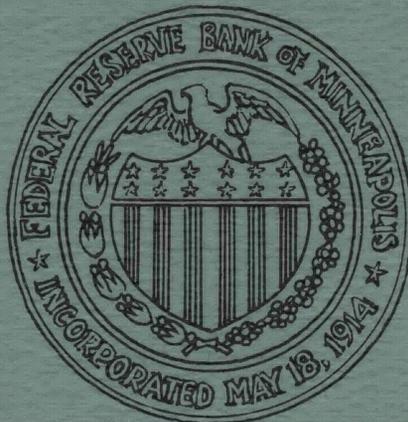
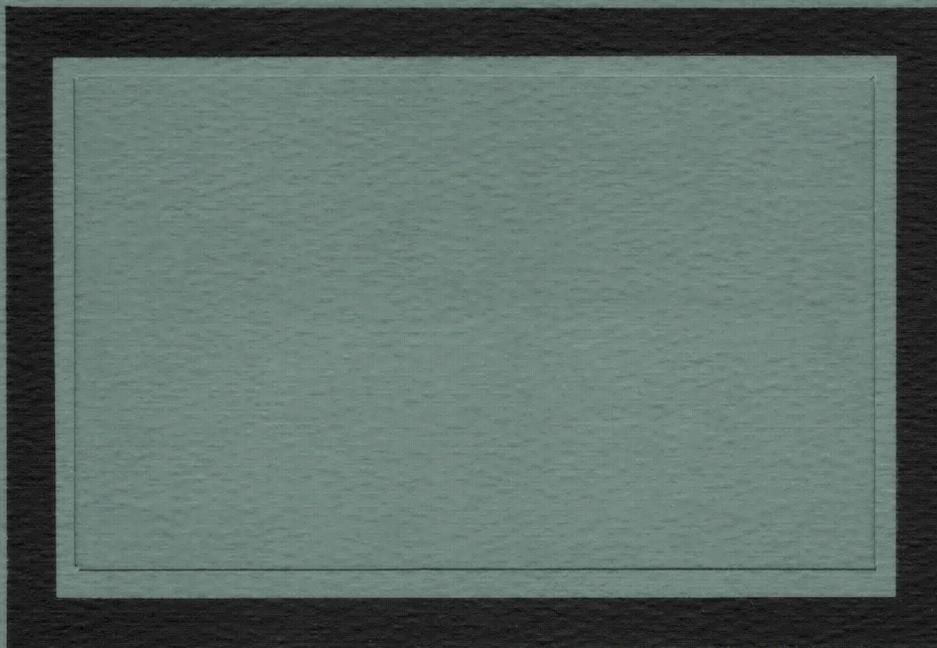


NOT FOR DISTRIBUTION

STAFF REPORT

working paper



**research
department**

federal reserve bank of minneapolis

Notes on Behavior Under Uncertainty

Thomas J. Sargent

March 1974

Working Paper #~~33~~ 32

Rsch. File #252.1

Thomas Sargent
March 1974

Notes on Behavior Under Uncertainty

We assume that the individual's preferences can be described by a utility function that makes utility depend on the amount of the one good that the individual consumes. So we write

$$(1) \quad \text{Utility} = U(C)$$

where C is the amount of the good consumed. We assume that utility is increasing in consumption, $U'(C) > 0$; that the marginal utility of consumption, though positive, decreases with increases in C , $U''(C) < 0$; and that $U(C)$ is bounded for $C \in [0, \infty)$.

We assume that the individual is making plans for the future, which we collapse to a single date in the future. To incorporate the existence of uncertainty, we assume that there are n mutually exclusive states of the world, indexed by $\theta = 1, 2, \dots, n$. The individual has a set of subjective probabilities $\pi(\theta) \geq 0$ giving the probability that he assigns to the occurrence of state θ .

An individual's claim to future consumption goods will in general depend on the state of nature that happens to occur. We let $C(\theta)$ denote his consumption if state θ occurs. The individual is assumed to maximize his expected utility

$$v = \sum_{\theta=1}^n \pi(\theta)U(C(\theta)),$$

subject to certain constraints.

Suppose that there are only two states of the world, so that $n = 2$. The individual's expected utility is then

$$v = \pi(1)U(C(1)) + \pi(2)U(C(2)).$$

Along lines of constant expected utility (indifference curves) we have

$$dv = 0 = \pi(1)U'(C(1))dC(1) + \pi(2)U'(C(2))dC(2)$$

which implies that

$$(2) \quad \left. \frac{dC(2)}{dC(1)} \right|_{dv=0} = - \frac{\pi(1)U'(C(1))}{\pi(2)U'(C(2))}.$$

Expression (2) gives the slope of the indifference curve. The concavity of the indifference curve is found by differentiating (2):

$$\frac{d^2C(2)}{dC(1)^2} = - \left[\frac{\pi(1)U''(C(1))}{\pi(2)U'(C(2))} \right] + \frac{\pi(1)U'(C(1))U''(C(2))\frac{dC(2)}{dC(1)}}{\pi(2)U'(C(2))^2} > 0$$

The slope of the indifference curves increases as $C(1)/C(2)$ increases, implying that they are convex.

Next notice that for $C(1) = C(2)$, we have

$$(3) \quad \left. \frac{dC(2)}{dC(1)} \right|_{\substack{dv=0 \\ C(1)=C(2)}} = - \frac{\pi(1)}{\pi(2)}.$$

Bundles of $C(1)$, $C(2)$ for which $C(1) = C(2)$ correspond to certain claims, since regardless of the state that occurs, the individual is able to consume the same amount. A 45° line through the origin in $C(1)$, $C(2)$ plane thus contains all certain bundles (see Figure 1), so that it is appropriately called the certainty line. Along the certainty line, the

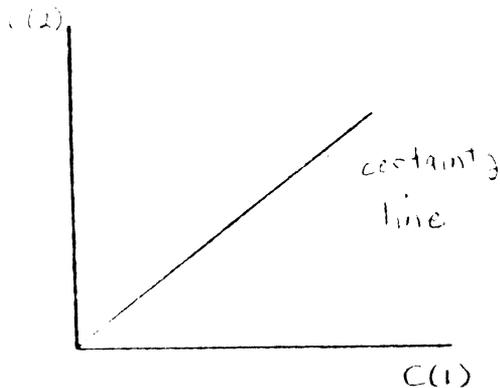


Figure 1

slope of the indifference curves equals $-\pi(1)/\pi(2)$, independently of the form of the utility function-- so long as the form of the utility function does not itself depend on the state of nature that occurs. The certainty line is thus an "expansion path."

Consider an individual whose initial endowment consists of a certain claim on Y_0 units of consumption goods. The individual is then confronted with a bet which he can undertake in any amount α so long as $\alpha \leq Y_0$. If α units of the bet are taken, the individual will receive an additional $\alpha X(1)$ units of the consumption good if state 1 occurs, but must sacrifice α units of output if state 2 occurs. Thus the payoff, cost $(dC(1), dC(2))$ associated with taking α goods worth of the bet is $(\alpha X(1), -\alpha)$. The bet is said to be "favorable" if its expected value in terms of goods is positive. The expected value of the bet's payoff stream is

$$\begin{aligned} \pi(1)\alpha X(1) - \pi(2)\alpha &= \pi(1)\alpha X(1) - (1-\pi(1))\alpha \\ &= \pi(1)\alpha(X(1) + 1) - \alpha. \end{aligned}$$

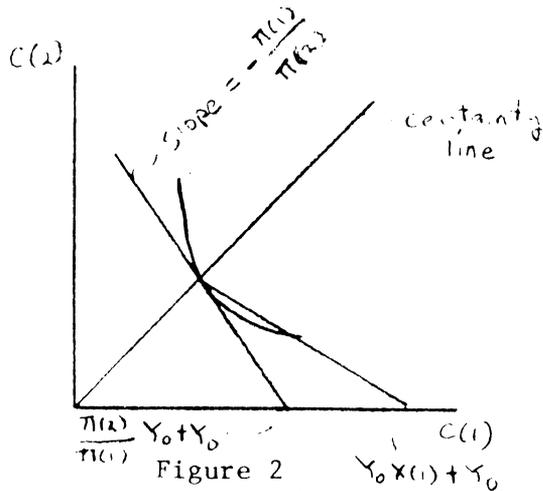
The bet is then favorable if

$$\pi(1)\alpha(X(1) + 1) - \alpha > 0$$

or

$$(4) \quad \pi(1) > \frac{1}{X(1) + 1}.$$

The bet is said to be "fair" if the above inequality is replaced by an equality.



If the individual undertakes α units of the bet, his claims to consumption across states of nature become

$$C(1) = Y_0 + \alpha X(1)$$

$$C(2) = Y_0 - \alpha,$$

from which we can deduce that by varying the amount of the bet taken, α , the individual can substitute $C(1)$ for $C(2)$ at the (constant) rate

$$(5) \quad \frac{dC(2)}{dC(1)} = - \frac{1}{X(1)} .$$

So $-1/X(1)$ is the slope of the "budget line" through (Y_0, Y_0) along which the individual can trade claims to consumption in state 1 for claims to consumption in state 2. If $\alpha = 0$, the individual's claims remain (Y_0, Y_0) . If $\alpha = Y_0$, the individual's claims become $(Y_0 X(1) + Y_0, 0)$. The straight line connecting these two points is the individual's budget line (see Figure 2).

As long as the slope of the budget line exceeds the slope of the individual's indifference curves at Y_0, Y_0 , the individual can increase his expected utility by undertaking at least a small part of the bet.

This requires that

$$- \frac{1}{X(1)} > - \frac{\pi(1)}{\pi(2)} = - \frac{\pi(1)}{1-\pi(1)} ,$$

from equations (3) and (5). The above inequality can be rearranged to read

$$\pi(1) > \frac{1}{X(1) + 1} ,$$

which is identical with inequality, (4), the condition that the bet be favorable. For our special case, we have thus proved Arrow's proposition that an individual will always take at least a small part of a favorable bet.¹

Within this framework, we now consider securities that entitle the individual to alternative patterns of consumption across our two states of nature. Consider a security, one unit of which entitles the owner to receive $X(1)$ units of the consumption good if state 1 occurs and $X(2)$ units if state 2 occurs. If the individual buys α units of the security, he is entitled to receive a pattern of returns $(\alpha X(1), \alpha X(2))$ across states of nature. In Figure 3, one unit of the security gives

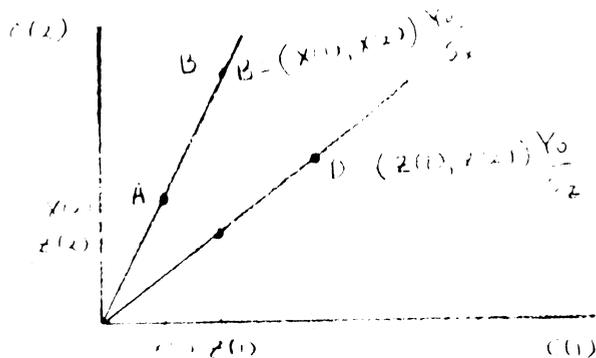


Figure 3

the returns labeled by point A.

Suppose that the security costs the individual S_X units of current output per unit of security.

If the individual has an investment

portfolio worth Y_0 units of current output, he could then buy Y_0/S_X units of the security and obtain a pattern of returns $(Y_0 X(1)/S_X, Y_0 X(2)/S_X)$ across states of nature. Point B in Figure 3 depicts such a pattern of returns.

Now suppose that there is a second security, one unit of which gives a pattern of returns $(Z(1), Z(2))$ across states of nature, where $Z(1)$ and $Z(2)$ are both denominated in consumption goods. If one unit of the security costs S_Z , the individual could purchase Y_0/S_Z units of the security if he put his whole portfolio of Y_0 current goods into that

security. Then his pattern of returns across states of nature would be $(Y_0 Z(1)/S_Z, Y_0 Z(2)/S_Z)$. Such a pattern of returns across states is labeled D in Figure 3.

Now suppose that the individual considers putting a percentage λ of his portfolio into security X, and $1-\lambda$ into security Z. He would then purchase $\lambda Y_0/S_X$ units of security X, and $(1-\lambda)Y_0/S_Z$ units of security Z. His pattern of returns across states of nature would then be

$$(6) \quad (C(1), C(2)) = Y_0 \left(\lambda \frac{X(1)}{S_X} + (1-\lambda) \frac{Z(1)}{S_Z}, \lambda \frac{X(2)}{S_X} + (1-\lambda) \frac{Z(2)}{S_Z} \right).$$

Such points are linear combinations of $(X(1), X(2))Y_0/S_X$ and $(Z(1), Z(2))Y_0/S_Z$, and so lie on the straight line connecting points D and B in Figure 3.

A change in λ brings changes in the consumption stream across states according to

$$dC(1) = \left[\frac{X(1)}{S_X} - \frac{Z(1)}{S_Z} \right] Y_0 d\lambda$$

$$dC(2) = \left[\frac{X(2)}{S_X} - \frac{Z(2)}{S_Z} \right] Y_0 d\lambda,$$

so that the "budget line" along which the consumer can alter the pattern of claims to the consumption good across states of nature has slope

$$\frac{dC(2)}{dC(1)} = \frac{\frac{X(2)}{S_X} - \frac{Z(2)}{S_Z}}{\frac{X(1)}{S_X} - \frac{Z(1)}{S_Z}}.$$

For this slope to be negative, the numerator and denominator must be of opposite sign, which means that one security must not dominate another.

That is, one unit of current output's worth of security X must offer more consumption in state 1 if it offers less in state 2 than does one unit of current output's worth of security 2.

By suitably choosing λ (which need not be between 0 and 1), the individual is able to obtain any combination of $C(1)$ and $C(2)$ in the nonnegative quadrant satisfying equation (6). A negative λ or one exceeding unity indicates that one security or the other is being sold short or being issued by the individual (see Figure 4). Notice that by choosing his

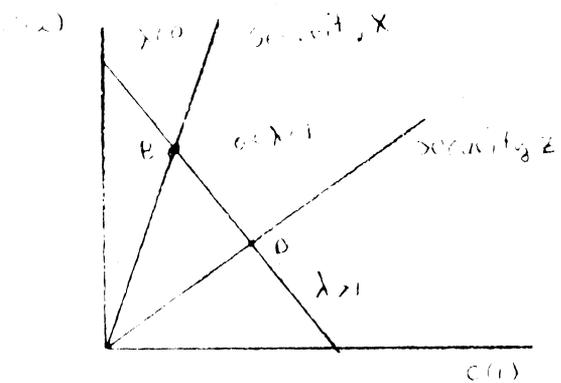


Figure 4

portfolio suitably, the individual can set $C(1)$ equal to $C(2)$, so that he need bear no risk, if that is his desire.

The individual chooses his portfolio so as to maximize

his expected utility subject to the budget constraint (6). Usually, this involves choosing λ so that it corresponds to a point of tangency between an indifference curve and the budget line. As always, however, corner solutions are possible.

Suppose now that a third security, security y , is added to our setup. The security has returns across states $(y(1)/S_y, y(2)/S_y)$ measured in consumption goods in states 1 and 2 per unit of current consumption good; S_y is the price of one unit of the security in terms of current consumption goods. Now unless the returns $(y(1)/S_y, y(2)/S_y)$ can be expressed as a linear combination of the returns of securities X and Z, one of the securities will not be held. To determine which of the securities will be held, simply plot the points $(X(1)/S_X, X(2)/S_X)$,

$(Z(1)/S_z, Z(2)/S_z)$, and $(y(1)/S_y, y(2)/S_y)$. Then determine which straight line through any two of the points provides the individual with the best opportunity locus or budget line in $C(1), C(2)$ space. These two securities will be held while the third will not (see Figure 5). The result will

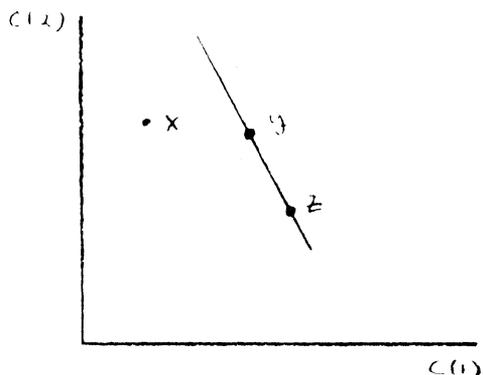


Figure 5

then be that the price of the security that isn't held will fall (or perhaps the prices of the others will rise) until all three points lie along the same line. Thus arbitrage requires that where

there are two states of the world there be at most two securities whose returns are linearly independent. Similar reasoning implies that where there are n states of the world, the returns on at most n securities can be linearly independent (i.e., the rank of the matrix of returns on securities across states is at most n).

It is sometimes analytically convenient to work with "pure" securities that pay off one unit of consumption in state i and nothing in any other state. Such securities were introduced by Arrow and are known as Arrow-Debreu contingent securities. The return vector for

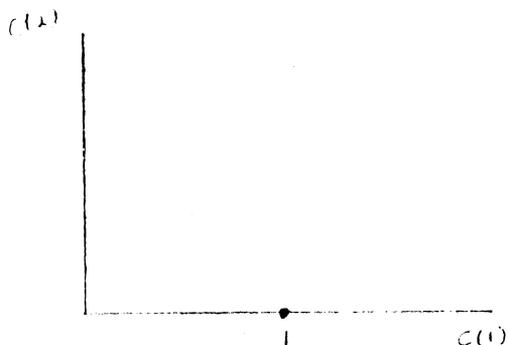


Figure 6

such a security lies along one of the axes. Thus, in our 2-state example, one unit of a state 1 contingent security offers a return vector $(1,0)$, while one unit of a state 2 security offers a return vector $(0,1)$ (see Figure 6).

Even where such contingent securities don't literally exist, it is possible effectively to "trade" them and to compute implicit prices for them where the number of ordinary securities equals the number of states of nature. For example, consider our 2-state example where securities X and Z exist. Security X derives its value from the value that consumers attach to the consumption stream the security delivers. Let $p(i)$ be the amount of current output an individual would sacrifice to obtain one more unit of consumption in state i . Then it must be so that

$$S_X = X(1)p(1) + X(2)p(2)$$

(7)

$$S_Z = Z(1)p(1) + Z(2)p(2),$$

i.e., the price of each real security must reflect the value of the consumption streams that the security represents a claim on. The above equations can be solved for $p(1)$, $p(2)$, the implicit prices of the contingent securities, so long as $X(1)Z(2) - Z(1)X(2) \neq 0$, i.e., so long as the returns on securities X and Z are not linearly dependent.

Liquidity Preference as Behavior Towards Risk²

Tobin's explanation of the demand for money as emerging partly as a result of wealthholders' desire to diversify their holdings can be viewed as an application of the theory just described. Suppose that there are two states of the world and that there are two assets: a risky asset that pays off $X(\theta)$ in state $\theta = 1, 2$ for each unit of current output's worth of the asset; and a riskless asset called "money" that pays off one unit of current output, regardless of state, for each unit of current output invested in it. From our preceding discussion, we know that the household will hold at least a little of the risky asset provided that the expected rate of return is positive, i.e., provided that holding the risky security amounts to undertaking a favorable bet. By investing one sure unit of output ("money") in the risky asset, the investor obtains an expected return of

$$\pi(1)X(1) + \pi(2)X(2)$$

which must exceed unity if the security is to offer the individual a favorable bet. Notice that for money to be held, it must be so that either $X(1) < 1$ or $X(2) < 1$, or else the risky asset would dominate money. The expected rate of return on the risky asset, denoted by r , is given by

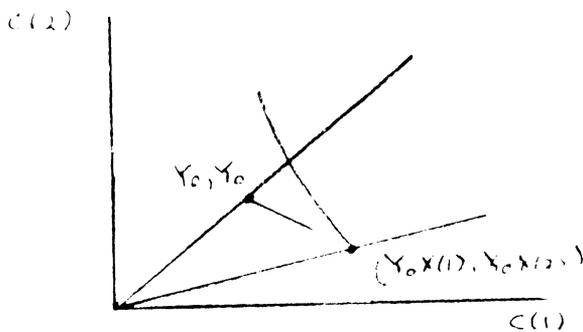
$$r = \pi(1)X(1) + \pi(2)X(2) - 1.$$

It is easy to show that if we start from a position in which $r = 0$, an increase in r , i.e., an increase in either $X(1)$ or $X(2)$, will cause the investor to increase his holdings of the risky asset and decrease his holdings of money. (This is an example of Arrow's proposition that at least a small part of a favorable bet will be undertaken. We leave it

to the reader to work out the details.) Clearly, by risk aversion, if $r = 0$, the investor will hold his entire portfolio in terms of money. Notice that we have established that at low enough interest rates, the investor's holdings of money will vary inversely with the interest rate on risky assets.

At higher interest rates, an increase in r (i.e., in $X(1)$ or $X(2)$) may or may not cause holdings of money to contract. As usual, there are two effects: a substitution effect inducing a movement along an indifference curve, an effect which leads to lower money holdings; and a wealth or income effect, which may or may not offset the substitution effect, depending on the shape of the investor's indifference curves.

Notice that it is possible that the investor will want to hold no money (though if $r > 0$, he will always want to hold some of the risky asset). This will occur if the rate of return on the risky asset is so high that the situation is as depicted in Figure 7, where the budget line is flatter than the indifference curve even where the investor's entire portfolio is in the risky asset.



(Y_0, Y_0) = returns vector if whole portfolio held in money.

$(Y_0 X(1), Y_0 X(2))$ = returns vector if whole portfolio held in risky asset.

Figure 7

The theory described above has often been embodied in a somewhat different form, set forth by Tobin. As above the individual is assumed to maximize expected utility

$$v = \sum_{\theta=1}^n \pi(\theta)U(C(\theta)).$$

If we know $C(\theta)$ and $\pi(\theta)$ for each θ , it is straightforward to deduce a probability distribution $g(C)$ which gives the probability that consumption will obtain the value C . In particular,

$$g(C) = \left[\sum_{\theta \in T} \pi(\theta), T = \left\{ \theta \mid C(\theta) = \tilde{C} \right\} \right].$$

In the finite-state case currently under discussion, $g(C)$ will obtain a nonzero, positive value at only a finite number of $n' \leq n$ values of consumption C . Denote these values of consumption as $C_1, C_2, \dots, C_{n'}$. Then expected utility v can be written as

$$(8) \quad v = \sum_{i=1}^{n'} U(C_i)g(C_i),$$

where $\sum_{i=1}^{n'} g(C_i) = 1$.

In a setup with a continuum of states of the world and where consequently C is allowed to take on any real value, the probability associated with consumption occurring in a neighborhood of width ϵ around C is given by $f(C;B)$ where $f(C;B)$ is the distribution function associated with C and B is a list of parameters determining that distribution. In this case expected utility v is

$$(9) \quad v = v(B) = \int_{-\infty}^{\infty} U(C)f(C;B)dC.$$

Here expected utility is a function only of the parameters B determining the distribution of consumption. If there is only one parameter in B, as would be true if C were distributed according to the Poisson distribution, then expected utility would depend only on the value of that one parameter. If there are p parameters in B, then expected utility depends on all p of them.

The theory has been developed for distributions $f(C;B)$ which can be characterized by two parameters--one measuring mean or central tendency, the other measuring variance. The normal distribution is an example of such a distribution, being completely characterized by the mean and variance of the distribution. Members of the class of stable distributions of Paul Levy are also characterized by two parameters.

Following Tobin, suppose B consists of the mean μ_C and standard deviation σ_C of consumption, so that

$$f(C;B) = f(C; \mu_C, \sigma_C).$$

It greatly facilitates the analysis also to assume that $f(C;B)$ is a "stable" distribution. A variate Z with density $f(Z;B_2)$ is said to be stable if when another variate y with the same form of density $f(y;B_y)$, perhaps with $B_Z \neq B_y$, is added to Z , the result is to produce a variate $X = Z + y$ obeying the same probability law $f(X;B_X)$. Assuming the distribution $f(C;B)$ is stable is natural because stable distributions are the only distributions that serve as the limiting distribution in central limit theorems. The normal distribution is the best known of stable distributions. The central limit property of stable distributions is useful here because the random variable C is often thought of as representing a linear combination of a large number of independently distributed returns on various investments, implying that it will approximately follow a stable distribution.

Assuming that $f(C;B)$ has two parameters, mean μ_C and standard deviation σ_C , and that it is also a stable distribution amounts to assuming that it is a normal distribution.³ That is because the normal distribution is the only (symmetric) stable distribution for which the standard deviation exists. Then expected utility v is

$$v(\mu_C, \sigma_C) = \int_{-\infty}^{\infty} U(C) f(C; \mu_C, \sigma_C) dC.$$

Defining the standardized variable Z as

$$Z = \frac{C - \mu_C}{\sigma_C},$$

we have that $C = \mu_C + \sigma_C Z$. Then

$$v = v(\mu_C, \sigma_C) = \int_{-\infty}^{\infty} U(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ,$$

where $f(Z; 0, 1)$ is the standard, unit variance normal distribution.

Since expected utility $v(\mu_C, \sigma_C)$ depends only on the two parameters μ_C, σ_C we can define indifference curves in the μ_C, σ_C plane, i.e., combinations of μ_C and σ_C that yield constant levels of expected utility. Along such curves, we have

$$\begin{aligned} dv = 0 &= d\mu_C \int_{-\infty}^{\infty} U'(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ \\ &+ d\sigma_C \int_{-\infty}^{\infty} Z U'(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ, \end{aligned}$$

so that the slope of the indifference curves in μ_C, σ_C plane is

$$(10) \quad \frac{d\mu_C}{d\sigma_C} = - \frac{\int_{-\infty}^{\infty} Z U'(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ}{\int_{-\infty}^{\infty} U'(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ}.$$

Since $U'' < 0$, while $f(Z; 0, 1)$ is symmetric the numerator on the right is negative so long as $\sigma_C > 0$ (negative Z 's being multiplied by larger U 's than positive Z 's); the denominator is positive since $U' > 0$. Thus the slope (10) is positive. To find the concavity of the indifference curves, we differentiate (10) with respect to σ_C to obtain

$$\frac{d^2 \mu_C}{d\sigma_C^2} = \frac{\int Z^2 U''(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ}{\int U'(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ}$$

$$- \frac{\frac{d\mu_C}{d\sigma_C} \int Z U''(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ}{\int U'(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ}$$

$$- \frac{d\mu_C}{d\sigma_C} \int Z U''(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ$$

$$- \frac{d\mu_C^2}{d\sigma_C^2} \int U''(\mu_C + \sigma_C Z) f(Z; 0, 1) dZ.$$

In each case, the limits of integration are $-\infty, \infty$. Since $Z^2 \geq 0$ and $U'' < 0$, the numerator of the first term is negative, making that term positive, since it is preceded by a negative sign. The numerator of the second term is also negative because in the large, U'' must be decreasing in absolute value as C increases in order for $U' > 0$ while $U'' < 0$ for all $C \in [0, \infty)$. This makes the second term positive. Likewise the third and fourth terms are also positive, taking into account the signs preceding them. Thus we have that

$$\frac{d^2 \mu_C}{d\sigma_C^2} > 0,$$

which shows that each indifference curve has a slope that increases as we move upward along a curve. An example of a map of such curves is depicted in Figure 8.

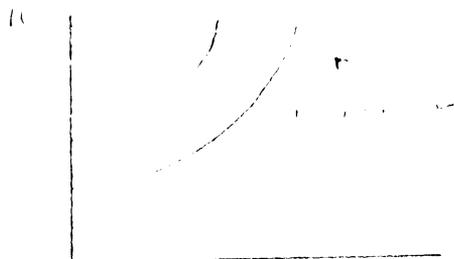


Figure 8

It is convenient to use equation (10) to compute the slope of the indifference curves at zero standard deviation. We have

$$\left. \frac{d\mu_c}{d\sigma_c} \right|_{\sigma_c = 0} = - \frac{\int ZU'(\mu_c) f(Z;0,1) dZ}{\int U'(\mu_c) f(Z;0,1) dZ} = \frac{\int Zf(Z;0,1) dZ}{\int f(Z;0,1) dZ} = \frac{-E(Z)}{1} = 0$$

The numerator equals zero, since the normal distribution is symmetric about $Z = 0$, and since $U'(\mu_c)$ is independent of Z . Thus the indifference curves have zero slope for $\sigma_c = 0$. This property of the indifference curves will be seen to reflect that an individual will always take at least a small part of a favorable risk.

To take a specific example, suppose

$$U(C) = -e^{-\lambda C} \quad \lambda > 0.$$

Notice that

$$U'(C) = \lambda e^{-\lambda C} > 0 \quad \text{for } C \in (-\infty, \infty)$$

$$U''(C) = -\lambda^2 e^{-\lambda C} < 0 \quad \text{for } C \in (-\infty, \infty).$$

The density function for the normal distribution is

$$f(C; \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(C-\mu)^2}{2\sigma^2}}.$$

Consequently expected utility is given by

$$\begin{aligned} E(U(C)) &= \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} -e^{-\lambda C} e^{-\frac{(C-\mu)^2}{2\sigma^2}} dC \\ (11) \quad &= \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} -e^{-(\lambda C + \frac{(C-\mu)^2}{2\sigma^2})} dC. \end{aligned}$$

Notice that

$$\begin{aligned} \lambda C + \frac{(C-\mu)^2}{2\sigma^2} &= \frac{2\lambda C\sigma^2 + C^2 - 2\mu C + \mu^2}{2\sigma^2} \\ &= \frac{[C-(\mu-\lambda\sigma^2)]^2 + 2\lambda\mu\sigma^2 - \lambda^2\sigma^4}{2\sigma^2} \\ &= \frac{[C-(\mu-\lambda\sigma^2)]^2}{2\sigma^2} + \lambda(\mu - (1/2)\lambda\sigma^2). \end{aligned}$$

Substituting the above expression into (11) gives

$$E[U(c)] = -e^{-\lambda(\mu - (1/2)\lambda\sigma^2)} \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{(C-(\mu-\lambda\sigma^2))^2}{2\sigma^2}} dC.$$

But we know that

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{(C-\mu')^2}{2\sigma^2}} dC = 1$$

for any μ' and $\sigma > 0$. So we have

$$E[U(C)] = -e^{-\lambda(\mu - \frac{\lambda\sigma^2}{2})}$$

Along curves of constant expected utility we require

$$\mu - (1/2)\lambda\sigma^2 = \text{constant}.$$

So the mean, standard deviation indifference curves satisfy

$$d\mu - \lambda\sigma d\sigma = 0$$

so that their slope is

$$\frac{d\mu}{d\sigma} = \lambda\sigma > 0 \quad \text{for } \sigma > 0.$$

and their concavity is

$$\frac{d^2\mu}{d\sigma^2} = \lambda > 0.$$

This concludes our analysis of our specific example for $U(C)$.

Having characterized the individual's preferences, we now describe his opportunities. Suppose that there is available to the individual a safe asset that has the property that if he puts his entire portfolio into this asset he will obtain a consumption stream characterized by mean C_0 and standard deviation zero. Suppose there is also an alternative asset (or maybe a portfolio of other assets) such that if the individual uses his entire portfolio to purchase this asset he obtains consumption goods in an amount C_0X where X is a normally distributed random variable with mean μ_X and variance σ_X^2 . If the individual invests a proportion $(1-\lambda)$ in the risky asset and λ in the safe asset, $0 \leq \lambda \leq 1$ he receives consumption

$$C = \lambda C_0 + (1-\lambda)C_0X.$$

Then the mean of his consumption would be

$$(12) \quad \mu_C = \lambda C_0 + (1-\lambda)C_0\mu_X.$$

Notice that

$$C - \mu_c = (1-\lambda)C_0(X - \mu_X),$$

so that

$$(13) \quad \sigma_c = (1-\lambda)C_0\sigma_X.$$

Solving (13) for λ , substituting into (12) and rearranging we obtain

$$(14) \quad \mu_c = C_0 + \left(\frac{\mu_X - 1}{\sigma_X}\right)\sigma_c, \quad 0 \leq \sigma_c \leq C_0\sigma_X$$

which gives the locus of combinations of μ_c, σ_c attainable by varying λ . The mean μ_c rises linearly with the standard deviation σ_c , the slope being $(\mu_X - 1)/\sigma_X$. Of course, the expected rate of return $(\mu_X - 1)$ must exceed zero for the opportunity locus to have a positive slope. This is the condition that the risk be favorable. Such an opportunity locus is depicted in Figure 9.




Figure 9

As we have seen, the

slope of the indifference curves at $\sigma_c = 0$ is zero. That means that if $\mu_X - 1 > 0$, the individual will always take at least a small part of the risk, since then the opportunity locus through $(0, C_0)$

has a positive slope, permitting the individual to move to a higher level of expected utility by taking some risk. It follows that beginning from a situation where $\mu_X - 1 = 0$, an increase in the rate of return on the risky asset, $\mu_X - 1$, will lead to a decrease in the amount held in the

safe asset and an increase in holdings of the risky asset. Thus, for a low enough rate of return on the risky asset, an increase in that rate does cause a decrease in the investor's demand for the safe asset. For higher values of the rate of return on the risky asset, however, an increase in that rate will not necessarily lead to a decrease in holdings of the safe asset, there being offsetting substitution and wealth effects. We leave it to the reader to study these offsetting effects in the context of the present graphical formulation of the theory. Needless to say, all of these features of the analysis have their counterparts in the state-preference version of the theory which we summarized above.

There are several unsatisfactory aspects of the theory that we have just sketched. For the formulation cast in terms of the mean and standard deviation of consumption, we have to assume that consumption is normally distributed, which requires that we act as if consumption can be an unbounded negative number. It is difficult to imagine negative consumption. If to circumvent this difficulty we restrict consumption to be nonnegative, we must pay for this by adopting a probability function for consumption that lacks the statistical property of stability, and so greatly weakens the appeal of the theory. But as we have seen above, the essence of the theory can be cast in terms of the state-preference analysis where the assumption that consumption has a normal distribution need play no role.

As a theory of the demand for money, the theory is certainly of limited applicability. For one thing, the occurrence of unforeseen price level changes makes money a risky asset in terms of goods, so that the "money" in the model above does not really correspond with the asset

called money in the real world. For another thing, money is dominated by assets like Treasury bills and savings deposits that are as risk-free as money but offer positive nominal yields. At best, the above theory is one about the demand for such assets, not money. To explain the demand for money it seems essential to take into account the presence of transactions costs.

The Modigliani-Miller Theorem

Throughout these pages, we have assumed that firms have no bonds outstanding, that they retain no earnings, and so they finance all of their investment by issuing equities. It is an implication of the "Modigliani-Miller theorem" that our assumptions about these matters are not restrictive.⁴ In particular, Modigliani and Miller's analysis implies that in the absence of a corporate income tax, the firm's cost of capital is independent of whether the firm raises the funds by retaining earnings, issuing bonds, or issuing equities. Moreover, Modigliani and Miller's theorem was proved in the context of a model that explicitly recognized the existence of uncertainty. These notes sketch the reasoning of Modigliani and Miller by using the state-preference presentation of Stiglitz.⁵

We collapse the entire future into a single point in the future. We assume that there is a finite number n of possible future states of the world, each state representing an entire constellation of possible outcomes of all sorts of events in the future. We let $\theta = 1, 2, \dots, n$ be an index over the possible states. In state $\theta = 1$, for example, it rains two inches in Eugene, Oregon, Ali defeats Foreman in the ring, Nixon wins a third term, and so on. States $\theta = 2, \dots, n$ correspond to different outcomes of this set of events. An individual's happiness, indexed by U , in the event that state θ prevails depends on the usual way on the amounts of n goods that he consumes:

$$U = U(q_1(\theta), \dots, q_m(\theta)) \quad \partial U / \partial q_i(\theta) > 0$$

$U(\)$ concave

where $q_i(\theta)$ is the amount of the i^{th} good consumed by the individual in state θ , $i = 1, \dots, m$. We have assumed that the form of the utility function $U(\cdot)$ is independent of the state θ .

The consumer's notions about the likelihood of various states of the world occurring are supposed to be summarized by a set of subjective probabilities $\pi(1), \pi(2), \dots, \pi(n)$ that obey

$$\sum_{\theta=1}^n \pi(\theta) = 1,$$

where $\pi(\theta)$ is the probability that the consumer assigns to state θ occurring. Individuals are assumed to maximize expected utility v :

$$(1) \quad v = \sum_{\theta=1}^n \pi(\theta) U(q_1(\theta), \dots, q_m(\theta)).$$

The consumer is assumed to come into a certain endowment q_i^0 of claims to goods $i = 1, \dots, m$, should state θ occur, $\theta = 1, \dots, n$. It is assumed that there exist competitive futures markets in which individuals trade claims to the i^{th} good in state θ prior to the occurrence of the state. The individual faces a price $p_i(\theta)$ at which he can buy or sell whatever claims he wishes on the i^{th} good contingent on state θ occurring. The value of the consumer's endowment is

$$\sum_{\theta=1}^n \sum_{i=1}^m p_i(\theta) q_i^0(\theta).$$

The consumer maximizes expected utility v subject to

$$(2) \quad \sum_{\theta=1}^n \sum_{i=1}^m p_i(\theta) q_i^0(\theta) = \sum_{\theta=1}^n \sum_{i=1}^m p_i(\theta) q_i(\theta),$$

which states that the market value of his endowment equals the market value of the bundle of contingent commodities that he purchases. Where λ is an undetermined Lagrange multiplier, the consumer's problem can be formulated as maximizing

$$J = \sum_{\theta=1}^n \left\{ \pi(\theta) U(q_1(\theta), \dots, q_m(\theta)) + \lambda \left(\sum_{i=1}^m p_i(\theta) (q_i(\theta) - q_i^0(\theta)) \right) \right\}.$$

The first order conditions are

$$(3) \quad \pi(\theta) \frac{\partial U}{\partial q_i(\theta)} + \lambda p_i(\theta) = 0 \quad i = 1, \dots, m$$

$$\frac{\partial J}{\partial \lambda} = 0 \quad \theta = 1, \dots, n.$$

Dividing (3) for θ and i by (3) for θ and j , we have

$$(4) \quad \frac{\frac{\partial U}{\partial q_i(\theta)}}{\frac{\partial U}{\partial q_j(\theta)}} = \frac{\tilde{\pi}(\theta) p_i(\theta)}{\pi(\theta) \tilde{p}_i(\theta)},$$

which is the analogue of the familiar static marginal equality for the household. From (4) and the budget constraint (2), demand curves for the nm contingent commodities can be derived. By aggregating these demand curves over the set of all consumers, market demand schedules can be obtained, which together with total market endowments permits computing a general equilibrium in which the prices $p_i(\theta)$, $i = 1, \dots, m$, $\theta = 1, \dots, n$ are determined.

Arrow⁶ has shown that consumers are just as well off where these nm markets in m commodities contingent on state $\theta (= 1, \dots, n)$ occurring are replaced by n markets in "contingent securities," with one security for each state. Each security promises to pay one dollar should state θ occur. Following the occurrence of a state, consumers then trade the m goods as described by the standard static model.

It is straightforward to add production to the model sketched above. We consider such a competitive model in which there exists a complete set of n markets for the n contingent securities, each promising to pay one dollar if state θ occurs in the future. The model is assumed to possess a general equilibrium in which the equilibrium present price of a claim to one dollar in state θ is $p(\theta)$. Notice that the price of a sure dollar next period is $\sum_{\theta=1}^n p(\theta)$, which can be interpreted as the reciprocal of one plus the risk-free rate of interest. The assumption that there exist perfect markets in the contingent securities for all n states of the world means that it is possible to insure against any risk. Individuals need bear no risks if that is their preference.

We will assume no taxes are present. Now consider a firm whose prospective returns, net of labor and materials costs, but gross of capital costs, are $X(\theta)$ dollars in state θ . Suppose that the firm issues an amount of B dollars worth of bonds. The firm now promises to pay $(r+1)B$ dollars to its bond holders next period, provided that it does not go bankrupt, i.e., provided that $X(\theta) \geq (r+1)B$. If the firm does go bankrupt, i.e., if $X(\theta) < (r+1)B$, then the bond holders receive only $X(\theta)$. Thus the realized rate of return on bonds $r(\theta)$ depends on the state of the world:

$$r(\theta) + 1 = \begin{cases} r + 1 & \text{if } X(\theta) \geq (r+1)B \\ \frac{X(\theta)}{B} & \text{if } X(\theta) < (r+1)B. \end{cases}$$

Only if $X(\theta) > (r+1)B$ for all θ , is $r(\theta)$ equal to the promised coupon rate r for all θ .

The value of the firm's bonds must equal the sum of the values of the contingent securities that the bond implicitly consists of. For each state in which the firm doesn't go bankrupt, the bonds will in total pay off $(r+1)B$. The present value of those returns is

$$(r+1)B \sum_{\theta \in S} p(\theta) \quad \text{where } S = \{\theta | X(\theta) \geq (r+1)B\}$$

For states θ in $S' = \{\theta | X(\theta) < (r+1)B\}$, in which the firm goes bankrupt, the bonds pay off $X(\theta)$. So the present value of payments in those states is

$$B \sum_{\theta \in S'} \frac{X(\theta)}{B} p(\theta).$$

The total present value of the firm's bonds B must thus satisfy

$$B = (r+1)B \sum_{\theta \in S} p(\theta) + B \sum_{\theta \in S'} \frac{X(\theta)}{B} p(\theta).$$

Dividing by B and solving for $(r+1)$, we obtain

$$(5) \quad (r+1) = \frac{1 - \sum_{\theta \in S'} \frac{X(\theta)}{B} p(\theta)}{\sum_{\theta \in S} p(\theta)},$$

which tells us that the rate of return a firm's bonds must bear depends on the firm's probability of defaulting, and so on the number of bonds it has issued. Notice that if there is zero probability of the firm's going bankrupt, S' being empty, r equals the risk-free rate of interest.

The firm's equities bear a payout stream across states of nature given by

$$\begin{aligned} X(\theta) - (r+1)B & \quad \text{if } X(\theta) \geq (r+1)B \\ 0 & \quad \text{if } X(\theta) < (r+1)B. \end{aligned}$$

As with bonds, the value of the firm's equities must equal the sum of the values of the contingent securities that the equities implicitly represent. So we have that the present value of equities E is

$$(6) \quad E = \sum_{\theta \in S} (X(\theta) - (r+1)B)p(\theta).$$

Substituting for $(r+1)$ from (5) in the above expression gives

$$\begin{aligned} E &= \sum_{\theta \in S} p(\theta)X(\theta) - B \left(\frac{1 - \sum_{\theta \in S'} \frac{X(\theta)}{B} p(\theta)}{\sum_{\theta \in S'} p(\theta)} \right) \sum_{\theta \in S} p(\theta) \\ &= \sum_{\theta \in S} p(\theta)X(\theta) - B + \sum_{\theta \in S'} X(\theta)p(\theta) \end{aligned}$$

$$(7) \quad E = \sum_{\theta \in S \cup S'} p(\theta)X(\theta) - B$$

or

$$(8) \quad E + B = \sum_{\theta \in S \cup S'} p(\theta)X(\theta).$$

Equation (8) states that the total value of the firm's debt plus equity equals the present value of the firm's return across state of nature, evaluated at the price of claims to one dollar contingent on the associated states of nature. The total value $E + B$ is therefore independent of the ratio of debt to equity.

Now assume that the firm is contemplating a project that costs C sure dollars today, and that will cause the firm's returns to change by $dX(\theta)$ in state θ . The value of stockholders' equity if the project

isn't undertaken is given by (7). If the project is undertaken, the value of the original stockholders' equity will be

$$E' = \int_{\theta} p(\theta)X(\theta) - B + \int_{\theta \in S} p(\theta)dX(\theta) - C.$$

The value of the original stockholders' equity is increased by undertaking the project so long as

$$\int_{\theta} p(\theta)dX(\theta) - C > 0;$$

the project ought to be undertaken by the firm so long as the above inequality is met because it will increase the value of the equity of initial stockholders. This is true regardless of whether the project is financed by issuing bonds or more equities. In particular, notice that the rate of interest r on the firm's bonds, which depends on the volume of bonds that the firm has outstanding, is not pertinent in helping the firm determine whether or not to undertake the project.

Effects of a Corporate Income Tax

We now suppose that the firm's profits net of interest payments to bond holders are taxed at a corporate profits tax rate t_K . The returns to stockholders then equal $(1-t_K)(X(\theta)-(r+1)B)$ for states in S , i.e., states satisfying $X(\theta) \geq (r+1)B$, and zero for states in which bankruptcy occurs. The interest rate r on the firm's bonds continues to obey (5). The value of the firm's equities is now given by

$$E = \sum_{\theta} (1-t_K)(X(\theta)-(r+1)B)p(\theta).$$

Substituting for r from (5) in the above equation and rearranging gives

$$(9) \quad E = (1-t_K) \sum_{\theta=1}^n X(\theta)p(\theta) - B + t_K B.$$

For $t_K > 0$, the value of the firm, $E + B$, varies directly with the stock of bonds outstanding. Equation (9) thus predicts that it is in stockholders' interest to have the firm levered an indefinitely large amount. The presence of the corporate income tax implies that there is an optimal debt-equity ratio for the firm (one indefinitely large) and thus causes the Modigliani-Miller theorem to fail to hold.

We should note that matters become much more complex when individual income taxes with different rates for interest income and capital gains are included in the analysis.

What it Means for There to be Markets in
Contingent Securities

Suppose that markets for the n contingent securities do not actually exist and that claims promising one dollar if state θ occurs never are traded. Instead, as in the real world, there are n' different companies each selling claims entitling the owner to share in the company's profits. It is easy to show that so long as there are more independent companies than states of nature, it is as if there existed markets in n contingent securities, since by buying and selling actual securities in the proper fashion, the individual can obtain any desired pattern of returns across states of nature.

We suppose that the i^{th} firm's returns across states of nature are given by $X_i(\theta)$, $\theta = 1, \dots, n$. Select n such firms each of whose patterns of returns across states of nature are not linearly dependent on the returns of the remaining $(n-1)$ firms. That is, for each $i = 1, \dots, n$ the vector $(X_i(1), X_i(2), \dots, X_i(n))$ cannot be written as a linear combination of the $(n-1)$ vectors $(X_j(1), X_j(2), \dots, X_j(n))$ for $j \neq i$. Let the market values of our n firms be V_1, V_2, \dots, V_n . If there were n contingent securities, each promising to pay one dollar in state θ and having price $p(\theta)$, the values of the n firms would have to obey

$$V_1 = \sum_{\theta=1}^n X_1(\theta)p(\theta)$$

$$V_2 = \sum_{\theta=1}^n X_2(\theta)p(\theta)$$

.

.

.

$$V_n = \sum_{\theta=1}^n X_n(\theta)p(\theta),$$

or in compact notation

$$(10) \quad V = Xp$$

$$\text{where } V = \begin{bmatrix} V_1 \\ \cdot \\ \cdot \\ \cdot \\ V_n \end{bmatrix}, \quad X = \begin{bmatrix} X_1(1)X_1(2)\dots X_1(n) \\ \cdot \\ \cdot \\ \cdot \\ X_n(1)X_n(2)\dots X_n(n) \end{bmatrix}, \quad p = \begin{bmatrix} p(1) \\ \cdot \\ \cdot \\ \cdot \\ p(n) \end{bmatrix}.$$

Since X is of full rank, (10) can be used to solve for p , giving

$$(11) \quad p = X^{-1}V.$$

Equation (11) tells us how to unscramble the implicit prices of the n implicit contingent securities from the market values of n firms and the patterns of their returns across states of nature.

Footnotes

1/ See Kenneth Arrow, Essays in the Theory of Risk Bearing, p. 100.

2/ James Tobin, "Liquidity Preference as Behavior Towards Risk," Review of Economic Studies, 1958.

3/ The density function for the normal distribution is

$$f(c; \mu_c, \sigma_c) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(c-\mu_c)^2}{2\sigma_c^2}} .$$

With this distribution, for expected utility to be defined, the utility function $U(C)$ must satisfy

$$|U(C)| \leq Ae^{BC^2}, \quad A > 0, \quad B > 0$$

for some A and B . See J. S. Chipman, "The Ordering of Portfolios in Terms of Mean and Variance," Review of Economic Studies, April, 1973.

4/ F. Modigliani and M. H. Miller, "The Cost of Capital, Corporation Finance, and the Theory of Investment," American Economic Review, June 1958.

5/ J. Stiglitz, "A Re-Examination of the Modigliani-Miller Theorem," American Economic Review, December 1969.

6/ K. Arrow, "The Role of Securities in the Optimal Allocation of Risk Bearing," Review Economic Studies, 1964.

