Economics 202A Lecture II Fall 2007

1. Check to see if books are available.

2. Check to see if reader is available.

3. Extra Reading lists, Data Assignments.

4. Bring Chalk.

5. Section this evening 6:10 - 8, 145 McCone Hall.

6. Next Tuesday's extra class will be in *141* McCone Hall, from 6:10 - 7:30.

There are five groups of people who are allowed to enroll.

1. Those who are already enrolled.

Then new enrollments will be restricted to four categories:

2. Economics graduate students for whom this is a required course.

3. ARE students for whom this is a required course.

4. Undergraduate majors in economics who are ready to take this course and who have my explicit permission.

5. One or two people who were earlier granted a promise that they could enroll in the class.

There will be no auditors.

I really do apologize to those of you who would like to take the class but are not allowed to enroll. I feel really very bad about this. I simply do not have a choice here. Last Spring, the University initially assigned us a room that was utterly terrible and also with fewer seats than this one. Phil Walz had to put up a real fight to get this room.

Let's turn to today's lecture.

I do have a suggestion for you. I want you to try to tell the story of the readings or of the lecture in your own words.

In seminars while the mathematics is being presented, that's what I try to do. I try

to tell a story about what lies behind it.

You should try to do that for yourself: to have your own story of what is behind the mathematics.

That allows you to do the manipulations on two levels:
Formally, in terms of the mathematics;
But also, informally, in terms of your story.

This is a way of being an active, rather than a passive, participant in the material.

What are we going to do today?

Today I will review Section I of the Reading List—the mathematical review.

Max will give a further review this evening in 145 McCone Hall.

And then next Tuesday's class will be from 6:10 to 7:30 in *141* McCone Hall.

Most of this lecture comes from what I learned in a Time Series course that I took in the statistics department taught by David Donaho, who has unfortunately left for Stanford.

I will probably mention Dave again in the class on Sargent next Tuesday evening.

The lecture will proceed in two parts.

1. I will review the solution to Difference Equations.

> **I will review the solution to the first and second order linear difference equations.**

2. I will give a short review of time series.

> **I will discuss, at least briefly, ARMA models.**
> **These will be mathematical tools that will recur throughout the course.**
> **It is very important that you understand them from the beginning.**
>
> > **So, you should definitely make it a first priority to read the sections of Harvey that are at the beginning of the Reader.**

Let me begin.

You should have a background that gives you the solution to *differential*

equations.

If that is the case then you should appreciate that it takes only minor adaptation to get the solution to difference equations.

(1) Let me remind you of the solution to the first order difference equation.

The solution to that is:

I think that you have already learned that before.
If you have missed it, you should review it, to see why this is the solution.

This solution makes sense.

You can check that the first terms of the series make sense. It follows the rule given by the difference equation for forming y_t given $y_{t-1.}$ So

$$y_0 = y_0$$

$$y_1 = \alpha y_0 + f_1$$

$$y_2 = \alpha(\alpha y_0 + f_1) + f_2 = \alpha^2 y_0 + \alpha f_1 + f_2$$

This corresponds to what we have in the formula for the solution.

For the moment I will put a BOX around this equation.

Put a BOX around both equations.

We put a box around it so that when we want this solution we can call on it—just the way a computer program will call upon a subroutine. We know it and we will not have to think about it.

You should play with this solution just long enough so that you see that it makes sense to you.

(2) similarly let's recall the solution to the second order linear difference equation.

That solution is

where x_1 and x_2 are the roots of the associated polynomial

That is the same a_1 and the same a_2 as in the original equation. <point>

That solution makes sense.
Why does it make sense?
If you plug in the solution here [POINT to LHS of equation] you will see that the LHS of the equation will equal the RHS of the equation.
The key to that is x_1 and x_2 are the roots of the associated polynomial.
That means, for example, that

So when you plug into the original equation *[point],* the terms with x_1 and the x_2 nicely vanish.

Let me continue describing the solution,

There is a slight problem if these roots are not *real.*
That occurs, if

$$a_1^2 - 4a_2 < 0.$$

In that case x_1 and x_2 are not real. Then a long and complicated argument shows that the *complex* elements cancel out and:

where θ = arc cos $((-a_1/2)/\sqrt{a_2}$.

We will use these *solutions*—so I want you to be familiar with them so we can call on them when needed.

It is probably not necessary for the purpose of this course to understand the deep reason why the complex elements cancel out, although it is almost always dangerous to use math that you do not completely understand.

****

Let's now turn to the mathematical innovation which forever changed macroeconomics.

In later classes—indeed, in the very next class—you will see why this is basic to macroeconomic systems.

An	**A**	**R**	**M**	**A**	**process is**
an	**Auto**	**Regressive**	**Moving**	**Average**	**process.**

Let me first talk about `	**A**	**R**	**processes:**
	Auto	**Regressive**	**processes.**

And then I will talk about

M	**A**	**processes:**
Moving	**Average**	**processes.**

And then I will talk about combining them.

The simplest example of an *Autoregressive* process, or an AR process, in the jargon, is:

where ϵ_t is i.i.d. $N(0, \sigma_\epsilon^2)$.

There is a *natural extension* of our *simplest autoregressive process.*

We could write:

where ϵ_t is independent, identically distributed, normal, or

i.i.d. $N(0, \sigma_\epsilon^2)$.

This is an autoregressive process of order *p.*
In abbreviation, it is an AR(p) process.

Alternatively, y_t might have been a

Moving Average Process.

The simplest MA process is:

This has a generalization:

In this case y_t is an MA process of order q.

In general we could combine *autoregression* with *moving averages* to have

Such a process is called an ARMA process of order (p,q) or an ARMA(p,q) process.

It is an *A*uto*R*egressive *M*oving *A*verage process of order p in the *autoregressive* part, and of order q in the *moving average* part.

I now want to show you that different ARMA processes have different *signatures.*

These *signatures* are the *covariances* between the dependent variable y_t and its realization *s* periods ago.

For notational convenience let's assume that all the y_t's have mean 0.

We can then consider the covariance of y_t with its value *s* periods ago.

$$E(y_t \cdot y_{t-s}).$$

This object has a *name.* It is called: *the autocovariance of order s,* and it has a conventional notation.

It turns out that different types of ARMA processes will have different patterns of *autocovariances.*

The autocovariance functions, which I will define later, are the *signatures of the processes.*
By that I mean that if you see an autocovariance function, you can tell what process caused it. You can *retrieve* the process.

Let me illustrate. I am going to show you that if I have the autocovariances for each value of *s* that I can retrieve the process.

I will show this in an example. That example is the simplest MA process.
Suppose y_t follows this process:

The ϵ_t's are assumed to be independent, identically distributed random variables with mean 0,

so that $E(\epsilon_t) = 0$

$$E(\epsilon_t^2) = \sigma_\epsilon^2 \quad \text{and}$$

$$E(\epsilon_t \cdot \epsilon_{t-s}) = 0, \qquad s \neq 0.$$

We shall use these properties to calculate the *autocovariance* of y_t for each lag value of *s.*

The important property that we will use, and that we will keep using over and over again, is that

$$E(\epsilon_t \cdot \epsilon_{t-s}) = 0, \qquad s \neq 0.$$

So I want to pause and I want you to note that. These cross terms are zero because ϵ_t is an innovation in information at t.

If ϵ_t is an innovation in information at t, it cannot be correlated with innovations in information in past periods, t - s.

According to the standard notation, I will denote the variance of y_t as γ_0. This is the zero-th order autocovariance of y_t.

With our MA(1) process:

The first-order autocovariance of y_t is $E(y_t \cdot y_{t-1})$ and it is denoted γ_1.

And for our MA(1) process we can calculate it as

The second order *autocovariance* is

$$\gamma_2 = \quad E(y_t \cdot y_{t-2}) = \quad E[(\epsilon_t + \theta\epsilon_{t-1})(\theta\epsilon_{t-2} + \epsilon_{t-3})]$$

$$= 0.$$

Footnote: Why? Because there are no cross-terms with matched time index. END FN

And you can see that all higher order autocovariances are zero because there will never be ϵ-crossproducts with matched time indexes.

****

Now we can plot the *autocovariance function.*

By definition the autocovariance function is γ_s as a function of s.

We can plot it for our MA(1) process. [NOTE: GRAPH GOES BELOW]

A GRAPH OF THE AUTOCOVARIANCE PROCESS FOR MA(1)

On the y-axis I will plot the values of γ.
On the x-axis I will plot the values of s.

For s = 0 the autocovariance is:

$\gamma_0 = (1 + \theta^2)\sigma_\epsilon^2$. **Let's plot it.** ***Mark this value with an x.***

For s=1

$\gamma_1 = \theta\, \sigma_\epsilon^2$. **Let's plot it.** ***Mark this value with an x.***

For s>1 the autocovariances are 0. They are zero for every value of s.

Mark these with x's.

Now I shall show that the autocovariance function is the *signature* of the MA function.

By that I mean that you can recover the MA process from the autocovariance function.

How do I know that?

First, the autocovariances die out after s=1.
This tells you that the process is MA(1).

If it were MA(2) or higher there would be nonzero autocovariances for s>1.

Second you can infer

$\sigma_\epsilon^2 (1 + \theta^2)$ **from** γ_0 **and**

$\sigma_\epsilon^2\theta$ **from** γ_1.

Thus you can infer what the value is individually of σ_ϵ and θ.
You have two equations in the two unknowns σ_ϵ and θ and so you can solve for them.

Thus, I have illustrated for this simplest case, that knowing the *autocovariances* and *plotting* them you can recover the whole process.
Indeed, you know the whole process exactly.

The method that I have just described is the core method used in (real) time series analysis.

****

Let me now give you *another definition*.
Instead of looking at the *autocovariance function,* it is common to look at the *autocorrelation function.*

To remind you, the autocovariance function is:

$$\gamma_s = E(y_t \cdot y_{t-s}).$$

The *autocorrelation function* is a normalized version of the autocovariance function.

It is:

In our example the autocorrelation function would be:

$$\rho_0 = 1$$

$$\rho_s = 0 \qquad \text{for all } s>1.$$

As you can see there is very little difference between the *autocovariance* function of an ARMA and the *autocorrelation* function.

Both are considered *signatures* of the process.

To repeat, by *signature* I mean that I can identify the process from looking at the

autocorrelation function.

Let's now look at the *autocovariance* functions of an AR process.
Let's first consider our simplest example of an AR process:

where ϵ_t is i.i.d. $N(0, \sigma_\epsilon^2)$.

Then *by definition*, the *s* th-order autocovariance is:

$$\begin{aligned}\gamma_s &= E(y_t \cdot y_{t-s}) \\ &= E[(\phi y_{t-1} + \epsilon_t) \cdot y_{t-s}] \\ &= \phi\, E(y_{t-1} \cdot y_{t-s}) + E(\epsilon_t \cdot y_{t-s}).\end{aligned}$$

ϵ_t and y_{t-s} are uncorrelated.

They are uncorrelated because ϵ_t is an innovation uncorrelated with past events.

So $E(\epsilon_t \cdot y_{t-s}) = 0$.

And $E(y_{t-1} \cdot y_{t-s})$ has a name.

It is just γ_{s-1}, or the (s-1)st order autocovariance.

Thus

$$\gamma_s = \phi\, \gamma_{s-1}.$$

So we get a difference equation of the γ_s's.

If we want to look at the *autocorrelation* function rather than the *autocovariance* function, we just divide γ_s by γ_0, a constant.

We would then find:

$$\rho_s = \phi\, \rho_{s-1}.$$

On inspection we can solve this difference equation:

We find that ρ_s is of the form:

$\rho_s \quad = \phi^s$.

FOOTNOTE: HOW DO YOU KNOW THAT?

Formally, for s=0, $\rho_0 = 1$.

And since ρ_s always obeys the first-order difference equation,

$\rho_s \quad = A\phi^s$

and A =1 since $\rho_0 = 1$. END FOOTNOTE

So the autocorrelation function of the AR(1) process has geometric decay, if $0 < \phi < 1$.

We can plot ρ_s as a function of s. [NOTE: GRAPH GOES BELOW]

GRAPH OF ρ_s AS A FUNCTION OF s.

Note that if $|\ \phi\ | > 1$ then the process is unstable and this procedure makes no sense.

FOOTNOTE: γ_0 **does not exist if** y_t **is generated by an unstable process. END FOOTNOTE**

In *general,* if you have a stable AR(p) process of the form:

you will find that its autocorrelation function will obey the difference equation:

The autocorrelation function will mirror the difference equation that generates the process.

In theory all we need to do is to look at the *autocorrelation function* and that will tell us what the process looks like.

If the autocorrelation function satisfies a difference equation of order *p*, then we will model the process as an AR(p) process with the coefficients of the difference equation being the coefficients of the AR(p) process.

That's what I mean when I say that the Autocorrelation function, or the Autocovariance function is the *signature of the process.* You can retrieve the process.

We can also express MA processes as AR processes and vice-versa.

Let me illustrate.

I will show that an MA(q) process is an AR(∞) process.

Let's consider the simplest example.

$$y_t = \epsilon_t + \theta\epsilon_{t-1}$$

So also: $$y_{t-1} = \epsilon_{t-1} + \theta\epsilon_{t-2}$$

We can multiply the RHS by θ, so

$$\theta y_{t-1} = \theta\epsilon_{t-1} + \theta^2\epsilon_{t-2}.$$

Subtracting

$y_t - \theta y_{t-1} \quad = \epsilon_t \qquad\qquad - \theta^2 \epsilon_{t-2}$.

We can now express ϵ_{t-2} in terms of y_{t-2} and ϵ_{t-3}.

With successive operations we find easily

$$y_t = - (\Sigma_{i=1}^{\infty} (-\theta)^i y_{t-i}) + \epsilon_t$$

So if y_t is an MA(1) process, we can express it also as an AR(∞) process if $|\theta|<1$.

****

Now let's see about converting an AR process to an MA process.

Suppose we have an AR(1) process.

We can write

$$y_t = \Sigma_{j=0}^{\infty} \phi^j \epsilon_{t-j}.$$

How do you know you can do that?

This is the solution to the first order difference equation that we put in the BOX at the beginning of class.

****

Let's now introduce some new terminology.

The condition that $|\phi|<1$ is necessary for the process to be stationary.

It may well be that

is *not stationary* because the largest root of the associated polynomial is *one or greater than one in absolute value.*

Let me give an example.

$$y_t = y_{t-1} + \epsilon_t.$$

But we may take the first difference and find then that

$$\Delta y_t = y_t - y_{t-1} = \epsilon_t$$

is *stationary.*

Let me give you *a new definition.*

An autoregressive *integrated* moving average process of order (p,d,q) is a process whose *d*th difference is an ARMA process of order (p,q).

It turns out that in *economics* this distinction may be important.

Quite often the *levels* of economic variables are not stationary, but the first differences will be stationary.

Thus an appropriate model is ARIMA(p,1,q);

ARIMA stands for AutoRegressive *Integrated* Moving Average.

I think that you will be later surprised how much this very short review of difference equations and ARMA processes simplifies our thinking about seemingly complex issues.

It is very important that you understand *very well and very thoroughly* what I have gone over in this brief review. The readings will go over what I have said a bit more; but for the most part you do not need to know much more than what I have gone over today for the use of ARMA processes in this course.

On Tuesday, we shall first cover the reading by Lucas and Sargent on the readings of Section II, and then we will cover the reading by Sargent in the evening class.

I want everyone to come to that class. It is probably the most important class of the term.