin{aligned} \pi_j \cdot \mathbb{P}(\text{leave } j) &= \sum_{i \neq j} \pi_i \cdot \mathbb{P}(\text{move from } i \text{ to } j) \\ \sum_{d \in D(j)} \pi_j q_j(d) (1 - p(j|j, d)) &= \sum_{i \neq j} \sum_{d \in D(i)} \pi_i q_i(d) p(j|i, d) \\ \sum_{d \in D(j)} \pi_{jd} (1 - p(j|j, d)) &= \sum_{i \neq j} \sum_{d \in D(i)} \pi_{id} p(j|i, d) \end{aligned}$$

Distribute and bring $-\sum \pi_{jd} p(j|j,d)$ over to the other side:

$$\sum_{d \in D(i)} \pi_{jd} = \sum_{i} \sum_{d \in D(i)} \pi_{id} p(j|i,d).$$

- That's a linear constant!
- Linear program:

$$\begin{split} \text{maximize} \quad & \sum_{i=1}^{N} \sum_{d \in D(i)} \pi_{id} r_{id} \\ \text{subject to} \quad & \sum_{i=1}^{N} \sum_{d \in D(i)} \pi_{id} = 1 \\ & \sum_{d \in D(i)} \pi_{jd} = \sum_{i} \sum_{d \in D(i)} \pi_{id} p(j|i,d) \\ & \pi_{id} \geq 0 \text{ for all } i, d \end{split}$$

Qotd

• There are four states, and two decisions for each state, so 8 decision variables

$$\max z = 100\pi_{E,NR} + 80\pi_{G,NR} + 50\pi A, NR + 10\pi_{B,NR} \\ - 100\pi_{E,R} - 100\pi_{G,R} - 100\pi A, R - 100\pi_{B,R} \\ \text{s.t.} \qquad \pi_{E,NR} + \pi_{G,NR} + \pi_{A,NR} + \pi_{B,NR} + \pi_{E,NR} + \pi_{G,NR} - (E \text{ state}) \qquad \pi_{E,NR} = 0.7(\pi_{E,NR} + \pi_{E,R} + \pi_{G,R} + \pi_{A,R} + \pi_{B,R}) \\ (G \text{ state}) \qquad \pi_{G,NR} + \pi_{G,R} = 0.3(\pi_{G,R} + \pi_{A,R} + \pi_{B,R} + \pi_{E,NR} + \pi_{C,NR} + \pi_{A,NR} \\ (B \text{ state}) \qquad \pi_{B,R} + \pi_{B,NR} = \pi_{B,NR} + 0.4\pi_{A,NR} \\ \pi_{i,d} \ge 0$$

• The optimal LP solution has z = 60 using

$$\pi_{E,NR} = 0.35, \pi_{G,NR} = 0.50, \pi_{A,R} = 0.15, \text{ all others } 0.50, \pi_{A,R} = 0.15, \text{ all others } 0.50, \pi_{A,R} = 0.15, \pi_{A,R} = 0.15$$

• Note for each state, only one decision has $\pi_{i,d} > 0!$

Fact 34

The LP for finding average reward always has an optimal solution where there is a unique d(i) such that $\pi_{i,d(i)} > 0$ or $\pi_i = 0$.

Back to qotd...

- $\pi_{E,NR} > 0$, $\pi_{E,R} = 0$. So never repair when in E!
- $\pi_{G,NR} > 0$, $\pi_{G,R} = 0$. So never repair when in *G*!
- $\pi_{A,R} > 0$, $\pi_{A,NR} = 0$. So always repair when in *A*!
- $\pi_{B,R} = \pi_{B,NR} = 0 \Rightarrow \pi_B = 0.$
- With this setup, never return to *B* once you repair!
- Optimal policy is stationary:

$$\delta(E) = NR, \ \delta(G) = NR, \ \delta(A) = R, \ \delta(B) = R.$$

Discount versus average

- Use average when you know t is large, situation stable (For instance machine working on factory floor)
- Use discount when future more uncertain

Value iteration approach

- This is another way to find optimal policies
- Let $V_t(i)$ denote the biggest average reward that you can acheive in t rounds starting from state i
- If at state i, you make decision d, get average reward in rounds $2, \ldots, t$ of:

$$r_{id} + \beta \sum_{i=1}^{n} p(j|i,d) V_{t-1}(j).$$

Note the factor of β discounts the future rewards.

• So the best you can do is to pick the decision that maximizes this:

$$V_t(i) = \max_{d \in D(i)} \left\{ r_{id} + \beta \sum_{i=1}^n p(j|i, d) V_{t-1}(j) \right\}$$

• Just start with $V_0(i) = 0$, and run this recursion forward to converge to the best solution!

Fact 35

Let $d_t^*(i)$ denote the optimal decision to make at state i with t steps, and $\delta^*(i)$ the optimal decision with an infinite horizon. Then

$$\lim_{t \to \infty} d_t^*(i) = \delta^*(i).$$

Eventually this method converges to the right solution...

Problem is, no way to know if convergence has occured!

Chapter 38 Variance Reduction

Question of the Day Suppose X_1, X_2, \ldots are iid Unif([0, 1]) and I want to estimate $\mathbb{E}(X_i^2)$ through simulation. Can I do better than just generating X_1, \ldots, X_n and using

$$\frac{X_1^2 + X_2^2 + \dots + X_n^2}{n}?$$

Today

- Variance Reduction
- Antithetic Variables

Antithetic Variables

- · Idea: try to introduct negative correlation between variables
- Recall:

$$\mathbb{V}(X+Y) = \mathbb{V}(X) + \mathbb{V}(Y) + 2\operatorname{Cov}(X,Y).$$

• The correlation has the same sign as the covariance:

$$\operatorname{Cor}(X,Y) = \frac{\operatorname{Cov}(X,Y)}{\operatorname{SD}(X)\operatorname{SD}(Y)}.$$

Example

- What is $\mathbb{E}(X)$, $X \sim \text{Unif}([0, 1])$?
 - 1. Draw $X_1, ..., X_n \sim \text{Unif}([0, 1])$ iid
 - 2. Let

$$\hat{\mu}_n = \frac{X_1 + \dots + X_n}{n}$$

- Recall: X,Y independent means $\mathrm{Cov}(X,Y)=0$ so $\mathbb{V}(X+Y)=\mathbb{V}(X)+\mathbb{V}(Y)$

167 184

Then

$$\mathbb{V}(\hat{\mu}_n) = \sum_{i=1}^n \frac{\mathbb{V}(X_i)}{n^2} = \frac{n\mathbb{V}(X_1)}{n^2} = \frac{\mathbb{V}(X_1)}{n}.$$

Better idea

- Draw $X_1 \leftarrow \mathsf{Unif}([0,1])$
- Then make $X'_1 \leftarrow 1 X_1$
- Estimate $\mathbb{E}(X)$ with

$$\frac{X_1 + X_1'}{2} = \frac{X_1 + 1 - X_1}{2} = \frac{1}{2}$$

- Right answer every time!
- Why?

$$Cov(X_1, X'_1) = \mathbb{E}(X_1 X'_1) - \mathbb{E}(X_1) \mathbb{E}(X'_1)$$

= $\mathbb{E}(X_1(1 - X_1)) - \mathbb{E}(X_1)^2$
= $\mathbb{E}(X_1) - \mathbb{E}(X_1^2) - \mathbb{E}(X_1)^2$
= $1/2 - 1/3 - (1/2)^2 = -1/12$
 $\mathbb{V}(X_1 + X'_1) = \mathbb{V}(X_1) + \mathbb{V}(X'_1) + 2 \operatorname{Cov}(X_1, X'_1)$
= $\frac{1}{12} + \frac{1}{12} - 2\frac{1}{12} = 0.$

- Negative correlation reduces the variance of the sum
- Idea; When f(U) monotonic in U, average U and 1 U to decrease variance of average

Definition 65 Call U and 1 - U **antithetic variables**.

Note

- Usually generation of random variables slowest part of simulation
- So can get 1 U from U without much extra effort

Qotd

- Want to estimate $\mathbb{E}[X_i^2]$
- X_i^2 monotonic in X_i for $X_i \in [0, 1]$
- First look at original variance:

$$\begin{aligned} \mathbb{V}(X_i^2) &= \mathbb{E}[X_i^4] - \mathbb{E}[X_i^2]^2 \\ &= \int_0^1 s^4 \, ds - \left[\int_0^1 s^2 \, ds\right]^2 \\ &= \frac{1}{5} - \frac{1}{3^2} = \frac{4}{45}. \end{aligned}$$

• Now for the variance of antithetic. Let

$$W = \frac{U^2 + (1 - U)^2}{2}.$$

$$\mathbb{V}(W) = \mathbb{V}(U^2) - \mathbb{V}(U)$$
$$= \frac{4}{45} - \frac{1}{12}$$
$$= \frac{1}{180}.$$

Note that

$$\frac{\mathbb{V}U^2}{\mathbb{V}(W)} = 16.$$

• So roughly speaking, you have to take 16 times as many samples to get the equivalent level of accuracy!

Exponentials

- Recall that for an M/M/1 queue, interarrival and service times are exponentially distributed.
- One way to generate an exponential of rate λ :

$$U \leftarrow \mathsf{Unif}([0,1]), \ A \leftarrow -\frac{1}{\lambda}\ln(U).$$

- So if one simulation uses U_1, U_2, \ldots
- A second simulation could use $1 U_1, 1 U_2, \ldots$ to decrease variance

• Suppose we want to estimate the mean exponential of rate 1

$$\mathbb{V}\left(\frac{-\ln(U) - \ln(1 - U)}{2}\right) = \left[\frac{1}{4}\mathbb{E}\left[\left(-\ln(U) - \ln(1 - U)\right)^2\right)\right] - \mathbb{E}\left[\ln(U(1 - U))^2\right] = \mathbb{E}\left[\ln(U)^2 + \ln(1 - U)^2 + 2\ln(U)\ln\left(\int_0^1 \ln(u)^2 du = 2 \text{ (integration by parts)}\right] - \int_0^1 \ln(u)\ln(1 - u) du = 2 - \frac{\pi^2}{6} \text{ (improper at both ends)}\right]$$
$$\mathbb{V}\left(\frac{-\ln(U) - \ln(1 - U)}{2}\right) = \frac{1}{4}\left(2 + 2 + 2\left(2 - \frac{\pi^2}{6}\right)\right) - 1$$
$$= 1 - \frac{\pi^2}{12}$$

Try it in R

```
x <- runif(10000)
> mean(-log(x))
[1] 0.9873291
> mean((-log(x)-log(1-x))/2)
> var((-log(x)-log(1-x))/2)
> var(-log(x))
```

Chapter A Probability review

A.1 Elementary facts

Combinatorics The number of ways to arrange *n* objects in order is *n* factorial:

$$n! = n(n-1)(n-2)\cdots 1,$$

where 0! = 1. The number of ways to choose r objects from n objects is:

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

For $n_1 + n_2 + ... n_r = n$, the number of ways to choose n_1 objects of type 1, n_2 objects of type 2, up to n_r objects of type r, is

$$\binom{n}{n_1, n_2, \dots, n_r} = \frac{n!}{n_1! n_2! \cdots n_r!}.$$

Definitions These are the basic definitions for talking about probability.

The set of outcomes is called the *sample space* or *outcome space*, and is usually denoted Ω .

An *event* is a subset E of Ω such that $\mathbb{P}(E)$ is defined (an event is also sometimes called a *measurable* subset). When A is an event, the complement of A is also an event. Also if A_1, A_2, \ldots is a sequence of events, then $\bigcup_{i=1}^{\infty} A_i$ is also an event. (Any set of events with these properties is called a σ -algebra or σ -field.)

 \mathbb{P} is a function that given an event A, outputs the probability that the outcome lies in A. The events A and B are *disjoint* or *mutually exclusive* if $A \cap B = \emptyset$.

Measures A probability is a special type of measure that obeys the following four rules:

1. For event $B, 0 \leq \mathbb{P}(B)$ (probabilities are nonnegative real numbers)

2. $\mathbb{P}(\emptyset) = 0$ (the probability nothing happens is zero).

3. For B_1, B_2, \ldots disjoint events,

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} B_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(B_i).$$

4. $\mathbb{P}(\Omega) = 1$ (the probability that something occurs is 1).

Simple facts Some basic facts follow from these rules.

Prop: $0 \leq \mathbb{P}(A) \leq 1$. **Prop:** $\mathbb{P}(A^C) = 1 - \mathbb{P}(A)$. **Prop:** $\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(AB)$ **Prop:** $\mathbb{P}(\emptyset) = 0$.

A word about intersection For sets *A* and *B*, the intersection of *A* an *B* can be denoted $A \cap B$, AB, or *A*, *B*. All of these notations mean the same thing:

$$A \cap B := \{ x : x \in A \text{ and } x \in B \}.$$

Conditional probabilities If $\mathbb{P}(B) > 0$, the conditional probability of A given B is

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(AB)}{\mathbb{P}(B)}$$

Bayes' Formula If F_1, \ldots, F_n are disjoint and $\bigcup_{i=1}^n F_i = \Omega$, then

$$\mathbb{P}(F_i|A) = \frac{\mathbb{P}(A|F_i)\mathbb{P}(F_i)}{\mathbb{P}(A|F_1)\mathbb{P}(F_1) + \dots \mathbb{P}(A|F_n)\mathbb{P}(F_n)}.$$

Random variables A *random variable* is a function of the outcome. The values the random variable can take on are called *states*, and lie in the *state space*. In other words, a random variable is a function from the sample space to the state space.

For a discrete random variable $X \in \{x_1, x_2, x_3, \ldots\}$, the expected value of X is

$$\mathbb{E}[X] = \sum_{i=1}^{\infty} x_i \mathbb{P}(X = x_i).$$

For a continuous random variable $X \in \mathbb{R}$ with density f_X , the expected value of X is

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} s f_X(s) \, ds.$$

For any two random variables X and Y,

$$\mathbb{E}[X+Y] = \mathbb{E}[X] + \mathbb{E}[Y].$$

For two random variables X and Y are *uncorrelated* if and only if

$$\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y].$$

Independent random variables (see below) are always uncorrelated, but uncorrelated random variables are not always independent!

Independence Two events *A* and *B* are *independent* if

$$\mathbb{P}(AB) = \mathbb{P}(A)\mathbb{P}(B) \Leftrightarrow \mathbb{P}(A \mid B) = \mathbb{P}(A).$$

Two random variables X and Y are independent if for any event $X \in A$ and $Y \in B$,

$$\mathbb{P}(X \in A, Y \in B) = \mathbb{P}(X \in A)\mathbb{P}(Y \in B).$$

A.2 A short guide to solving probability problems Equally likely outcomes. If all outcomes are equally likely,

$$\mathbb{P}(E) = \frac{\text{number of outcomes in } E}{\text{total number of outcomes}}.$$

Trick #1: Use complements. It is often easier to find $\mathbb{P}(A^C)$ then $\mathbb{P}(A)$, remember

$$\mathbb{P}(A) = 1 - \mathbb{P}(A^C).$$

Trick #2: Use independence to turn intersections into products. If we want the probability of the intersection of A_1, \ldots, A_n , then we can break it apart only when the events are independent:

$$\mathbb{P}(A_1 \cdots A_n) = \mathbb{P}(A_1)\mathbb{P}(A_2) \cdots \mathbb{P}(A_n).$$

Trick #3: Use disjointness to turn unions into sums. If the events A_1, \ldots, A_n are disjoint,

$$\mathbb{P}(A_1 \cup \cdots \cup A_n) = \mathbb{P}(A_1) + \mathbb{P}(A_2) + \dots \mathbb{P}(A_n).$$

Trick #4: Use Principle of In/Ex to deal with any union. We can *always* break apart unions of events $A_1 \ldots A_n$ using the Principle of Inclusion/Exclusion, which we use most often when n = 2:

$$\mathbb{P}(A_1 \cup A_2) = \mathbb{P}(A_1) + \mathbb{P}(A_2) - \mathbb{P}(A_1 A_2).$$

Its easier to say the Principle of Inclusion/Exclusion in words than symbols: the probability of any event occurring is the sum of the probabilities that one event occurs minus the sum of the probabilities that 2 events occur plus the sum of the probabilities that 3 events occur etcetera until we reach the probability that all events occur.

Trick #5: Use De Morgan's Laws to covert unions and intersections. Convert back and forth between union and intersection using De Morgan's Laws:

$$(A_1A_2\cdots A_n)^C = A_1^C \cup A_2^C \cdots \cup A_n^C,$$
$$(A_1 \cup A_2 \cup \cdots \cup A_n)^C = A_1^C A_2^C \cdots A_n^C.$$

173 184

Trick #6: Use Bayes' Formula to reverse conditional probabilities. If you know $\mathbb{P}(A \mid F_i)$ for all *i* as well as $\mathbb{P}(F_i)$, and want $\mathbb{P}(F_i \mid A)$, then use Bayes' Formula.

Trick #7: Acceptance/Rejection Method 1 Suppose that we perform a trial which if successful, has outcomes A_1, \ldots, A_n . If we fail, then we try again until one of A_1 through A_n occur. Then

 $\mathbb{P}(A_i \text{ occurs on final trial}) = \mathbb{P}(A_i \text{ on first trial}) | \text{first trial a success}) = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial})}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial a success)}}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial a success)}}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial a success)}}{\mathbb{P}(text first trial a success)} = \frac{\mathbb{P}(A_i \text{ on first trial a success)}}{\mathbb{P}(text first trial a success)} = \frac{\mathbb$

Trick #8: Acceptance/Rejection Method 2 The other way to tackle acceptance rejection problem is using infinite series. Remember, when |r| < 1,

$$\sum_{i=0}^{\infty} r^i = \frac{1}{1-r}.$$

Common errors Some things to watch out for! Events use complements, unions, and intersections. A statement like $\mathbb{P}(A)^C$ doesn't make sense, since $\mathbb{P}(A)$ is a number. What was probably meant was $\mathbb{P}(A^C)$. Similarly, use +, - and times for numbers like probabilities, and never for sets. We haven't defined A + B, what was probably intended was $\mathbb{P}(A) + \mathbb{P}(B)$.

Steps to a problem: If you don't know how to get started on a problem, the following steps usually can get you going:

(1) Write down the sample space. Even if you can't write down the whole sample space, write down some of the outcomes. Make up symbols, like H for head or T for tails or W for win and L for a loss to make writing outcomes easier.

(2) Write down the events that you are given probabilities for, and the event that you are trying to find the probability of (the target event).

(3) See if you can express the target event in terms of union, intersection, or complements of the events that you are given (here is where the five tricks come into play).

Simple checks on an answer: Make sure that your final probabilities lie between o and 1. If you know that a set of probabilities must add to 1, then check by actually adding them. If you have a simple intuitive reason to believe that A is more likely than B, check that $\mathbb{P}(A) > \mathbb{P}(B)$.

A.3 A short guide to counting

Order matters When order matters, then there are n! ways to order n objects.

Thinking about *n* **choose** *k*. There are several ways of thinking about $\binom{n}{k}$, all of which are equivalent.

1. It's the number of the ways to choose a subset of size k from a set of size n.
- 2. It's the number of ways to order a group of letters $A \dots AB \dots B$ where A appears k times and B appears n k times.
- 3. Given n spaces, it's the number of ways to mark k of those spaces in some way.
- 4. It's the number of ways of choosing k out of n trials to be successful.

Multichoosing Now $\binom{n}{n_1,\ldots,n_r}$ is similar, in that it generalizes $\binom{n}{k}$. This is because $\binom{n}{k} = \binom{n}{n-k}$. The number *n* multichoose n_1, n_2, \ldots, n_r counts the following.

(1) It's the number of the ways to choose a partition of a set of size n where the first subset has size n_1 , the second n_2 , etcetera.

(2) It's the number of ways to order a group of letters $A_1 \dots A_1 A_2 \dots A_2 \dots A_r \dots A_r$ where A_i appears n_i times.

(3) Given n spaces, it's the number of ways to mark n_1 of those spaces with a 1, n_2 spaces with a 2, up to n_r spaces with n_r .

(4) Suppose each trial has r different outcomes. Then its the number of ways of choosing n_1 trials to have outcome 1, n_2 trials to have outcome 2, up to n_r trials having outcome r.

When all else fails. Almost any problem can be written as a problem with ordering. If you are uncomfortable with n choose r or can't figure out what should be ordered and what shouldn't then give everything in your problem a number and order everything.

For example, what's the probability of choosing a given five card hand from a set of 52 cards? One way: number of outcomes is 1, total number of outcomes is $\binom{52}{5}$, so

$$\mathbb{P}(\text{hand}) = \frac{1}{\binom{52}{5}}.$$

Another way: number all the cards $1, \ldots, 52$ and order them in any one of 52! ways. Then any outcome where the five cards we are interested in appear first in the ordering of cards works. There are 5! ways to order these cards and (52 - 5)! ways to order the remaining 47 cards, so the total number of outcomes is 5!(47!), so

$$\mathbb{P}(\mathsf{hand}) = \frac{5!47!}{52!},$$

which is the same answer as the other way.

Another example: given a random ordering of the letters MIIIISSSSPP, what's the probability that it spells MISSISSIPPI? Think about numbering every symbol, so we are ordering $x_1x_2x_3x_4x_5x_6x_7x_8x_9x_{10}x_{11}$, where $x_1 = M$, x_2 through x_5 equal I, etc. Then the total number of outcomes is 11!. The number of outcomes that are successful? Well x_1 has to be in first position, x_2 , x_3 , x_4 and x_5 have to occupy positions 2, 5, 8, and 10 (which they can do in 4! ways, there are 4! ways to order the x_i that equal S and 2! ways to order the x_i that equal P. So

$$\mathbb{P}(\text{MISSISSIPPI}) = \frac{1!4!4!2!}{11!}.$$

175 184

A.4 How to find $\mathbb{E}[X]$

Step 1 Find the values that *X* can take on with positive probability (this is called the *positive support* of *X*). If *X* is discrete, this will be either a finite number of values $\{x_1, \ldots, x_n\}$ or a countable number of values $\{x_1, x_2, \ldots\}$. If *X* is continuous, it could be an interval or union of intervals, like $(0, \infty)$ or $(3, 4) \cup [10, 15)$.

Step 2 Use the right formula. If X is discrete, then $\mathbb{E}[X]$ is the sum over all values of x such that $\mathbb{P}(X = x) > 0$ of the outcome times the probability. So if $X \in \{x_0, x_1, \ldots\}$, then

$$\mathbb{E}[X] = \sum_{x:p(x)>0} xp(x) = \sum_{i=1}^{\infty} x_i \mathbb{P}(X = x_i).$$

If X is continuous with density f_X then

$$\mathbb{E}[X] = \int_{\mathbb{R}} x f_X(x) \, dx.$$

If $X \in \{0, 1, 2, 3, ...\}$, then the *Tail Sum Formula* gives an alternate way to find the expected value:

$$\mathbb{E}[X] = \sum_{i=0}^{\infty} \mathbb{P}(X > i).$$

If X is continuous and $\mathbb{P}(X\geq 0)=1,$ then the Tail Sum Formula is

$$\mathbb{E}[X] = \int_{x=0}^{\infty} \mathbb{P}(X > x) \, dx.$$

Conditional expectation To find $\mathbb{E}[A|B]$, treat *B* as a constant and calculate the probability in the exact same way as above. For all random variables *A* and *B*:

$$\mathbb{E}[\mathbb{E}[A|B]] = \mathbb{E}[A].$$

Note: If we wish to find $\mathbb{E}[g(X)]$ then use

$$\mathbb{E}[g(X)] = \sum_{x:\mathbb{P}(X=x)>0} g(x)\mathbb{P}(X=x) = \sum_{i=1}^{\infty} g(x_i)\mathbb{P}(X=x_i),$$

and

$$\mathbb{E}[g(X)] = \int_{\mathbb{R}} g(s) f_X(s) ds.$$

For uncorrelated random variables, $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$. Independent random variables are uncorrelated, but uncorrelated random variables might not be independent.

Some properties of expected value:

• For any two random variables (correlated or uncorrelated) $\mathbb{E}[X+Y] = \mathbb{E}[X] + \mathbb{E}[Y]$.

A.5 How to find $\mathbb{V}(X)$

Method 1: Use

$$\mathbb{V}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2.$$

Method 2: Use

 $\mathbb{V}(X) = \mathbb{E}[(X - \mathbb{E}[X])^2].$

Some properties

- For uncorrelated random variables, $\mathbb{V}(X+Y) = \mathbb{V}(X) + \mathbb{V}(Y)$.
- For random variable X and constant $\alpha \in \mathbb{R}$, $\mathbb{V}(\alpha X) = \alpha^2 \mathbb{V}(X)$, $SD(\alpha X) = \alpha SD(X)$.

A.6 Distributions

The *distribution* of a random variable is a complete listing of $\mathbb{P}(X \in A)$ for all sets A of interest. The distribution also referred to as the law of X, and denoted $\mathcal{L}(X)$. When X and Y have the same distribution, this is denoted

$$X \sim Y$$
, or $\mathcal{L}(X) = \mathcal{L}(Y)$.

The *distribution function* of a random variable X (also known as the cumulative distribution function) is

$$F(a) = \mathbb{P}(X \le a).$$

This is a function that is bounded, that is, it always lies between 0 and 1. It is also right continuous, that is if a_1, a_2, a_3, \ldots decrease and their limit is a, then limit of $F(a_1), F(a_2), \ldots$ equals F(a).

Because of a theorem from measure theory called the Carathéodory Extension Theorem, knowing F allows computation of $\mathbb{P}(X \in A)$ for any A of interest. In particular, if A = (a, b], then $\mathbb{P}(X \in A) = F(b) - F(a)$. (Looks a bit like the fundamental theorem of calculus, which is one reason why F is always capitalized when used for the distribution function.)

More precisely, if F_X is the distribution function of X and F_Y is the distribution function of Y, then

$$\mathcal{L}(X) = \mathcal{L}(Y) \iff F_X(a) = F_Y(a) \ \forall a.$$

If X is discrete then the graph of F(a) will have jumps, if X is continuous then F(a) will be continuous. Some more formulas that come in handy:

$$\begin{split} \mathbb{P}(a < X \le b) &= F(b) - F(a) \\ \mathbb{P}(a < X < b) &= F(b) - F(a) - \mathbb{P}(X = b) \\ \mathbb{P}(a \le X < b) &= F(b) - F(a) - \mathbb{P}(X = b) + \mathbb{P}(X = a) \\ \mathbb{P}(a \le X \le b) &= F(b) - F(a) + \mathbb{P}(X = a). \end{split}$$

Remember that for continuous random variables $\mathbb{P}(X = s) = 0$ for any s, so the right hand side of these formula just becomes F(b) - F(a). Also for continuous X,

$$f(a) = \frac{dF(a)}{da}$$

and

$$F(a) = \int_{-\infty}^{a} f(a) da,$$

where f(x) is the probability density function (sometimes just called the density) of X.

Finally, say that X_1, X_2, \ldots are independent identically distributed, or iid, if they are independent and all have the same distribution.

A.7 Discrete distributions

A random variable is *discrete* if it only takes on a finite or countably infinite number of values. The distribution of a discrete random variable is also called discrete in this instance.

Uniform Written: $Unif(\{1, ..., n\})$. The story: roll a fair die with *n* sides.

$$\mathbb{P}(X=i) = \frac{1}{n}\mathbb{1}(i \in \{1,\dots,n\})$$
$$\mathbb{E}[X] = \frac{n+1}{2}$$
$$\mathbb{V}(X) = \frac{(n-1)(n+1)}{12}$$

Bernoulli Written: Bern(p). The story: flip a coin that comes up heads with probability p, and count the number of heads on the single coin flip. Also, the number of successes in a single trial where the trial is a success with probability p.

$$\mathbb{P}(X=1) = p, \ \mathbb{P}(X=0) = 1 - p$$
$$\mathbb{E}[X] = p$$
$$\mathbb{V}(X) = p(1-p).$$

Binomial Written: Bin(n, p). The story: flip iid coins n times where the probability of heads is p and count the number of heads. Also, the number of successes in a single trial where the trial is a success with probability p. Also if X_1, \ldots, X_n are iid Bern(p), then $X = X_1 + X_2 + \ldots + X_n \sim Bin(n, p)$.

$$\mathbb{P}(X=i) = \binom{n}{i} p^i (1-p)^{n-i} \mathbb{1}(i \in \{\{0,\dots,n\})$$
$$\mathbb{E}[X] = np$$
$$\mathbb{V}(X) = np(1-p).$$

178 184

Geometric Written: Geo(p). The story: flip iid coins with probability p of heads and counting the number of flips needed for one head. Also, the number of trials needed for 1 success when the probability of success at each trial is p and each trial is independent.

$$\mathbb{P}(X=i) = (1-p)^{i-1}p\mathbb{1}(\{0,1,\ldots\})$$
$$\mathbb{E}[X] = \frac{1}{p}$$
$$\mathbb{V}(X) = \frac{1-p}{p^2}.$$

Negative Binomial Written: NegBin(r, p). The story: flipping iid coins with probability p of heads and counting the number of flips needed for r heads to arrive. Also, the number of trials needed for r successes when the probability of success at each trial is p and each trial is independent.

Also $X = X_1 + X_2 + ... X_r$, where X_i are iid and distributed as Geo(p).

$$\mathbb{P}(X=i) = \binom{i-1}{r-1} p^r (1-p)^{i-r} \mathbb{1}(\{0,1,\ldots\})$$
$$\mathbb{E}[X] = \frac{r}{p}$$
$$\mathbb{V}(X) = r \frac{1-p}{p^2}.$$

Hypergeometric Written: Hypergeo(N, m, n). The story: drawing n balls from an urn holding m green balls and N - m blue balls and counting the number of green balls chosen.

$$\mathbb{P}(X=i) = \frac{\binom{m}{i}\binom{N-m}{n-i}}{\binom{N}{n}}\mathbb{1}(\{0,1,\ldots,n\})$$
$$\mathbb{E}[X] = \frac{nm}{N}$$
$$\mathbb{V}(X) = \frac{N-n}{N-1}np(1-p).$$

Zeta Written: $Zeta(\alpha)$. A.k.a. Zipf or power law. The story: things like city sizes and incomes have Zeta distributions.

$$\begin{split} \mathbb{P}(X=i) &= \frac{C}{i^{\alpha+1}}\mathbb{1}(\{1,2,\ldots\})\\ \mathbb{E}[X] &= \text{no closed form}\\ \mathbb{V}(X) &= \text{no closed form.} \end{split}$$

Special notes: Except for special values of α like 1, we do not have a closed form solution for the value of C, the normalizing constant. Choose C so that $\sum_{i=1}^{\infty} \mathbb{P}(X = i) = 1$. Similarly, there are no closed form solutions for $\mathbb{E}[X]$ or $\mathbb{V}(X)$. These must be evaluated numerically. When $\alpha < 1$, $\mathbb{E}[X]$ does not exist (or is considered infinite). Similarly, when $\alpha < 2$, $\operatorname{Var}(X)$ does not exist (or can be considered infinite).

Poisson Written: $\text{Pois}(\mu)$. The story: given that the chance of an arrival in time t to t+dt is λdt , and $\mu = \lambda T$, then this is the number of arrivals in the interval [0, T]. X_1, X_2, \ldots , it is

$$\max_i X_1 + X_2 + \ldots + X_i < 1.$$

$$\mathbb{P}(X = i) = e^{-\mu} \frac{\mu^i}{i!} \mathbb{1}(\{0, 1, \ldots\})$$
$$\mathbb{E}[X] = \mu$$
$$\mathbb{V}(X) = \mu.$$

A.8 Continuous Distributions

A random variable is *continuous* if $\mathbb{P}(X = a) = 0$ for all a. The distribution of a continuous random variable is also called continuous.

Uniform (continuous) Written: Unif(A). The story: a point is uniform over A if for all $B \subseteq A$, the chance the point falls in B is the Lebesgue measure of B divided by the Lebesgue measure of A. For general A:

$$f(x) = \frac{1}{\text{Lebesgue measure of } A} \mathbb{1}(x \in A)$$

When A = [a, b], more specifically:

$$f(x) = \frac{1}{b-a} \mathbb{1}(x \in (a, b])$$

$$F(x) = \frac{x-a}{b-a} \mathbb{1}(x \in [a, b]) + \mathbb{1}(x > b)$$

$$\mathbb{E}[X] = \frac{b+a}{2}$$

$$\mathbb{V}(X) = \frac{(b-a)^2}{12}$$

Normal Written: $N(\mu, \sigma^2)$. The story: when you sum variables with finite mean and standard deviation together, they are well approximated by a normal distribution.

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$$
$$F(x) = \Phi\left(\frac{x-\mu}{\sigma}\right)$$
$$\mathbb{E}[X] = \mu$$
$$\mathbb{V}(X) = \sigma^2$$

Addition of normals. Adding independent normal random variables gives back another normal random variable. If $X_i \sim N(\mu_i, \sigma_i^2)$, and $X = X_1 + X_2 + \ldots + X_n$, then

$$X \sim \mathsf{N}\left(\sum_{i} \mu_{i}, \sum_{i} \sigma_{i}^{2}\right).$$

For X, Y independent N(0, 1) random variables, the joint distribution of (X, Y) is rotationally invariant.

Normal random variables are symmetric around μ , and so $\Phi(x) = 1 - \Phi(-x)$.

Exponential Written: $Exp(\lambda)$. What it is: when events occur continuously over time at rate λ , this is the time you have to wait for the first event to occur.

$$\begin{split} f(t) &= \lambda e^{-\lambda t} \mathbbm{1}(t \in (0,\infty)) \\ F(t) &= \begin{cases} 1 - e^{-\lambda t} & a \ge 0 \\ 0 & a < 0 \end{cases} \\ \mathbb{E}[X] &= \frac{1}{\lambda} \\ \mathbb{V}(X) &= \frac{1}{\lambda^2} \end{split}$$

A.9 How to use the Central Limit Theorem (CLT)

The CLT says that if X_1, X_2, \ldots are identically distributed random variables and $Z_n = X_1 + \ldots + X_n$, then

$$\lim_{n \to \infty} \mathbb{P}\left(\frac{Z_n - \mathbb{E}[Z_n]}{\sqrt{\mathbb{V}(Z)}} \le a\right) = \Phi(a).$$

We use it as an approximation tool for $Z = X_1 + \ldots X_n$:

$$\mathbb{P}\left(\frac{Z - \mathbb{E}[Z]}{\sqrt{\mathbb{V}(Z)}} \le a\right) \approx \Phi(a).$$

Often we are interested in approximating the probability of things like $\mathbb{P}(Z \leq b)$ where $Z = X_1 + \ldots X_n$. This takes two steps.

Step 1 If Z is integral, apply the half integer correction. So instead of $\mathbb{P}(Z \le i)$ we write $\mathbb{P}(Z \le i + 1/2)$.

Step 2 Subtract off $\mathbb{E}[Z]$ and divide by the square root of Var(Z). So

$$\mathbb{P}(Z \le b + 0.5) = \mathbb{P}\left(\frac{Z - \mathbb{E}[Z]}{\sqrt{\mathbb{V}(Z)}} \le \frac{b + 0.5 - \mathbb{E}[Z]}{\sqrt{\mathbb{V}(Z)}}\right).$$

Step 3 Apply the CLT and say

$$\mathbb{P}(Z \le b) \approx \Phi\left(\frac{b + 0.5 - \mathbb{E}[Z]}{\sqrt{\mathbb{V}(Z)}}\right)$$

Chapter B Proofs of Theorems

B.1 The Mixture Space Theorem

Fact 36 (Mixture Space Theorem, Herstein and Milnor) A preference relation \preceq on a convex set Π is independent and continuous if and only if there exists an affine utility representation $U : \Pi \to \mathbb{R}$ of \preceq .

Moreover, if $U : \Pi \to \mathbb{R}$ is an affine representation of \leq , then $U' : \Pi \to \mathbb{R}$ is an affine representation of \leq iff there exist a > 0 and $b \in \mathbb{R}$ such that U' = aU + b.

This proof comes from lecture notes of Roee Teper.

Proof: (When a maximum and minimum element exists). Let L^* be the largest element of Π and L_* the smallest. $L_* \prec L^*$ (otherwise all elements are equal and the theorem is trivially true.)

Step 1 $\alpha \leq \beta$ if and only if $\alpha L^* + (1 - \alpha)L_* \preceq \beta L^* + (1 - \beta)L_*$.

Start with $\alpha < \beta$. Note that $\beta > 0$. So set $\gamma = \alpha/\beta \in [0, 1]$ and independence gives us

$$L_* = \beta L_* + (1 - \beta) L_* \prec \beta L^* + (1 - \beta) L_*.$$

Which in turn gives us:

$$\beta L^* + (1-\beta)L_* = (1-\gamma)[\beta L^* + (1-\beta)L_*] + \gamma[\beta L^* + (1-\beta)L_*]$$

$$\succ (1-\gamma)L_* + \gamma[\beta L^* + (1-\beta)L_*] \text{ (independence)}$$

Now

$$(1-\gamma) + \gamma(1-\beta) = 1 - \gamma + \gamma - \alpha = 1 - \alpha,$$

so

$$\beta L^* + (1 - \beta)L_* \succ (1 - \alpha)L_* + \alpha L^*$$

Now suppose $\alpha = \beta$. Then the statement is trivially true. Since the argument above also applies to $\beta < \alpha$, the if and only if holds.

Step 2 Given $L \in \Pi$, there exists a unique α^* such that $\alpha^* L^* = (1 - \alpha^*)L_* = L$. The fact that such an $\alpha^* \in [0, 1]$ exists is continuity, the previous step shows that it must be unique.

Step 3 Create U. For $L \in \Pi$, let U(L) be the unique α^* from the previous step. For any two L and M in Π , the following holds

$$L \preceq M \Leftrightarrow U(L)L^* + (1 - U(L))L_* \qquad \preceq U(M)L^* + (1 - U(M))L_* \Leftrightarrow U(L) \leq U(M).$$

The first equivalence is just the definition of U(L) and U(M), the second equivalence is just Step 1 again.

Step 4 This U is affine. Consider two $L, M \in \Pi$ and $\alpha \in [0, 1]$. Let $N = \alpha L + (1 - \alpha)M$. Then by the way U was defined:

$$U(N)L^{*} + [1 - U(N)]L_{*} =$$

= $N = \alpha L + (1 - \alpha)M$
= $\alpha [U(L) \cdot L^{*} + (1 - U(L))L_{*}] + (1 - \alpha)[U(M) \cdot L^{*} + (1 - U(M))L_{*}]$
= $[\alpha U(L) + (1 - \alpha)U(M)]L^{*} + [1 - (\alpha U(L) + (1 - \alpha)U(M))]L_{*}$

By uniqueness of the representation

$$\alpha U(L) + (1-\alpha)U(M) = U(N) = U(\alpha L + (1-\alpha)M).$$

Note: Mixture space theorem actually holds over any convex subset of \mathbb{R}^n , not just probability simplex.

Step 6 Suppose U is an affine representation of \leq . Then U' is also an affine representation if and only if there exists a > 0 and $b \in \mathbb{R}$ such that

$$U' = aU + b.$$

One direction is straightforward: if U' = aU + b, then U' is an affine, strictly increasing transformation of U, and so must also represent \preceq .

For the other direction, given U and U', let

$$a = \frac{U(L^*) - U(L_*)}{U'(L^*) - U'(L_*)}.$$

This is positive since $L_* \prec L^*$. Let

$$b = U'(L_*) - aU(L_*)$$

 \square

Then it is just algebra to show that U' = aU + b.

184 184