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.

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!

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

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.

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 $s_y = 0$ 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

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

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.

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.

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] 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.

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.