

MATH 56A SPRING 2008
STOCHASTIC PROCESSES

KIYOSHI IGUSA

CONTENTS

1. Finite Markov Chains	22
1.1. Concept and examples	22
1.2. Long term behavior	27
1.3. Invariant probability distribution	31
1.4. Transient classes	36
1.5. Canonical form of P	40
1.6. The substochastic matrix Q	45
1.7. Transient and recurrent	53
Homework 1a: Finite Markov Chains	54
Homework 1b: Leontief Model	55

1. FINITE MARKOV CHAINS

1.1. **Concept and examples.** On the first day I explained the concept behind finite Markov chains, gave the definition and two examples. But first I explained how you convert a higher order difference equation into a first order matrix equation. When we randomize this process we get a finite Markov chain.

1.1.1. *reduction to first order.* I used the Fibonacci sequence as an example to illustrate how higher order equations can be reduced to first order equations in more variables. The Fibonacci sequence is the sequence

$$1, 1, 2, 3, 5, 8, 13, \dots$$

given by the second order difference equation

$$f(n) = f(n-1) + f(n-2).$$

To convert this to first order you let

$$g(n) := f(n-1).$$

Then $f(n-2) = g(n-1)$ and the original equation becomes:

$$f(n) = f(n-1) + g(n-1).$$

Thus $(f(n), g(n))$ depends only on $(f(n-1), g(n-1))$ and the relation is given by the matrix equation:

$$(f(n), g(n)) = (f(n-1), g(n-1)) \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$$

I explained it like this: You have to make your decision about what to do tomorrow based on the information you have today. You can only use the information that you had yesterday if you *recorded it*. Thus, every day, you need to record the important information, either on paper or in your computer, otherwise it is lost and won't be available tomorrow.

The Fibonacci sequence, in this first order form, looks like this:

n	0	1	2	3	4	5
$f(n)$	1	1	2	3	5	8
$g(n)$	0	1	1	2	3	5

So, on Day 4, the information you have is today's number 5 and the record you kept of yesterday's number 3. You add these to get tomorrow's number 8 and you record the number 5 so that you have still have it tomorrow. Each day you look at the information you get that day and the information that was recorded from the past. So, this process is realistic and makes sense.

1.1.2. *concept.* “A stochastic process is a random process which evolves with time.” This definition is too broad for a careful, complete mathematical analysis, especially at the beginning.


We want to start with simple models that we can analyze and understand completely. Then we will go to more and more general models adding complexities one step at a time.

A Markov chain is a stochastic process which has four simplifying assumptions:

- (1) There are only finitely many states. For example, in the Kermack-McKendrick model, there were only 3 states: S, I, R . I also used the example of the Brandeis campus. If we made the movement of people on campus into a Markov process then the set of states would be the buildings (plus one for the outside). Your exact location, for example which room you were in, is disregarded.
- (2) Time is discrete. Time is a nonnegative integer (starting at $t = 0$). For example, for movement of people on campus, people are only allowed to move from building to building on the hour. Or, we only record or notice which building people are in at 1pm, 2pm, 3pm, etc.
- (3) You forget the past. What happens at time $n + 1$ depends only on the situation at time n . Which building you are in at 2pm depends only on which building you were in at 1pm. If you add more states (more variables), you can keep track of information from the past and still satisfy the “forget the past” rule.
- (4) Rules of movement do not change with time. If, at 2pm, everyone in building 2 move to building 5 then the same thing will happen at 3pm, 4pm, etc. The Fibonacci sequence or any first order recurrence has this property.

I used two examples to illustrate these principles.

1.1.3. *mouse example.*¹



¹from “Markov Chains ... ” by Pierre Brémaud

In this example, a mouse is randomly moving from room to room. The cat and cheese do not move. But, if the mouse goes into the cat's room, he never comes out. If he reaches the cheese he also does not come out.

This will be a Markov chain with the following details.

- (1) There are 5 states (the five rooms). I numbered them: 1,2,3,4,5.
- (2) The mouse moves in integer time, say every minute.
- (3) The mouse does not remember which room he was in before. Every minute he picks an adjacent room at random, possibly going back to the room he was just in.
- (4) The probabilities do not change with time. For example, whenever the mouse is in room 3 he will go next to room 2,4 or 5 with equal probability.

The mouse moves according to the *transition probabilities*

$$p(i, j) = \mathbb{P}(\text{the mouse goes to room } j \text{ when he is in room } i).$$

These probabilities form a matrix called the *transition matrix*:

$$P = (p(i, j)) = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{matrix} & \begin{pmatrix} 0 & 1/2 & 0 & 1/2 & 0 \\ 1/2 & 0 & 1/2 & 0 & 0 \\ 0 & 1/3 & 0 & 1/3 & 1/3 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}$$

I pointed out the two important properties of this matrix:

- (1) Every row adds up to 1. This is because the mouse has to go somewhere or stay where he is. When all of the possibilities are listed and they are mutually exclusive, the probabilities must add up to 1.
- (2) The entries are nonnegative and at most 1 (because they are probabilities).

1.1.4. *students example.* ²

Each year, the students at a certain college either flunk out, repeat the year or go on to the next year with the following probabilities:

$$\begin{aligned} p &= \mathbb{P}(\text{flunking out of school}) \\ q &= \mathbb{P}(\text{repeating a year}) \\ r &= \mathbb{P}(\text{passing to the next year}) \end{aligned}$$

The first step is to determine what are the states. Then find the transition matrix. Later we can answer other questions, such as: What is

²from "Finite Markov Chains" by Kemeny and Snell

the probability that a Freshman will eventually graduate? and how long will it take?

There are 6 states: the student is either

- (1) Freshman
- (2) Sophomore
- (3) Junior
- (4) Senior
- (5) graduated
- (6) flunked out

The transition matrix is

$$P = (p(i, j)) = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{matrix} & \begin{pmatrix} q & r & 0 & 0 & 0 & p \\ 0 & q & r & 0 & 0 & p \\ 0 & 0 & q & r & 0 & p \\ 0 & 0 & 0 & q & r & p \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{matrix}$$

It is important to notice that the rows add up to 1:

$$p + q + r = 1.$$

This means there are no other possibilities except for the three that were listed.

1.1.5. *definition.* Here is the precise mathematical definition.

Definition 1.1. A *finite Markov chain* is a sequence of random variables X_0, X_1, \dots which take values in a finite set S called the *state space* so that, for all $n \geq 0$ and all values of x_0, x_1, \dots, x_n , we have:

$$\mathbb{P}(X_{n+1} = x \mid X_0 = x_0, X_1 = x_1, \dots, X_n = x_n) = \mathbb{P}(X_1 = x \mid X_0 = x_n)$$

The $S \times S$ matrix P with entries

$$p(x, y) := \mathbb{P}(X_1 = y \mid X_0 = x)$$

is called the *transition matrix*.


The probability equation can be broken up into two steps:

$$\mathbb{P}(X_{n+1} = x \mid X_0 = x_0, X_1 = x_1, \dots, X_n = x_n) = \mathbb{P}(X_{n+1} = x \mid X_n = x_n)$$

$$\mathbb{P}(X_{n+1} = x \mid X_n = x_n) = \mathbb{P}(X_1 = x \mid X_0 = x_n)$$

The first equation says that what happens at time $n + 1$ depends only on the state at time n and not on the state at previous times. The second equation says that the transition probabilities are the same at time n as they were at time 0.

1.1.6. *graphic notation.* I first discussed the graphic representation of a Markov chain. Here is an example.



This diagram is a (*directed*) *graph*. It has vertices representing states and arrows representing possible movement in one unit time. The numbers on the arrows are the transition probabilities. Thus the transition matrix is:


$$P = \begin{pmatrix} 0 & 1 & 0 \\ 1/2 & 0 & 1/2 \\ 0 & 1 & 0 \end{pmatrix}.$$

This example is called a *random walk with reflecting walls* on the graph:



The term “random walk” refers to the fact that, at any internal vertex (of which there is only one), you move in any direction with equal probability and at any endpoint (“leaf”) of the graph you move inward with probability one.

The Mouse-Cat-Cheese example becomes the following diagram in this notation:



At this point I explained that there are *implied* or *implicit* loops at the vertices 4,5 with probability 1. These two states also happen to be *absorbing states* which means that you can enter but you cannot leave. In the transition matrix there is a 1 in the (i, i) position if i is an absorbing state.

1.2. Long term behavior. I did an overview of many of the concepts in this chapter, centered around the following question: What is the long term behavior of the Markov chain? Of course Markov chains are random. So, we can only talk about probabilities. So, a long term probability question might be:


$$\mathbb{P}(X_{1000} = 2 \mid X_0 = 1) = ?$$

for the random walk with reflecting walls. I.e., what is the probability that you will end up at state 2 after 1000 steps if you start in state 1?

The answer is that this probability is 0. This is because the Markov chain is *periodic* of period 2. The states are divided into even and odd numbers and you always move from even to odd and odd to even at every step. So, it is not possible to move from an odd position such as 1 to an even position such as 2 in an even number of steps, such as $n = 1000$. Thus, if you start at vertex 1, this Markov chain will forever oscillate between two sets of possibilities:

- (1) You are at vertex 2 for all odd n .
- (2) You are at one of the two odd vertices 1, 3 with equal probability at all even times $n \geq 2$.

I also gave another example of a hexagon:



In the long run this tends to oscillate between two probability distributions. We will study this more carefully on Wednesday.

The theorem is that the distribution stabilizes if the Markov chain is aperiodic and “irreducible” which means that it forms one “communication class.”

1.2.1. Communication classes. I made the following definitions. Since students were confused, I will write the key point first: The *communication class* of a state x , by definition, consists of x and all those y which communicate with x :

Definition 1.2. I write $x \rightarrow y$ if it is possible to move from x to y in a finite number of steps:

$$x \rightarrow y \quad \Leftrightarrow \quad \mathbb{P}(X_n = y \mid X_0 = x) > 0 \text{ for some } n > 0$$

$$\Leftrightarrow \sum_{n=1}^{\infty} \mathbb{P}(X_n = y | X_0 = x) > 0$$

I also write $x \leftrightarrow y$ if $x \rightarrow y$ and $y \rightarrow x$, i.e., if you can go from x to y and from y back to x . The communication class of x is

$$C(x) := \{x\} \cup \{y \in S | x \leftrightarrow y\}$$

In class I wrote this as:

$$C(x) = \{y \in S | x \leftrightarrow y \text{ or } x = y\}$$


The point is that $x \in C(x)$ by definition.

1.2.2. *Transient and recurrent classes.* There are only two kinds of communication classes: transient and recurrent.

- (1) A *transient class* is a communication class from which is possible to leave.
- (2) A *recurrent class* (also called a *absorbing class*) is a communication class from which it is not possible to leave.

I pointed out that, although a recurrent class is the same as an absorbing class, the adjectives “recurrent” and “absorbing” are not interchangeable. They mean different things.

It is very easy from the diagrams to determine the transient and recurrent classes. I did some examples in class and students did Worksheet 1.



In this example, there are two recurrent classes which you can see at the bottom: 4, 5 form one recurrent class and 6, 7 form the other recurrent class. There are three transient classes: 1, 2, 3. Each of these forms

a transient communication class. There are also several implied loops which are not drawn.

In the Cat-Mouse-Cheese example there is one transient class consisting of states 1,2,3 and there are two recurrent classes which are the absorbing states. I pointed out that absorbing states always form one point recurrent communication classes.


1.2.3. *Powers of P.* At the end I started to write down the formulas. We want to know the long range probabilities such as:

$$p_n(i, j) := \mathbb{P}(X_n = j \mid X_0 = i).$$

This is the probability of going from state i to state j in exactly n steps. The theorem is:

Theorem 1.3. *The number $p_n(i, j)$ is equal to the (i, j) entry of the matrix P^n .*

I didn't prove this but I did an example for $n = 2$. I took the first example in the Worksheet.



The question is: What is $p_2(1, 3)$? This is the probability of going from 1 to 3 in two steps. But there are two ways to do this. You can either go across to 2 and then down to 3 or you can go down to 4 and then across to 3. The probabilities for these two paths are added. Each path has two segments and the probabilities are multiplied since you must go through both segments. (The probability of A or B is the sum of the probabilities, assuming A, B are exclusive, and the probability of A and B is the product of the probabilities, assuming A, B are independent.) So, the probability is

$$p_2(1, 3) = \left(\frac{1}{2}\right) \left(\frac{2}{3}\right) + \left(\frac{1}{2}\right) \left(\frac{1}{4}\right) = \frac{1}{3} + \frac{1}{8} = \frac{11}{24}$$

If you look at the $(1, 3)$ entry of P^2 you see the same thing:

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

So, to compute $(P^2)_{13}$ you multiply the blue numbers:

$$P^2 = \begin{pmatrix} 0 & 1/2 & 0 & 1/2 \\ 0 & 1/3 & 2/3 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/3 & 3/4 \end{pmatrix} \begin{pmatrix} 0 & 1/2 & 0 & 1/2 \\ 0 & 1/3 & 2/3 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/3 & 3/4 \end{pmatrix}$$

$$(P^2)_{13} = 0 \cdot 0 + \left(\frac{1}{2}\right) \left(\frac{2}{3}\right) + 0 \cdot 0 + \left(\frac{1}{2}\right) \left(\frac{1}{4}\right)$$

Or, more generally:

$$p_2(i, j) = \sum_k p(i, k)p(k, j) = (P^2)_{ij}.$$

1.3. Invariant probability distribution.

Definition 1.4. A *probability distribution* is a function

$$\pi : S \rightarrow [0, 1]$$

from the set of states S to the closed unit interval $[0, 1]$ so that

$$\sum_{i \in S} \pi(i) = 1.$$

When the set of states is equal to $S = \{1, 2, \dots, s\}$ then the condition is:


$$\sum_{i=1}^s \pi(i) = 1.$$

Definition 1.5. A probability distribution π is called *invariant* if

$$\pi P = P.$$

I.e., π is a left eigenvector for P with eigenvalue 1.

1.3.1. *probability distribution of X_n .* Each X_n has a probability distribution. I used the following example to illustrate this.



The numbers $1/3$ and $1/4$ are transition probabilities. They say nothing about X_0 . But we need to start in a random state X_0 . This is because we need to understand how the transition from one random state to another works so that we can go from X_n to X_{n+1} .

X_0 will be equal to either 0 or 1 with probability:

$$\alpha_1 = \mathbb{P}(X_0 = 1),$$

$$\alpha_2 = \mathbb{P}(X_0 = 2).$$

These two numbers are between 0 and 1 (inclusive) and add up to 1:

$$\alpha_1 + \alpha_2 = 1.$$

So, $\alpha = (\alpha_1, \alpha_2)$ is a probability distribution. α is the probability distribution of X_0 and is called the *initial (probability) distribution*.

Once the distribution of X_0 is given, the probability distribution of every X_n is determined by the transition matrix:

Theorem 1.6. *The probability distribution of X_n is the vector αP^n .*

So, in the example,

$$\begin{aligned} \mathbb{P}(X_2 = 2) &= \sum_{i=1}^2 \sum_{j=1}^2 \underbrace{\mathbb{P}(X_0 = i)}_{\alpha_i} \underbrace{\mathbb{P}(X_1 = j | X_0 = i)}_{p(i,j)} \underbrace{\mathbb{P}(X_2 = 2 | X_1 = j)}_{p(j,2)} \\ &= \sum_{i,j} \alpha_i p(i,j) p(j,2) = (\alpha P^2)_2. \end{aligned}$$

This is the sum of the probabilities of all possible ways that you can end up at state 2 at time 2.

To prove this in general, I used the following probability formula:


Lemma 1.7. *Suppose that our sample space is a disjoint union*

$$\Omega = \coprod B_i$$

of events B_i . Then

$$(1.1) \quad \mathbb{P}(A) = \sum_i \mathbb{P}(B_i) \mathbb{P}(A | B_i)$$

I drew this picture to illustrate this basic concept that you should already know.



Proof of theorem 1.6. By induction on n . If $n = 0$ then $P^n = I$ is the identity matrix. So,

$$\alpha P^n = \alpha P^0 = \alpha I = \alpha.$$

This is the distribution of $X_n = X_0$ by definition. So, the theorem holds for $n = 0$.

Suppose the theorem holds for n . Then, by Equation (1.1)

$$\begin{aligned} \mathbb{P}(X_{n+1} = 1) &= \\ &= \underbrace{\mathbb{P}(X_n = 1)}_{B_1} \underbrace{\mathbb{P}(X_{n+1} = 1 | X_n = 1)}_{p(1,1)} + \underbrace{\mathbb{P}(X_n = 2)}_{B_2} \underbrace{\mathbb{P}(X_{n+1} = 1 | X_n = 2)}_{p(2,1)} \\ &= (\alpha P^n)_1 p(1,1) + (\alpha P^n)_2 p(2,1) = [(\alpha P^n)P]_1 = (\alpha P^{n+1})_1. \end{aligned}$$

And similarly, $\mathbb{P}(X_{n+1} = 2) = (\alpha P^{n+1})_2$. So, the theorem holds for $n + 1$. So, it holds for all $n \geq 0$. \square

Corollary 1.8. *If the initial distribution $\alpha = \pi$ is invariant, then X_n has probability distribution π for all n .*

Proof. The distribution of X_n is

$$\alpha P^n = \pi P^n = \underbrace{\pi P}_{=\pi} P \cdots P = \pi$$

since every time you multiply P you always have π . \square

1.3.2. *Perron-Frobenius Theorem.* I stated this very important theorem without proof. However, the proof is outlined in Exercise 1.20 in the book.

Theorem 1.9 (Perron-Frobenius). *Suppose that A is a square matrix all of whose entries are positive real numbers. Then, A has a left eigenvector π , all of whose coordinates are positive real numbers. I.e.,*

$$\pi P = \lambda \pi.$$

Furthermore,

- (a) π is unique up to a scalar multiple. (If α is another left eigenvector of A with positive real entries then $\alpha = C\pi$ for some scalar C .)
- (b) $\lambda_1 = \lambda$ is a positive real number.
- (c) The eigenvector λ_1 is larger in absolute value than any other eigenvalue of A : $|\lambda_2|, |\lambda_3|, \dots < \lambda_1$. (So, λ_1 is called the maximal eigenvalue of A .)
- (d)

$$\lim_{n \rightarrow \infty} \frac{1}{\lambda_1^n} P^n = \begin{pmatrix} \pi \\ \pi \\ \vdots \\ \pi \end{pmatrix}$$

assuming that π is a probability distribution, i.e., $\sum \pi_i = 1$.

I didn't prove this. However, I tried to explain the last statement. When we raise P to the power n , it tends to look like multiplication by λ_1^n . So, we should divide by λ_1^n . If we know that the rows of the

matrix $\frac{1}{\lambda_1^n} P^n$ are all the same, what is it?

$$\pi \frac{1}{\lambda_1^n} P^n = (\pi_1, \pi_2, \dots) \begin{pmatrix} \alpha \\ \alpha \\ \vdots \\ \alpha \end{pmatrix} = \left(\sum \pi_i \right) \alpha$$

But $\pi \frac{1}{\lambda_1} P = \pi$. So,

$$\pi \frac{1}{\lambda_1^n} P^n = \pi = \left(\sum \pi_i \right) \alpha$$

and

$$\alpha = \frac{1}{\sum \pi_i} \pi.$$

This theorem applies to Markov chains but with some conditions. First, I stated without proof the fact:

Theorem 1.10. *The maximal eigenvalue of P is 1. More precisely, all eigenvalues of P have $|\lambda| \leq 1$.*

Proof. Suppose that P has an eigenvalue λ with absolute value greater than 1. Then, there is an eigenvector x so that $xP = \lambda x$. Then $xP^n = \lambda^n x$ diverges as n goes to infinity. But this is not possible since the entries of the matrix P^n are all between 0 and 1 since P^n is the probability transition matrix from X_0 to X_n by Theorem 1.6. \square

I'll explain this proof later. What I did explain in class is that 1 is always an eigenvalue of P . This follows from the fact that the rows of P add up to 1:

$$\sum_j p(i, j) = 1.$$

This implies that the column vector with all entries 1 is a right eigenvector of P with eigenvalue 1:

$$P \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix}$$

For example, if

$$P = \begin{pmatrix} 2/3 & 1/3 \\ 1/4 & 3/4 \end{pmatrix}$$

then

$$P \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 2/3 & 1/3 \\ 1/4 & 3/4 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$


The invariant distribution π is a left eigenvector with eigenvalue 1. The unique invariant distribution is:

$$\pi = \left(\frac{3}{7}, \frac{4}{7} \right)$$

This means:

$$\pi P = \left(\frac{3}{7}, \frac{4}{7} \right) \begin{pmatrix} 2/3 & 1/3 \\ 1/4 & 3/4 \end{pmatrix} = \left(\frac{3}{7}, \frac{4}{7} \right) = \pi.$$

You can find π using linear algebra or a computer. You can also use intuition. In the Markov chain:



we can use the *Law of large numbers* which says that, if there are a large number of people moving randomly, then the proportion who move will be approximately equal to the probability. So, if there are a large number of people in states 1 and 2 then one third of those at 1 will move to 2 and one fourth of those in 2 will move to 1. If you want the distribution to be stable, the numbers should have a 3:4 ratio. If there are 3 guys at point 1 and 1/3 of them move, then one guy moves to 2. If there are 4 guys at 2 and 1/4 of them move then one guy moves from 2 to 1 and the distribution is the same. To make it a probability distribution, the vector $(3, 4)$ needs to be divided by 7.

Theorem 1.11. *If the Markov chain is aperiodic and irreducible then it satisfies the conclusions of the Perron-Frobenius theorem.*

Proof. These conditions imply that $A = P^n$ has all positive entries for some finite n . Then, the Perron-Frobenius eigenvector for A is the invariant distribution for P . \square

The Perron-Frobenius theorem tells us that the distribution of X_n will reach an equilibrium (the invariant distribution) for large n (assuming aperiodic). Then next question is: How long does it take?

1.4. **Transient classes.** I asked the question:

How long does it take to escape from a transient class? I started with a really simple example:

$$\bullet_1 \xrightarrow{p} \bullet_2$$

This is a Markov chain with one transient class $\{1\}$ and one absorbing class $\{2\}$. The question is: How long can you stay in the transient class? I was glad to see that students know basic probability:

$$\mathbb{P}(X_n = 1 \mid X_0 = 1) = (1 - p)^n.$$

What happens when n goes to infinity?

$$\lim_{n \rightarrow \infty} (1 - p)^n = 0 \quad \text{if } p > 0.$$

Proof. And you guys helped me with this proof:

$$\begin{aligned} L &= \lim_{n \rightarrow \infty} (1 - p)^n \\ \ln L &= \lim_{n \rightarrow \infty} \underbrace{n}_{\rightarrow \infty} \underbrace{\ln(1 - p)}_{\substack{< 1 \\ < 0}} = -\infty. \end{aligned}$$

So, $L = 0$. □

So, the probability of remaining indefinitely in state 1 is zero. In other words, you will eventually escape the transient class with probability one (at least in this example). For future reference I recorded this conclusion as follows.

Theorem 1.12. *If the probability of success is $p > 0$ and if you try infinitely many times, then you will eventually succeed with probability one.*

But how long does it take? Let

$$T := \text{smallest } n \text{ so that } X_n = 2.$$

Then

$$\mathbb{P}(T = n \mid X_0 = 1) = p(1 - p)^{n-1}.$$

For example, if $n = 3$ then, $T = 3$ which means we have:

$$\begin{array}{ccccccc} X_0 & & X_1 & & X_2 & & X_3 \\ \bullet & \xrightarrow{1-p} & \bullet & \xrightarrow{1-p} & \bullet & \xrightarrow{p} & \bullet \\ 1 & & 1 & & 1 & & 2 \end{array}$$

$$\mathbb{P}(T = 3 \mid X_0 = 1) = p(1 - p)^2.$$

The numbers $p(1-p)^n$ add up to 1 and give what is called the *geometric distribution* on the nonnegative integers.

From this formula we can calculate the conditional expected value of T :

$$\mathbb{E}(T | X_0 = 1) = \sum_{n=0}^{\infty} np(1-p)^{n-1}$$

and, yes, the $n = 0$ term is zero. This is easy to calculate using a little calculus. First we start with the geometric series

$$g(x) := \sum_{n=0}^{\infty} x^n = 1 + x + x^2 + x^3 + \dots = \frac{1}{1-x}.$$


Then differentiate:

$$g'(x) = \sum_{n=0}^{\infty} nx^{n-1} = \frac{1}{(1-x)^2}.$$

Applying this formula to the expected value problem, $x = 1 - p$ and we get:

$$\mathbb{E}(T | X_0 = 1) = p \sum n(1-p)^{n-1} = p \left(\frac{1}{(1-(1-p))^2} \right) = p \frac{1}{p^2} = \boxed{\frac{1}{p}}.$$

This was actually intuitively obvious from the beginning. For example:



When $p = 1/5$ then you expect it to happen in 5 trials. So, $\mathbb{E}(T) = 5 = 1/p$.

1.4.1. *larger transient classes.* Last time I explained (Theorem 1.12) that, if

$$\mathbb{P}(\text{success in one trial}) = p > 0$$

then

$$\mathbb{P}(\text{success with } \infty \text{ many trials}) = 1.$$

But you can say more:

Corollary 1.13. *Furthermore, you will almost surely succeed an infinite number of times.*

Proof. Suppose that you succeed only finitely many times, say 5 times:

$$n_1, n_2, n_3, n_4, n_5.$$

If n_5 is the last time that you succeed, it means that, after that point in time, you try over and over infinitely many times and fail each time. This has probability zero by the theorem. So,

$$\mathbb{P}(\text{only finitely many successes}) = 0.$$

But, the number of successes is either finite or infinite. So,

$$\mathbb{P}(\text{infinitely many successes}) = 1.$$

□

Apply this to Markov chains:

$$X_0, X_1, X_2, \dots$$

These locations are random states in the finite set S of all states. This means that there is at least one state that is visited infinitely many times. Let

$$I := \{i \in S \mid X_n = i \text{ for infinitely many } n\}$$

This is the set of those states that the random path goes to infinitely many times.

Theorem 1.14. *(A.s.) I is one recurrent class.*


At this point we had a discussion about the meaning of this. The set I is a random set. Since a general finite Markov chain has several recurrent classes, which one you end up in is a matter of chance. The probability distribution of X_n for large n will include a linear combination or “superposition” of several possible futures. So, several recurrent classes have positive probability at the beginning. However, when you actually go into the future, you pick one path and you get stuck in one recurrent class from which you cannot escape. This theorem says

that you will wander around and visit every site in that recurrent class infinitely many times.

Proof. In order to prove this theorem I first proved:

Lemma (a) If $i \in I$ and $i \rightarrow j$ then $j \in I$.

This means: if it is possible to go from i to j then $j \in I$.



Proof of Lemma (a): It is given that $i \in I$. I.e., we go to i infinitely many times. Each time we go to i we have a probability $p > 0$ of going to j . Theorem 1.12 says that, with probability one, we will eventually go to j . But then (b) we have to eventually go back to i because, we are going to i infinitely many times. So, by Corollary 1.13, with probability one, you cross that bridge infinitely many times. So, $j \in I$. (The picture is a little deceptive. The path from i to j can have more than one step.)


This proof also says: **(b)** $j \rightarrow i$ since you need to return to i infinitely many times. Therefore, I is one communication class. We just need to show that this class is recurrent.

But (a) implies that I is recurrent. Otherwise, there would be a j not in I so that $i \rightarrow j$ for some $i \in I$ and this would contradict (a). \square

Corollary 1.15. *The probability is zero that you remain in a transient class indefinitely.*

1.5. **Canonical form of P .** Next, I talked about the canonical form of P which is given on page 20 of our book.

1.5.1. *definition.* Suppose that R_1, R_2, \dots, R_r are the recurrent classes of a Markov chain and T_1, T_2, \dots, T_s are the transient classes. I drew a picture similar to the following to illustrate this.



Then the *canonical form* of the transition matrix P is given by the following “block” form of the matrix: (In the book, all transient classes are combined. So, I will do the same here.)

$$P = \begin{array}{c} \\ R_1 \\ R_2 \\ T \end{array} \begin{array}{ccc} R_1 & R_2 & T \\ \left(\begin{array}{ccc} P_1 & 0 & 0 \\ 0 & P_2 & 0 \\ S_1 & S_2 & Q \end{array} \right) \end{array}$$

If you start in the recurrent class R_1 then you can't go anywhere else. So, there is only P_1 in the first row. In the example, it is a 2×2 matrix. Similarly, the second row has only P_2 since, if you start in R_2 you can't get out.

The matrices P_1 and P_2 are *stochastic matrices*. Their rows add up to one since, in the entire matrix P , there are no other numbers in those rows. This also reflects the fact that the recurrent classes R_1 and R_2 are, in themselves, (irreducible) Markov chains.

The transient class T is not a Markov chain. Why not? There are several reasons. If you look at the picture, you see that you can leave the transient class out of the bottom. So, it is not a “closed system.” Another reason is that the matrix Q is not stochastic. Its rows do not add up to one. So, Q does not define a Markov chain.

The bottom row in the canonical form describes what happens if you start in any transient class. You either go to another transient state or you go to a recurrent state. The matrix Q is the *transient-to-transient* matrix. The matrix

$$S = (S_1, S_2)$$

is the *transient-to-recurrent* matrix. It has one block S_i for every recurrent state R_i .

Since each recurrent state R_i is an irreducible Markov chain, it has a unique invariant distribution π_i .

Theorem 1.16. *If π_i is the invariant distribution for P_i then the invariant distributions for P are the positive linear combinations of the π_i (with coefficients adding to 1). In other words,*

$$\pi = \sum t_i \pi_i$$

where $t_i \geq 0$ and $\sum t_i = 1$. In the case of two recurrent states, this is:

$$\pi = t\pi_1 + (1-t)\pi_2$$

where $0 \leq t \leq 1$.

Proof. Suppose that π_1, π_2 are invariant distributions for P_1, P_2 . Then they are row vectors of the same size as P_1, P_2 , respectively, and

$$\pi_1 P_1 = \pi_1, \quad \pi_2 P_2 = \pi_2.$$

When $t = 1/3$ we get the invariant distribution:

$$\pi = \left(\frac{1}{3}\pi_1, \frac{2}{3}\pi_2, 0\right).$$

You need to multiply by $1/3$ and $2/3$ (or some other numbers ≥ 0 which add to 1) so that the entries of π add up to 1. Block matrix multiplication show that this is an invariant distribution:

$$\begin{aligned} \pi P &= \left(\frac{1}{3}\pi_1, \frac{2}{3}\pi_2, 0\right) \left(\begin{array}{cc|c} P_1 & 0 & 0 \\ 0 & P_2 & 0 \\ \hline S_1 & S_2 & Q \end{array} \right) \\ &= \left(\frac{1}{3}\pi_1 P_1, \frac{2}{3}\pi_2 P_2, 0\right) = \left(\frac{1}{3}\pi_1, \frac{2}{3}\pi_2, 0\right) = \pi \end{aligned}$$

This shows that the positive linear combinations of the invariant distributions π_i are invariant distributions for P .

The converse, which I did not prove in class is easy: Suppose that $\pi = (\alpha, \beta, \gamma)$ is an invariant distribution. Then we must have $\gamma = 0$, since otherwise

$$(\alpha, \beta, \gamma)P^n = (\alpha, \beta, \gamma)$$

indicating that we have a positive probability of remaining in a transient state indefinitely, a contradiction to what we just proved. So, $\pi = (\alpha, \beta, 0)$ and

$$(\alpha, \beta, 0) \left(\begin{array}{cc|c} P_1 & 0 & 0 \\ 0 & P_2 & 0 \\ \hline S_1 & S_2 & Q \end{array} \right) = (\alpha P_1, \beta P_2, 0) = (\alpha, \beta, 0)$$

which means that $\alpha P_1 = \alpha$ and $\beta P_2 = \beta$. So, α, β are scalar multiples of invariant distributions for P_1, P_2 . \square

The next two pages are what I handed out in class, although the page numbers have shifted.

1.5.2. *example.* The problem is to find all invariant distributions of the following transition matrix.

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

An invariant distribution π is the solution of:

$$\pi P = \pi.$$

This equation can be rewritten as:

$$\pi(P - I) = 0$$

where $I = I_4$ is the identity matrix. In other words π is a left null vector of

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

Corollary 1.17. *The dimension of the null space of $P - I$ is equal to the number of recurrent classes. A basis is given by the invariant distributions of each recurrent class.*

Note that the numbers in each row of $P - I$ adds up to zero. This is the same as saying that the column vectors of $P - I$ add up to the zero vector.

In order to find the left null space of $P - I$ we have to do *column operations* on $P - I$ to reduce it to *column echelon form*! This is not such a terrible thing. For example, you can always eliminate the last column using column operations, namely, add the first three columns to the last column. It becomes all zero! So we have:

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

Now, multiply the fourth column by 4 then, using column operations, clear the 2nd row:

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

This is not quite in column echelon form. But it is good enough to answer all the questions because:

Every row has at most one nonzero entry.

- (1) The rank of $P - I$ is 2, the number of nonzero columns.
- (2) The dimension of the null space of $P - I$ is 2 since

$$\dim \text{Null space} = \text{size} - \text{rank} = 4 - 2 = 2.$$

Therefore, there are 2 recurrent classes.

- (3) A basis for the null space is given by
 - (a) $(0, 0, 1, 0)$
 - (b) $(1, 0, 0, 2)$
- (4) If we normalize these two vectors (divide by the sum of the coordinates), we get the *basic invariant distributions*:
 - (a) $\beta = (0, 0, 1, 0)$
 - (b) $\gamma = (1/3, 0, 0, 2/3)$
- (5) These are the unique invariant distributions for the two recurrent classes. So, their *supports* $\{3\}$ and $\{1, 4\}$ are the recurrent classes.
- (6) Now we can find all invariant distributions. They are given by

$$\pi = t\gamma + (1 - t)\beta = \left(\frac{t}{3}, 0, 1 - t, \frac{2t}{3} \right)$$

for $0 \leq t \leq 1$.

- (7) This represents the long term distribution where t is the probability of ending up in the recurrent class $\{1, 4\}$ and $1 - t$ is the probability of ending up in the other recurrent class $\{3\}$. For example, if the initial distribution is

$$\alpha = (1/4, 1/4, 1/4, 1/4)$$

then


$$t = \frac{1}{4} + \frac{2}{3} \cdot \frac{1}{4} + 0 + \frac{1}{4} = \frac{2}{3}.$$

So, in the long run (as $n \rightarrow \infty$) we get:

$$\lim_{n \rightarrow \infty} \alpha P^n = \left(\frac{t}{3}, 0, 1 - t, \frac{2t}{3} \right) = \left(\frac{2}{9}, 0, \frac{1}{3}, \frac{4}{9} \right)$$

1.6. **The substochastic matrix Q .** On Monday, I used the Mouse-Cat-Cheese problem to explain the use of the matrix Q . On Wednesday, I explained this further and applied it to the Leontief model.

1.6.1. *mouse-cat-cheese.* Here is the diagram for the problem (from page 26 with 4,5 switched).



The recurrent classes are the absorbing states: $R_1 = \{4\}, R_2 = \{5\}$. In canonical form, we put these first:

$$\begin{array}{c}
 \begin{array}{cc}
 & \begin{array}{cc} 4 & 5 \end{array} \\
 \begin{array}{c} 4 \\ 5 \\ 1 \\ 2 \\ 3 \end{array} & \left(\begin{array}{cc|ccc}
 1 & 0 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 & 0 \\
 \hline
 1/2 & 0 & 0 & 1/2 & 0 \\
 0 & 0 & 1/2 & 0 & 1/2 \\
 1/3 & 1/3 & 0 & 1/3 & 0
 \end{array} \right)
 \end{array}
 \end{array}$$

$\begin{array}{cc} S_1 & S_2 \\ Q \end{array}$

The *transient-to-transient* matrix is called Q . This will be a square matrix whose rows add up to ≤ 1 .

Now we want to calculate the probability that the mouse will be eaten by the cat if he starts in rooms 1,2,3. We get three numbers t_1, t_2, t_3 :

$$t_i := \mathbb{P}(X_\infty = 4 \mid X_0 = i)$$

which form a column vector t . The theorem is:

Theorem 1.18. *The probability of ending in the first recurrent state R_1 is the vector*

$$\vec{t} = \begin{pmatrix} t_1 \\ t_2 \\ t_3 \end{pmatrix} = (I - Q)^{-1}S_1.$$

Proof. We need to consider how long the mouse is going to be moving around in the transient class $T_1 = \{1, 2, 3\}$. So, let T be the time it takes for the mouse to reach a recurrent state:

$$T := \text{smallest } n \text{ so that } X_n = 4 \text{ or } 5.$$

Then:

$$t_i = \mathbb{P}(X_\infty = 4 \mid X_0 = i) = \sum_{n=0}^{\infty} \mathbb{P}(T = n \text{ and } X_n = 4 \mid X_0 = i).$$

But, the probability that $T = n$ and $X_n = 4$ given $X_0 = i$ is the i th coordinate of the vector

$$Q^{n-1}S_1.$$

The reason is that, in order to get to a recurrent class at time $T = n$, the mouse needs to move around in the transient states for exactly $n - 1$ turns. This is given by the matrix Q^{n-1} . Then S_1 gives the probability of moving to the recurrent class 4. The product $Q^{n-1}S_1$ gives the probability of doing one then the other. This is a column vector where the row number indicates the starting point.

For example,

$$QS_1 = \begin{pmatrix} 1 & 1/2 & 0 \\ 1/2 & 0 & 1/2 \\ 0 & 1/3 & 0 \end{pmatrix} \begin{pmatrix} 1/2 \\ 0 \\ 1/3 \end{pmatrix} = \begin{pmatrix} 1/2 \\ 5/12 \\ 0 \end{pmatrix}.$$

The number $5/12$ represents the probability of getting from room 2 to room 4 in exactly two steps:

$$\begin{aligned} \mathbb{P}(T = 2 \text{ and } X_2 = 4 \mid X_0 = 2) &= \\ p(2, 1)p(1, 4) + p(2, 3)p(3, 4) &= \frac{1}{2} \frac{1}{2} + \frac{1}{2} \frac{1}{3} = \frac{5}{12} \end{aligned}$$

Therefore, the vector \vec{t} is the sum of these vectors for all n :

$$\vec{t} = S_1 + QS_1 + Q^2S_1 + \cdots = (I + Q + Q^2 + \cdots)S_1.$$

This is equal to

$$(I - Q)^{-1}S_1$$

by the following lemma. □

Lemma 1.19. *The matrix $I - Q$ is invertible and its inverse is given by*

$$(I - Q)^{-1} = I + Q + Q^2 + Q^3 + \cdots$$

Also, this series converges.

Proof. I assumed that the series converges. From this it follows that the limit is the inverse of $I - Q$. The proof is deceptively simple:

$$\begin{aligned} (I - Q)(I + Q + Q^2 + Q^3 + \cdots) &= \\ = I + Q + Q^2 + Q^3 + \cdots &= \\ -Q - Q^2 - Q^3 - \cdots + I &= I. \end{aligned}$$

□

Using the formula in the theorem we get:

$$\vec{t} = (I - Q)^{-1}S_1 = \frac{1}{7} \begin{pmatrix} 10 & 6 & 3 \\ 6 & 12 & 6 \\ 2 & 4 & 9 \end{pmatrix} \begin{pmatrix} 1/2 \\ 0 \\ 1/3 \end{pmatrix} = \begin{pmatrix} 6/7 \\ 5/7 \\ 4/7 \end{pmatrix}.$$

So, e.g., the probability that the mouse will be eaten by the cat if he starts in room 2 is

$$t_2 = 5/7.$$

Next, I talked very quickly about the time that it takes the mouse to reach the cat. I explained this better on Wednesday.

1.6.2. *expected time.* Today I did a more thorough explanation of the expected time until we reach a recurrent state. I started by reviewing the basics of substochastic matrices.

Definition 1.20. A *substochastic matrix* is a square matrix Q with nonnegative entries so that every row adds up to at most 1.

For example,

$$Q = \begin{pmatrix} 1/2 & 1/3 \\ 1/4 & 3/4 \end{pmatrix}$$

is substochastic.

Given any subset C of the set of states S , the C -to- C transition matrix will always be substochastic.

Lemma 1.21. *If C contains no recurrent class then*

- a) $I + Q + Q^2 + Q^3 + \dots$ converges to $(I - Q)^{-1}$.
- b) $I + 2Q + 3Q^2 + 4Q^3 + \dots$ converges to $(I - Q)^{-2}$.

Proof. (a) follows from the computation

$$(I - Q)(I + Q + Q^2 + Q^3 + \dots) = I.$$

For (b), the argument is:


$$\begin{aligned} (I - Q)(I + 2Q + 3Q^2 + 4Q^3 + \dots) &= \\ &= I + 2Q + 3Q^2 + 4Q^3 + \dots \\ &\quad - Q - 2Q^2 - 3Q^3 - \dots \\ &= I + Q + Q^2 + Q^3 \dots = (I - Q)^{-1} \end{aligned}$$

by (a). Therefore,

$$I + 2Q + 3Q^2 + 4Q^3 + \dots = \frac{(I - Q)^{-1}}{I - Q} = (I - Q)^{-2}$$

□

I gave another example to illustrate both of these formulas.



In this example, there is only one recurrent class $R_1 = \{2, 3, 4\}$. The set $C = \{1, 2\}$ contains the recurrent state 2 but it does not contain a recurrent class (R_1 is not contained in C). Therefore, the lemma applies.

Inserting the implied loop at vertex 2, we saw that the substochastic matrix is

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

Then

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

So,

$$(I - Q)^{-1} = \begin{pmatrix} 1 & 3/2 \\ 0 & 2 \end{pmatrix}, \quad (I - Q)^{-2} = \begin{pmatrix} 1 & 9/2 \\ 0 & 4 \end{pmatrix}.$$

These numbers are used to answer questions such as the following.

Question 1: If you start at 1, what is the probability that you reach 4 before you reach 3?

The first step in answering this question is to realize that: It does not matter what happens after you reach 3 or 4 because, at that point, the question has been answered. So, the numbers

$$p(3, 2) = 1/3, p(4, 3) = 1/2$$

are irrelevant and we can replace them with 0. In other words, we can make 3 and 4 into absorbing states. This simplification process gives a new probability transition matrix which we put into canonical form:

$$\tilde{P} = \begin{array}{c} \begin{array}{cc} 3 & 4 & 1 & 2 \\ \begin{array}{c} 3 \\ 4 \\ 1 \\ 2 \end{array} & \left(\begin{array}{cc|cc} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \hline 1/4 & 0 & 0 & 3/4 \\ 1/4 & 1/4 & 0 & 1/2 \end{array} \right) \end{array} \\ \begin{array}{ccc} S_1 & S_2 & Q \end{array} \end{array}$$

S_1 is the C -to-3 transition matrix and S_2 is the C -to-4 transition matrix. We also need their total:

$$S_T = S_1 + S_2 = \begin{pmatrix} 1/4 \\ 1/2 \end{pmatrix}.$$

Answer: The answer to the question is given by the first coordinate of the vector

$$(I - Q)^{-1} S_1 = \begin{pmatrix} 1 & 3/2 \\ 0 & 2 \end{pmatrix} \begin{pmatrix} 1/4 \\ 1/4 \end{pmatrix} = \begin{pmatrix} \frac{1}{4} + \frac{3}{8} \\ \frac{2}{4} \end{pmatrix} = \begin{pmatrix} 5/8 \\ 1/2 \end{pmatrix}$$

which is $5/8$. (If we started at state 2, the answer would be the second coordinate which is $1/2$.) This was proved on Monday (Theorem 1.18).

The matrix $(I - Q)^{-2}$ is used to answer a different question.

Question 2: If we start at $X_0 = 1$, how long will it take to reach 3 or 4? In other words, we want to calculate the conditional expected value

$$\mathbb{E}(T | X_0 = 1) = ?$$

of

$$T := \text{smallest } n \text{ so that } X_n = 3 \text{ or } 4.$$

Answer: Just take the definition of expected value:

$$\mathbb{E}(T | X_0 = 1) := \sum_{n=0}^{\infty} n \mathbb{P}(T = n | X_0 = 1).$$

But

$$\mathbb{P}(T = n | X_0 = 1) = (Q^{n-1} S_T)_1$$

is the first coordinate of

$$Q^{n-1} S_T$$

because $T = n$ means we stay in the set C for $n - 1$ turns and then go to one of the recurrent states 3, 4 (according to the modified transition matrix \tilde{P}). Q^{n-1} gives the probability of stay in C for $n - 1$ turns and S_T gives the probability of moving on the n th turn from C to 3 or 4.

$$\begin{aligned} \mathbb{E}(T | X_0 = 1) &:= \sum_{n=0}^{\infty} n \mathbb{P}(T = n | X_0 = 1) \\ &= \underbrace{\sum_{n=0}^{\infty} n (Q^{n-1} S_T)_1}_{(I-Q)^{-2}} = ((I - Q)^{-2} S_T)_1 \\ &= \begin{pmatrix} 1 & 9/2 \\ 0 & 4 \end{pmatrix} \begin{pmatrix} 1/4 \\ 1/2 \end{pmatrix} = \begin{pmatrix} \boxed{5/2} \\ 2 \end{pmatrix} \end{aligned}$$

So, the answer is

$$\mathbb{E}(T | X_0 = 1) = 5/2.$$

1.6.3. *Leontief model.* The final example is the Leontief economic model. In this model, the numbers in a substochastic matrix are interpreted as being the input requirements of several industries or factories. I started with the substochastic matrix:

$$Q = \begin{pmatrix} 0.3 & 0.5 \\ 0.4 & 0.2 \end{pmatrix} = \begin{pmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{pmatrix}$$

Q being *substochastic* means

- (1) $q_{ij} \geq 0$,
- (2) $q_{11} + q_{12} \leq 1$ and
- (3) $q_{21} + q_{22} \leq 1$.

The states are factories A and B plus one more C (the bank). q_{ij} is the amount of output of factory j which factory i needs for each dollar's worth of output. For example, the numbers in the first row mean that for each dollar of product A , factory A needs

30¢ worth of product A

50¢ worth of product B

which leaves:

$$20¢ \longrightarrow \text{profit!} \longrightarrow C(\text{bank})$$

Similarly, Factory B makes 40¢ worth of profit for each dollar of output. Putting in the third state C which is recurrent. (Assume the factories use the same bank. Or think of C as representing all banks put together.) Then we get the following probability transition function:

$$P = \begin{array}{c} C \\ A \\ B \end{array} \left(\begin{array}{c|cc} C & A & B \\ \hline 1 & 0 & 0 \\ .2 & .3 & .5 \\ .4 & .4 & .2 \end{array} \right)$$

The rows add to 1. So, this gives a Markov chain. The question is: What are we measuring the probability of ?

This matrix is keeping track of the *money* as it is being passed back and forth between the factories and the bank. Since the bank is recurrent, all of the money eventually ends up in the bank. To make it random, we think of the money as a pile of one dollar bills. Then

$$X_n = \text{location of one random dollar at time } n.$$

For example, suppose you get 100\$.

You put 50\$ in the bank C

You buy 30\$ of product A

You buy 20\$ of product B .

Then

$$X_0 = (.5, .3, .2)$$

because, if you “mark” one of the dollars, the probability that the marked dollar will go to the bank is 0.5, the probability of that marked dollar going to Factory A is 0.3 and for Factory B it is 0.2.

Suppose that each factory keeps a stockpile of supplies. After filling your order, each factory will have used up a certain amount of its inventory. It will order supplies to replenish its stockpile and the money will move:

$$X_1 = X_0P = (.64, .17, .19)$$

represents the location of the money after one day. For example, Factory A always puts $1/5$ of its income into the bank. So, it puts $30/5 = 6\$$ in the bank. Factory B puts 40% of its income into the bank. So, it puts $20 * 4/10 = 8\$$ into the bank. So

$$50 + 6 + 8 = 64\$$$

will be in the bank after one day. This is the first coordinate of X_1 times 100\$.

$$X_n = X_0P^n$$

gives the distribution of the money you put into the system after n days. Notice that, in this model, the total amount of money never changes! But the amount of goods produced can be very large.

Using what we know about Markov chains we can answer questions about the output of the factories.

Question: How much does Factory A need to produce in total?

The answer is the A coordinate of the vector

$$100\$ \cdot (X_0 + X_0P + X_0P^2 + \dots)$$

We can ignore the first (C) coordinate since the money in the bank just sits there and doesn't do anything. So, the answer is equal to

$$(30, 20)(I + Q + Q^2 + \dots)_1$$

where the $()_1$ means 1st coordinate. The vector $(30, 20)$ represents $100X_0$. Using the formula that the series $I + Q + Q^2 \dots$ converges to $(I - Q)^{-1}$ we get

$$(30, 20)(I - Q)_1^{-1} = (\$88.888\dots, \$80.555\dots)_1 = \$88\frac{8}{9}$$

Students figured out that what I wrote on the board did not make sense because the numbers must be greater than 30 and 20. Here, I used Excel to calculate the matrix inverse and product more accurately. (I used “= *index(minverse(...), i, j)*”).

1.7. Transient and recurrent. In a finite Markov chain, every state is either transient or recurrent depending on whether it is in a transient or recurrent communication class. Here is a summary of the differences.

1.7.1. *transient classes.*

- (1) If i is a transient class then you will a.s. visit i only finitely many times.
- (2) If i, j are transient states then the expected number of visits to j , starting at i

$$\mathbb{E}(\text{visits to } j \mid X_0 = i) = \sum_{n=0}^{\infty} \mathbb{P}(X_n = j \mid X_0 = i)$$

is equal to the (i, j) coordinate of the matrix

$$I + Q + Q^2 + Q^3 + \cdots = (I - Q)^{-1}.$$

- (3) Starting at a transient state i , the expected value of $T =$ the number of steps to reach a recurrent state is the i th coordinate of

$$(I - Q)^{-2} S_T$$

- (4) The expected length of time to return to i given that $X_0 = i$ is $\mathbb{E}(\text{smallest } n > 0 \text{ so that } X_n = i \mid X_0 = i) = \infty.$

The reason is that there is a nonzero chance p that you could be waiting forever. Then $\mathbb{E} = (\cdots) + p \cdot \infty = \infty$

1.7.2. *recurrent classes.*

- (1) If j is recurrent then you will a.s. return to j an infinite number of times if $X_0 = j$.
- (2) Every recurrent class R_i has an invariant distribution π_i .
- (3) The long term probability of being in state $j \in R_i$ given that $X_0 \in R_i$ is equal to the j -coordinate $\pi_i(j)$ of π_i .
- (4) The expected number of visits to j is infinite:

$$\mathbb{E}(\text{number of visits to } j \mid X_0 = j) = \infty.$$

- (5) The expected length of time between visits to j is

$$\mathbb{E}(\text{smallest } n > 0 \text{ so that } X_n = j \mid X_0 = j) = \frac{1}{\pi_i(j)}.$$

(For example if $\pi_i(j) = 1/3$ then you spend $1/3$ of the time at j and the average time between visits is 3. Since this is obvious I won't prove it.)

HOMEWORK 1A: FINITE MARKOV CHAINS

Three problems (corrected) due next Monday, Feb 4:

0.1. Do 1.5 in the book.

0.2.

A man is playing two slot machines (call them A and B). Machine A gives a payoff with a probability of $1/6$ and machine B gives a payoff with probability $1/16$. The man starts by picking a machine at random. Then he plays the machine until he has lost twice (not in a row, just in total). Then he switches machines and continues. For example, suppose that his winning (1) and losing (0) sequence might be:

$${}_01{}_10{}_20{}_31{}_40{}_51{}_60{}_71{}_80{}_91{}_100{}_110{}_12$$

Then he will switch machines after $n = 2$ since he lost twice. (He switches in the time between $n = 2$ and $n = 3$). He switches back after $n = 6$ and then again after $n = 10$.

- Make this into a Markov chain with 4 states: A_0, A_1, B_0, B_1 where the subscript keeps track of the number of losses. [This is an example of recording information to convert a stochastic process to first order.]
- What is the probability that the man will be playing machine A at $n = 4$ if he starts at machine A ? What about if he starts at a machine picked at random?
- Find the invariant distribution.
- In the long run, how much of the time will the man be playing the better machine?

0.3. Suppose that

$$P = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ .4 & .4 & 0 & .2 \\ .7 & 0 & 0 & .3 \end{pmatrix}$$

- Find the unique invariant distribution and explain why it is unique.
- Draw the diagram and find the communication classes.
- What is the probability that X_{100} is in the transient class given that you start in the transient class? What about if you start at a random location?

HOMEWORK 1B: LEONTIEF MODEL

(This is the same as the homework handed out in class.)

This homework project is due next Thursday, Feb 14:

Your assignment is to analyze the Leontief economic model and work out one example. Don't do the calculations by hand.

In the Leontief model, there are "factories" which require the output of other factories to make their product. For each dollar of output, Factory i requires q_{ij} dollars worth of output of factory j . The total amount that factory i needs to spend for each dollar of output is

$$q_{i1} + q_{i2} + \cdots + q_{ir} \leq 1.$$

We always assume the sum is ≤ 1 . (But it is allowed to be equal to 1.)

Work out the following example and answer the questions (in complete sentences so that your kid brother can read it!)

Example: We have 4 factories:

S = Steel

W = Water

E = Coal/Gas/Oil

P = Plastic

To produce \$1 worth of steel, the steel factory needs 50¢ worth of energy and 25¢ worth of water (and no plastic). This goes into the matrix Q in the first line.

$$Q = \begin{pmatrix} 0 & .25 & .5 & 0 \\ .1 & 0 & .4 & .2 \\ .2 & .1 & .3 & .1 \\ .1 & .2 & .2 & .1 \end{pmatrix}$$

Each factory keeps a stockpile of material, say 10\$ worth of each item. When it get an order for goods, the factory uses its inventory and then orders replacements. So, the steel factory, after filling out an order for 1\$ worth of steel will order 25¢ worth of water and 50¢ worth of energy.

(*) Write in words: What do the numbers in the fourth row of the matrix mean?

The consumer wants 1\$ of steel, 2\$ of water, 10\$ of energy and 2\$ of plastic.

- (a) How much does each factory need to make?
- (b) Follow the money: Where do the 15\$ go after 4 rounds?
- (c) How long does it take for all factories to regain 99% of their original inventory assuming that they keep 10\$ worth of each commodity in stock.
- (d) Follow the energy. Take the total amount of energy (your answer to part (a)) that is needed. Where does it go?

Theory: Assume we have a more general matrix Q representing the requirements of each industry in the Leontief model.

- (e) The rows of Q may not add up to less than 1. If row i adds up to 1, what does it mean about state i ? Give an example.
- (f) Suppose that every row of Q adds up to at most $p = 0.9$. Then prove that each row of Q^n adds up to

$$p^n = (0.9)^n$$

or less.

Why does this imply that the sequence

$$I + Q + Q^2 + Q^3 + \dots$$

converges? [Hint: a series (infinite sum) of matrices converges if and only if, for each i and j the sum of the (i, j) entries converges. Use the comparison test, comparing these entries to a geometric series to show that the series converges.]