Topic 4: Rational Expectations Notes

Eddie Shore
Corporate Finance – ECONS4280

July 12, 2022

1 Introducing Uncertainty

Throughout the course, we haven't mentioned the fact that some things are uncertain. Given that we typically care about how choices today influence options tomorrow, we should be concerned about the fact that **the future is uncertain**.

How do we solve this problem? Well, people form beliefs! In order to think about how people do that, we have to implement a bit of structure on how we think some variable, y_t, might evolve in the future. In order to see this, let's think about a few real world examples:

- When you think about how hot or cold it will be tomorrow, how do you do it?
- When thinking about which library to study in based on how busy you think it will be, how do you choose?
- How do you pick your lottery numbers?

Each one of these systems of thinking can be formalised. I'm going to argue that each of these questions have answers that can be roughly modelled by the following three equations, where ϵ_{t+1} is a **random shock**, i.e. something that we can't predict, that affects our variable of interest at time $t+1$:

$$y_{t+1} = \rho y_t + \epsilon_{t+1} \tag{1}$$

$$y_{t+1} = \bar{y} + \epsilon_{t+1} \tag{2}$$

$$y_{t+1} = \epsilon_{t+1} \tag{3}$$

Let's consider each of these in turn, **working with the assumption that** ϵ_t is an iid random variable, i.e. it's distribution is fixed over time, and we draw from out randomly, without reference to previous draws.

Equation 1

Equation 1 answers the first question; what the temperature is tomorrow is probably going to be pretty similar to what it was today. When we think some variable Y_t has this relationship, where Y_{t+1} is a function of ρY_t and some random shock, we call it an **AR(1)** process, which stands for **autoregressive(1)**; i.e. we can regress Y_{t+1} on itself one period back. As you might expect, an AR(2) takes the form:

$$Y_{t+1} = \rho Y_t + \phi Y_{t-1} + \epsilon_{t+1}$$

For simplicity, we'll stick to AR(1)s for now. Remember that at time t, we don't know what Y_{t+1} is going to be; why is that? Because ϵ_{t+1} is **not something we can predict**. It is an iid random variable. What is the expected value of Y_{t+1}? Well if we take the expectation conditional on all the information available at time t, which I will denote using the notation $\mathbb{E}_t$, then:

$$\begin{aligned}\mathbb{E}_t[Y_{t+1}] &= \mathbb{E}_t[\rho Y_t + \epsilon_{t+1}] \\ &= \mathbb{E}_t[\rho Y_t] + \mathbb{E}_t[\epsilon_{t+1}] \\ &= \rho \mathbb{E}_t[Y_t] + \mathbb{E}_t[\epsilon_{t+1}]\end{aligned}$$

But if we're conditioning on all the information available at time t, then it must be the case that we know what Y_t is, right? The value of Y_t has been realised! It is a non-random constant now. Therefore:

$$\mathbb{E}_t[Y_{t+1}] = \rho Y_t + \mathbb{E}_t[\epsilon_{t+1}]$$

What value do you think $\mathbb{E}_t[\epsilon_{t+1}]$ ought to take? Well if it is a random shock, that we cannot predict, then we must think that in expectation it is 0. If instead we expected it to be some positive constant, then **that wouldn't be a random shock**; note that if we thought for example that there was some trend in our process, i.e. $\mathbb{E}_t[\epsilon_{t+1}] = a$, then we could equally describe the process described in Equation

14 as:

$$y_{t+1} = a + \rho y_t + \tilde{\epsilon}_{t+1}$$

And here it should be obvious that the expected value of $\tilde{\epsilon}_{t+1} = 0$. So when we talk about **random shocks**, we're always talking about something that is **mean equal to 0**. So, our best prediction of tomorrow's Y?

$$\mathbb{E}_t[Y_{t+1}] = \rho Y_t$$

What about Y_{t+j}? Given our information at time t, what should our expectation of Y_{t+j} be? Recall that our process tells us that $Y_{t+j} = \rho Y_{t+j-1} + \epsilon_{t+j}$; now consider the following:

$$\begin{aligned} \mathbb{E}_t[Y_{t+j}] &= \mathbb{E}_t[\rho Y_{t+j-1} + \epsilon_{t+j}] \\ &= \rho \mathbb{E}_t[Y_{t+j-1}] \end{aligned}$$

By identical logic we can keep going, substituting in our governing process as outlined in Equation 14, until we get:

$$\mathbb{E}_t[Y_{t+j}] = \rho^j Y_t \tag{4}$$

Typically in economics, we assume that $\rho \leq 1$; why is that? Suppose $\rho > 1$; what would happen to Y_t? It would explode; put mathematically, we can see that if $\rho > 1$, then:

$$\lim_{j \to \infty} \mathbb{E}_t[Y_{t+j}] = \lim_{j \to \infty} \rho^j Y_t = \infty$$

In economics, nothing goes to infinity, ever! Therefore any process that we're using to model economic behaviour should not lead to the conclusion that such behaviour will diverge to infinity!

Equation 2

The way you might think about how busy a given library is going to be is by assuming that it will be as busy as it has been on average, plus some random shock. If we let $\bar{Y}_t = \frac{1}{t} \sum_{i=0}^{t} Y_i$, then we get the expression in Equation 2. What should we expect in terms of how busy the library is tomorrow? Using the same procedure outlined above:

$$\mathbb{E}_t[Y_{t+1}] = \mathbb{E}_t[\bar{Y} + \epsilon_{t+1}]$$
$$= \bar{Y}$$

Clearly,

$$\mathbb{E}_t[Y_{t+2}] = \mathbb{E}_t[\bar{Y} + \epsilon_{t+1} + \epsilon_{t+2}]$$
$$= \bar{Y}$$

And so on, for as far forward as you care to look.

Equation 3

This is just a simplified version of Equation 2; here, we have no idea what tomorrow's Y_{t+1} is going to be. So, as with the lottery, we just have to guess.

A Side Note; what do we mean by conditional on all information at time t?

I mentioned above that $\mathbb{E}_t$ denotes the expectation conditional on all information at time t; I could just have easily said, let Ω_t denote the set of all information available at time t, and then written $\mathbb{E}[.|\Omega_t]$. But what constitutes the set Ω_t?

This is a good and important question. In the world at large, identifying the set of all available information is an impossible task; for example, when taking the expectation of tomorrow's prices, we have information on the price of copper in futures markets in Bulgaria, but do we really make use of this information? It is available! But processing all the available information in the world pertinent to our expectations is too difficult.

In our restricted context, i.e. in our model, we don't need to worry about all of this extra baggage. We have written the expression for Y_t as determined solely by past observations of Y_t and shocks ϵ_t. Claiming that we're using all information available today is simply to say that we observe and take into account Y_t when forecasting Y_{t+1}.

2 Rational Expectations

What are Rational Expectations?

Whilst it may not be immediately obvious, what we just did above in forecasting the expected future values of Y_{t+1} was an exercise in **Rational Expectations**. What do we mean by Rational Expectations? We can summarise this concept using two distinct notions; expectations are rational if **and only if**:

- The belief framework is consistent with the model framework.
- All information available is used to form beliefs.

Consider the first component; what this is telling us is that if in our model, Y_t evolves according to the process outlined in Equation 1, **then agents in our model believe that it evolves according to Equation 1**. This means they are using the correct system to forecast what will happen to Y_t in the future.

The second component builds on our previous discussion by emphasising it. Rational expectations requires that agents **do not ignore information that is pertinent to them**. In our model, that means they make use of all available information to forecast the future.

It should be clear to you that in all the examples above, we obeyed these two principles, and hence we solved the problem as though we had rational expectations.

What are *not* Rational Expectations?

A common mistake is to assume that rational expectations means that the agents are always 'correct'. Let's go back to our previous examples, and suppose that our random shock is distributed continuously (for example, suppose it is a normal random variable); remember that we found that our best guess of Y_{t+1} according to Equation 1 was ρY_t? Well the probability that $Y_{t+1} = \rho Y_t$ is zero. Why? Given ϵ_{t+1} is continuous, although in expectation it is zero, the probability that it *is* zero is zero. In formal language, we say that $Y_{t+1} \neq \rho Y_t$ *almost surely*. So rational expectations doesn't require that agents can predict the future!

Why are Rational Expectation reasonable:

They are **enormously** useful. Under the framework of Rational Expectations, we can produce consistent and coherent results from our models. To see this, suppose that we do away with the first component of rational expectations, i.e. we allow agents to have false systems of beliefs. How would you model this? How would you select an alternative belief system? What justification could you give for this belief system?

Crucially, how could you claim that agents would **persistently believe in a false system**? It is possible to produce Rational Expectations using a **learning model**, where agents are initially agnostic about the system governing a process, but through observation come to understand the process. If you're going to claim that they're not doing that, then why aren't they? Are people stupid? Do people not respond to incentives?

Suppose we relax the second term, so that we allow agents to ignore some information; which information do we allow them to ignore? Why that information and not some other? Is the information that agents do observe something you're imposing? If so, what justification can you give for that? Why wouldn't agents use all available information?

The fact remains that **even at the research frontier, we (as in, humanity, not just economists) do not know how people form expectations**. It is **very difficult** to credibly offer an alternative mechanism that is robust to the learning argument.

Why are Rational Expectations *not* reasonable:

Despite all of this, it should be clear to all of you that people don't have rational expectations. It is clearly the case that people don't perfectly understand the systems that govern random processes relevant to them, and it's also clearly the case that we don't think about the future price of futures markets of copper in Bulgaria, even though it might be relevant and available information to us. **Don't worry**, economists know this. Frontier work in economics focuses extensively on what happens when we relax both of these features of Rational Expectations.

A Small but 'Important' rant:

Something that really bothers me when I hear it in the news is when people say, 'These stupid economists believe that people form rational expectations! What idiots!' Whilst there are a few eccentrics out there who really do believe in rational expectations, the vast majority of the economics profession are aware that people don't actually have rational expectations; RE is a model. It's also worth noting that it's a model that does **extremely well** in most cases, in that it produces results that have been enormously useful, not just in understanding the world but in producing more effective policy that has improved the lives of millions.

Challenging a model by offering a credible alternative framework is both a noble and valuable pursuit. Shouting that rational expectations aren't true is just a waste of air; there is nothing impressive or valuable in decrying a model framework because of its implausibility, nothing creative in that, nothing worthy of praise. The reason why Tversky and Kahneman, and Thaler won Nobel Prizes in this area is **not** because they said 'Rational Expectations is nonsense'. They won Nobel Prizes because they were able to make statements of the following form; 'expectations are not rational, because of **this consistent and modellable characteristic**, that governs **broad human behaviour** across people, over time.' By identifying a **mechanism that differs from the rational paradigm**, these economists were able to provide a **coherent and credible system** that can challenge rational expectations, and the whole rational framework. This is an extraordinary achievement, and **is** worthy of praise.

For those of you unconvinced of the progress economics has made since the rational expectations revolution (which occurred nearly 50 years ago, at a time when **less than 1% of the US population owned a computer**), I would draw your attention to work by Mike Woodford, here at Columbia. Or Hassan Afrouzi, or Mark Dean, or Alessandra Casella. So there are **four** Professors **just at Columbia**, whose research deals almost exclusively with deviations from Rational Expectations.

I want to reiterate that although Rational Expectations are clearly false, **that does not mean they are not useful**. In most settings, rational expectations are a very good proxy for real world behaviour. They are also **extremely tractable**; this feature may not appear hugely important to you now, but this is an enormously non-trivial advantage of this modelling framework. Without rational expectations, many models become **literally impossible** to solve. Okay, that's my rant over!

3 Rational Expectations and Project Valuation

Up to now, we have assumed that the payoffs associated with a given project are known with certainty. This is clearly unreasonable. Suppose instead that the stream of payments associated with a project, $\{y_0, y_1, ...\}$, is random. How do we assess the value of this project? A slight modification to our formula for PV will allow us to extend the concept to cover uncertain payoff returns. Assuming a fixed discount rate, $(1+r)^{-1}$, then:

$$\begin{aligned} PV(Y) &= \mathbb{E}_0\left[\sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)y_t\right] \\ &= \sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)\mathbb{E}_0\left[y_t\right] \end{aligned}$$

This formula tells us that, conditional on a set of *beliefs* over the path of $\{y_t\}$, then we can calculate the present value in a similar as before. Under the assumption of *rational expectations*, then these beliefs should be *as good as possible*. In simple, toy examples, establishing the present value of a stream of random payments is simple. For example, suppose that:

$$y_t = \bar{y} + \epsilon_t$$

$$\epsilon_t \sim N(0,1)$$

where $\bar{y}$ is known. Then:

$$\begin{aligned} \mathbb{E}_0[y_t] &= \mathbb{E}_0[\bar{y} + \epsilon_t] \\ &= \bar{y} + \mathbb{E}_0[\epsilon_t] \\ &= \bar{y} \end{aligned}$$

Therefore, present value is given by:

$$\begin{aligned} PV(Y) &= \sum_{t=0}^{\infty} \left(\frac{1}{1+r} \right) \mathbb{E}_0 \left[y_t \right] \\ &= \sum_{t=0}^{\infty} \left(\frac{1}{1+r} \right) \bar{y} \\ &= \bar{y} \sum_{t=0}^{\infty} \left(\frac{1}{1+r} \right) \\ &= \frac{1+r}{r} \bar{y} \end{aligned}$$

3.1 A problem...

Compare the present value of the following two projects, projects A and B, where the random process governing $\{y_t\}$ differs in the following way:

$$\begin{aligned} A: \; y_t &= \bar{y} + \epsilon_t \\ \epsilon_t &\sim N(0,1) \\ B: \; y_t &= \bar{y} + \epsilon_t \\ \epsilon_t &\sim N(0,1000000) \end{aligned}$$

What is the difference in their present value? There is *no difference*. Which project would you prefer? Surely Project A! What this captures is the fact that present value, *in its current form*, tells us very little about the **riskiness** of projects when returns are uncertain.

3.1.1 An application to economics

The idea that the riskiness of a project should affect our assessment of it has a very natural parallel with the notion of expected value vs. expected utility in economics.

To illustrate with an example, suppose I offered you the following game: I will flip a coin, and if it lands on tails, I will give you \$2, and the game ends. If I flip heads, then I will flip again. This time, if I land on tails, I will give you \$4. If it lands heads, I flip again. Again, I double the reward from tails (now \$8), and if I land on heads, we keep going. Note that the returns here are *risky*. How much would you be willing to pay for this game? Think about it for a couple of seconds.

Now that you've established an intuitive impression of how much the game is worth, let's establish what it is worth in *expected value*. Recall that the expected value is just the sum of all the potential payouts, scaled by their probability. The probability that I get a payout of \$2 is 0.5. The probability of \$4 is 0.25. The probability of \$8 is 0.125, and so on. Putting this together:

$$\begin{aligned} EV &= 0.5 * 2 + 0.25 * 4 + 0.125 * 8 + ... \\ &= 1 + 1 + 1 + ... \\ &= \infty \end{aligned}$$

So the expected value of this game is ∞. Did any of you value the game that high? I doubt it! This result is known as the **St Petersburg paradox**, and it formed one of the bases for establishing *expected utility* theory. That theory tells us that the raw payoffs are not the only thing that matter to us. In mathematical terms, *we care about more than just the first moment*, and that's because we dislike 'risky' prospects. Let's show how expected utility resolves the paradox.

3.1.2 Resolving the St Petersburg Paradox

Suppose that instead of taking the value in absolute terms of the payoffs, we instead take the *utility* of the expected payoff instead. For simplicity, suppose that:

$$u(y_t) = \log(y_t)$$

Note that assuming a concave utility function is the same things as assuming risk aversion. To see this, let's start by defining risk aversion.

> **Definition 1:** *Risk aversion* denotes a preference for the expected value of a risky payoff rather than taking the risky payoff itself. For example, if a coin toss delivers \$1 for heads, and \$0 for tails, then a risk averse agent would not accept this payoff if asked to pay 50 cents for it. More formally, suppose WTP is the maximum amount that an agent is willing to pay for a risky payoff, then the *Risk Premium* (RP) is defined as:
>
> $$RP = EV - WTP$$
>
> The agent is risk averse(loving) if RP is greater(less) than 0. Risk neu-

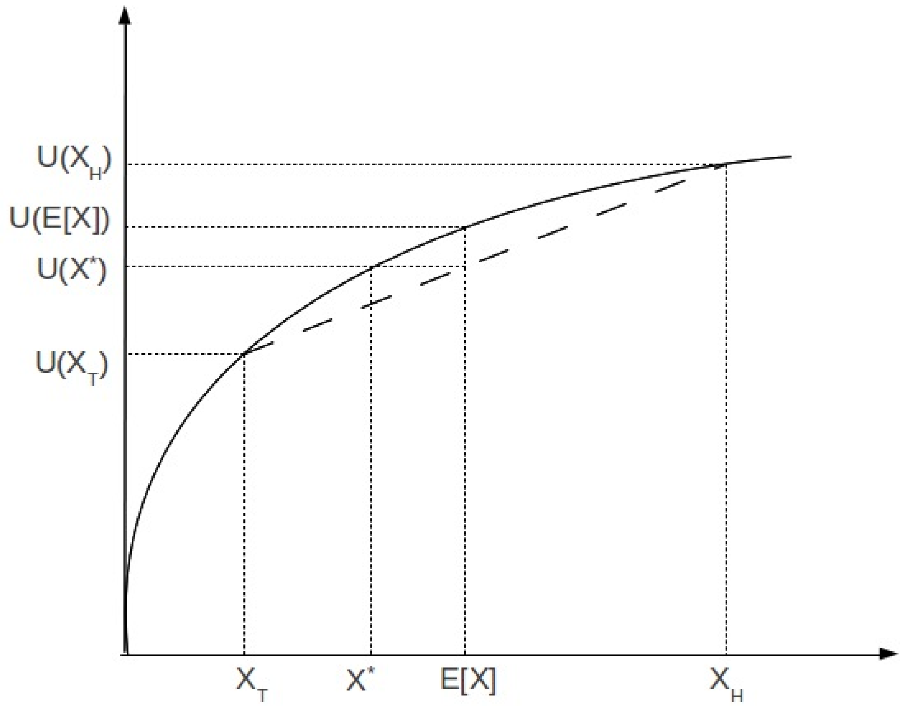

Figure 1: Concavity in the utility function implies risk averse behavior.

trality implies that $RP = 0$.

Why does a concave utility function imply this? Perhaps the graph in Figure 1 will illustrate. In the graph, we are supposing a risky scenario in which there are two potential payoffs, $\{X_H, X_T\}$, that are (roughly) equally likely. The expected *utility* of this payoff is then given by:

$$EU(Risky) = 0.5 * U(X_T) + 0.5 * U(X_H) = U(X^*)$$

where $U(X*)$ is shown on the graph. Suppose that instead we gave the agent the *expected value* of the gamble for sure: that would be a payoff of $0.5X_T + 0.5 * X_H$. The expected utility of this payoff would be:

$$EU(Certain) = U(0.5X_T + 0.5 * X_H) = U(E[X])$$

where $U(E[X])$ is also shown on the graph, *and is higher than* $U(X*)$. In other words, *if asked to pay* $E[X]$ *for this gamble, the expected utility would be less than* *the cost.* In other words, the agent is risk averse. This also leads to another definition of risk aversion:

> **Definition 2:** For a given gamble X, an agent is *risk-averse* if and only if:
> $$EU(X) < U(E[X])$$

Okay, now we've got that over with, what does the application of expected *utility* do to the St Petersburg problem? Well, if we calculate the expected utility of the game presented in the problem, we will find something like the following:

$$\begin{aligned}
\therefore EU =& 0.5 * \log(2) + 0.25 * \log(4) + 0.125 * \log(8) + ... \\
=& \sum_{j=1}^{\infty} 2^{-j}. \log(2^j) \\
=& \sum_{j=0}^{\infty} 2^{-j}. \log(2^j) - \underbrace{2^{-0}. \log(2^0)}_{=0} \\
=& \log(2) \sum_{j=0}^{\infty} j2^{-j} = \log(2) \sum_{j=0}^{\infty} j(0.5)^j \\
& \textit{here noting that } \sum_{j=0}^{\infty} ja^j = \frac{a}{(1-a)^2} \\
=& \log(2). \frac{0.5}{(1-0.5)^2} \\
=& \log(2).2 = \log(4)
\end{aligned}$$

So the expected utility is $\log(4)$, which by the definition of the utility function means that the value in cash is just \$4! Probably not far off how much you'd pay for it!

3.2 Application to Finance

In economics, risk aversion in the utility function is typically assumed ex ante as an invariant property of human preferences. In finance, we can go better by considering the fact that *resources are not infinite.*

In the example between project A and B above, there is a distinct possibility of enormous losses. For example, there is a roughly 15% chance of a \$1,000,000 loss *or worse* in every period of this project. There is roughly a 2% chance that this happens consecutively. A sequence of bad losses in the real world can cause a firm to collapse, meaning they will be unable to continue with the project. As of right now, we've allowed no space for such a concern.

Risk is **always** a factor in financial decision making, so we have to expand our toolset to allow us to incorporate the riskiness of a project when assessing its value. How do we do that? Well, we do it **through the discount rate.**

3.2.1 Modifying the discount rate

In the example above, I specified that the discount rate was fixed and given by r. But recall what r is intended to capture: namely, the *opportunity cost* of the project —the returns that we *could* have achieved by investing in something else *similar* to the project in question. Clearly Projects A and B are *not* similar. So, the appropriate discount rate should also be different.

Next class we will discuss how the riskiness and the return of a project are intimately linked. This result will allow us to conclude that the r we use to discount riskier projects should be *higher* than for safer projects. Thus, with this in place, we can assign discount rates to projects A and B such that $r^A < r^B$. The proof below shows that the PV of A is now *higher* than that of B, even though the *expected* payoffs are the same.

Theorem 1. *For two projects, A and B, such that both have the same expected payoffs, , the variance of payoffs of A, σ^A, is less than the variance of payoffs of B, σ^B, and the discount rate for A is less than that of B, $0 < r^A < r^B$: then the PV of A is greater than that of B.*

Proof. Note that:

$$PV(A) = \frac{1 + r^A}{r^A}\bar{y}$$

$$PV(B) = \frac{1 + r^B}{r^B}\bar{y}$$

Therefore, it is enough to show that $\frac{1+r^A}{r^A} > \frac{1+r^B}{r^B}$. Suppose this is not so:

$$
\begin{aligned}
\frac{1+r^A}{r^A} &< \frac{1+r^B}{r^B} \\
\therefore (1+r^A)r^B &< (1+r^B)r^A \\
\implies r^B + r^A r^B &< r^A + r^A r^B \\
\implies r^B &< r^A
\end{aligned}
$$

Which is a contradiction. $\square$