Lecture notes on

FINANCIAL MARKETS

Marco Li Calzi

Dipartimento di Matematica Applicata
Università "Ca' Foscari" di Venezia

February 2002

M. Li Calzi, *Lecture notes on Financial Markets*

Introduction

In 2000, when Bocconi University launched its *Master in Quantitative Finance*, I was asked to develop a course touching upon the many aspects of decision theory, financial economics, and microstructure that could not otherwise fit in the tight schedule of the program. Reflecting this heterogeneity, the course was dubbed "Topics in Economics" and I was given a fair amount of leeway in its development.

My only constraint was that I had to choose what I thought best and then compress it in 15 lectures. These notes detail my choices after two years of teaching "Topics in Economics" at the *Master in Quantitative Finance* of Bocconi University and a similar class more aptly named "Microeconomics of financial markets" at the *Master of Economics and Finance* of the Venice International University.

The material is arranged into 15 units, upon whose contents I make no claim of originality. Each unit corresponds to a 90-minutes session. Some units (most notably, Unit 5) contain too much stuff, reflecting either the accretion of different choices or my desire to offer a more complete view. Unit 7 requires less time than the standard session: I usually take advantage of the time left to begin exploring Unit 8.

I have constantly kept revising my choices (and my notes), and I will most likely do so in the future as well, posting updates on my website at `http://helios.unive.it/~licalzi`.

Contents

1. Expected utility and stochastic dominance

1.1 Introduction

Most decisions in finance are taken under a cloud of uncertainty. When you plan to invest your money in a long-term portfolio, you do not know how much will its price be at the time of disinvesting it. Therefore, you face a problem in choosing the "right" portfolio mix. Decision theory is that branch of economic theory which works on models to help you sort out this kind of decisions.

There are two basic sorts of models. The first class is concerned with what is known as *decisions under risk* and the second class with *decisions under uncertainty*.

1.2 Decisions under risk

Here is a typical decision under risk. Your investment horizon is one year. There is a family of investment funds. You must invest all of your wealth in a single fund. The return on each fund is not known with certainty, but you know its distribution of past returns. For lack of better information, you have decided to use this distribution as a proxy for the probability distribution of future returns.[1]

Let us model this situation. There is a set C of consequences, typified by the one-year returns you will be able to attain. There is a set A of alternatives (i.e., the funds) out of which you must choose one. Each alternative in A is associated with a probability distribution over the consequences. For instance, assuming there are only three funds, your choice problem may be summarized by the following table.

Fund α		Fund β		Fund γ	
return	prob.ty	return	prob.ty	return	prob.ty
-1%	20%	-3%	55%	2.5%	100%
+2%	40%	+10%	45%		
+5%	40%				

Having described the problem, the next step is to develop a systematic way to make a choice.

Def. 1.1 [Expected utility under risk] Define a real-valued utility function u over consequences. Compute the expected value of utility for each alternative. Choose an alternative which maximizes the expected utility.

[1] The law requires an investment fund to warn you that past returns are not guaranteed. Trusting the distribution of past returns is a choice you make at your own peril.

How would this work in practice? Suppose that your utility function over a return of $r\%$ in the previous example is $u(r) = r$. The expected utility of Fund α is

$$U(\alpha) = -1 \cdot 0.2 + 2 \cdot 0.4 + 5 \cdot 0.4 = 2.6.$$

Similarly, the expected utility of Fund β and γ are respectively $U(\beta) = 2.85$ and $U(\gamma) = 2.5$. According to the expected utility criterion, you should go for Fund β and rank α and γ respectively second and third.

If you had a different utility function, the ranking and your final choice might change. For instance, if $u(r) = \sqrt{r+3}$, we find $U(\alpha) \approx 2.31$, $U(\beta) \approx 1.62$ and $U(\gamma) \approx 2.35$. The best choice is now γ, which however was third under the previous utility function.

All of this sounds fine in class, but let us look a bit more into it. Before you can get her to use this, there are a few questions that your CEO would certainly like you to answer.

Is expected utility the "right" way to decide? Thank God (or free will), nobody can pretend to answer this. Each one of us is free to develop his own way to reach a decision. However, if you want to consider what expected utility has in it, mathematicians have developed a partial answer. Using expected utility is equivalent to taking decisions that satisfy three criteria: 1) consistency; 2) continuity; 3) independence.

Consistency means that your choices do not contradict each other. If you pick α over β and β over γ, then you will pick α over γ as well. If you pick α over β, you do not pick β over α.

Continuity means that your preferences do not change abruptly if you slightly change the probabilities affecting your decision. If you pick α over β, it must be possible to generate a third alternative α' by perturbing slightly the probabilities of α and still like α' better than β.

Independence is the most demanding criterion. Let α and β be two alternatives. Choose a third alternative γ. Consider two lotteries: α' gets you α or γ with equal probability, while β' gets you β or γ with equal probability. If you'd pick α over β, then you should also pick α' over β'.

If you are willing to subscribe these three criteria simultaneously, using expected utility guarantees that you will fulfill them. On the other hand, if you adopt expected utility as your decision making tool, you will be (knowingly or not) obeying these criteria. The answer I'd offer to your CEO is: "if you wish consistency, continuity and independence, expected utility is right".

Caveat emptor! There is plenty of examples where very reasonable people do not want to fulfill one of the three criteria above. The most famous example originated with Allais who, among other things, got the Nobel prize in Economics in 1988. Suppose the consequences are given as payoffs in millions of Euro. Between the two alternatives

α		β	
payoff	prob.ty	payoff	prob.ty
0	1%	1	100%
1	89%		
5	10%		

Allais would have picked β. Between the two alternatives

γ		δ	
payoff	prob.ty	payoff	prob.ty
0	90%	0	89%
5	10%	1	11%

he would have picked γ. You can easily check (yes, do it!) that these two choices cannot simultaneously be made by someone who is willing to use the expected utility criterion.

Economists and financial economics, untroubled by this, assume that all agents abide by expected utility. This is partly for the theoretical reasons sketched above, but mostly for convenience. To describe the choices of an expected utility maximizer, an economist needs only to know the consequences, the probability distribution over consequences for each alternative, how to compute the expected value and the utility function over the consequences. When theorizing, we'll do as economists do: we assume knowledge of consequences, alternatives and utility functions and we compute the expected utility maximizing choice.

For the moment, however, let us go back to your CEO waiting for your hard-earned wisdom to enlighten her.

What is the "right" utility function? The utility function embeds the agent's preferences under risk. In the example above, when the utility function was $u(r) = r$, the optimal choice is Fund β which looks a lot like a risky stock fund. When the utility function was $u(r) = \sqrt{r+3}$, the optimal choice was Fund γ, not much different from a standard 12-month Treasury bill. It is the utility function which makes you prefer one over another. Picking the right utility function is a matter of describing how comfortable we feel about taking (or leaving) risks. This is a tricky issue, but I'll say more about it in Lecture 4.

Sometimes, we are lucky enough that we can make our choice without even knowing what exactly is our utility function. Suppose that consequences are monetary payoffs and assume (as it is reasonable) that the utility function is increasing. Are there pairs of alternatives α and β such that α is (at least, weakly) preferred by all sorts of expected utility maximizers?

In mathematical terms, let F and G be the cumulative probability distributions respectively for α and β. What is the sufficient condition such that

$$\int u(x)\,\mathrm{d}F(x) \geq \int u(x)\,\mathrm{d}G(x)$$

for all increasing utility functions u?

Def. 1.2 [Stochastic dominance] Given two random variables α and β with respective cumulative probability distributions F and G, we say that α stochastically dominates β if $F(x) \leq G(x)$ for all x.

Stochastic dominance of α over β means that $F(x) = P(\alpha \leq x) \leq P(\beta \leq x) = G(x)$ for all x. That is, α is less likely than β to be smaller than x. In this sense, α is less likely to be small.

If you happen to compare alternatives such that one stochastically dominates the other ones and you believe in expected utility, you can safely pick the dominating one without even worrying to find out what your "right" utility function should be. This may not happen often, but let us try not to overlook checking for this clearcut comparison.

Isn't this "expected utility business" too artificial? Well, it might be. But we are not asking you to use expected utility to take your decisions. Expected utility is what economists use to model your behavior under risk. If you happen to use a different route which fulfills the three criteria of consistency, continuity and independence, an economist will be able to fit your past choices to an expected utility model and guess your future choices perfectly.

We can put it down to a matter of decision procedures. Expected utility is one: it is simple to apply but it requires to you to swallow the idea of a utility function. There are other procedures which lead you to choices that are compatible with expected utility maximization in a possibly more natural way.

Here is an example[2] of an alternative procedure. Suppose that you are a fund manager and that your compensation depends on a benchmark. Your alternatives are the different investing strategies you may follow. Each strategy will lead to a payoff at the end of the year which is to be compared against the benchmark. If you beat the benchmark, you'll get a fixed bonus; otherwise, you will not. The performance of the benchmark is a random variable B. Using past returns, you estimate its cumulative probability distribution H. Moreover, since you are only one of many fund managers, you assume that the performance of the benchmark is independent of which investing strategy you follow.

Your best bet is to maximize the probability of getting the bonus. If your investing strategy leads to a random performance α with c.d.f. F, the probability of getting the bonus is simply

$$P(\alpha \geq B) = \int P(x \geq B)\,\mathrm{d}F(x) = \int H(x)\,\mathrm{d}F(x).$$

While (naturally) trying to maximize your chances of getting your bonus, you will be behaving as if (artificially) trying to maximize a utility function $u(x) = H(x)$.

What is the "right" probability distribution for an alternative? Ah, that's a good question. You might have read it already but, thank God (or free will), nobody can answer this. Each one of us is free to develop his own way to assess the probabilities. In the example above, I mischievously assumed that you were willing to use the distribution of past returns but this may sometimes be plainly wrong.

Economists have since long recognized the importance of this question. Their first-cut answer is to isolate the problem by assuming that the probabilities have already been estimated by someone else and are known to the agent. Whenever this assumption holds, we speak of decision making under risk. Whenever we do not assume that probabilities are already known, we enter the realm of decisions under uncertainty.

[2] This is drawn from research developed in Bocconi and elsewhere. See Bordley and LiCalzi (2000) for the most recent summary.

1.3 Decisions under uncertainty

Here is a typical decision under uncertainty. Your investment horizon is one year. There is a family of investment funds. You must invest all of your wealth in a single fund. The return on each fund is not known with certainty and you do not think that the distribution of past returns is a good proxy. However, there is a list of scenarios upon which the return on your investment depends.

Let us model this situation. There is a set C of consequences, typified by the one-year returns you will be able to attain. There is a list S of possible future scenarios. There is a set A of alternatives (i.e., the funds) out of which you must choose one. Each alternative α in A is a function which tells you which consequence c you will be able to attain under scenario s: that is, $\alpha(s) = c$. Assuming there are only three funds, your choice problem may be summarized by the following table.

	Fund α	Fund β	Fund γ
scenario	return	return	return
s_1	-1%	-3%	2.5%
s_2	+2%	-3%	2.5%
s_3	+2%	-3%	2.5%
s_4	+2%	+10%	2.5%
s_5	+5%	+10%	2.5%

Def. 1.3 [Expected utility under uncertainty] Assess probabilities for each scenario. Define a utility function over consequences. Compute the expected value of utility for each alternative over the scenarios. Choose an alternative which maximizes the expected utility.

After you assess probabilities for each scenario, you fall back to the case of decision under risk. For instance, assessing $P(s_1) = 20\%$, $P(s_2) = 30\%$, $P(s_3) = P(s_4) = 5\%$, and $P(s_5) = 40\%$ gets you back to the case studied above. If your utility function were $u(r) = r$, the optimal choice would again be β. However, staying with the same utility function, if you'd happen to assess $P(s_1) = P(s_2) = P(s_4) = P(s_5) = 5\%$, and $P(s_3) = 80\%$, the optimal choice would be γ.

Under uncertainty, the analysis is more refined. What matters is not only your attitude to risk (as embedded in your choice of u), but your beliefs as well (as embedded in your probability assessment).

Expliciting scenarios may matter in a surprising way, as it was noted in Castagnoli (1984). Suppose the consequences are given as payoffs in millions of Euro. Consider the following decision problem under uncertainty.

	Fund α	Fund β
scenario	payoff	payoff
s_1	0	4
s_2	1	0
s_3	2	1
s_4	3	2
s_5	4	3

Suppose that you assess probabilities $P(s_1) = 1/3$, and $P(s_2) = P(s_3) = P(s_4) = P(s_5) = 1/6$. Then β would stochastically dominate α even though the probability that α beats β is $P(\alpha \geq \beta) = 2/3$. Any expected utility maximizer (if using an increasing utility function) would pick β over α. However, if you are interested only in choosing whichever alternative pays more between the two, you should go for α.

References

[1] R. Bordley e M. Li Calzi (2000), "Decision analysis using targets instead of utility functions", *Decisions in Economics and Finance* **23**, 2000, 53–74.

[2] E. Castagnoli (1984), "Some remarks on stochastic dominance",*Rivista di matematica per le Scienze Economiche e Sociali* **7**, 15–28.

2. Irreversible investments and flexibility

2.1 Introduction

Under no uncertainty, NPV is the common way to assess an investment. (In spite of contrary advice from most academics, consultants and hence practitioners use the payback time and the IRR as well.) If you have to decide whether to undertake an investment, do so only if its NPV is positive. If you have to pick one among many possible investments, pick the one with the greatest NPV.

When uncertainty enters the picture, the easy way out is to keep doing NPV calculations using expected payoffs instead of the actual payffos, which are not known for sure. This might work as a first rough cut, but it could easily led you astray. The aim of this lecture is to alert you about what you could be missing. Most of the material is drawn from Chapter 2 in Dixit and Pyndick (1994).

2.2 Price uncertainty

Consider a firm that must decide whether to invest in a widget factory. The investment is irreversible: the factory can only be used to produce widgets and, if the markets for widgets should close down, the firm could not be scrapped down and sold to someone else. The firm can be built at a cost of $c = 1600$ and will produce one widget per year forever, with zero operating cost. The current price of a widget is $p_0 = 200$, but next year this will rise to $p_1 = 300$ with probability $q = 1/2$ or drop to $p_1 = 100$ with probability $1 - q = 1/2$. After this, the price will not change anymore. The risk over the future price of widgets is fully diversifiable and therefore we use the risk-free rate of interest $r = 10\%$.

Presented with this problem, a naive CFO would compute the expected price $p = 200$ from the next year on. Using the expected price, the NPV of the project is

$$\text{NPV} = -1600 + \sum_{t=0}^{\infty} \frac{200}{(1.1)^t} \approx 600.$$

Since the NPV is positive, the project gets the green light and the firm invests right away.

A clever CFO would consider also the possibility of waiting one year. At the cost of giving up a profit of 2000 in year $t = 0$, one gains the option to invest if the price rises and to not invest otherwise. The NPV for this investment policy is

$$\text{NPV} = \frac{1}{2}\left[\frac{-1600}{1.1} + \sum_{t=1}^{\infty} \frac{300}{(1.1)^t}\right] \approx 773.$$

This is higher than 600, and therefore it is better to wait than to invest right away. The value of the flexibility to postpone the investment is $773 - 600 = 173$.

Ex. 2.1 For a different way to assess the value of flexibility, check that the opportunity to build a widget factory *now and only now* at a cost of $c = 1600$ yields the same NPV as the opportunity of building a widget factory *now or next year* at a cost of (about) 1980.

Ex. 2.2 Suppose that there exists a futures market for widgets, with the futures price for delivery one year from now equal to the expected future spot price of 200. Would this make us anticipate the investment decision? The answer is no. To see why, check that you could hedge away price risk by selling short futures for 11 widgets, ending up with a sure NPV of 2200. Subtract a cost of 1600 for building the factory, and you are left exactly with an NPV of 600 as before. The futures market allows the firm to get rid of the risk but does not improve the NPV of investing now.

2.3 Real options

We can view the decision to invest now or next year as the analog of an american option. An american option gives the right to buy a security any time before expiration and receive a random payoff. Here, we have the right to make an investment expenditure now or next year and receive a random NPV. Our investment option begins "in the money" (if it were exercised today, it would yield a positive NPV), but waiting is better than exercising now. This sort of situations, where the underlying security is a real investment, are known as real options. The use of real options is getting increasingly popular in the assessment of projects under uncertainty.

Let us compute the value of our investment opportunity using the real options approach. Denote by F_0 the value the option today, and by F_1 the value next year. Then F_1 is a random variable, which can take value

$$\left[\sum_{t=0}^{\infty} \frac{300}{(1.1)^t}\right] - 1600 \approx 1700$$

with probability 1/2 (if the widget price goes up to 300) and value 0 with probability 1/2 (if it goes down). We want to find out what is F_0.

Using a standard trick in arbitrage theory, consider a portfolio in which one holds the investment opportunity and sells short n widgets at a price of P_0. The value of this portfolio today is $\Pi_0 = F_0 - nP_0 = F_0 - 200n$. The value of the portfolio next year is $\Pi_1 = F_1 - nP_1$. Since $P_1 = 300$ or 100, the possible values of Π_1 are $1700 - 300n$ or $-100n$. We can choose n and make the portfolio risk-free by solving $1700 - 300n = -100n$, which gives $n = 8.5$. This number of widgets gives a sure value $\Pi_1 = -850$ for the portfolio.

The return from holding this portfolio is the capital gain $\Pi_1 - \Pi_0$ minus the cost of shorting the widgets; that is, $\Pi_1 - \Pi_0 - 170 = -850 - (F_0 - 1700) - 170 = 680 - F_0$. Since this portfolio is risk-free, it must earn the risk-free rate of $r = 10\%$; that is, $680 - F_0 = (0.1)\Pi_0 = 0.1(F_0 - 1700)$, which gives $F_0 = 773$. This is of course the same value we have already found.

2.4 Assessing your real option

Once we view an investment opportunity as a real option, we can compute its dependence on various parameters and get a better understanding. In particular, let us determine how the value of the option — and the decision to invest — depend on the cost c of the investment, on the initial price P_0 of the widgets, on the magnitudes of the up and down movements in price next period, and on the probability q that the price will rise next period.

a) Cost of the investment. Using the arbitrage argument, we find (please, do it) that the short position on widgets needed to obtain a risk-free portfolio is $n = 16.5 - 0.005c$. Hence, $\Pi_1 = F_1 - nP_1 = 0.5c - 1650$ and $\Pi_0 = F_0 - 3300 + c$. Imposing a risk-free rate of $r = 10\%$ yields

$$F_0 = 1500 - 0.455c, \tag{1}$$

which gives the value of the investment opportunity as a function of the cost c of the investment.

We can use this relationship to find out for what values of c investing today is better than investing next year. Investing today is better as long as the value V_0 from investing is greater than the direct cost c plus the opportunity cost F_0. Since the NPV of the payoffs from investing today is 2200, we should invest today if $2200 > c + F_0$. Substituting from (1), we should invest as long as $c < 1284$.

In the terminology of financial options, for low values of c the option is "deep in the money" and immediate exercise is preferable, because the cost of waiting (the sacrifice of the immediate profit) outweighs the benefit of waiting (the ability to decide optimally after observing whether the price has gone up or down).

b) Initial price. Fix again $c = 1600$ and let us now vary P_0. Assume that with equal probability the price P_1 will be twice or half the current price P_0 (and remain at this level ever after). Suppose that we want to invest when the price goes up and we do not want to invest if it goes down (we will consider other options momentarily). Set up the usual portfolio and check (yes, do it) that its value is $\Pi_1 = 16.5P_0 - 1600 - 1.5nP_0$ if the price goes up and $\Pi_1 = -0.5nP_0$ is the price goes down. Equating these two values, we find that $n = 16.5 - (1600/P_0)$ is the number of widgets that we need to short to make the portfolio risk-free, in which case $\Pi_1 = 800 - 8.25P_0$ whether the price goes up or down.

Recall that the short position requires a payment of $0.1nP_0 = 1.65P_0 - 160$ and compute the return on this portfolio. Imposing a risk-free rate of $r = 10\%$, we have $\Pi_1 - \Pi_0 - [1.65P_0 - 160] = 0.1\Pi_0$ which yields

$$F_0 = 7.5P_0 - 727. \tag{2}$$

This value of the option to invest has been calculated assuming that we would only want to invest if the price goes up next year. However, if P_0 is low enough we might never want to invest, and if P_0 is high enough it might be better to invest now rather than waiting.

Let us find for which price we would never invest. From (2), we see that $F_0 = 0$ when $P_0 \approx 97$. Below this level, there is no way to recoup the cost of the investment even if the price rises by 50% next year. Analogously, let us compute for which price we would always invest today. We should invest now if the NPV of current payoffs (i.e., $11P_0$) exceeds the total cost $1600+F_0$ of investing now. The critical price $\hat{P}_0$ satisfies $11\hat{P}_0 - 1600 = F_0$ which, after substituting from (2), gives $\hat{P}_0 = 249$. Summarizing, the investment rule is:

Price region	Option value	Investment rule
if $P_0 \leq 97$	then $F_0 = 0$	and you never invest;
if $97 < P_0 \leq 249$	then $F_0 = 7.5P_0 - 727$	and you invest next year if price goes up;
if $P_0 > 249$	then $F_0 = 11P_0 - 1600$	and you invest today.

c) Probabilities. Fix an arbitrary P_0 and let us vary q. In our standard portfolio, the number of widgets needed to construct a risk-free position is $n = 8.5$ and is independent of q (yes, check it). The expected price of widgets next year is $E(P_1) = q(1.5P_0) + (1-q)(0.5P_0) = (q+0.5P_0)$; therefore the expected capital gain on widgets is $[E(P_1)-P_0]/P_0 = q-0.5$. Since the long owner of a widget demands a riskless return of $r = 10\%$ but gets already a capital gain of $q-0.5$, she will ask a payment of $[0.1-(q-0.5)]P_0 = (0.6-q)P_0$ per widget. Setting $\Pi_1 - \Pi_0 - (0.6-q)nP_0 = 0.1\Pi_0$ with $n = 8.5$ we find (for $P_0 > 97$) that the value of the option is

$$F_0 = (15P_0 - 1455)q. \tag{3}$$

For $P_0 \leq 97$ we would never invest and $F_0 = 0$.

What about the decision to invest? It is better to wait than to invest today as long as $F_0 > V_0 - c$. Since $V_0 = P_0 + \sum_1^\infty [(q+0.5)P_0]/(1.1)^t = (6+10q)P_0$, it is better to wait as long as $P_0 < \hat{P}_0 = (1600-1455q)/(6-5q)$ — yes, check this. Note that $\hat{P}_0$ decreases as q increases: a higher probability of a price increase makes the firm more willing to invest today. Why?

d) Magnitudes. Fix $q = 0.5$ and let us change the magnitudes of variations in price from 50% to 75%. This leaves $E(P_1) = P_0$ but increases the variance of P_1. As usual, we construct a risk-free portfolio by shorting n widgets. The two possible values for Π_1 are $19.25P_0 - 1600 - 1.75nP_0$ if the price goes up and $-0.25nP_0$ if the price goes down — yes, check this. Equating these two values and solving for n gives $n = 12.83 - (1067/P_0)$, which makes $\Pi_1 = 267 - 3.21P_0$ irrespective of P_1. Imposing a risk-free rate of $r = 10\%$ (please fill in the missing steps) yields

$$F_0 = 8.75P_0 - 727. \tag{4}$$

At a price $P_0 = 200$, this gives a value $F_0 = 1023$ for the option to invest significantly higher than the 773 we found earlier. Why does an increase in uncertainty increase the value of this option?

Ex. 2.3 Show that the critical initial price sufficient to warrant investing now instead rather than waiting is $\hat{P}_0 \approx 388$, much larger than the 249 found before. Can you explain why?

References

[1] A.K. Dixit and R.S. Pyndick (1994), *Investment under Uncertainty*, Princeton (NJ): Princeton University Press.

3. Optimal growth and repeated investments

3.1 Introduction

Standard portfolio theory treats investments as a single-period problem. You choose your investment horizon, evaluate consequences and probabilities for each investment opportunity and pick a portfolio which, for a given level of risk, maximizes the expected return over the investment horizon. The basic lesson from this static approach is that volatility is "bad" and diversification is "good".

The implicit assumptions are that you know your investment horizon and that you plan to make your investment choice once and for all. However, when we begin working over multiperiod investment problems, some of the lessons of the static approach take a whole new flavour. The aim of this lecture is to alert you about some of the subtleties involved. Most of the material is drawn from Chapter 15 in Luenberger (1998).

3.2 An example

At each period, you are offered three investment opportunities. The following table reports their payoffs to an investment of Euro 100. An identical but independently distributed selection is offered each period, so that the payoffs to each investment are correlated within each period, but not across time.

	α	β	γ	
scenario	payoff	payoff	payoff	prob.ty
s_1	300	0	0	1/2
s_2	0	200	0	1/3
s_3	0	0	600	1/6

You start with Euro 100 and can invest part or all of your money repeatedly, reinvesting your winnings at later periods. You are not allowed to go short, but you can apportion your investment over different opportunities. What should you do?

Consider the static choice over a single period. There is an obvious trade-off between pursuing the growth of the capital and avoiding the risk of losing it all. More precisely, for an investment of 100, the first lottery has an expected value of 150 and a 50% probability of losing the capital. The second lottery has an expected value of (about) 67 and a 66.6% probability of losing the capital. The third lottery has an expected value of 100 and a 83.3% probability of losing the capital. Comparing β against γ, Lottery β minimizes the risk of being ruined while γ offers the highest expected return. However, note that α dominates β and γ under both respects. If you want to maximize your expected gain, investing 100 in α is the best choice.

This intuition does not carry over to the case of a multiperiod investment. If you always invest all the current capital in α, sooner or later this investment will yield 0 and therefore you are guaranteed to lose all of your money. Instead of maximizing your return, repeatedly betting the whole capital on α guarantees your ruin.

Let us consider instead the policy of reinvesting your capital each period in a fixed-proportion portfolio $(\alpha_1, \alpha_2, \alpha_3)$, with $\alpha_i \geq 0$ for $i = 1, 2, 3$ and $\alpha_1 + \alpha_2 + \alpha_3 \leq 1$. Each of these portfolios leads to a series of (random) multiplicative factors that govern the growth of capital.

For instance, suppose that you invest Euro 100 using the $(1/2, 0, 0)$ portfolio. With probability 50%, you obtain a favorable outcome and double your capital; with probability 50%, you obtain an unfavorable outcome and your capital is halved. Therefore, the multiplicative factors for one period are 2 and 1/2, each with probability 50%. Over a long series of investments following this strategy, the initial capital will be multiplied by a multiple of the form

$$\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(2)\left(\frac{1}{2}\right)(2)\,(2)\left(\frac{1}{2}\right)\cdots\left(\frac{1}{2}\right)\left(\frac{1}{2}\right)(2)$$

with about an equal number of 2's and (1/2)'s. The overall factor is likely to be about 1. This means that over time the capital will tend to fluctuate up and down, but is unlikely to grow appreciably.

Suppose now to invest using the $(1/4, 0, 0)$ portfolio. In the case of a favorable outcome, the capital grows by a multiplicative factor 3/2; in the case of an unfavorable outcome, the multiplicative factor is 3/4. Since the two outcomes are equally likely, the average multiplicative factor over two periods is $(3/2)\,(3/4) = 9/8$. Therefore, the average multiplicative factor over one period is $\sqrt{(9/8)} \approx 1.06066$. With this strategy, your money will grow, on average, by over 6% per period.

Ex. 3.4 Prove that this is the highest rate of growth that you can attain using a $(k, 0, 0)$ portfolio with k in $[0, 1]$.

Ex. 3.5 Prove that a fixed-proportions strategy investing in a portfolio $(\alpha_1, \alpha_2, \alpha_3)$ with $\min_i \alpha_i > 0$ and $\max_i \alpha_i < 1$ guarantees that ruin cannot occur in finite time.

3.3 The log-optimal growth strategy

The example is representative of a large class of investment situations where a given strategy leads to a random growth process. For each period $t = 1, 2, \ldots$, let X_t denote the capital at period t. The capital evolves according to the equation

$$X_t = R_t X_{t-1}, \tag{5}$$

where R_t is the random return on the capital. We assume that the random returns R_t are independent and identically distributed.

In the general capital growth process, the capital at the end of n trials is

$$X_n = (R_n R_{n-1} \ldots R_2 R_1)\, X_0.$$

After a bit of manipulation, this gives

$$\log\left(\frac{X_n}{X_0}\right)^{1/n} = \frac{1}{n}\sum_{t=1}^{n}\log R_t.$$

Let $m = E(\log R_1)$. Since all R_t's are independent and identically distributed, the law of large numbers states that the right-hand side of this expression converges to m as $n \to +\infty$ and therefore

$$\log\left(\frac{X_n}{X_0}\right)^{1/n} \to m$$

as well. That is, for large values of t, X_t is asymptotic to $X_0 e^{mt}$. Roughly speaking, the capital tends to grow exponentially at rate m.

It is easy to check (please, do it) that $m + \log X_0 = E(\log X_1)$. Thus, if we choose the utility function $U(x) = \log x$, the problem of maximizing the growth rate m is equivalent to finding the strategy that maximizes the expected value of $EU(X_1)$ and applying this same strategy in every trial. Using the logarithm as a utility function, we can treat the problem as if it were a single-period problem and this single-step view guarantees the maximum growth rate in the long-run.

3.4 Applications

a) The Kelly criterion. Suppose that you have the opportunity to invest in a prospect that will either double your investment or return nothing. The probability of the favorable outcome is $p > 1/2$. Suppose that you have an initial capital of X_0 and that you can repeat this investment many times. How much should you invest each time to maximize the rate of growth of the capital?

Let α be the proportion of capital invested in each period. If the outcome is favorable, the capital grows by a factor $1+\alpha$; if it is unfavorable, the factor is $1-\alpha$. In order to maximize the growth rate of his capital, you just need to maximize $m = p\log(1+\alpha)+(1-p)\log(1-\alpha)$ to find the log-optimal value $\alpha^* = 2p - 1$.

This situation resembles the game of blackjack, where a player who mentally keeps track of the cards played can adjust his strategy to ensure (on average) a 50.75% chance of winning a hand. With $p = .5057$, $\alpha^* = 1.5\%$ and thus $e^m \approx 1.01125$, which gives an (expected) .00125% gain each round.

b) Volatility pumping. Suppose that there are only two assets available for investment. One is a stock that in each period doubles or halves your capital with equal probability. The other is a risk-free bond that just retains value — like putting money under the mattress. Neither of these investments is very exciting. An investment left in the stock will have a value that fluctuates a lot but has no overall growth rate. The bond clearly has no growth rate. Nevertheless, by using these two investments in combination, growth can be achieved!

Suppose that we invest α of our capital in the stock and $(1-\alpha)$ in the bond, with α in $[0,1]$. The expected growth of this strategy is

$$m = \frac{1}{2}\log(1+\alpha) + \frac{1}{2}\log\left(1 - \frac{1}{2}\alpha\right),$$

which is maximized at $\alpha^* = 1/2$. For this choice of α, $e^m \approx 1.0607$ and the growth rate of the portfolio is about 6% per period which significantly outperforms the palsy 0% average growth of each of the two assets.

The gain is achieved by using the volatility of the stock in a pumping action. Remember that the strategy says that no more (and no less) than 50% of the stock should go in the stock in each period. When the stock goes up in certain period, some of its capital gains are reinvested in the bond; when it goes down, additional capital is shifted from the bond to the stock. Capital is pumped back and forth between the two assets in order to achieve growth greater than either could provide alone. Note that this strategy follows automatically the dictum "buy low and sell high" by the process of rebalancing the investment in each period.

Ex. 3.6 Suppose that the two assets are stocks that in each period double or halve the capital with equal probability. Assume that each asset moves independently of the other. Prove that the optimal portfolio has $\alpha^* = 50\%$ and that the growth rate per period is about 11.8%.

c) Optimal growth portfolio. Let us go back to the example of Section 3.2. Consider all portfolio strategies $(\alpha_1, \alpha_2, \alpha_3)$, with $\alpha_i \geq 0$ for $i = 1, 2, 3$ and $\alpha_1 + \alpha_2 + \alpha_3 \leq 1$. What is the portfolio which achieves the maximum growth?

Under scenario s_1, the return is $R = 1 + 2\alpha_1 - \alpha_2 - \alpha_3$. Under scenario s_2, the return is $R = 1 - \alpha_1 + \alpha_2 - \alpha_3$. Under scenario s_3, the return is $R = 1 - \alpha_1 - \alpha_2 + 5\alpha_3$. To find an optimal portfolio, it suffices to maximize

$$m = \frac{1}{2}\log(1 + 2\alpha_1 - \alpha_2 - \alpha_3) + \frac{1}{3}\log(1 - \alpha_1 + \alpha_2 - \alpha_3) + \frac{1}{2}\log(1 - \alpha_1 - \alpha_2 + 5\alpha_3).$$

Ex. 3.7 Show that $(1/2, 1/3, 1/6)$ is an optimal portfolio. Prove that it is not unique, by checking that $(5/18, 0, 1/18)$ is also optimal.

Any optimal portfolio has a growth rate of about 6.99%, higher than what was found above.

d) Continuous lotteries. There is one stock that can be purchased at a price of Euro 100 per share. Its anticipated price in one year is uniformly distributed over the interval $[30, 200]$. What is the log-optimal strategy over this period?

The log-optimal strategy is found by maximizing

$$m = \int_{.3}^{2} [\log(1 - \alpha + \alpha r)] \left(\frac{10}{17}\right) \mathrm{d}r.$$

We find that the log-optimal investment in the stock is $\alpha^* \approx .63$, which gives a growth rate of about 4.82% per period.

3.5 Excursions

An extensive list of several properties of the log-optimal strategy is given in MacLean et alii (1992). This paper and Li (1998) discuss the trade-offs between growth and security in multiperiod investment analysis.

References

[1] Y. Li (1993), "Growth-security investment strategy for long and short runs", *Management Science* **39**, 915–924.

[2] D.G. Luenberger (1998), *Investment Science*, New York: Oxford University Press.

[3] L.C. MacLean, W.T. Ziemba and G. Blazenko (1992), "Growth versus security in dynamic investment analysis", *Management Science* **38**, 1562–1585.

[4] L.M. Rotando and E.O. Thorp (1992), "The Kelly criterion and the stock market", *American Mathematical Monthly* **99**, 922–931.

4. Risk aversion and mean-variance preferences

4.1 Risk attitude

Consider a choice problem among the following three lotteries, whose expected values are written by their name.

α (480)		β (525)		γ (500)	
payoff	prob.ty	return	prob.ty	return	prob.ty
480	100%	850	50%	1000	50%
		200	50%	0	50%

If we were to base our choices on the expected value, β would be our preferred choice. However, there are people who would rather pick α (which on average pays less) on the ground that is less "risky" or maybe γ on the ground that it is an even riskier choice.

What can we say about the elusive notion of "risk"? While it is hard to define what exactly "risk" means, certainly a sure outcome like α should be deemed riskless. (And avoiding risk justifies its choice in the example above.) By contrast, let us call risky any lottery which does not yield a sure outcome. Although this is a poor distinction, it recognizes that risk lies in the unpredictability of the resulting outcomes.

However, avoiding risk cannot be the only reason driving a choice. If we use the expected value of a lottery as a benchmark to measure its "return", we should argue that β is a more "profitable" choice (on average). An agent who chooses α is avoiding risk at the cost of accepting a lower "return". On the contrary, an agent who always pick the lottery with the highest expected value is not affected by risk considerations.

Def. 4.4 An agent is risk neutral if he evaluates lotteries by their expected value.

If an agent is not risk neutral, he may be repelled by or attracted to risk: think of people buying insurance or respectively playing lotto. How do we tell which case is which? There is a simple question we may ask the agent: suppose you possess lottery β; how much sure money would you ask for selling it? We name this "price" $c(\beta)$ called by the agent the *certainty equivalent* of a lottery.

If the agent does not care about risk, he should be willing to exchange β for a sum equivalent to its expected value. That is, we should have $c(\beta) = E(\beta)$. We can deviate from equality in two directions. If $c(\beta) < E(\beta)$, the agent values β less than his expected value: therefore, the risk in β reduces the value to him of holding it — we say that he is *risk averse*. Vice versa, if $c(\beta) > E(\beta)$, he is *risk seeking*.

For a different way to say exactly the same thing, define the *risk premium* $r(\beta)$ of a lottery β as the difference between its expected value and its certainty equivalent. A risk averse agent is willing to forfeit the risk premium $r(\beta)$ in order to replace the "risky" lottery β by the sure outcome $c(\beta) < E(\beta)$. That is, he is willing to accept for sure a payment which is less than the average payoff of β in order to get rid of the risk associated with β.

4.2 Risk attitude and expected utility

All of this holds in general, even if the agent is not an expected utility maximizer. However, in the special case of expected utility maximizers, there exists a simple criterion to recognize whether an agent is risk averse, neutral or seeking.

THM. 4.5 An expected utility maximizer is risk neutral (resp., averse or seeking) if his utility function is linear (resp., concave or convex).

Thus, while the increasing monotonicity of the utility function speaks about the greediness of the agent, its curvature tells us something about his attitude to risk.

EX. 4.8 Check that expected utility can rationalize any of the three choices in the example above using different utility functions. If an expected utility maximizer has a utility function $u_1(x) = x$ he prefers β; if it is $u_2(x) = \sqrt{x}$ he prefers α; and if it is $u_3(x) = x^2$ he prefers γ. This is evidence of the flexibility of the expected utility model.

Here is a simple application. There are two assets. One is a riskless bond that just retains its value and pays 1 per euro invested. The other is a risky stock that has a random return of R per euro invested; we assume that $E(R) > 1$ so that on average the stock is more profitable than the riskless bond. Suppose that an agent is risk-averse and maximizes the expected value of a (concave and strictly increasing) utility function u over returns. The agent must select a portfolio and invest a fraction α of his wealth in the risky asset and a fraction $1 - \alpha$ in the riskless bond. Short-selling is not allowed and thus α is in $[0, 1]$.

The maximization problem is $\max_\alpha Eu(\alpha R + 1 - \alpha)$. Risk aversion implies that the objective function is concave in α (can you prove it?). Therefore, the optimal portfolio satisfies the first-order Kuhn-Tucker condition:

$$E\left[(R-1)u'(\alpha R + 1 - \alpha)\right] \begin{cases} = 0 & \text{if } 0 < \alpha < 1 \\ \leq 0 & \text{if } \alpha = 0 \\ \geq 0 & \text{if } \alpha = 1 \end{cases} .$$

Since $E(R) > 1$, the first-order condition is never satisfied for $\alpha = 0$. Therefore, we conclude that the optimal portfolio has $\alpha^* > 0$. That is, if a risk is actuarially favorable, then a risk averter will always accept at least a small amount of it.

4.3 Mean-variance preferences

There exist alternative approaches to the formalization of risk. One that is very common relies on the use of indices of location and dispersion, like mean and standard deviation. The expected value is taken as a measure of the (average) payoff of a lottery. Risk, instead, is present if the standard deviation (or some other measure of dispersion) is positive. The preferences of the agent are represented by a functional $V(\mu, \sigma)$, where μ and σ are respectively the expected value and the standard deviation of the lottery.

If offered several lotteries with the same standard deviation, a (greedy) agent prefers the one with the highest expected value. If offered several lotteries with the same expected value, a risk averse agent prefers the one with the lowest variance. Thus, a greedy and risk averse agent has preferences represented by a functional $V(\mu, \sigma)$ which is increasing in μ and decreasing in σ.

While intuitively appealing, this approach postulates that the agent dislikes any kind of positive standard deviation. It turns out that this is not consistent with the definition of risk aversion given above. Therefore, the so-called "mean-variance" preferences are in general incompatible both with the standard definition of risk aversion and in particular with the expected utility model.

Ex. 4.9 Suppose that a risk averse agent is an expected utility maximizer with utility function $u(x) = x$ for $x \leq 1$ and $u(x) = 1 + 0.1(x - 1)$ for $x > 1$. Compare a lottery α offering a payoff of 1 with probability 1 versus another lottery β offering a payoff of 1.1 with probability 10/11 and 0 with probability 1/11. While $\mu(\alpha) = \mu(\beta)$ and $\sigma(\alpha) = 0 < \sigma(\beta)$, the agent strictly prefers β to α.

If one wishes to relate this approach with the expected utility model, the best she can do is to view it as a crude approximation. Given an arbitrary value x^*, consider the Taylor expansion of u around x^*:

$$u(x) = u(x^*) + \frac{(x - x^*)}{1!}u'(x^*) + \frac{(x - x^*)^2}{2!}u''(x^*) + \frac{(x - x^*)^3}{3!}u'''(x^*) + \dots \tag{6}$$

The "mean-variance" approach ignores the third and all successive terms in the Taylor expansion of the utility function. Or, in more statistical terms, it looks only at the first two moments of the probability distribution of the lotteries.

Consider the two lotteries

α		β	
payoff	prob.ty	return	prob.ty
-1	0.999	1	0.999
999	0.001	-999	0.001

Since they have the same mean and the same standard deviation, the second-order approximation cannot tell them apart. However, most people is not indifferent between the two.

If we want to distinguish α from β, we need to reach the third term in (6), which represents skewness. Skewness is zero for a symmetric distribution; it is positive if there is a hump on the left and a long thin tail on the right; and negative in the opposite case. So α is positively skewed and β is negatively skewed. Most commercial lotteries and games of chance are positively skewed: if people like them because of this, the second-order approximation cannot capture their preferences.

Even if crude, the second-order approximation may be justified by two different kinds of assumptions: either (i) the utility function has a special form (namely, it is quadratic and therefore the approximation is correct); or, (ii) the probability distributions belong to a particular family which is completely characterized by the first two moments (for instance, they are normal).

4.4 Risk attitude and wealth

The risk attitude is especially studied with reference to choices involving the wealth of an agent. To keep things simple, assume in the following that all risky decisions concern monetary payoffs and that agents maximize expected utility. The utility functions are defined over the positive reals and, whenever necessary, they are twice differentiable with a strictly positive first derivative.

We consider one way in which wealth affects the risk attitude. Suppose that an agent with current wealth w must choose between a risky lottery α and a sure outcome b. If he chooses α, his future wealth will be $w + \alpha$; if he chooses b, $w + b$. Suppose that at the current level of wealth the agent prefers the risky α to the sure b. If he is an expected utility maximizer, this implies that $Eu(w + \alpha) > u(w + b)$.

DEF. 4.6 An agent has decreasing risk aversion (with respect to wealth) whenever, for arbitrarily given α and b, $Eu(w + \alpha) > u(w + b)$ implies $Eu(w' + \alpha) > u(w' + b)$ if $w' > w$.

Similar definitions holds for constant and increasing risk aversion. There exists a simple criterion to recognize whether an agent has decreasing, constant or increasing risk aversion.

THM. 4.7 An agent has decreasing risk aversion (resp., constant or increasing) if and only if his coefficient of (absolute) risk aversion

$$\lambda(x) = -\frac{u''(x)}{u'(x)}$$

is decreasing (resp., constant or increasing) in x.

The standard assumption in economic theory is that agents are risk averse and have decreasing risk aversion. However, in applications it is very common to postulate that they are risk neutral or that they have constant risk aversion because this greatly simplifies the choice of their utility function.

THM. 4.8 The only utility functions with constant absolute risk aversion are the linear utility function $u(x) = x$ which has $\lambda(x) = 0$, and the exponential utility function $u(x) = -\text{sgn}(k)e^{-kx}$ which has $\lambda(x) = k$.

EX. 4.10 Suppose that the agent is an expected utility maximizer with a constant coefficient of absolute risk aversion $k > 0$. The choice set contains only lotteries normally distributed. Given a lottery $X \sim N(\mu, \sigma)$, check that the preferences of the agent can be represented by the functional $V(\mu, \sigma) = \mu - (1/2)k\sigma^2$.

4.5 Risk bearing over contingent outcomes

Suppose that there is a finite number of states of the worlds (or scenarios). Each state s_i ($i = 1, 2, \ldots, n$) occurs with probability π_i. There exists a single commodity, which has a price p_i in state s_i. The agent is endowed with the same initial income y in each scenario and he derives a differentiable utility $u(c_i)$ from consuming a quantity c_i of the commodity in the scenario s_i. When the agent is an expected utility maximizer, he chooses his state contingent consumption by solving

$$\max \sum_i \pi_i u(c_i) \qquad \text{s.t.} \sum_i p_i c_i = y.$$

Assuming an interior solution, the first-order condition requires $[\pi_i u'(c_i)]/p_i$ to be constant for all i. This result states that the risk-bearing optimum has the same expected marginal utility per dollar of income in each and every state and is known as the fundamental theorem of risk-bearing.

References

[1] J. Hirshleifer and J.G. Riley (1992), *The Analytics of Uncertainty and Information*, Cambridge University Press.

[2] D. Kreps (1988), *Notes on the Theory of Choice*, Westview Press.

[3] A. Mas-Colell, M.D. Whinston and J.R. Green (1995), *Microeconomic Theory*, Oxford University Press.

5. Information structures and no-trade theorems

5.1 Introduction

One traditional view about trading in financial markets is that this has two components: liquidity and speculation. Some people trade because they need the liquidity (or have other pressing demands from the real economy); others trade because they have asymmetric information and hope to profit from it. According to this view, high volume trading should be explained mostly by differences in information among traders. See for instance Ross (1989):

> "It is difficult to imagine that the volume of trade in security markets has very much to do with the modest amount of trading required to accomplish the continual and gradual portfolio balancing inherent in our current intertemporal models. It seems clear that the only way to explain the volume of trade is with a model that is at one and at the same time appealingly rational and yet permits divergent and changing opinions in a fashion that is other than ad hoc."

This lecture illustrates a few theoretical results showing that in fact asymmetry in information alone is not sufficient to stimulate additional trade. In fact, there are even cases where it might reduce the trading volume and lead to market breakdowns. This family of results are known as "no-trade" or "no-speculation" theorems.

5.2 The Dutch book

The common wisdom about trading motivated by asymmetric information is that people with different beliefs can concoct mutually beneficial trades. In fact, even more is true: when two or more risk-neutral agents have different beliefs about the probability of some events and are willing to bet among them on these events, a bookie can arrange a set of bets that each of the agents is willing to take and still guarantee himself a strictly positive expected profit. This phenomenon is known as a Dutch book.

Here is a simple illustration. There are two agents: Ann and Bob. Suppose that Ann believes that in a month the Mib30 will be up with probability p while Bob believes that this will happen with probability $q < p$. Neither one can be absolutely certain about the direction of the market, so we assume $q \neq 0$ and $p \neq 1$.

The bookie can make a sure profit by offering to each player a specific bet customized for his beliefs. The following table gives the bets' payoffs for each of the two possible events: the index goes up or down. Given any $x > 0$, the bookie can cash a sure (and strictly positive) profit of $x - \varepsilon$, where $0 < \varepsilon < x$ is the "sweteener" that induces agents to take the bets. While reading the table, recall that the bookie's payoffs are the opposite of the sum of agents' payoffs.

	Bull	Bear
Ann's bet	$x\dfrac{1-p}{p-q}+\varepsilon$	$-x\dfrac{p}{p-q}$
Bob's bet	$-x\dfrac{1-q}{p-q}$	$x\dfrac{q}{p-q}+\varepsilon$
bookie's bet	$x-\varepsilon$	$x-\varepsilon$

Ann's uses her own beliefs to assess her bet. She estimates that the bet has an expected value of

$$p\left[x\frac{1-p}{p-q}+\varepsilon\right]+(1-p)\left[-x\frac{p}{p-q}\right]=p\varepsilon>0$$

and therefore is willing to take it. Bob, using his own beliefs, reaches a similar conclusion. The bookie, regardless of his own beliefs, makes money in each and every state. Everybody is happy (a priori) but the bookie is setting up an arbitrage to his advantage. Hence, he must be exploiting somehow at least one of the two agents. Who is being exploited?

Well, if two agents hold different beliefs about the probability of an event, at least one of the two must be "wrong". The bookie, even though he himself may have no idea who is wrong, is simply exploiting this fact. Differing beliefs imply that there is an arbitrage opportunity lurking somewhere, and agents who recognize this fact will forfeit trading and revise their beliefs for fear of being exploited. This is the basic intuition behind the no-trade theorems. It also motivates the wide-spread buzzword that there is only one "market probability" to which agents' beliefs should conform to; otherwise, so goes the argument, they could be exploited.

Ex. 5.11 There are three risk-neutral agents: Ann, Bob and Carol. Consider three events: in a month the Mib30 will go down of more than 5%, or it will be up of more than 5%, or it will move less than 5%. Call these three events respectively E_1, E_2 and E_3. Ann believes that E_i occurs with probability p_i, for $i = 1, 2, 3$. Similarly, denote Bob' beliefs with q and Carol's beliefs with r. Assume $0 < p_3 = q_1 = r_2 < p_2 = q_3 = r_1 < p_1 = q_2 = r_3 < 1$. Find a Dutch book. (Warning: solving this is longer than usual.)

5.3 The red hats puzzle

We should turn the intuition above into a formal argument. The first step is to appreciate that there exist different levels of knowledge of a fact. We illustrate this with a puzzle. An evil King decides to grant an amnesty to a group of prisoners who are kept incommunicado in the royal dungeons. Summoning the prisoners to his presence and commanding them under penalty of death not to look upward, the King has a hat placed upon the head of each one. Two hats are red; the rest are white. Each prisoner can see the hat of his fellow mates, but not his own.

The King speaks to all summoned prisoners: "Most of you wear white hats. *At least one of you is wearing a red had.* Every day, at dawn, you will be summoned here and any of you can attempt to guess the colour of your own hat. If you guess correctly, you will go free; if you guess incorrectly, you will be instantly beheaded."

How many days would it take the two red-hatted prisoners to infer the colour of their hats? On the first day, each prisoner with a white hat sees two red hats and each prisoner with a read hat sees only one. Since they all know what the King has told them, none can conclude anything on the colour of his own hat. Thus, none will dare to guess.

On the second day, the white-hatted prisoners still see two red hats and the two red-hatted prisoners still see one red hat. Now, the two red-hatted prisoners (named Ann and Bob) learn a new piece of information. One day is passed and neither Ann nor Bob have spoken and left. Because they were both present at the King's speech, Ann knows that Bob knows that there is at least one red hat. If he had not already left, it is because he must have seen a red hat. For if he had not, Bob would have seen only white hats and thus would have correctly inferred that the only possible red hat was perched atop his own head. "Hence, — Ann reasons, — since the only red hat *I* can see is on Bob's head, that which Bob must have seen on the first day must have been on mine." Bob, having the same information set and abilities as Ann, will arrive at the same conclusion. Both will infer their hats to be red on the second day, and will go free.

This puzzle establishes that there are different levels of knowledge of a specific fact. Consider the proposition "everyone sees a red hat". On the first day, Ann sees Bob's red hat (and vice versa) but does not know whether Bob also sees one (and vice versa). Thus on the first day the proposition is actually true but not everyone knows that it is true. It takes another day to make the proposition "everyone knows that everybody sees a red hat" become true.

The highest-level proposition is "everybody knows that everybody knows that everybody knowns ... that [everyone sees a red hat]". Independently of which event you put within brackets, this is an instance of what is called "common knowledge" of an event. Not only everybody knows the event, but he knows also that others know it, and so on.

The puzzle shows that sometimes the difference between staying in prison and going free rests with understanding the difference. Similarly, sometimes making profits in a financial market rests on the capacity to understand the difference between the proposition "Autostrade trades at a price of 7" and "everybody knows that Autostrade trades at a price of 7". Since trading markets are public, both propositions are true; but the second proposition conveys more information than the first one.

Ex. 5.12 Convince yourself that, if the red hats had been three, the proposition "everyone knows that everybody sees two red hats" becomes true on the third day.

Ex. 5.13 Prove that, if the red hats had been k, then k prisoners would have gone free on the k-th day.

5.4 Different degrees of knowledge

There is a set Ω of states. Each agent has an *information map* P that associates with each state ω in Ω a nonempty subset $P(\omega)$ of Ω. When the true state is ω, the agent thinks that all (and only all) the states in $P(\omega)$ may be true; in other words, the agent knows only that the true state belongs to the set $P(\omega)$. An event E is a subset of Ω. An agent for whom $P(\omega) \subseteq E$ knows, in state ω, that some state in the event E has occurred or, more simply, he knows that E must be true.

Let us consider an example. There are two risk-neutral agents, Ann and Bob. Momentarily, we are going to consider the following bet, based on the value of θ: if $\theta = 0$, Bob must pay Ann Euro 1; if $\theta = 1$, Ann must pay Bob Euro 1. In case of indifference, assume that either agent prefers to turn down the bet. Hence, Ann bets only if she thinks that the probability of $\theta = 0$ is strictly greater than 1/2 and Bob bets only if she thinks that it is strictly lower than 1/2. Trade is possible only if agents have different beliefs.

There are five states $\omega_1, \ldots, \omega_5$ and the pairs (θ, ω) are drawn from the following prior joint distribution, which is known to both agents:

	$\theta = 0$	$\theta = 1$
ω_1	0.20	0.05
ω_2	0.05	0.15
ω_3	0.05	0.05
ω_4	0.15	0.05
ω_5	0.05	0.20

Ann and Bob have access to different information, which is represented by their information maps. We assume that the information maps are *partitional*; that is: (i) $\omega \in P(\omega)$; and (ii) $\omega' \in P(\omega)$ implies $P(\omega') = P(\omega)$. A partitional information map may be specified by the information partition it induces.

Ex. 5.14 Convince yourself that (i) makes sure that the agent is never certain that the state is different from the true state. Show that, if (ii) were false, in state ω the agent could infer that the state is not ω': therefore, ω' could not belong to $P(\omega)$.

Continuing the example, Ann's information partition is $\mathcal{P}_A = (\{\omega_1, \omega_2\},\ \{\omega_3, \omega_4\},\ \{\omega_5\})$. Bob's information partition is $\mathcal{P}_B = (\{\omega_1\},\ \{\omega_2, \omega_3\},\ \{\omega_4, \omega_5\})$. Since the two information maps are different, there may be states where agents entertain different posterior beliefs and are thus both willing to undertake the bet.

For instance, consider state ω_2. By her map, Ann knows only that either ω_1 or ω_2 are true. Using Bayes' rule, she attributes probability 5/9 to $\theta = 0$. By a similar reasoning, Bob knows only that either ω_2 or ω_3 are true and he attributes probability 1/3 to $\theta = 0$. Therefore, in state ω_2, agents have different beliefs and are willing to subscribe the bet. Trade can take place. Consider now state ω_5. Now, Ann attributes probability 1/5 to $\theta = 0$ and is not willing to bet. In this state, no trade can take place.

Ex. 5.15 Compute agents' posterior probabilities in all states.

Using the results of Exercise 5.15, it is easy to show that players' disposition to bet in each state is the following.

State	ω_1	ω_2	ω_3	ω_4	ω_5
Ann	Bet	Bet	Bet	Bet	No bet
Bob	No bet	Bet	Bet	Bet	Bet

It seems that trade should occur in states ω_2, ω_3 and ω_4. But this would be true only if Ann and Bob are naive and have not digested the meaning of the red hats puzzle. To see why, consider again state ω_2.

As before, Ann initially knows only that either ω_1 or ω_2 are true. However, by computing the table above, she knows only that Bob is not willing to bet in state ω_1. Therefore, if she hears Bob proposing to take his part of the bet, she knows that the state cannot be ω_1. But then it must be exactly ω_2, in which case the probability that $\theta = 0$ is $1/4$ and therefore Ann should not accept the bet. The mere fact that Bob is willing to trade suffices to alert Ann to something that make her revise her beliefs and lead her to refuse the bet.

Ex. 5.16 By a similar argument, prove that no trade can occur in state ω_4. Having established that no trade should occur in states $\{\omega_1, \omega_2, \omega_4, \omega_5\}$, apply another round of rationality and prove that no trade occurs in state ω_3 either.

Trade is impossible when everybody behaves rationally, even if Ann and Bob have different information maps! The public information revealed by their willingness of trade realigns their probabilities and make them give up trade.

More intuitively, here is the gist of the argument. If I know you are so smart, why should I buy a piece of risky stock from you? If you are willing to sell it at a price p, you must think that its value is less than p. Since you are no fool, I better take this into account and revise my estimate of the stock downwards. As far as you are willing to sell, you must know something that I do not know: then I better be careful and avoid being a sucker, so I won't buy.

5.5 Can we agree to disagree?

In the example above, only two rounds of deductive thinking are needed to reach the conclusion that no trade is preferable. Similar techniques can be used to obtain similar results after one formally introduces the notion of common knowledge.

Restricting for simplicity to the case of two individuals, here is how to formalize it. Let P_A and P_B the information maps of Ann and Bob. An event F is *self-evident* if for all ω in F we have $P_i(\omega) \subseteq F$ for $i = A, B$. An event E is *common knowledge* in state ω if there exists a self-evident event F for which $\omega \in F \subseteq E$.

Ex. 5.17 Let $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$. Suppose that the information partitions induced by P_A and P_B are respectively

$$\begin{aligned} \mathcal{P}_A &= \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4, \omega_5\}, \{\omega_6\}\}, \\ \mathcal{P}_B &= \{\{\omega_1\}, \{\omega_2, \omega_3, \omega_4\}, \{\omega_5\}, \{\omega_6\}\}. \end{aligned}$$

Prove that in any state the event $E = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ is not common knowledge. Prove that the event $E = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5\}$ is common knowledge in any state but ω_6.

Ok, enough with abstract theory. People do not actually talk about information maps and states. But traders use probabilities. So let us relate these notions. Suppose that an agent has a prior probability distribution $p(\cdot)$ on Ω: for instance, he attaches probability $p(E)$ to an event E. Given his information map, in state ω the agent will know that only the states in $P(\omega)$ can occur and therefore will update his probability distribution to a posterior $p(\cdot|P(\omega))$. If $E \cap P(\omega) \neq \emptyset$, it will be assigned a probability $p(E|P(\omega))$; otherwise, it will be deemed impossible.

Thus information maps and states can be used to describe how people update their beliefs in the wake of new information. It turns out that a common prior and common knowledge of the posteriors force people to have the same posterior probabilities for that event.

THM. 5.9 Suppose that Ω is finite and that Ann and Bob share the same prior. If the information maps are partitional and in state ω^* it is common knowledge that each agent $i = 1, 2$ assigns probability π_i to some event E, then $\pi_1 = \pi_2$.

Proof: The assumptions imply that there is a self-evident event F containing ω^* which is a subset of $\{\omega : p(E|P_i(\omega)) = \pi_i\}$ for $i = 1, 2$. The event F must be a union of members of i's information partition. Since Ω is finite, so is the number of sets in each union; thus let $F = \cup_k A_k = \cup_k B_k$. Given two nonempty disjoint sets C and D with $p(E|C) = p(E|D) = \pi_i$, we have $p(E|C \cup D) = \pi$. Since for each A_k we have $p(E|A_k) = \pi_1$, it follows that $p(E|F) = \pi_1$; similarly for $p(E|F) = \pi_2$. Hence $\pi_1 = \pi_2$. □

Following the same logic, one can show that if two agents with partitional information maps have the same prior it can be common knowledge that they assign different probabilities to some event. However, it *cannot* be common knowledge that the probability assigned by Ann exceed that assigned by Bob. This makes it impossible to stipulate some purely speculative bets.

For instance, suppose that Ann and Bob are strictly risk averse. If Ann accepts an even-odds bet that a coin will come up heads, she reveals that she is assigning "heads" a probability greater than 1/2. If Bob accepts to bet on "tails", he reveals that he is assigning "heads" a probability lower than 1/2. The theorem says that it cannot be common knowledge that Ann' posterior probability for heads is different from Bob's posterior.

Let us work out a second example where common knowledge is needed to reach the no-trade outcome. Ann and Bob are considering to set up the same bet as before. However, this time there is a countable number of states $\omega_0, \omega_1, \omega_2, \ldots$ and the information partitions are

$$\mathcal{P}_A = \Big(\{\omega_0\},\ \{\omega_1, \omega_2\},\ \{\omega_3, \omega_4\},\ \{\omega_5, \omega_6\},\ \{\omega_7, \omega_8\},\ \ldots\Big)$$

and

$$\mathcal{P}_B = \Big(\{\omega_0\},\ \{\omega_1\},\ \{\omega_2, \omega_3\},\ \{\omega_4, \omega_5\},\ \{\omega_6, \omega_7\},\ \ldots\Big).$$

Note that in any state each player (at worst) is unsure which of two states hold. Therefore, ruling out one of the two candidates enables an agent to correctly infer the true state.

For some ε in $(0, 1/32)$, the prior joint distribution on (θ, ω) is the following.

	$\theta = 0$	$\theta = 1$
ω_0	$1/4$	$(1/4) + (2/3)\varepsilon$
ω_1	$1/8$	$(1/8) - \varepsilon$
ω_2	$1/16$	$(1/16) + (1/2)\varepsilon$
ω_3	$1/32$	$(1/32) - (1/4)\varepsilon$
ω_4	$1/64$	$(1/64) + (1/8)\varepsilon$
$\vdots$	$\vdots$	$\vdots$
ω_n	$\frac{1}{2^{n+2}}$	$\frac{1}{2^{n+2}} + (-1)^n \frac{1}{2^{n-1}}\varepsilon$
$\vdots$	$\vdots$	$\vdots$

Ex. 5.18 Verify that

$$\sum_{n=0}^{+\infty} \frac{1}{2^{n+2}} = \frac{1}{2} \quad \text{and} \quad \sum_{n=1}^{+\infty} (-1)^n \frac{1}{2^{n-1}} = -\frac{2}{3}$$

and check that this is a well-defined probability distribution.

In state ω_0, even a naive Ann does not want to trade; similarly, in state ω_1, even a naive Bob does not want to trade. In state ω_2, Ann needs a second round of deductive thinking to know that the state is not ω_1 (therefore it must be ω_2) and hence to decline trade. The argument is identical to the previous example: the crucial step is that Ann can deduce that Bob know that the state is not ω_1. For short, Ann knows that Bob knows that the state is not ω_1.

Going further, in state ω_3 Bob needs a third round of deductive thinking to conclude that the state is not ω_2 (therefore it must be ω_3) and to decline trade. The crucial step is that Bob must be able to deduce that Ann knows that the state is not ω_2. To do so, he needs to replicate the reasoning above and therefore he needs to know that in state ω_2 Ann knows that Bob knows that the state is not ω_1.

Continuing like this, it turns out that it takes n rounds for Ann (if n is even) or Bob (if n is odd) to deduce that the state is unfavorable and hence decline to trade. For no trade to occur in state n, the agents must undertake a chain of n deductions. Since the number of states is infinite, for no trade to occur in any state, they must run an infinite chain of deductions. This infinite chain requires the assumption of common knowledge.

5.6 No trade under heterogenous priors

The no-speculation results are fairly robust. Here is an example with heterogenous priors. Since priors are different, common knowledge is not an impediment to trade. However, this time trade fails because of information asimmetry.

Ann and Bob are risk-neutral agents who are considering a bet which depends on whether the interest rate will go up or down tomorrow. Ann gets a binary signal which she can reveal to Bob. They have different priors, so either player interprets the signal differently. When Ann gets signal u, she assigns probability p_1^u to a rise in the interest rate; when she gets signal d, she assigns probability p_1^d to a decrease in the interest rate. If Bob is revealed signal u (respectively, d), he assigns probability p_2^u (respectively, p_2^d) to the corresponding event.

Gains from trade are possible if $p_1^u \neq p_2^u$ or $p_1^d \neq p_2^d$. For example, suppose $p_1^u = 2/3$, $p_1^d = 1/3$, $p_2^u = 3/4$ and $p_2^d = 1/4$. Then Ann could agree to reveal the signal to Bob with the understanding that she would pay her 2 if the signal is correct and be paid 5 from him if it is incorrect. For either possible signal, Ann expects an average gain of $(2/3)(-2) + (1/3)(+5) = 1/3$.

On the other hand, if Ann reveals him her signal, Bob can expect for either signal an average gain of $(3/4)(+2) + (1/4)(-5) = 1/4$ if the signal is truthful and $(3/4)(-5) + (1/4)(+2) = -13/4$ if it is not. Since Bob can secure 0 by not trading, he is willing to buy Ann's signal only if she reports truthfully.

Unfortunately, Ann can gain more by misreporting his signal. If she lies and always report the wrong signal, she has an expected gain of $(2/3)(+5) + (1/3)(-2) = 8/3$, which is greater than her gain from reporting truthfully. Even if Ann prefers being truthful and sell her signal to no trade, once the deal is signed she gains more by lying. Knowing this, Bob cannot trust Ann and does not buy her signal. So trade breaks down due to informational asimmetry: if Bob could monitor the truthfulness of Ann's report, trade would resume.

References

[1] P. Milgrom and N. Stokey (1982), "Information, trade and common knowledge," *Journal of Economic Theory* **26**, 17–27.

[2] S. Morris (1994), "Trade with heterogenous beliefs and asymmetric information," *Econometrica* **62**, 1327–1347.

[3] M.J. Osborne and A. Rubinstein (1994), *A course in game theory*, Cambridge (MA): The MIT Press.

[4] S. Ross (1989), "Intertemporal asset pricing", in: S. Bhattacharya and G. Constantinides (eds.), *Theory of valuation: Frontiers in modern financial theory*, Totowa (NJ): Rowman and Littlefield.

6. Herding and informational cascades

6.1 Introduction

Consider a situation where 100 people have to choose one of two movies shown at the same theater. The prior probability for A (respectively, B) being better is 51% (respectively, 49%). People arrive at the movie theater in sequence, observe the choices made by the people arrived before them, and decide for one or the other of the movies. Each agent knows the prior probability and, before going to the theater, receives a private (independent) signal about which movie is better. All signals are of the same quality.

Suppose that 99 people receive a signal favoring B and only one a signal favoring A. It is obvious that, if the agents could aggregate their information, everybody would choose B. Consider what could happen when aggregation of information is not possible because people arrive in sequence and do not talk to each other.

In particular, suppose that the only person who got an A-signal arrives first. Clearly, Primus will go to A. Secunda will now know that Primus got an A-signal, while she had a B-signal. Since the two signals are of equal quality, they cancel out, and the rational choice is to go by the prior probabilities and pick A. Note that if Secunda had got an A-signal, she would also have chosen A.

Therefore, Secunda chooses A regardless of her signal. But then her choice provides no information to the third agent in line. Facing exactly the same choice as Secunda, he will also choose A and so on. Everybody ends up watching movie A even though the aggregate information strongly suggests that B is better. (Food for thought: have you ever wondered why producers heavily advertise a movie only in its opening days?)

A group dynamics where most people coordinate on the same choice is called "herding". We may view some buying or selling frenzies as instances of herding behavior. This lecture illustrates what may rationally start and maintain this kind of imitative behavior, as well as which consequences herding may have on the valuation of financial assets.

6.2 Some terminology

Herding behavior is associated with people blindly following others' choices. The most likely cause is probably human stupidity (something never too rare). For our purposes, however, we are interested in situations where herding is the outcome of a rational choice.

Imitation can be rational if the predecessors' actions affects either the successors' payoffs or their beliefs (or both). In the first case, we say that there is a *payoff externality*; for example, think of the your last decision to upgrade to the ump-th version of Word (packed with many new frills absolutely worthless to you) only because the old version cannot read files prepared with the new version. In the second case, we say that there is an *information externality* as in the movies' example above.

In many cases, both types of externalities are present and they may oppose or reinforce each other. In the movies' example, for instance, I could have added the assumption that people dislike crowded shows. Then the information externality pushing for A-herding would have been opposed by a payoff externality driving later customers to choose the (seemingly slightly worse but) less crowded movie show.

Sequencing can be *exogenous*, and then the order of moves is given, or *endogenous*, in which case agents are allowed to decide when to move and whether to move first or to herd. The state space can be *discrete* or *continuous*, as well as the choice set. The movies examples had a discrete state space (A is better or B is better), a discrete choice set (watching A or B) and exogenous sequencing.

6.3 Public offerings and informational cascades

Consider a situation in which a seller is making an IPO through an underwriter. Suppose that market participants hold perfectly accurate information when aggregated and that the offer price must be uniform. Then we should expect overpriced offerings to fail and underpriced offerings to succeed.

However, if distribution channels are limited, the underwriter needs time to approach interested investors. If later investors can observe how well the offering has sold to date, herding based on information externalities is possible and the conclusion above may be upset.

Suppose that the issuer and the n investors are risk-neutral. Each share of the firm has a value θ, which is uniformly distributed in $[0, 1]$. Neither the issuer nor the investors know θ. The issuer is willing to accept any price $p \geq 0$ but wants to maximize revenue from the IPO. The issue has sufficient shares to offer one to each investor; each investor can only afford one share. Therefore, the investor's decision is simply whether to buy or abstain.

Each investor receives an independent binary signal, which is *H*igh with probability θ and *L*ow with probability $1 - \theta$. It can be shown that the posterior expected value of θ when a number k of H-signals have been observed out of n signals is

$$E(\theta | k \text{ out of } n) = \frac{k+1}{n+2}, \tag{7}$$

while the probability of observing k or more H-signals out of n signals is

$$P(\text{at least } k \text{ H-signals}) = 1 - [k/(n+1)].$$

We examine three different scenarios.

Perfect communication. In the first scenario investors can freely communicate. This is the benchmark case. By (7), the highest price that the issuer can charge (if k out of n investors receive an H-signal) is

$$p(k) = \frac{k+1}{n+2}. \tag{8}$$

Since the issuer wants to choose the price which maximizes his expected gain, he must choose a price equal to $p(k)$ for some $k = 0, 1, \ldots, n$. The higher the price, the higher the number of H-signals that must be observed to convince all investors to buy; therefore, the higher the price, the lower the probability that the IPO will go through.

When choosing a price $p(k)$, the issuer is implicitly betting on the event that at least k investors will get an H-signal. If the issuer charges the price $p(0) = 1/(n+2)$, the price is so low that all investors will buy independently of the number of H-signals observed. Hence the IPO will certainly go through and yield a sure profit of $n/(n+2)$.

On the other hand, if he charges $p(k)$ with $k > 0$, his expected gain is

$$P(\text{at least } k \text{ H-signals}) \cdot p(k) = \left(1 - \frac{k}{n+1}\right) \frac{k+1}{n+2}.$$

This is maximized for $k^* = n/2$. By (8), the optimal price is $p(k^*) = 1/2$ which yields an expected profit of $[n(n+2)]/[4(n+1)]$. Since this is greater than the profit $n/(n+2)$, $p^* = 1/2$ is the optimal price.

Ex. 6.19 Prove that: (i) the issuer never prices above the expected value of θ; (ii) no successful offerings are overpriced; (iii) offerings are at least as likely to succeed as they are to fail; and (iv) for large n, the expected profit of the issuer is at most 1/4 per share.

Perfect communication from early to late investors. In the second scenario each investor observes his own and the private information of the investors who were previously approached by the underwriter. Now the profits to the seller are path-dependent.

For example, suppose that $p = 2/5$ and $n = 2$. If the first investor gets an H-signal and the second investor an L-signal, the issuer gets a profit of $2 \cdot (2/5)$. Vice versa, if the first investor gets an L-signal and the second investor an H-signal, the issuer gets only 2/5. When the circulation of information is restricted, the profits to the seller may depend on the order in which signals are revealed to the investors.

Only actions are observable. In the third scenario, late investors knows only whether early investors purchased or not; that is, only actions are observable. This can start informational cascades. For instance, suppose that at some point investor i has an H-signal and finds that it is not in his interest to invest. Then each subsequent investor realizes that i's decision not to invest should not be interpreted as i having actually received an L-signal. Therefore, she will not be able to learn anything from i's action. This implies that investor $i+1$ faces the same choice (and has the same information) as investor i; so she will not invest regardless of her signal. By induction, so will all later investors.

Here is an example. Suppose that $p = 1/2 - \varepsilon$ and the first three investors observe (in this order) an H-signal and two L-signals. Primus believes that the value of the project is 2/3 so he purchases. Consequently, Secunda infers that the value of the project is 1/2 and she purchases as well. The third investor, regardless of his signal, cannot infer the signal of Secunda. Like her, he will purchase even with an L-signal. This produces a positive cascade: everybody purchases and the IPO goes through.

Consider now what happens if the first three investors observe (in this order) L, L and H. They infer share values of 1/3, 1/4 and 2/5 respectively. Regardless of his signal, the third investor does not purchase. Then all subsequent investors abstain as well, and a negative cascade ensues: nobody purchases and the IPO fails.

The example shows that the order in which people with different signals move can determine the success or the failure of the IPO. This can lead to striking situations. For instance, imagine the case where there are 100 investors: 99 have observed an L-signal and only one has had an H-signal. Under perfect communication, nobody should buy and the IPO should fail. However, if it so happens that the sequencing of moves has the investor with the H-signal in first place, everybody will purchase and an overpriced offering will be entirely sold.

THM. 6.10 Even with an infinite number of investors, for $p > 1/3$ there is positive probability that an offering of ultimate value $\theta < 1$ (and in particular an underpriced issue) fails completely (no investor purchases) because of a negative cascade. For $p < 2/3$, there is positive probability that an offering of ultimate value $\theta > 0$ (and in particular an overpriced issue) succeeds perfectly (all investors buy) because of a positive cascade.

Because of cascades, a well-informed market fails to aggregate information efficiently. Other phenomena that can be explained by this sort of models include: clumping on an incorrect decision, low consensus in opinions (when polled) but high uniformity in actions, fragility (a little bit of information can reverse a long-standing cascade), and strong dependence on initial conditions.

6.4 Excursions

Herding can arise in different ways. Welch (2000) distinguishes six major causes: i) utility interactions, ii) sanctions on deviants, iii) direct payoff externalities, iv) principal-agent externalities, v) informational externalities, and vi) irrational behavior. Devenow and Welch (1996) surveys models based on iii), iv) and v). Welch (2000) updates the references therein and provides empirical evidence about herding among financial analysts.

References

[1] A.V. Banerjee (1992), “A simple model of herd behavior,” *Quarterly Journal of Economics* **107**, 797–817.

[2] A. Devenow and I. Welch (1996), “Rational herding in financial economics,” *European Economic Review* **40**, 603–615.

[3] I. Welch (1992), “Sequential sales, learning, and cascades,” *Journal of Finance* **47**, 695–732.

[4] I. Welch (2000), “Herding among security analysts,” *Journal of Financial Economics* **58**, 369–396.

7. Normal-CARA markets

7.1 Introduction

Suppose that Primus is interested in taking actions that depend on an unknown parameter y. For instance, he has a utility function $u(y)$ which depends on the value of y. If Primus is a (subjective) expected utility maximizer, he should pick the action which maximizes the expected value of his utility computed with respect to his (subjective) probability distribution for y.

Suppose now that, before making any decision, Primus receives additional information about y. Then he should update his prior probability distribution to a *posterior* probability distribution and use this latter one to choose an optimal action. Updating beliefs when new information is available, therefore, is crucial for taking informed decisions.

This basic format for taking decisions under uncertainty when new information is revealed is the tenet of Bayesian rationality. Out of the many possible variations, the literature especially insists on a model where Primus has an exponential utility function and normally distributed random variables (both prior and posterior). This is made for analytical convenience. This lecture is meant to collect the basic mathematics necessary to deal with the *cara–normal* model.

7.2 Updating normal beliefs

Suppose that the parameter Y is normally distributed with mean m and standard deviation $s_y > 0$. For short, we write $Y \sim N(m, s_y)$. Imagine that Primus can receive signals about Y. Each signal x is independently and identically distributed according to a normal distribution with mean y and standard deviation $s_x > 0$; that is, a signal is an (iid) draw from $X \sim N(y, s_x)$.

We wonder what should be Primus' posterior distribution for Y after having observed one signal X. If we denote by $g(y)$ the prior distribution for Y and by $f(x|y)$ the conditional distribution of the signal, we have respectively

$$g(y) = \frac{1}{\sqrt{2\pi} s_y} \exp\left[-\frac{1}{2s_y^2}(y-m)^2\right]$$

$$f(x|y) = \frac{1}{\sqrt{2\pi} s_x} \exp\left[-\frac{1}{2s_x^2}(x-y)^2\right]$$

By Bayes' rule, the posterior density function for Y given a signal $X = x$ is given by

$$g(y|x) = \frac{f(x|y) \cdot g(y)}{\int f(x|y) \cdot g(y)\,\mathrm{d}y}. \tag{9}$$

Carrying out substitutions, we can find the posterior density of $Y|x$ and check that

$$Y|x \sim N\left(\frac{\frac{m}{s_y^2} + \frac{x}{s_x^2}}{\frac{1}{s_y^2} + \frac{1}{s_x^2}} \, , \, \frac{1}{\frac{1}{s_y^2} + \frac{1}{s_x^2}} \right). \tag{10}$$

Three properties are worth being noted. First, the posterior is a normal as well. If we begin with a normal prior and the signal is normally distributed, the posterior remains normal. This feature is extensively used in models with rational expectations.

Second, we can simplify (10) by defining the *precision* of a normally distributed signal as the inverse of its variance. In particular, let $\tau_y = (1/s_y^2)$ and $\tau_x = (1/s_x^2)$ respectively the precisions of Y and X. Then (10) can be written as

$$Y|x \sim N\left(\frac{m\tau_y + x\tau_x}{\tau_y + \tau_x} \, , \, \frac{1}{\tau_y + \tau_x} \right). \tag{11}$$

Thus, the posterior mean of $Y|x$ can be written more simply as the average of the prior mean and of the signal weighted by their respective precisions. In the following, we make frequent use of this simple method for computing the expected value of a posterior belief.

Ex. 7.20 After having observed $X = x_1$, $Y|x_1$ is distributed according to (11). Suppose that Primus receives a second (iid) signal $X = x_2$ and derive his new posterior distribution for $Y|\{x_1, x_2\}$. Extend your answer to the case where Primus got n (iid) signals $x_1, x_2, \ldots, x_n$.

Ex. 7.21 If the signal about y has infinite precision, we have $s_x^2 = 0$ and (10) is no longer valid. What is the distribution of $Y|x$ when the signal X has infinite precision?

Third, note that the Bayesian posterior beliefs converge to the truth as the number of signals increase. After n (iid) draws $x_1, x_2, \ldots, x_n$, the variance of the posterior goes to zero while the Strong Law of Large Numbers implies that the posterior mean converges to m.

7.3 Cara preferences in a normal world

If Primus is an expected utility maximizer with constant absolute risk aversion, his utility function must be linear or exponential. In particular, if we also assume that he is strictly risk averse, his utility function over the wealth w must be a negative exponential

$$u(w) = -e^{-kw} \tag{12}$$

where $k > 0$ is his coefficient of (absolute) risk aversion.

Suppose that Primus has preferences which satisfy these assumptions and that his beliefs are normally distributed so that $W \sim N(\mu, \sigma)$. You were asked in Exercise 4.3 in Lecture 4 to check that his expected utility can be written

$$Eu(W) = \int \left[-e^{-kw}\right] \cdot \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{1}{2\sigma^2}(w-\mu)^2\right] \mathrm{d}w = -\exp\left\{-\left[k\mu - \frac{1}{2}k^2\sigma^2\right]\right\}.$$

This expression is a function only of the current mean μ and the current variance σ^2 of W. Since $-\exp(-kw)$ is a strictly increasing function of w, Primus' preferences in a cara–normal world are more simply expressed by the functional

$$V(\mu, \sigma) = \mu - \frac{1}{2} k\sigma^2$$

and we can characterize them using simply the two statistics μ and σ^2 and the coefficient of absolute risk aversion k.

Ex. 7.22 Suppose that $W = Y$ and that the agent has received a signal x such that his posterior beliefs about $Y|x$ are given by (10). Derive his (posterior) expected utility.

7.4 Demand for a risky asset

Suppose that Primus is an expected utility maximizer with constant absolute risk aversion $k > 0$. Assume that there are two assets, one of which is a risky stock and the other is a riskless bond. The bond has a current price normalized to 1 and will pay a riskless amount $(1 + r)$ at the end of the period. The stock has a current price of p and will pay a risky amount $Y \sim N(m, s)$ at the end of the period. Primus' current endowment is w. What is Primus' optimal portfolio?

Primus is interested in maximizing the expected value of his terminal wealth (at the end of the period). Assuming that short sales are allowed, he can invest in any portfolio of α stocks and β bonds such that $\alpha p + \beta = w$. Therefore, $\beta = w - \alpha p$. Note that α and β are unrestricted in sign and may not add to 1.

The terminal value of such a portfolio is normally distributed with mean $\alpha m + \beta(1+r) = \alpha m + (w - \alpha p)(1+r)$ and variance $\alpha^2 s^2$. Hence, the expected utility of Primus of a portfolio with α stocks is

$$\alpha m + (w - \alpha p)(1 + r) - \frac{1}{2} k\alpha^2 s^2.$$

Maximizing with respect to α, this yields Primus' demand function for stock:

$$\alpha(p, r; m, s^2; k) = \frac{m - p(1 + r)}{ks^2}. \tag{13}$$

Note that the demand for the risky stock is separately monotone in each of its arguments; for instance, it is increasing in the mean m and decreasing in the variance s^2.

Ex. 7.23 Provide an economic justification for each of these monotone relationships.

8. Transmission of information and rational expectations

8.1 Introduction

The main question addressed by rational expectations models is what happens when people with different information decide to trade. How market prices are affected by traders' information affects how the traders can infer information from market prices. The fundamental insight is that prices serve two purposes: they clear markets and they aggregate information. This dual role can make the behavior of prices and markets much more complex than assumed in simple models of asset behavior.

Let us begin with an example. Suppose that there are two agents in the market for q widgets. Primus receives a binary signal about the true value of widgets: if the signal is High, his demand for widgets is $p = 5 - q$; if the signal is Low, his demand is $p = 3 - q$. We say that Primus is *informed* because his demand depends on which signal he receives. Secunda receives no signal and offers an unconditional supply of widgets $p = 1+q$. Morevoer, assume that, if she could receive signals, Secunda would change her supply to $p = 1 + 3q$ with an H-signal and to $p = 1$ with an L-signal.

When Secunda is sufficiently naive, the following situation occurs. If Primus receives an H-signal, the demand from the informed Primus equates the supply from an uninformed Secunda at a price of $p^H = 3$ (and $q = 2$ widgets are exchanged). If he receives an L-signal, his demand equates the supply from Secunda at a price of $p^L = 2$ (and $q = 1$ widget is exchanged). Different prices clear markets for different signals: $p = 3$ when the signal is H and $p = 2$ when it is L.

This outcomes, however, presumes that Secunda does not understand that prices also convey information. The market-clearing price is $p = 3$ if (and only if) the signal is H. Thus, if Secunda sees that markets clear at a price of $p = 3$ she can infer that Primus has received an H-signal and this suffices to let her change the supply function to $p = 1+3q$. But in this case the market must clear at a price such that $5-q = 1+3q$, that is $p = 4$. Similarly, if the market-clearing price would be $p^L = 2$, Secunda would understand that Primus got an L-signal and her supply would switch to $p = 1$, making this the market-clearing price.

In other words, if Secunda passively lets the prices clear the market, the prices are $p^H = 3$ and $p^L = 2$. If she exploits the information embedded in different prices, the prices will be $p^H = 4$ and $p^L = 1$. The first case ($p^H = 3$ and $p^L = 2$) can be an equilibrium only if we assume that Secunda is not sufficiently rational to understand that prices reveal information, or to use the information which is revealed by prices. Market-clearing equilibria with rational agents require that the information embedded in prices is fully exploited, and this is what the notion of rational expectations equilibrium is about.

8.2 An example

We consider a two-asset one-period economy in which all random variables are independent and normally distributed, with strictly positive standard deviations. The two available assets are a risky stock and a riskless bond. The bond has a current price normalized to 1 and will pay a riskless amount $(1+r)$ at the end of the period. The stock has a current price of p and will pay a risky amount $Y \sim N(m, s_y)$ at the end of the period. For convenience, denote by $\tau_y = 1/s_y^2$ the precision of Y.

There are two traders in the economy. They have identical cara preferences expressed by the utility function (12), defined over terminal wealth, with $k = 1$. Primus has an initial endowment $w = m + Z_1$, where m is the amount of bond he holds and $Z_1 \sim N(0, s_z)$ is a random endowment of stock that he receives at the beginning of the period. Secunda has a similar endowment $w = m + Z_2$, where Z_2 is independent but identically distributed to Z_1. Because of this uncertainty about initial endowments, neither agent knows exactly what the overall supply $Z = Z_1 + Z_2$ of risky asset is.

Primus is an informed trader, while Secunda is uninformed. In particular, Primus receives a signal $X \sim N(y, s_x)$ about the value of the stock; denote by $\tau_x = 1/s_x^2$ the precision of this signal. By (11), Primus' posterior distribution is

$$Y|x \sim N\left(\frac{m\tau_y + x\tau_x}{\tau_y + \tau_x}, \frac{1}{\tau_y + \tau_x}\right).$$

Therefore, by (13), Primus' demand for the risky asset is

$$\alpha_1 = \frac{\frac{m\tau_y + x\tau_x}{\tau_y + \tau_x} - p(1+r)}{\frac{1}{\tau_y + \tau_x}} = m\tau_y + x\tau_x - p(1+r)(\tau_y + \tau_x).$$

The uninformed agent receives no signal. However, she knows that trading from the informed agent affects the price. If she knew how this trading affects the market-clearing prices, she could extract information about the informed trader's signal from the market price.

A rational expectations equilibrium postulates that there is a specific *pricing rule* $P(\cdot)$ which links Primus' information with the market price and let Secunda use this rule to extract information. In equilibrium, the pricing rule must be correct; that is, Secunda must extract information that is consistent with Primus' signal.

Suppose that the pricing rule is linear (we will check this in a moment). That is, assume

$$p = am + bx - cZ \tag{14}$$

for some appropriate coefficients a, b, c to be determined as part of the equilibrium. Given this pricing rule, Secunda (who is the uninformed trader) can construct the observable random variable

$$\eta := \frac{p - am}{b} = x - \frac{c}{b}Z.$$

Since $Z \sim N(0, \sqrt{2}s_z)$, this implies that η is an unbiased (but garbled) estimate of the signal x actually received by Primus. This estimate is garbled by the additional zero-mean noise associated with Z. Indeed, as we see from the right-hand side, $\eta \sim N(m, s_x + (c/b)\sqrt{2}s_z)$. Therefore, η is a signal about m that Secunda can use to obtain information about Y.

Let τ_η be the precision associated with $s_\eta = s_x + (c/b)\sqrt{2}s_z$; of course, $\tau_\eta < \tau_x$. By (11), Secunda's posterior distribution for Y is

$$Y|\{P(\cdot), p\} \sim N\left(\frac{m\tau_y + \eta\tau_\eta}{\tau_y + \tau_\eta}, \frac{1}{\tau_y + \tau_\eta} \right).$$

Again by (13), Secunda's demand for the risky asset is

$$\alpha_2 = \frac{\frac{m\tau_y + \eta\tau_\eta}{\tau_y + \tau_\eta} - p(1+r)}{\frac{1}{\tau_y + \tau_\eta}} = m\tau_y + \eta\tau_\eta - p(1+r)(\tau_y + \tau_\eta).$$

The equilibrium price can be found by equating the aggregate demand $\alpha_1 + \alpha_2$ with the aggregate supply Z. This yields (after substituting for η)

$$p = \frac{2m\tau_y + x(\tau_x + \tau_\eta) - Z[1 + (c/b)\tau_\eta]}{(1+r)(2\tau_y + \tau_x + \tau_\eta)}.$$

This expression can be rewritten as $p = am + bx - cZ$ by appropriately choosing a, b, c such that

$$a = \frac{2\tau_y}{(1+r)(2\tau_y + \tau_x + \tau_\eta)},$$

$$b = \frac{\tau_x + \tau_\eta}{(1+r)(2\tau_y + \tau_x + \tau_\eta)},$$

$$c = \frac{1 + (c/b)\tau_\eta}{(1+r)(2\tau_y + \tau_x + \tau_\eta)}.$$

This confirms that the pricing rule conjectured in (14) is linear and concludes the example.

If the aggregate supply Z were not noisy, Secunda could use η to infer exactly what is the signal x received by Primus. We say that the rational expectations equilibrium is *fully revealing* if prices can be used to infer exactly what are the signals. In this case, prices are sufficient statistics for all the available information and the market is "efficient" in the strong sense (even private information is embedded in prices).

The rational expectations equilibrium is *partially revealing* if only a partial inference is possible, as in the example discussed, where prices reflect both private information and exogenous noise. Now, if p goes up, Secunda cannot tell whether the cause is more positive private information (i.e., a higher signal for Primus) or smaller asset supply (i.e., a lower realization of Z).

Ex. 8.24 Assume that the aggregate supply is known to be z. Note that $\tau_\eta = \tau_x$ and compute the fully revealing rational expectations equilibrium.

8.3 Computing a rational expectations equilibrium

Suppose that there are n traders and j assets. Each trader observes a signal X_i ($i = 1, 2, \ldots, n$) about one or more assets. The construction of a rational expectations equilibrium (REE) can be outlined in five steps.

1. Specify each trader's prior beliefs and propose a pricing rule (which for the moment is only a conjecture) P^c mapping the traders' information to the prices of the assets. The pricing rule $P^c(X_1, X_2, \ldots, X_n, \varepsilon)$ may incorporate some noise ε. The traders takes this mapping as given. The pricing rule must be determined in equilibrium; at this stage, it is parameterized by undetermined coefficients because the true equilibrium price is not known yet.

2. Derive each trader's posterior beliefs, given the parameterized price conjectures and the important assumption that all traders draw inferences from prices. The posterior beliefs depend on the proposed pricing rule (e.g., from the undetermined coefficients).

3. Derive each trader's optimal demand, based on his (parameterized) beliefs and his preferences.

4. Impose the market clearing conditions for all markets and compute the endogenous market clearing prices. Since individual demands depend on traders' beliefs, so do prices. This gives the actual pricing rule $P^a(X_1, X_2, \ldots, X_n, \varepsilon)$ which provides the actual relationship between traders' signals and the prices.

5. Impose rational expectations; that is, make sure that the conjectured pricing rule P^c coincides with the actual pricing rule P^a. This can be achieved by equating the undetermined coefficients of P^c with the actual P^a.

8.4 An assessment of the rational expectations model

The REE provides a few key insights on which the following literature has built upon.

1. Prices play two roles: they clear markets and they convey information.

2. In a fully-revealing equilibrium, individual asset demands depend only on price, not on the trader's private information.

3. Therefore, in a fully-revealing equilibrium there is no incentive to invest in costly information: this incentive is restored in a partially-revealing equilibrium.

On the other sides, there are three important difficulties with the notion of a REE. First, there is the issue of its existence. If the number of possible signals is finite, then for a generic set of economies there exists a REE. This result, of course, does not apply to signals drawn from a normal distribution. Similarly, for j the number of assets in the economy, if the dimension S of the space of signals is lower than $j - 1$, then for a generic set of economies there exists a REE. Intuitively, if there are more prices than signals, there is sufficient flexibility to both clear markets and aggregate information. On the other hand, if $S = j - 1$, there is an open set of economies for which no REE exists. For instance, in a model with two assets (a risky stock and a riskless bond) and a one-dimensional signal, the existence of a REE, while possible, is a fragile result. Finally, for $S > j - 1$, all weird things can happen.

Second, there is the issue of how the pricing rule is discovered. In a REE, traders must know the pricing rule that specifies the equilibrium prices. How such knowledge comes about, however, is left unspecified. One possible explanation is that it is learned over time, but learning to form rational expectations is not an easy task. During the learning process, traders must act according to "wrong" conjectures; their behavior, then, is likely to upset the emergence of the "correct" conjectures.

Third, there is the issue of price formation. In a REE, it is implicitly assumed that prices are set by a Walrasian auctioneer which (magically) equates demand and supply. This implicit auctioneer collects the "preliminary orders" and uses them to find the market-clearing prices. This avoids the difficulty of actually specifying the actual trading mechanism but hides the effect of the market microstructure on the process of price formation. In the next lecture, we will see that the microstructure literature improves on the rational expectations model by making explicit the process of price determination.

All in all, models based on rational expectations are difficult to construct and difficult to interpret. The common approach circumvents these difficulties by using specific examples. While this makes things tractable, the approach is special in the sense that even smaller deviations from the assumptions in the example may change the equilibrium drastically, or make it disappear.

References

[1] S.J. Grossman and J.E. Stiglitz (1980), "On the impossibility of informationally efficient markets," *American Economic Review* **70**, 393–408.

[2] R.K. Lyons (2001), *The microstructure approach to exchange rates*, Cambridge (MA): The MIT Press, Chap. 4.

[3] M. O'Hara (1995), *Market microstructure theory*, Oxford: Basil Blackwell.

9. Market microstructure: Kyle's model

9.1 Introduction

The models on market microstructure differ on the assumptions made on how the best available price is set. In auction models (also known as "order-driven"), the best available price is defined by the submitted orders. In dealer models (also known as "quote-driven"), it is defined by dealer quotes.

The rational expectations model can be seen as an example of an (implicit) auction model. Kyle's (1985) model improves on this by making it explicit the process by which prices are formed, while keeping the auction structure and the use of an expectations consistent pricing rule. In particular, Kyle assume batch-clearing; that is, all orders are fulfilled simultaneously at the same price.

The model assumes that there is a *market-maker* (named Secunda) who set prices and thus acts as an auctioneer. Moreover, the market-maker can take trading positions and has privileged access to information on the order flow. This changes the nature of the pricing rule because the act of price setting is assigned to a player within the model. The market-maker must set prices using only the information which is available to him, which is determined by the trading protocol. This generates a relationship between the price and the trading protocol.

Besides the market-maker, Kyle assumes that there is one informed agent (named Primus) and a number of liquidity traders in the market. The market maker aggregates the orders and clears all trades at a single price. The informed trader chooses those transactions which maximize the value of his private information. This provides a relationship between the price and the strategic use of information by the informed trader. Thus price reflects both the trading protocol and the strategic behavior of the informed trader.

9.2 The model

We consider a one-asset one-period economy, with a zero riskless interest rate. There are three types of agents: one informed trader (Primus), one uninformed market-maker (Secunda), and many noise traders who trade only for liquidity or hedging reasons. Primus and Secunda are risk-neutral expected utility maximizers. All random variables are independent and normally distributed, with strictly positive standard deviations. The only available asset is a risky stock which will pay a risky amount $Y \sim N(m, s_y)$ at the end of the period.

Primus is the informed trader, who receives a perfectly precise signal about Y and tus learns what is the realization $Y = y$. Secunda is the uninformed market-maker, who knows only the prior distribution of Y. After Primus learns $Y = y$, market orders from Primus and the noise traders are submitted to Secunda to be executed at a uniform market-clearing price p. The noise traders submit a random demand $Q^u \sim N(0, s_u)$; if this is negative, they are on balance selling. Primus submits a demand Q^i without observing the realization of Q^u. The market-maker receives an aggregate demand $Q = Q^i + Q^u$; she knows the sum but not who demanded what.

The market-maker's pricing rule mandates that she earns zero profits. This is consistent with free entry of competing market-makers, which impairs any monopoly power of the single market-maker. This implies that the market-maker sets prices such that $p = E(Y|Q)$.

Primus places an order Q^i which maximizes his profit $E[(Y - p)Q^i|Y = y] = (y - p)Q^i$. Primus' order is influenced by the price quoted by the market maker. At the same time, his demand Q^i affects the price quoted. This strategic interaction between the market-maker and the informed trader is what makes the model tick: Secunda's choice of p depends on Q^i and Primus' choice of Q^i depend on p.

Kyle proves that there exists a linear equilibrium for this model such that the market-maker pricing rule is

$$P(Q) = m + \alpha Q \tag{15}$$

and Primus' trading rule is

$$Q^i = \beta(Y - m), \tag{16}$$

where

$$\alpha = \frac{1}{2}\frac{s_y}{s_u} \qquad \text{and} \qquad \beta = \frac{s_u}{s_y}. \tag{17}$$

We will show momentarily a piece of the argument which establishes this result. Before that, a quick commentary may be useful.

Note that both the pricing and the trading rule depend on the same parameters (although their ratio is inverted). When α is high and orders have a significant price impact, then β is low because Primus trades less aggressively (to avoid the impact of his own trades). When s_y is high, Primus' information is more likely to be substantial and therefore Secunda adjusts price more aggressively. When s_u is high, Primus' order is less likely to be a conspicuous component of the total order flow, and therefore he can afford to trade more aggressively.

To show that (15) and (16) constitute an equilibrium, we need to prove two facts. First, that the best reply of Secunda to an insider using the trading rule (16) is precisely (15). Second, that the best reply of Primus to a market-maker using the pricing rule (15) is precisely (16). We prove only the first fact, and leave the second as an exercise.

Thus, assume (16). Secunda's prior for Y is that $Y \sim N(m, s_y)$. Since she observes the aggregate demand $Q = Q^i + Q^u$, she can use this information to update her prior. As $Q^i = \beta(Y - m)$, the total order flow can be written $Q = \beta(Y - m) + Q^u$, with $Q^u \sim N(0, s_u)$. Dividing Q by β (which is known in equilibrium) and adding m (which is known), Secunda can construct the observable random variable

$$\eta := m + \frac{Q^i + Q^u}{\beta} = Y + \frac{Q^u}{\beta},$$

which is normally distributed with mean Y and standard deviation $s_\eta = s_y + s_u/\beta$. This can be used to make inferences about Y. Exploiting known results from Lecture 7 (with precisions instead of variances), we obtain

$$Y|\eta \sim N\left(\frac{m\tau_y + \eta\tau_\eta}{\tau_y + \tau_\eta}, \frac{1}{\tau_y + \tau_\eta}\right).$$

Looking at $E(Y|\eta)$ and substituting for τ_η and τ_y, we find

$$E(Y \mid \eta) = \frac{m\tau_y + \eta\tau_\eta}{\tau_y + \tau_\eta} = \frac{ms_u^2 + \eta s_y^2\beta^2}{s_u^2 + s_y^2\beta^2}.$$

Replacing the value of β from (17), this becomes

$$E(Y \mid \eta) = \frac{1}{2}(\eta + m).$$

Secunda sets her price equal to her best estimate of Y; that is, $P = E(Y|\eta)$. By definition, $\eta = m + (Q/\beta)$. Hence, the price set by Secunda is:

$$P = E(Y \mid \eta) = \frac{1}{2}\left(m + \frac{Q}{\beta} + m\right) = m + \frac{1}{2\beta}Q.$$

Using again the value of β from (17), it is easy to check that this pricing rule indeed matches (15), as it was to be shown.

9.3 Lessons learned

There are a few major insights to be gained from Kyle's model.

1. When setting the price, the market-maker explicitly updates the analysis of fundamentals (corresponding to her prior) with the information embedded in the order flow. She uses information gathered by the traders as transmitted by their demands. Order flow communicates information about fundamentals because it contains the trade of those who analyze/observe fundamentals.

We can actually measure how much of the trader's information is revealed in Kyle's model by looking at the variance of Secunda's posterior distribution for Y. Before she observes the order flow, this variance is s_y. After she observes Q and updates her prior, the variance becomes (check the steps!)

$$V(Y|\eta) = \frac{1}{\tau_y + \tau_\eta} = \frac{s_y^2 s_u^2}{\beta^2 s_y^2 + s_u^2} = \frac{s_y^2}{2}.$$

This is exactly half of the prior variance s_y^2. Regardless of the exact value, the important message is that the updated variance is somewhere in between the prior variance and a zero variance. If the updated variance were to remain s_y^2, Secunda would learn nothing and Primus could make infinite profits. If it were to drop to zero, Primus would make no profit because all his trades would clear at the perfectly revealing price Y.

2. Since the market-maker cannot separate informative from uninformative trade, the transmission of information is noisy and the informed trader can use this to his advantage. Primus can partially hide his trade from the market-maker in the order flow from the noise traders, reducing the amount by which price moves against him.

3. Liquidity and market efficiency are deeply related. Efficient markets tend to gravitate towards constant liquidity, defined as the *price impact* of orders. To see why, suppose that liquidity is not constant: for example, suppose it is known that the market will be fairly illiquid in the next month and then will revert to standard liquidity. Then if Primus buys 10's worth of stock each day in the next month, each purchase will push the price of the stock progressively upward. This is because if, Primus' trades communicate information about fundamentals, these price increases induced by the order flow should persist. Then, when at the end of next month liquidity has returned to normality, Primus could suddenly sell the 300's worth of stock purchased during the month with almost no price impact and make a riskless excessive profit similar to an arbitrage opportunity. An efficient market should prevent making excessive profits.

Kyle (1985) considers a multiple-period extension of his model which generates constant liquidity in equilibrium. Intuitively, this follows because the informed trader must balance the effects of his current trades on his future trading opportunities: if he trades too much too soon, the price will adjust rapidly and his profits will be smaller.

9.4 Excursions

Huddart et alii (2001) modifies Kyle's model to take into account the case of regulation requiring insiders to disclose their stock trades after the fact. This accelerates price discovery and reduces insider profits. Luo (2001) shows that public information is detrimental for the insider, because it increases the information revealed by the price.

References

[1] S. Huddart, J.S. Hughes and C.B. Levine (2001), "Public disclosure and dissimulation of insider trades," *Econometrica* **69**, 665–681.

[2] A. Kyle (1985), "Continuous auctions and insider trading," *Econometrica* **53**, 1315–1335.

[3] S. Luo (2001), "The impact of public information on insider trading," *Economics Letters* **70**, 59–68.

[4] R.K. Lyons (2001), *The microstructure approach to exchange rates*, Cambridge (MA): The MIT Press, Chap. 4.

[5] M. O'Hara (1995), *Market microstructure theory*, Oxford: Basil Blackwell.

10. Market microstructure: Glosten and Milgrom's model

10.1 Introduction

The market in Kyle's model is "order-driven": the market-maker sees the order flow and sets a price which clears the market in a single batch. Demand and supply meet at a single price and simultaneously. This lecture considers a different trading protocol, possibly closer to reality, which is "quote-driven" and involves sequential trades. Unlike Kyle's auction, this sequential-trade model describes a dealership market.

All trades involve a dealer, who posts bid and ask prices. Traders arrive sequentially and can trade at the current bid-ask prices. Thus orders are fulfilled sequentially at (possibly) different prices. After each trade, the dealer updates his bid and ask prices to reflect the information he has learned from his privileged position. At each trade, the current trader may be informed or uninformed so the model can accomodate more than one informed trader.

This trading protocol implies some important differences from Kyle's model. First, there is an explicit bid-ask spread, as opposed to the single market-clearing price of Kyle's model. Second, the spread occurs even if the dealer is risk-neutral and behaves optimally incorporating in the price all the information he can extract from the order flow. Third, since trades occur sequentially, it is possible to analyze explicitly how the information content of each trade affects the bid-ask spread.

10.2 The model

We consider a one-asset one-period economy, with a zero riskless interest rate. There are three types of agents: a single dealer (Primus), several informed traders and many uninformed traders. Both the dealer and the informed traders are risk-neutral expected utility maximizers, while the uninformed traders trade only for liquidity or hedging reasons.

The only available asset is a risky stock which will pay a risky amount Y at the end of the period. For simplicity, we assume that it can take only the high value $Y = 1$ or the low value $Y = 0$. The prior distribution is $P(Y = 1) = p$ and $P(Y = 0) = 1 - p$. All the informed traders know the realization of Y (because they have received a perfectly informative signal about it), while Primus does not.

Trading is organized as a sequence of bilateral trading opportunies, taking places one after the other but always within the end of the period. The pool of potential traders is given and the dealer knows that q of it are informed traders and $(1 - q)$ are uninformed traders. At each trading opportunity, one trader is randomly chosen from this pool (with replacement) and is offered the chance to buy or sell one unit of the stock at the current bid or ask price. The informed trader can buy, sell, or pass at her discretion. As for the uninformed traders, we assume for simplicity that they buy with probability r or sell with probability $1 - r$.

As in Kyle's model, the dealer sets prices such that the expected profit on any trade is zero. (This can be justified by assuming a competitive dealers' market.) This implies that Primus must set prices equal to his conditional expectation of the asset's value given the type of transaction taking place.

More specifically, before trading opportunity t occurs, Primus must offer a bid price b_t and an ask price a_t such that

$$b_t = E(Y \mid \text{next trader sells}) \qquad \text{and} \qquad a_t = E(Y \mid \text{next trader buys}). \tag{18}$$

This rule takes explicitly into account the effect that the sale/purchase of one unit would have on Primus' expectations. This makes sure that his prices are "regret-free", in the sense that — given the trade that actually occurs — the dealer believes that the price is fair.

Such "regret-free" price-setting behavior makes sure that prices incorporate the information revealed by a trade. Due to the signal value of each trade, as trading goes on, the dealer keeps revising his beliefs and sets new trading prices. This generates a sequence of bid-ask prices $\{b_t, a_t\}$ that change over time, paralleling the evolution of Primus' beliefs. Let us work out an example and see what happens.

10.3 An example

Suppose $p = q = r = 1/2$. Denote by B_t and S_t respectively the event that at the trading opportunity t there is a buy or a sale. By (18), the ask price at the first trading opportunity should be

$$a_1 = E(Y \mid B_1) = 1 \cdot P(Y = 1 \mid B_1) + 0 \cdot P(Y = 0 \mid B_1), \tag{19}$$

while the bid price should be

$$b_1 = E(Y \mid S_1) = 1 \cdot P(Y = 1 \mid S_1) + 0 \cdot P(Y = 0 \mid S_1). \tag{20}$$

In order to find what a_1 should be, we need to compute $P(Y = 1 \mid B_1)$. By Bayes' rule,

$$P(Y = 1 \mid B_1) = \frac{P(Y = 1) \cdot P(B_1 \mid Y = 1)}{P(Y = 1) \cdot P(B_1 \mid Y = 1) + P(Y = 0) \cdot P(B_1 \mid Y = 0)}. \tag{21}$$

Since we know by assumption that $P(Y = 1) = P(Y = 0) = 1/2$, it suffices to determine $P(B_1 \mid Y = 1)$ and $P(B_1 \mid Y = 0)$. We know that the uninformed traders buy always with probability 1/2. On the other hand, the informed trader know Y and therefore buy only if Y is high and sell only if Y is low. That is, they buy with probability 1 if $Y = 1$ and sell with probability 1 if $Y = 0$.

Therefore, conditional on $Y = 1$, the probability of a buy is 1/2 if it comes from an uninformed trader and 1 if it comes from an informed trader. Since uninformed and informed traders are equally likely to come up for trade, the overall probability is $P(B_1 \mid Y = 1) = (1/2)\cdot(1/2)+(1/2)\cdot 1 = 3/4$. By a similar reasoning, $P(B_1 \mid Y = 0) = (1/2)\cdot(1/2)+(1/2)\cdot 0 = 1/4$. Substituting in (21), we find

$$P(Y = 1 \mid B_1) = \frac{(1/2) \cdot (3/4)}{(1/2) \cdot (3/4) + (1/2) \cdot (1/4)} = \frac{3}{4}.$$

Ex. 10.25 By a similar reasoning, compute $P(S_1 \mid Y = 1)$ and $P(S_1 \mid Y = 0)$ and deduce that $P(Y = 1 \mid S_1) = 1/4$.

By (19) and (20), we obtain that the dealers sets ask and bid prices respectively equal to

$$a_1 = \frac{3}{4} \qquad \text{and} \qquad b_1 = \frac{1}{4}.$$

Suppose that there actually occurs a buy at a price of 3/4. What will be the new bid and ask prices? The probability distribution for Y after a buy is $P(Y = 1 \mid B_1) = 3/4$ and $P(Y = 0 \mid B_1) = 1/4$. This acts as a new prior for the next trading opportunity. Then

$$a_2 = E(Y \mid B_1, B_2) = 1 \cdot P(Y = 1 \mid B_1, B_2) + 0 \cdot P(Y = 0 \mid B_1, B_2)$$

and

$$b_2 = E(Y \mid B_1, S_2) = 1 \cdot P(Y = 1 \mid B_1, S_2) + 0 \cdot P(Y = 0 \mid B_1, S_2).$$

We need to compute $P(Y = 1 \mid B_1, B_2)$ and $P(Y = 1 \mid B_1, S_2)$. By Bayes' rule,

$$P(Y = 1 \mid B_1, B_2) = \frac{(3/4) \cdot (3/4)}{(3/4) \cdot (3/4) + (1/4) \cdot (1/4)} = \frac{9}{10},$$

which in turn implies $a_2 = 9/10$.

Ex. 10.26 Check that $P(Y = 1 \mid B_1, S_2) = 1/2$ and deduce that $b_2 = 1/2$.

Note that, if a sale had occurred instead of a buy, the price would have been set to $a_2 = 1/2$ and $b_2 = 1/10$. Therefore, the fact that the first transaction is a buy or a sale reveals information. On the other hand, since $(a_t - 1)(b_t - 0) \neq 0$ for all t, information is never fully revealed.

The exercise can be repeated. The exact sequence of bid-ask prices will depend on the actual trading events. For instance, if the first four events are B_1, S_2, B_3, and B_4, the sequence will be

t	b_t	a_t	trade
1	1/4	3/4	B_1
2	1/2	9/10	S_2
3	1/4	3/4	B_3
4	1/2	9/10	B_4
5	3/4	27/28	etc.

A sufficient statistics for the current bid and ask prices is the difference between the number β_t of buys occurred before t and the number σ_t of sales occurred before t. For instance, if $\beta_t - \sigma_t = 0$, then $b_t = 1/4$ and $a_t = 3/4$.

Ex. 10.27 Suppose that $p = r = 1/2$ but $q = 1/4$. Prove that $b_1 = 3/8$ and $a_1 = 5/8$.

10.4 Comments on the model

The assumptions about the trading protocol are crucial. Informed traders profit from trading if prices do not yet reflect all the available information. An informed trader prefers to trade as much (and as often) as possibile. By so doing, informed traders quickly reveal information which is immediately incorporated in prices.

This cannot occur in the model because the only trader who is allowed to trade is chosen randomly and she can only buy or sell one unit of stock. Thus, if an informed trader desires to trade further, she must return to the pool of traders and wait to be selected again.

The probabilistic selection process dictates that the population of traders facing the dealer is always the same as the population of potential traders. This makes it possible for the dealer to know the probability that he is trading with an informed trader. Moreover, it implies that plausible trading scenarios are ruled out. For instance, whenever information is likely to become more dispersed over time, the fraction of informed traders should increase with time (and the dealer would need to learn another parameter yet). This cannot occur here.

An important result of the model is that transaction prices form a martingale. That is, the best predictor for the transaction price in $t+1$ (given the information I_t available after trade t) is the transaction price in t: $E(p_{t+1}|I_t) = p_t$.

For instance, in the example above, consider the situation after having observed the first buy. The current price is $p_t = a_t = 3/4$. The next transaction may come from an uninformed trader with probability 1/2 or from an informed trader with probability 1/2. The probability that it will be a buy is 1/2 if the next trader is uninformed and 3/4 if he is informed. Therefore, the probability that the next transaction will be a buy is $(1/2) \cdot (1/2) + (1/2)(3/4) = 5/8$. Hence, the next price will be 9/10 with probability 5/8 and 1/2 with probability 3/8. Then $E(p_{t+1}|I_t) = (9/10) \cdot (5/8) + (1/2) \cdot (3/8) = 3/4 = p_t$.

The martingale property dictates that prices respect semi-strong efficiency, in the sense that they reflect all the information available to the dealer. It can be shown that, in the limit, all information is revealed and $a_t - b_t \to 0$ with both a_t and b_t converging to 0 or 1 depending on whether, respectively, $Y = 0$ or $Y = 1$.

10.5 Lessons learned

There are a few major insights to be gained from Glosten and Milgrom's model.

1. Information alone is sufficient to induce spreads, independently of the risk attitude of the dealer or of his inventory costs. The equilibrium spread in this model is such that when the dealer trades with an informed trader he loses money (as in Kyle's model, the informed trader knows exactly the value of the asset). To prevent overall losses, the dealer must offset them with gains from trading with uninformed traders. The equilibrium spread balances these losses and gains exactly so that expected profits are zero.

2. Learning takes place over time, as it involves the sequential arrival of distinct orders. The dealer does not know whether behind a single trade there is an informed trader who knows something that the dealer does not, or an uninformed trader who must trade for reasons unrelated to fundamentals. However, if a preponderance of sales takes place over time, the dealer adjusts her beliefs and prices downward. The private information gradually finds its way in the dealer's prices.

3. There is a process of price discovery behind the (dynamic) adjustment to market efficiency. The dealer must uncover the private information hidden behind trades before market prices can be efficient.

References

[1] L. Glosten and P. Milgrom (1985), "Bid, ask, and transaction prices in a specialist market with heterogenously informed traders," *Journal of Financial Economics* **13**, 71–100.

[2] R.K. Lyons (2001), *The microstructure approach to exchange rates*, Cambridge (MA): The MIT Press, Chap. 4.

[3] M. O'Hara (1995), *Market microstructure theory*, Oxford: Basil Blackwell.

11. Market microstructure: market viability

11.1 Introduction

In Glosten and Milgrom's (1985) model, spreads arise to protect the dealer from the risk of trading with an informed trader. However, if the threat of information-based trading is too high, the necessary spread may be too high for the uninformed traders to keep staying in the market. If so, they might quit the market, exposing the dealer to a sure loss against the remaining informed traders. If it is not possible to sustain a spread large enough to let the dealer break even, the only option is a trading halt in which the market ceases to function.

11.2 An example

Consider a one-asset one-period economy, with a zero riskless interest rate. There are three types of agents: an informed trader (Primus), a single market maker (Secunda), and nine uninformed traders. Among the uninformed traders, some are willing to place orders (for buying or selling) of 5 units, while the others always place orders only for 1 unit. We call these two types of uninformed traders respectively "big" and "little".

The only available asset is a risky stock which will pay a risky amount Y at the end of the period. We assume that Y is randomly distributed over the interval $[0, 4]$ with unconditional mean $m = 2$. The informed trader knows the realization of Y (because he has received a perfectly informative signal about it), while Secunda does not.

Finally, we assume that both the uninformed traders are price sensitive. Any uninformed trader placing an order of 1 is willing to pay up to a price of 2, equal to the unconditional mean. On the other hand, the four big uninformed traders are willing to accept a higher ask price if they decide to place a big order. In particular, U_1 is willing to buy his batch of 5 units up to a price of $p = 2.3$ each; U_2 up is willing to buy it up to a price of $p = 2.45$ each; U_3 up to a price of $p = 2.6$ each; and U_4 up to a price of $p = 2.9$ each.

For simplicity, suppose that Primus knows that the true value of the asset is 4 and therefore wants to buy the asset. Consider Secunda's decision about the ask price. She must a set a price for little orders of 1 unit and a (possibly different) price for big orders of 5.

We prove that, if Secunda acts competitively, there exists no ask price which can keep the market open. First, assume that Primus trades big. If Secunda acts competitively, she must sets a price for a big trade of 5 such that her expected losses to Primus equal the expected gains from the big uninformed traders. Since there is a one in five chance that the big trade comes from Primus, the price p must solve the equation $(1/5)(p - 4) + (4/5)(p - 2) = 0$, from which $p = 2.4$.

However, if Secunda were to quote a price $p = 2.4$, U_1 would not pass his order because the price is too high for him. This would make the number of uninformed big traders drop to three, in which case the competitive equilibrium price should solve the equation $(1/4)(p-4)+(3/4)(p-2)=0$, from which $p=2.5$. At such a price, U_2 would drop out rising the new competitive price (with only two big uninformed traders left) to $p=2.67$.

At such a price, U_3 would drop out and the new competitive price (with only one big uninformed trader left) should be $p=3$. This would make U_4 drop out and the final competitive price would shoot up to $p=4$ (because only the informed trader would be willing to trade). At this price, however, Primus' profit would be zero and therefore he would rather trade a small quantity and make a positive profit. Therefore, the large trade market would cease to exist.

Thus, assume now that Primus trades a small quantity. Since Primus is informed, Secunda must set her competitive ask price for a small trade of 1 such that $(1/10)(p-4)+(9/10)(p-2)=0$, from which $p=2.1$. However, this is higher than the willingness of pay of any uninformed trader. Therefore, they would quit the market leaving Secunda with no uninformed customer to offset the sure loss caused by Primus' order. Then Secunda would rather close the market (and make a zero profit) than face a sure loss.

11.3 Competitive versus monopolistic market making

The example suggests that asymmetric information may cause trading halts. However, while trading halts occasionally occur, they are not usual. The question, then, is what prevents such information-induced difficulties from arising more frequently. One possible answer is that some characteristics of the trading mechanism make it less responsive to these difficulties. Glosten (1989) provides a model where this stabilizing feature stems from the monopoly position of the specialist.

The basic idea is to compare a perfectly competitive dealer market against the "specialist system" of the New York Stock Exchange, where each investor has to trade through a single monopolist. In the competitive case, each (risk-neutral) dealer must set the trade price equal to the asset's expected value given the trade. This implies that the expected profit on each and every trade must be zero. As shown in the example, this condition may be too restrictive and could lead to a market failure.

In the monopolist case, the specialist does not face competition and can use a different pricing rule. In particular, she can set price so as that they maximize profits *on average*. Secunda may willingly accept to be losing money over large trades provided that she can make enough positive profits on small trades to offset these losses. (This cannot occur in a competitive equilibrium, because competitors would compete over small trades driving ask prices down and bid prices up.)

Under asymmetric information, because of the greater protection afforded to the market maker, a monopolistic regime can lead to a greater social welfare than a competitive regime.

11.4 The basic steps in the model

We consider a one-asset one-period economy in which all random variables are independent (unless otherwise mentioned) and normally distributed, with strictly positive standard deviations. The risky stock will pay a risky amount $Y \sim N(m, s_y)$ at the end of the period.

Let τ_y denote the precision of Y.

There are many traders and a single market maker in the economy. Each trader is a risk-averse expected utility maximizer with the same constant coefficient of (absolute) risk aversion $k > 0$. That is, each trader has a utility function $u(w) = -e^{-kw}$ defined over his terminal wealth w. The market maker (named Secunda) is a risk-neutral expected utility maximizer.

At the beginning of the period, each trader starts with a known amount of cash c and a random endowment $W_0 \sim N(0, s_w)$ of the risky asset. Finally, each trader receives a noisy signal $X = Y + \varepsilon$, where $\varepsilon \sim N(0, s_x)$. Let τ_x be the precision of the signal X. A trader may seek a trade for speculative reasons (taking advantage of his information) or (regardless of his signal) for liquidity or hedging reasons, such as rebalancing his portfolio.

Let us dub Primus the generic trader and denote by Q the size of his order (with Q being negative for sales). Then, given the pricing rule $p(Q)$ chosen by the market maker, Primus's decision problem is to choose Q so as to maximize the expected utility of his final wealth

$$W_1 = [c - Q \cdot p(Q)] + (W_0 + Q)Y.$$

By the assumption of normality and cara preferences, and taking into account the information associated with Primus' signal X, we can solve this problem by maximizing

$$E(W_1 \,|\, X) - \frac{1}{2}kV(W_1 \,|\, X) = c - Q \cdot p(Q) + (W_0 + Q)E(Y|X) - \frac{1}{2}k(W_0 + Q)^2 V(Y|X).$$

Assuming a differentiable pricing rule $p(Q)$ (which is yet to be determined), this leads to the first-order condition

$$\frac{[Q \cdot p'(Q) + p(Q) - m]\,(\tau_y + \tau_x) + kQ}{\tau_x} = X - \frac{kW_0}{\tau_x} - m. \tag{22}$$

Now, suppose that the market maker acts in a competitive market. Since Secunda sees Q, she can calculate the left-hand side of (22) and add m to obtain the signal

$$\eta = X - \frac{kW_0}{\tau_x},$$

which is normally distributed with mean X and standard deviation $s_x + s_w$. This is equivalent to getting a noisy version of the same signal X observed by Primus. Using this information (revealed by Primus' order flow), Secunda sets her pricing rule $p(Q)$. Assuming differentiability of $p(Q)$, Glosten (1989) proves that in the competitive market it has the form

$$p(Q) = m + A\left(\frac{Q}{2\alpha - 1}\right) + k\ [\mathrm{sgn}(Q)]\ |Q|^{\beta},$$

where $\alpha, \beta = \alpha/(1-\alpha)$ and A are strictly positive and depend on the precisions of Y, X, and W_0 and on the coefficient of risk aversion k. Moreover, and more importantly, a necessary condition for a solution to exist is $\alpha > (1/2)$. This leads to two important observations.

11.5 Lessons learned

First, the pricing rule is upward sloping. Higher orders are filled at higher (and, hence, worse) prices. This may cause a welfare loss due to the risk of information-based trading. When a trader transacts for portfolio-rebalancing reasons, the greater the slope, the more costly is for him to achieve his optimal holding. Instead, with a uniform pricing schedule, he could move costlessly to his optimal holdings. Glosten (1989) proves that *ex ante* utility is strictly inferior under asymmetric information.

Second, a necessary condition for a (differentiable) pricing rule and hence a solution to exist is $\alpha > (1/2)$. This condition requires that the risk of informed trading is not so large as to overwhelm the market maker's ability to set market-clearing prices which let her break-even. Since α depends on the precisions and the traders' risk aversion, it can be lower than (1/2) and therefore there is no guarantee that a trading halt can be avoided.

On the other hand, if we suppose that the market maker acts as a monopolist, she can choose a pricing rule $p(Q)$ to maximize her expected profits. Glosten shows that there exists one such rule which guarantees that the market never shuts down. With this pricing strategy, the specialist quotes prices such that she loses money on large trades but makes money on small trades. Since by normality extreme trades are unusual, Secunda is essentially betting on the greater frequency of profits from small trades to offset the few occasional losses from large trades.

This implies that prices for small trades are higher than they would be in a competitive market. From the viewpoint of a liquidity trader, this is clearly undesirable. On the other hand, these monopolistic prices makes trade possible where it would otherwise be not. When the market would otherwise fail, Glosten proves that welfare is higher in a monopolistic market.

References

[1] M.K. Brunnermeier (2001), *Asset pricing under asymmetric information: Bubbles, crashes, technical analysis and herding*, Oxford: Oxford University Press.

[2] L. Glosten (1989), "Insider trading, liquidity, and the role of the monopolist specialist," *Journal of Business* **62**, 211–236.

[3] M. O'Hara (1995), *Market microstructure theory*, Oxford: Basil Blackwell.

12. Noise trading: limits to arbitrage

12.1 Introduction

Many models in financial economics rely on the existence of traders who are unable to fully exploit available information or to correctly maximize their utility. Somehow, these people fail to behave optimally. They are called "noise traders".

When there are many noise traders, the market is more liquid in the sense of having frequent trades that allow people to observe prices. On the other hand, the price of the stock reflects both the correct information on the part of informed agents and the distortions generated by the noise traders. The more noise trading, the less information gets into prices, and therefore the less efficient the market.

As price information gets blurred by the effects of noise trading, speculators and arbitrageurs cannot be sure that they are trading on information rather than noise. This limits their willingness to take large positions and slows down the process by which adjustment to fundamental values takes place. In effect, it is even possible that noise traders may profit more than "rational" traders if the amount of noise trading is large enough to sustain an imbalance that in the short-term favors noise traders.

In the words of Black (1986), "noise creates the opportunity to trade profitably, but at the same time makes it difficult to trade profitably".

12.2 A model with noise trading

We begin with a known result. Suppose that Primus is an expected utility maximizer with a coefficient of constant absolute risk aversion $k > 0$. Assume that there are two assets, one of which is a risky stock with a current price of p and a random payoff $Y \sim N(m, s)$ and the other is a riskless bond with a current price of 1 and a riskless payoff of $(1 + r)$. Recall from Lecture 8 that Primus' demand function for the stock is

$$\alpha = \frac{m - p(1 + r)}{ks^2} \tag{23}$$

Consider now an economy with a riskless bond and a risky stock, where agents live two periods. Each agent chooses a portfolio of two assets in the first period (when young) and consumes the payoff from his portfolio in the second period (when old). As time goes on, generations overlap. In each period, there are old timers who sell the assets in their portfolios to youngsters.

The riskless asset in each period is a one-period bond which costs 1 in period t and pays a sure amount $(1 + r)$ in period $t + 1$. The risky asset is a stock that costs p_t in period t and pays a fixed dividend r in the next period. Therefore, the payoff to one unit of stock in the next period is $(p_{t+1} + r)$, where (based on the information available in period t) p_{t+1} is normally distributed with mean $E_t(p)$ and standard deviation s_t. Hence, in period t, youngsters should regard the stock as normally distributed with payoff $Y \sim N(E_t(p)+r, s_t)$.

All agents in the economy have the same preferences as Primus. However, there are two types of agents in the economy with respect to beliefs. A fraction $(1-\lambda)$ are sophisticated investors who have rational expectations and correctly believe that the payoff to the stock is $Y \sim N(E_t(p)+r, s_t)$. The remaining λ are noise traders, who in period t misperceive the expected price of the stock $E_t(p)$ by an identical amount ε_t and wrongly believe that the payoff to the stock is $Y + \varepsilon_t \sim N(E_t(p) + r + \varepsilon_t, s_t)$. Note that all noise traders share the same misperception: this makes the risk associated with their presence nondiversifiable.

We assume that the noise traders' misperception is a normal random variable $\varepsilon_t \sim N(\mu, \sigma)$ independent and identically distributed over time. The mean misperception μ is a measure of the average "bullishness" of the noise traders.

Each agent maximizes his expected utility. By the result recalled above, in each period t a sophisticated investor demands a quantity

$$\alpha_t^i = \frac{E_t(p) + r - p_t(1+r)}{ks_t^2}$$

of the stock, while a noise trader demands a quantity

$$\alpha_t^n = \frac{E_t(p) + r + \varepsilon_t - p_t(1+r)}{ks_t^2} = \alpha_t^i + \frac{\varepsilon_t}{ks_t^2}.$$

When noise traders overestimate expected returns, they demand more of the risky asset than sophisticated investors; when they underestimate, they demand less. Sophisticated investors have a stabilizing influence because they offset the volatile position of noise traders.

To calculate the equilibrium prices, note that the old sell to the young and so in each period $\lambda\alpha_t^n + (1-\lambda)\alpha_t^i = 1$. This implies

$$p_t = \frac{E_t(p) + r + \lambda\varepsilon_t - ks_t^2}{1+r}, \tag{24}$$

which gives the price of the stock in period t as a function of period t misperception, of the technological (r) and behavioral (k) parameters of the model, and of the first two moments of the distributions of p_{t+1} (conditional on the information available in t).

We consider only steady-state equilibria by imposing the requirement that the unconditional distribution of p_{t+1} be identical to the distribution of p_t. We can then eliminate $E_t(p)$ by recursive substitution to obtain

$$p_t = 1 + \frac{\lambda(\varepsilon_t - \mu)}{1+r} + \frac{\lambda\mu}{r} - \frac{ks_t^2}{r}. \tag{25}$$

Since all terms but the second one are constant, the one-step-ahead variance of Y is

$$s_t^2 = \frac{\lambda^2\sigma^2}{(1+r)^2}.$$

Substituting back into (25), we obtain the final form of the pricing rule

$$p_t = 1 + \frac{\lambda(\varepsilon_t - \mu)}{1+r} + \frac{\lambda\mu}{r} - \frac{k\lambda^2\sigma^2}{r(1+r)^2}. \tag{26}$$

Ex. 12.28 Explicitly carry out the recursive substitution leading to Equation (25).

Interpretation. The last three terms in (26) represent the impact of noise trading on the price of the stock. As the distribution of the misperception converges to a point mass of zero (and thus $\mu \to 0$ and $\sigma \to 0$), the equilibrium pricing function for the stock approaches its fundamental value of one.

The second term captures the fluctuations in price due to variations in noise traders' misperceptions. The higher their bullish beliefs, the higher the positive difference between the current price and the fundamental value. Moreover, the higher the fraction of noise traders, the higher the volatility of the price of the stock.

The third term captures the permanent effect on price due to the fact that the average misperception of traders is not zero. If noise traders are bullish on average, there is a positive "pressure" which raises the price above its fundamental value.

The fourth term shows that there is a systematic underpricing of the stock due to the uncertainty about noise traders' beliefs in the next period. Both noise traders and sophisticated investors in period t believe that the stock is mispriced, but the uncertainty about p_{t+1} makes them unwilling to bet too much on this mispricing. In fact, if it were not for traders' misperceptions, the stock would not be risky: its dividend is fixed in advance and the only uncertainty about its payoff comes from p_{t+1}, which is not affected by any fundamental risk but depends on the noise traders' misperceptions. In a sense, noise traders "create their own space", driving the price of the stock down and its return up.

12.3 Relative returns

The model can also be used to show that the common belief that noise traders earn lower returns than sophisticated investors and therefore are eventually doomed to disappear may not be true. All agents earn the same return on the riskless bond. Hence, assuming equal initial wealth, the difference between noise traders' and sophisticated investors' total returns is the product of the difference in their holdings of the stock and the excess return paid by a unit of the stock

$$\Delta R = (\alpha_t^n - \alpha_t^i)\left[p_{t+1} + r - p_t(1+r)\right].$$

The difference between noise traders' and sophisticated investors' demands for the stock is

$$(\alpha_t^n - \alpha_t^i) = \frac{\varepsilon_t}{ks_t^2} = \frac{(1+r)^2\varepsilon_t}{k\lambda^2\sigma^2}.$$

Note that this difference becomes very large as λ becomes small. Noise traders and sophisticated investors take enourmous positions of opposite signs because the small amount of noise trader risk makes each group think that it has an almost riskless arbitrage opportunity.

Substituting from (24), the expected value of the excess return (conditional on the information available in t) is

$$E_t\left[p_{t+1} + r - p_t(1+r)\right] = ks_t^2 - \lambda\varepsilon_t = \frac{k\lambda^2\sigma^2}{(1+r)^2} - \lambda\varepsilon_t.$$

Therefore, the (conditional) expected value of the excess total return of noise traders is

$$E_t\left[\Delta R\right] = \varepsilon_t - \frac{(1+r)^2\varepsilon_t^2}{k\lambda\sigma^2}.$$

Taking the unconditional expectation, we have

$$E\left[\Delta R\right] = \mu - \frac{(1+r)^2\mu^2 + (1+r)^2\sigma^2}{k\lambda\sigma^2}, \tag{27}$$

which can be positive for intermediate (and positive) values of μ. Thus, noise traders with the right amount of "bullish bias" can earn positive excess returns with respect to sophisticated investors.

According to De Long et alii (1990), Equation (27) may be interpreted in terms of four effects. The first μ on the right-hand side of (27) represents the "hold more" effect. For $\mu > 0$, this increases noise traders' expected returns relative to sophisticated investors because noise traders on average hold more of the risky asset and thus earn a larger share of the rewards to risk bearing.

The first term in the numerator of (27) represents the "price pressure" effect. As noise traders become more bullish, they demand more of the risky asset and drive up its price. This reduces the return to risk bearing and thus contrasts the "hold more" effect on the return differential.

The second term in the numerator represents the "buy high–sell low" effect. Because of their misperceptions, noise traders have a bad market timing. The more variable their beliefs are, the more damage their poor market timing does to their returns.

The denominator represents the "create space" effect. A higher variability in noise traders' beliefs increases the price risk. To take advantage of noise traders' misperceptions, sophisticated investors must bear this greater risk. Being risk averse, however, they reduce the extent to which they bet against noise traders in response to this increased risk. If this effect is sufficiently large, it may overcome the "price pressure" and the "buy high–sell low" effects.

Ex. 12.29 Suppose $\sigma^2 = 1$. Prove that $k\lambda \geq 2(1+r)^2$ is a necessary and sufficient condition for $E\left[\Delta R\right] \geq 0$.

12.4 An appraisal

There is a traditional claim that speculation in asset markets is a price stabilizer. According to Friedman (1953), "people who argue that speculation is generally destabilizing seldom realize that this is largely equivalent to saying that speculators lose money, since speculation can be destabilizing in general only if speculators on average sell when currency is low in price and buy when it is high."

This line of reasoning combines two arguments: first, that rational arbitrageurs push asset prices towards their fundamental value; second, that over time irrational traders will have bad results, and therefore be driven out of the market.

De Long et alii (1991) shows that, when noise traders follow positive feedback strategies, sophisticated investors may prefer to anticipate the bandwagon. This forward–looking speculators destabilize prices by increasing the volatility of prices. Shleifer and Vishny (1997), shows that the existence of noise traders is enough to limit the positions of sophisticated investors in an agency context where professional arbitrageurs must attract outside funding. Therefore, the first of Friedman's arguments alone does not suffice to imply the stabilizing effect.

The model by De Long et alii (1990) discussed here, instead, shows that under reasonable conditions noise traders can dominate rational traders. A similar point is made in Blume and Easley (1992) where it is shown that the fitness (in the sense of higher wealth share) and the rationality (in the sense of utility-based choice) of investment rules may often depart. Thus, the second argument alone is not conclusive either.

12.5 Excursions

De Long (1990) shows that noise traders may earn higher expected returns than rational investors. However, since they are risk averse and must also bear more risk, this does not imply that they enjoy a higher expected utility. Palomino (1996) shows that in an imperfectly competitive market even a higher utility may occur.

References

[1] F. Black (1986), "Noise," *Journal of Finance* **41**, 529–543.

[2] E. Blume and D. Easley (1992), "Evolution and market behavior", *Journal of Economic Theory* **58**, 9–40.

[3] E. Blume and D. Easley (1993), "Economic natural selection", *Economics Letters* **42**, 281–289.

[4] J.B. De Long, A. Shleifer, L.H. Summers, and R.J. Waldmann (1990), "Noise trader risk in financial markets", *Journal of Political Economy* **98**, 703–738.

[5] J.B. De Long, A. Shleifer, L.H. Summers, and R.J. Waldmann (1991), "Positive feedback investment strategies and destabilizing rational speculation", *Journal of Finance* **45**, 379–395.

[6] M. Friedman (1953), *Essays in positive economics*, Chicago: University of Chicago Press.

[7] F. Palomino (1996), "Noise trading in small markets", *Journal of Finance* **51**, 1537–1550.

[8] A. Shleifer and R.W. Vishny (1997), "The limits of arbitrage", *Journal of Finance* **52**, 35–55.

13. Noise trading: simulations

13.1 Introduction

Models which explicitly include noise trading among their assumptions can be of a few different kinds. For convenience, we can distinguish three main cases.

First, there are the agents who are (correctly) solving a problem other from the one under study. For instance, if a trader cares about a longer horizon than the one studied in a one-period model, his choices may appear nonoptimal in the framework of one-period, even though they are so over the longer horizon. Similarly, if a trader is solving her personal hedging or portfolio balancing problem, she may appear to trade for reasons unrelated to speculation. In some of the earlier lectures, we have implicitly relied on this interpretation when considering a few examples in which noise trading affects informed agents' plans.

Second, there are agents who are (incorrectly) solving the problem under study. This may occur for many reasons, but it is a reasonable approximation that in order to do so they must be using "wrong" preferences or "wrong" beliefs. They might have heterogenous priors, or they might be unable to correctly follow Bayes' rule, or they may not understand the pricing rule of a rational expectations model, or they might just exchange noise for valuable information. Whatever the reason impairing the rationality of their choices, it remains true that non-noise agents must formulate their plans by taking into account the actions of the noise traders. In the previous lecture, we have relied on this interpretation.

Third, agents' choices may be rule-driven. When rules are not perfectly rational, they may lead to noise trading. Consider feedback traders, who base their trading on a rule driven by the past price dynamics. In particular, when the price of an asset goes up, positive feedback traders tend to buy it while negative feedback traders tend to sell it. Of course, this might also be the outcome of a deliberate rational choice. However, whenever the feedback rule is mechanically imposed, their behavior is ascribed to noise trading. This lecture discusses two models which fall into this latter category.

13.2 A simple dynamic model

Consider a one-asset economy, with a zero riskless interest rate. Let p_t the current price of the asset and y its long-term expected value. We assume that the value of the asset is constrained to lie between a minimum m and a maximum M.

There are three types of agents: fundamentalists, noise traders and a single market maker. For convenience, we assume that all traders of the same type behave identically and that they are all risk-neutral. Then we can dub Primus a (generic) fundamentalist and Secunda a (generic) uninformed trader.

Besides the current price p_t, Primus uses all the additional information I_t available to reach his best estimate u_t of the long-term value of the asset. To makes his investment decisions, Primus takes into account the difference between his best estimate u_t and the current price p_t. Dropping subscripts for simplicity, when $u > p$ a capital gain is expected and hence Primus buys; similarly, he sells if $u - p < 0$. Moreover, the larger $|u - p|$, the higher is the expected payoff to this strategy. To model this second effect, we assume that there is a bimodal positive function π over $[m, M]$ achieving its minimum in $p = u$; the higher π, the higher the (expected) profitability and hence Primus' willingness to trade.

Thus, the excess demand from fundamentalists who base their transactions on the spread $u - p$ can be modelled as

$$D_1(p) = a \cdot (u - p) \cdot \pi(p)$$

for p in $[m, M]$ and 0 otherwise. The parameter $a > 0$ measures the strength of fundamentalists' demand. (Note: $D_1(p)$ is decreasing over $[m, M]$, positive and convex over $[m, u]$, negative and concave over $[u, M]$.) This matches Black's (1986) hypothesis that "the farther the price of a stock is from its (investment) value, the more aggressive the information traders become. More of them will come in, and they will take a larger position."

While Primus is a sophisticated investor who closely monitors the market, Secunda is a more simple-minded operator who follows simpler rules. Instead of assessing u_t using all the available information, she uses a simpler adaptive rule such as

$$u_t = p_t + s \cdot (p_t - y),$$

where $s > 0$ measures the speed of the adaptive rule. Using a simple extrapolative rule, Secunda responds only to the spread between the current price and the (long-term) investment value.

Unlike Primus, Secunda does not have a well-planned strategy for trading, but simply enters the market when the price is high (expecting it to go up) and exits when it is low. Given that noise traders chase prices, their excess demand can be modelled as

$$D_2(p) = k \cdot (u - p) = b \cdot (p - y)$$

for p in $[m, M]$ and 0 otherwise. The parameter $b = ks > 0$ measures the strength of noise traders' reaction to price movements. Note how Secunda "chases" prices up and down, following a pattern which may potentially sustain both bull and bear markets.

The market maker is a specialist who equilibrates the market, matching excess demand out of his inventory. He announces a price p_t in each period, and then executes all orders received at that price filling out excess demand or absorbing excess supply. The market maker adjusts prices over time to balance his holdings over time while at the same time moderating his response to avoid destabilizing the market. Given that the aggregate demand is $D_1 + D_2$, we assume that he sets prices such that

$$p_{t+1} = p_t + c[D_1 + D_2]$$

where $c > 0$ measures the rapidity of the price's adjustments.

Simulations and hard results

Substituting numeric values for the parameters, we can run simulations. Day and Huang (1990) do so for $u = y$ and find out that (for reasonable value of the parameters) the model generates a time series matching a regime of irregular fluctuations around shifting levels. That is, we obtain the appearance of randomly switching bear and bull markets.

The intuition is the following. Suppose that p_0 is just above y. Secunda enters the market, while Primus is not much willing to sell. Given the aggregate excess demand, the market maker must sell from his inventory. This drives the prices up, initiating a bull market until the price reaches a level at which Primus begins to sell consistent amounts and creates an excess supply. Then the price is pulled back, yielding a temporary respite or initiating a bearish regime.

This market admits two types of equilibria. In the (unique) full equilibrium, both Primus' and Secunda's demand is zero; this occurs when $p = u$ and Secunda is in equilibrium when $p = y$. In a temporary equilibrium, the aggregate demand is $D_1 + D_2 = 0$: Primus' and Secunda's demands exactly offset each other.

Depending on the parameters, different kinds of behavior may emerge. For a pictorial representation, draw the phase diagram of p_{t+1} versus p_t on the interval $[m, M]$, with the 45-degree line and the price adjustment function which has a local minimum near m and a local maximum near M (the function is initially convex and then concave). There are four cases: a) bullish market, when the price adjustment function crosses the 45-degree line (from above) only once near m; b) bearish market when this happens (from above) only once near M; c) stable market when it crosses (from above) only once in y; and d) "bear and bull" if it crosses thrice, one (from above) at p_m (near m), one from below) at y and one (from above) at p_M (near M). For $u = y$, only Cases c) and d) may occur.

For the case we are most interested in, a useful distinction concerns the sign of $[D_1'(y)+b]$. If this is negative, the demand from Primus at $p = y$ locally overwhelms the demand from Secunda: we say that flocking is weak; otherwise, we say that it is strong.

Thm. 13.11 Suppose $D_1'(y) < 0$. If flocking is weak, prices converge to y for $c < c^* = -2/[D_1'(y) + b]$ and locally unstable 2-period cycles around y arise for $c > c^*$. If flocking is strong, the full equilibrium is unstable.

Case d) above occurs under strong flocking where, for instance, there are prices high enough between the equilibrium price p^M and M to make excess demand from Primus fall so precipitously that the price is pulled below y. Then Secunda interprets this as a signal of a further fall and there is a negative feedback effect initiating a dramatic fall in prices. A similar fluctuation arises on the other side.

Thm. 13.12 For appropriate (and robust) values of the parameters, the following may occur

1. Chaos: perpetual, erratic speculative fluctuations.
2. Switching regimes: stock prices switch between bull and bear markets at random intervals with irregular fluctuations around a low and a high level respectively.
3. Ergodicity: the frequency of observed prices converges.

4. Appearance of deceptive order: the price trajectory may pass close to cycles of varying periodicities but, as these are unstable, it will eventually be driven away.

Lessons learned

The simulations and the theorems are valid for $u = y$, but the conclusions do not change much if $u \neq y$. For instance, if $u > y$, Primus has bullish expectations and we obtain the bullish market described as Case a). This market will have fluctuations with an occasional long run-up followed by a sharp pullback. Of course, as u changes over time due to information arrival, the switch between bullish and bearish regimes may be driven by exogenous shocks.

As u may change over time, Secunda may occasionally sell at prices above her purchase price. However, in general, she will buy high and sell low, acting as the sheep to be sheared by Primus. How can her behavior survive? A piece of the answer is that she is most often right than not. When Secunda buys in, the market usually goes up (except for a few occasional sharp drops). So, when she buys near the peak and the price falls, she might attribute this to a timing problem: she just bought too late.

The possibility of cycles makes technical trading possible. However, if chaos occurs, cycles of all orders are possible. Therefore, assuming that there are enough feedback traders, technical trading may be temporarily effective, but sooner or later it is likely to go astray.

13.3 An artificial stock market

A financial market involves a group of interacting agents, who adapt to the arrival of new information and update their model of the world. Modelling this aspect is not easy, but simulations may be used to gain some insight. LeBaron et alii (1999) is a good example. They simulate an artificial stock market where agents can add variable selection to their forecasting problem.

Consider an economy with a riskless bond and a risky stock, extending over several periods. In each period, the bond pays a constant rate of interest r and the stock pays a stochastic dividend d_t which follows the autoregressive process $d_t = \delta + \rho(d_{t-1} - \delta) + \varepsilon_t$, where ε_t is a zero-mean normally distributed random variable.

There are 25 shares of the stock and 25 agents trading, each of which is an expected utility maximizer with the same (constant) coefficient of absolute risk aversion k. In each period, an agent acts myopically to maximize his expected utility with respect to his current beliefs. If stock prices were normally distributed, there would be a rational expectations equilibrium where each agent holds one share and the price of the stock in one period is a linear function of the dividend paid in that period: $p_t = \alpha d_t + \beta$.

However, the interaction of agents using different forecasting models can upset this equilibrium and change the distribution of stock prices. In this simulated market, each agent is endowed with the capacity of using his own several forecasting rules. Some are based on fundamental information such as the dividend-price ratio and others on technical information such as moving averages. These rules may apply at all times, or in certain specific cases. They are monitored for accuracy and only the best rules are actually used in forecasting. After a chunk of trading periods, each agent reexamines the rules, eliminates the worst ones and generate new ones.

In each period, each agent uses his best available forecasting rule to determine his demand for shares. A (fictitious) auctioneer clears the market fixing a price which balances the demand and the fixed supply of shares. At the end of each period, each agent updates the accuracy of his forecasting rules against the realized stock price.

Simulations are run under the assumption of slow or fast learning, which refers to the frequency with which agents revise their forecasting rules. We mention only the most significant departures from the benchmark case of the rational expectations equilibrium. First, the stock price has a significantly greater kurtosis, especially under fast learning. Under fast learning there is some evidence of technical trading predictability, although technical trading is not consistently useful and its impact undergoes large swings over time. Trading volume series are highly persistent and there are significant cross correlations between volume and volatility, as is the case in actual stock markets.

References

[1] F. Black (1986), "Noise," *Journal of Finance* **41**, 529–543.

[2] R.H. Day and W. Huang (1990), "Bulls, bears and market sheep," *Journal of Economic Behavior and Organization* **14**, 299–329.

[3] B. LeBaron, W.B. Arthur, and R. Palmer (1999), "Time series properties of an artificial stock market," *Journal of Economic Dynamics and Control* **23**, 1487–1516.

14. Behavioral finance: evidence from psychology

14.1 Introduction

Most economic theory deals with an idealized *homo œconomicus* who embodies principles of rational decision-making such as expected utility maximization, Bayesian learning and rational expectations. However, psychologists have amassed a great deal of evidence that human beings depart from the idealized picture in many ways.

The building blocks of the decision-making process of the *homo œconomicus* are beliefs and preferences. He makes judgements about the probabilities and assigns values (or utilities) to outcomes. By combining these together, he derives preferences about uncertain options.

The point repeatedly made by psychologists is that actual people's judgements are systematically wrong in many ways. These systematic errors of judgement are called biases. They are further aggravated by distortions in evaluating preferences or by improper combinations of probabilities and values.

For instance, people tend to use heuristics and shortcuts to make their probabilistic judgements and their preferences are often influenced by how a problem is framed. The purpose of this lecture is to make you aware of the most important findings and some of their consequences for predicting investors' behavior in financial markets.

14.2 Judgement biases

Overconfidence. Most people cannot correctly calibrate their probabilistic beliefs. Suppose you are asked:

State your 98% confidence interval for the SP500 one month from today.

and then your prediction is compared against the actual outcome. Repeat this exercise many times. If you are good at probabilistic judgements, you should expect to encounter about 98% of outcomes inside the confidence interval and thus be "surprised" only 2% of the times. Most people, instead, experience surprise rates between 15 and 20%. Thus, beware the investor who is 99% sure.

Hindsight. After an event has occurred, people cannot properly reconstruct their state of uncertainty before the event. Suppose you are asked:

On the day before the event, what was your probability of a 5% drop in the SP500?

After the facts, financial pundits and common investors believe that they have an exact explanation for what happened. It almost seems that the event was so inevitable that it could have been easily predicted. Hindsight has two important consequences. First, it tends to promote overconfidence, by fostering the illusion that the world is more predictable than it is. Second, it often turns (in investors' eyes) reasonable gambles into foolish mistakes.

Optimism. Most people entertain beliefs biased in the direction of optimism. Suppose you are asked:

Is your ability as a driver above or below the median?

and that the same question is posed to the members of a group. Usually, about 80% of the people believe that they are above the median. The combination of optimism and overconfidence leads investors to overestimate their knowledge and underestimate risks, leaving them vulnerable to statistical surprises. It also fosters the illusion of control, which makes people believe that risk can be managed by knowledge and trading skill. In a survey admistered to 45 investors, De Bondt (1998) reports that 89% agree with the statement "I would rather have in my stock portfolio just a few companies that I know well than many companies that I know little about" whereas only 7% agree with the statement "Because most investors do not like risk, risky stocks sell at lower market prices" and just 18% agree with the statement that "The risk of a stock depends on whether its price typically moves with or against the market."

Spurious regularities. People tend to spot regularities even where there is none. Suppose you are asked:

Which sequence is more likely to occur if a coin is tossed — HHHTTT or HTHTTH?

Although the two sequences are equally likely, most people wrongly believe that the second one is more likely. Odean (1998) reports that, when individual investors sold a stock and quickly bought another, the stock they sold on average outperformed the stock they bought by 3.4% in the first year (excluding transactions costs). This costly overtrading may be explained in terms of spurious regularities and overconfidence. A very frequent error is the extrapolation bias which makes people optimistic in bullish markets and pessimistic in bear markets. De Bondt (1993) reports that the average gap between the percentage of investors who is bullish and the percentage who is bearish increases by 1.3% for every percentage point that the *Dow Jones* raises in the previous week.

14.3 Distortions in deriving preferences

Changes matter more than states. People's preferences are especially sensitive to changes. Suppose you are asked two questions:

A: Imagine than you are richer by Euro 20,000 than you are today. Would you prefer an additional gain of 5,000 for sure or a 50–50 chance for a gain of 10,000?

B: Imagine than you are richer by Euro 30,000 than you are today. Would you prefer an additional loss of 5,000 for sure or a 50–50 chance for a loss of 10,000?

Although the final outcomes in the two problems are exactly the same, most people choose the gamble in Question A and the sure loss in Question B. Apparently, they tend to favor the narrow framing based on gains and losses rather than the broader (and more relevant) framing based on the final wealth.

Loss aversion. People's sensitivity to losses is higher than their sensitivity to gains. Suppose you are asked the question:

Consider a bet on the toss of the coin. If heads, you lose Euro 100. What is the minimum gain if tails that would make you accept the gamble?

Most answers typically fall in the range from 200 to 250, which reflects a sharp asymmetry between the values that people attach to gains and losses.

Probability weighting. With respect to the expected utility paradigm, people tend to weigh probabilities differently. Suppose you are asked:

Given a chance for a gain of Euro 20,000, would you pay more to raise the probability of gain from 0 to 1%, from 41 to 42%, or from 99 to 100%?

While expected utility predicts that the answer should be the same, most people would pay significantly less for raising the probability to 42%. In particular, low probabilities are overweighted: people tend to find a 1% chance of winning Euro 1,000 preferable to a sure Euro 10.

Prospect theory. The distortions in this section are often modelled using the prospect theory proposed in Kahneman and Tversky (1979), which suitably modifies the expected utility formulation. The major modifications are three. First, the utility function is defined over changes in wealth rather than wealth levels. Second, the slope of the utility function over changes in wealth is greater for losses than for gains. Third, the agent generates decision weights w_i from the probability distribution and maximizes $\sum w_i u(x_i)$. Note that the decision weights are not necessarily interpretable as probability weights. This basic model is often enriched by the assumption that the utility function is concave over gains and convex over losses: people are risk averse when dealing with gains and risk prone when dealing with losses.

Shape and attractiveness. The following table list eight gambles in Euros as ranked by financial analysts according to their attractiveness to investors. They have the same expected value and the same number of possible outcomes.

Gamble	Payoff 1	Prob.ty (%)	Payoff 2	Prob.ty (%)
A	5,000	95	105,000	5
B	5,000	50	15,000	50
C	1,000	10	11,000	90
D	1,000	90	91,000	10
E	2,000	50	18,000	50
F	0	50	20,000	50
G	$-2,000$	90	118,000	10
H	$-5,000$	50	25,000	50

The ideal gamble combines a high probability of a moderate gain and a small probability of a very large gain. Its preferability may be explained by a combination of the overweighting of the small probability of a large gain and of the different risk attitudes over gains and losses.

14.4 Framing effects

Reference point. People's trading decisions are affected by reference points which act as performance benchamrks. Suppose you are asked the question:

Primus owns 100 shares of an asset, originally bought at 100. Secunda owns 100 shares of the same asset, originally bought at 200. The price of the share was 160 yesterday and it is 150 today. Who is more upset?

It is the initial price who determines how the stock is doing. A consequence of this is the disposition effect: an investor who needs cash and owns two stocks is much more likely to sell the stock whose price has gone up; see Odean (1998a).

Compartmentalization. People often fail to take into account the joint effect of their decisions. Suppose you are asked the question:

You face two concurrent decisions. Examine both decisions and state the options you prefer for each decision.

Decision A: choose between a sure gain of 2,400 and a 25% chance of a gain of 10,000.

Decision B: choose between a sure loss of 7,500 and a 75% chance of a loss of 10,000.

In accord with loss aversion, most people tend to choose the sure gain in A and the gamble in B. Now, consider the following question:

Choose between a 25% chance to win 2,400 and a 75% chance to lose 7,600 versus a 25% chance to win 2,500 and a 75% chance to lose 7,500.

There is no doubt about the right choice: the second option dominates the first one. However, note that the pair chosen in the first question amounts to the inferior option, while the pair rejected is equivalent to the dominating choice. People find it easier to deal with decisions one at a time and tend to forget the overall picture. A common effect is the use of a "mental accounting" strategy by which people classify their savings (or expenses) in different segments and deal with them in different ways; for instance, a windfall rise in income rarely finds its way in the (mentally separated category of) retirement savings or health insurance.

References

[1] W.F.M. De Bondt (1993), "Betting on trends: Intuitive forecasts of financial risk and return", *International Journal of Forecasting* **9**, 355–371.

[2] W.F.M. De Bondt (1998), "A portrait of the individual investor", *European Economic Review* **42**, 831–844.

[3] D. Kahneman and M. Riepe (1998), "Aspects of investor psychology", *Journal of Portfolio Management* **24**, 52–65.

[4] D. Kahneman and A. Tversky (1979), "Prospect theory: An analysis of decision under risk", *Econometrica* **47**, 263–291.

[5] T. Odean (1998), "Do investors trade too much?", *American Economic Review* **89**, 1279–1298.

[6] T. Odean (1998a), "Are investors reluctant to realize their losses?", *Journal of Finance* **53**, 1775–1798.

15. Behavioral finance: asset pricing

15.1 Introduction

There is an enourmous discrepancy between returns on stocks and fixed income securities. Between 1926 and 1990, for instance, the annual real rate of return on U.S. stocks has been about 7%, while the real return on U.S. treasury bills has been less than 1% percent. Mehra and Prescott (1985) shows that the combination of a high equity premium, a low risk-free rate, and smooth consumption cannot be reconciled with plausible levels of investors' risk aversion within a standard paradigm of expected utility maximization. This has become known as the equity premium puzzle.

Mehra and Prescott estimate that the coefficient of relative risk aversion should be higher than 30 to explain the historical equity premium, whereas previous estimates and theoretical arguments suggest that the actual figure should be close to 1. A coefficient of relative risk aversion close to 30 corresponds to a situation like the following one: suppose that Primus is offered a 50–50 gamble between USD 100,000 and USD 50,000; then his certainty equivalent should be around USD 51,209. Certainly this is an extreme risk aversion.

Reitz (1988) has argued that the equity premium may be a rational response to a time-varying risk of economic catastrophe, as bonds protect capital investment better than equity. This explanation is not testable. Moreover, since the data from 1926 contain the 1929 crash, the catastrophe in question must be of greater magnitude. Also, the hypothetical catastrophe should affect equity but not bonds: thus, hyperinflation would not qualify.

A different line of research has managed to explain part of the equity premium introducing nonexpected utility preferences. In particular, Constantinides (1990) suggested a habit-formation model in which the utility of consumption depends on past levels of consumption as well. People become averse to reductions in their consumption and this may be used to explain the equity risk premium. Campbell and Cochrane (1999) perfect this intuition with a carefully crafted model. While this sort of models are probably on the right track, emphasizing only consumption-based habit-forming neglects the weighty role of pension funds, endowments, and very wealthy individuals with long horizons.

This leaves us with two questions that are still open: why is the equity premium so large, and why is anyone willing to hold bonds? This lecture is meant to show you how this sort of questions can be approached using insights from behavioral finance.

15.2 Myopic loss aversion

Benartzi and Thaler (1995) puts forth a behavioral explanation of the equity premium puzzle based on a partial equilibrium model. They exploit two ideas from the psychological evidence about of decision-making.

The first notion is *loss aversion*, which refers to the tendency for individuals to be more sensitive to reductions in their levels of well-being than to increases. This translates into a slope of the utility function which is greater over wealth decrements than over increments.

The second notion is *mental accounting*, which refers to the practice of implicitly earmarking financial outcomes as belonging to different accounts. This bears relevance on how outcomes are aggregated: because of loss aversion, aggregation rules may not be neutral.

Here is an example drawn from Barberis and Huang (2001). An investor named Primus exhibits loss aversion, modelled by a utility function over wealth increments such as $u(x) = x$ if $x \geq 0$ and $u(x) = 2x$ if $x < 0$. Primus is thinking about buying a portfolio of one share each of two stocks, which are currently trading at 100. Primus believes that their values one year from now will be either 70 or 150 with equal probability and that the two distributions are stochastically independent.

If Primus is loss averse over portfolio fluctuations, the expected utility of the investment is

$$\frac{1}{4}u(100) + \frac{1}{2}u(20) + \frac{1}{4}u(-60) = 5;$$

while if he is loss averse over individual stock fluctuations, it is

$$2\left[\frac{1}{2}u(50) + \frac{1}{2}u(-30)\right] = -10,$$

which is obviously not as attractive.

The relevance of mental accounting for the equity premium puzzle can be seen by confronting Primus with the choice between a risky asset paying an expected 7% per year with standard deviation of 20% and a safe asset yielding a sure 1%. The attractiveness of the risky asset depends on the time horizon of the investor, which is another form of mental accounting. The longer Primus' investment horizon, the more attractive the risky asset, *provided that the investment is not evaluated frequently.* It is the combination of loss aversion and a short evaluation period (which we call *myopic loss aversion*) that makes the investor unwilling to bear the risks associated with holding equities.

15.3 A partial equilibrium model

Let us go back to the equity premium puzzle. Suppose that Primus' preferences conform to prospect theory. With respect to the standard expected utility formulation, this makes three important modifications to take into account observed behavior: i) utility is defined over wealth increments; ii) it exhibits loss aversion; ii) the expected utility $\sum w_i u(x_i)$ is computed using decision weights w_i which distort the actual probabilities.

The second element necessary to build a multi-period model is a specification of the length of time over which an investor aggregates returns; that is, his *evaluation period.* This is not his investment horizon. Consider a young investor saving for retirement 30 years off in the future who gets a newsletter from his mutual fund every quarter. If he experiences the utility associated with his wealth increments every quarter (when he gets the newsletter), this agent has an evaluation period of three months and an investment horizon of thirty years. Accordingly, he will behave as someone else whose investment horizon is just one quarter.

Mehra and Prescott (1985) investigates the equity premium by asking how risk averse should be the representative investor to explain historical evidence. Benartzi and Thaler (1995) approaches the puzzle by asking how long should be the evaluation period of an investor with prospect theory preferences to explain the equity premium.

An answer is obtained using simulations based on the historical (1926–1990) monthly returns on stocks, bonds, and treasury bills. The stock index is compared both with treasury bills returns and with five-year bond returns, and these comparisons are done both in real and nominal terms. It is argued that the use of bonds is preferable because they are more profitable substitutes for long-term investors. And it is argued that nominal terms are preferable because they are used in most annual reports (and because evaluation in real terms would yield negative prospective utility over most periods). However, the results remain robust under any of the four possible specifications.

It is found that the evaluation period that makes a portfolio of 100% stock indifferent to a portfolio of 100% bonds in nominal terms is about 13 months. (If the comparison is made in real terms, the equilibrium period is between 10 and 11 months. If bills are used in place of bonds, this period is one month shorter.) This suggests that an evaluation period of about 12 months may lead people to consider bonds as feasible alternative to stocks.

An obvious criticism to this findings is that most people prefer to invest in portfolios containing both stocks and bonds. A second simulation is thus run, checking (under 10% increments) which mix of bonds and stocks would maximize prospective utility. Portfolios carrying between 30 and 55% of stocks all yield approximately the same prospective value. This result is consistent with observed behavior. For instance, the most frequent allocation in TIIA-CREF (a very large defined benefit retirement plan in U.S.) is 50-50.

As the evaluation period lenghtens, stocks become more attractive. The actual equity premium in the data used was 6.5% per year, and this is consistent with an evaluation period of one year. What happens if the evaluation period lengthens? With a two-year evaluation period, the premium falls to 4.65; with a five-year period, it falls to 3.0%, and with 20 years to 1.4%. Therefore, assuming 20 years as the benchmark case, we can say that the price of excessive vigilance is about 5.1%.

A common asset allocation for pension funds has about 60% in stocks and 40% in bonds. Given that it is reasonable to assume that pension funds have an infinite investment horizon, they should favor stocks much more. A possible explanation links myopic loss aversion with an agency problem. Although the pension fund has an infinite investment horizon, its managers must report annually on the performance of their investments and cannot afford negative returns over long periods. Their choice of a short horizon creates a conflict of interest between the manager and the stockholders.

Another source of a conflict of interest is the rule adopted in foundations and trusts that only a fixed percentage of an n-year moving average of the value of the endowment (usually, $n \leq 5$) can be spent every year. The goals of maximizing the present value of spending over an infinite horizon versus maintaining a steady operating budget compete against each other.

15.4 An equilibrium pricing model

Barberis et alii (2001) builds on Benartzi and Thaler (1995) to provide a behavioral explanation of the equity premium puzzle based on an equilibrium pricing model. The new twist they add is that prior outcomes influence the way gains and losses in wealth are experienced.

Thaler and Johnson (1990) finds that a loss is *less* painful to people when it comes after substantial earlier increases in wealth: the earlier gains "cushion" the subsequent loss and make it more bearable. In financial markets, this translates into a behavior such that investors care less for a market dip that follows substantial prior gains because they are "still up, relative to a year ago."

Starting from an underlying consumption growth process with low variance, the combination of prospect theory and the effect of prior outcomes can generate stock returns with high mean, high volatility and significant predictability, while maintaining a riskless interest rate with low mean and volatility. This is obtained in a Lucas-style consumption-based asset pricing model with a continuum of identical infinitely-lived agents and two assets: a riskfree asset in zero net supply and a risky asset with a total supply of one unit. Except for the modifications in investors' preferences, the (representative) investor (named Primus) is fully rational and dynamically consistent.

The driving force in the model is a story of changing risk aversion. After a run-up in prices, Primus is less risk averse because those gains will cushion any subsequent loss. After a fall in stock prices, he is more wary of further losses and hence more risk averse. This variation in risk aversion allows returns to be much more volatile than the underlying dividends: an unusually good dividend raises prices, but this price increase also makes Primus less risk averse, driving prices still higher.

This process generates a predictability in returns very similar to what is empirically observed: following a significant rise in price, the investor is less risk averse and subsequent returns are therefore on average lower. Moreover, since the high volatility of returns leads to frequent losses for stocks, the loss averse investor requires a high equity premium to be willing to hold stocks.

References

[1] N. Barberis and M. Huang (2001), "Mental accounting, loss aversion, and individual stock returns", *Journal of Finance* **56**, 1247–1292.

[2] N. Barberis, M. Huang and T. Santos (2001), "Prospect theory and asset prices", *Quarterly Journal of Economics* **116**, 1–53.

[3] S. Benartzi and R.H. Thaler (1995), "Myopic loss aversion and the equity premium puzzle," *Quarterly Journal of Economics* **110**, 73–92.

[4] J.Y. Campbell and J.H. Cochrane (1999), "By force of habit: A consumption-based explanation of aggregate stock market behavior," *Journal of Political Economy* **107**, 205–251.

[5] G.M. Constantinides (1990), “Habit formation: A resolution of the equity premium puzzle,” *Journal of Political Economy* **98**, 519–543.

[6] R. Mehra and E.C. Prescott (1985), “The equity premium puzzle,” *Journal of Monetary Economics* **15**, 145–161.

[7] T. Reitz (1988), “The equity risk premium: A solution?,” *Journal of Monetary Economics* **21**, 117–132.

[8] R.H. Thaler and E.J. Johnson (1990), “Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice”, *Management Science* **36**, 643–660.