

NBER WORKING PAPER SERIES

INFORMATION AND CAPITAL MARKETS

Joseph E. Stiglitz

Working Paper No. 678

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge MA 02138

May 1981

The research reported here is part of the NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

Information and Capital Markets

ABSTRACT

This paper provides the foundations of a general theory of information and the capital market.

We show that in a pure gambling market, even with asymmetric information, there cannot exist an equilibrium with trade with rational individuals. We argue that although a pure exchange stock market is not a pure gambling market, most of the trade on the stock market arises from irrationality on the part of some investors and the rational response on the part of other investors to take advantage of that irrationality.

We show that the private returns to information acquisition and dissemination differ markedly from social returns and as a result the market equilibrium is not a (constrained) Pareto optimum. Moreover, we show how firms' actions, e.g. the fraction of shares retained by the original entrepreneurs, the debt equity ratio, and the level of investment, may convey information about firm characteristics. This in turn affects the behavior of firms. As a result, the original owners of firms will be incompletely diversified, firms will not take actions which maximize their stock market value, and, in particular, they may behave in a risk averse manner, paying attention to own risk (which traditional theory suggests that the only risk firms should care about is the correlation with the market).

Professor Joseph E. Stiglitz
Department of Economics
Dickinson Hall
Princeton University
Princeton, N.J. 08544

(609) 452-4014

INFORMATION AND CAPITAL MARKETS*

by

Joseph E. Stiglitz

I. Introduction

Is the stock market just a gambling casino? Are too many resources being spent on obtaining information by individuals who are attempting to beat the market? Is it in fact possible to beat the market?

The answer to these questions is important for an evaluation of the role of the stock market in our economy. In neoclassical economic theory, the stock market provides a crucial link between consumers and producers. In theory, the stock market is supposed to provide the signals for firms to make the correct investment decisions. If, at the extreme, individuals had no "information" about the relative probabilities of different events and the relative merits of different investment opportunities, prices would simply be random, and the stock market would be unable to perform its role in allocating investment.

There are two contrasting views of information in the stock market: one that if the market worked well, so prices reflected values, there would be no incentive to obtain information and therefore, there cannot exist an informed market equilibrium. Even if prices did not reflect true values, to the extent that insiders' gains are at the expense of outsiders, rational outsiders will

*This is a revised version of Part II of a paper presented at the New Orleans meetings of the Econometric Society, December, 1971. Since that time, the literature on the subjects discussed here has grown enormously. I have not attempted to reference all of these more recent contributions. I am indebted to M. Rothschild, C. von Weizsacker, R. Lindsey, V. Krishna, S. Salop, A. Dixit, and A. Weiss for extremely helpful discussions. Forthcoming in Financial Economics: Essays in Honor of Paul H. Cootner, edited by Cathryn Cootner and William F. Sharpe, Prentice-Hall. My interest in the questions discussed here was greatly stimulated by Paul Cootner, both directly and through his writings. His influence on my thinking, and my indebtedness to him, should be apparent.

refuse to "play"; and if that happens, the insiders have no one from whom they can make money.

In contrast, Hirschleifer has argued that in a pure exchange market, there are strong incentives for information acquisition; for any individual who finds out information about which state of nature is about to occur, e.g. which securities are going to rise in price and which are going to fall, will be able to make a profit out of the capital gain.

A simple example along the lines of that employed by Hirschleifer might be instructive. Let there be two factories. If state 1 occurs factory 1 has an output of Q , factory 2 an output of $\frac{1}{2}Q$; if state 2 occurs, factory 1 has an output of $\frac{1}{2}Q$, factory 2 an output of Q . Initially, with individuals assuming the two states equally probable if they choose a portfolio to maximize their expected utility then the relative price of the two factories will be unity. If an individual knows that state 1 will occur, he will demand shares of factory 1 and sell short shares of factory 2 so long as the relative price of factory 1 shares is less than 2. When it then becomes publicly known that state 1 will occur, the price of the shares of firm 1 will rise, those of firm 2 will fall, and the informed speculator will thus make a large profit.

Indeed, Hirschleifer argues that because the gains from information acquisition are purely redistributive, any information acquisition has no social value: some gain at the expense of others.

These two hypotheses appear to be in direct contradiction of one another: the one suggesting that even when there is a social return to investing in information, there may be no investment in information, the other

suggesting that even when there is no social return to investing in information, there will be some investment. We argue that both positions are wrong--the latter because it implicitly assumes irrationality and/or noncompetitive behavior, the former because it fails to recognize the role of the seller of securities in providing information, because it fails to take account of the presence of "noise" which results in prices imperfectly conveying information from the informed to the uninformed,¹ and because it fails to note that, even with free entry into the "information" industry, equilibrium may be characterized by partial monopolies, in which particular individuals may have a monopoly of particular pieces of information.²

Both views suffer from failure to identify fully the nature of the social and private gains to information.

It is clear that central to an understanding of the functioning of capital market is an analysis of

- (a) the incentives, within a market economy, for individuals to acquire information;
- (b) the extent to which market prices reflect the information of informed individuals; and
- (c) the role that market prices play in determining the behavior of managers of firms.

We argue that within a competitive economy there are only limited incentives to acquire information which is of general value;

¹This aspect of the argument is developed more fully in subsequent work by Grossman and Stiglitz (1976, 1980a) and Grossman (1977).

²This aspect of the argument is developed briefly in Stiglitz (1975a).

indeed, in one central, but limiting case, of a pure exchange economy
the only equilibrium entails zero expenditure
on information (and no trade). In more general cases, where
 individuals differ, there will be some incentives for acquiring
 information, but, if information is costly, market prices reflect
the information of the informed individuals only imperfectly.

The fact that market prices/^{do not}reflect the true value of firms has
 two important consequences:

Owner-managers who know that their firms are undervalued will
 diversify imperfectly. Firms which they control will behave in a
 risk averse manner, even though their output (profits) have a zero
 correlation with the market, and even though the market values
 them in a risk neutral manner. Firms do not, we argue, pursue a
 market value maximizing policy but are not, at the same time,
 subjected to the threat of take-over.

Secondly, good firms may attempt to convey to the market information
 about their quality. Since direct statements (such as that their
 prospects of returns are very good) have only limited credibility,
 indirect screening devices may be employed. Since owner-managers
 of good firms will be more willing to retain a larger proportion
 of their own firms than are owner-managers of poor firms, the
 fraction of shares owned by insiders may serve as an effective
 screening device.¹ In order to do this, however, owner-managers
 of good firms have to retain a larger fraction of their own shares

¹In the more recent literature on imperfect information, this
 is sometimes referred to as an example of a self-selection device;
 and the fraction of shares owned is said to serve as a signal.
 See Stiglitz (1982), Spence (1973) and Salop and Salop (1976).

than they otherwise would, so that even in situations where the market prices do accurately reflect the returns, the firm acts in a risk averse manner. There are other situations, however, where the only equilibrium is one in which the market does not differentiate between "good" and "bad" firms; in order for the fraction of shares retained by the owner-manager not to convey information about quality, the poor quality firms have to imitate the good firms, i.e. they too have to retain significant fractions of their shares.

Other policies, such as the debt-equity ratio and dividend policy may also serve as screening devices. We do not pursue the implications of our analysis for these other important aspects of financial policy here.¹

II. Can There Be an Informed Securities Market?

In this section, we consider in somewhat more detail the paradox discussed in the introduction; alternative resolutions of the paradox are presented, of which only one will be pursued in subsequent sections of the paper.

We limit ourselves to a pure exchange economy. Individuals are endowed with ownership claims on different "factories." The output of the factories is distributed at the end of the period to different individuals in proportion to their ownership claims at the time.

¹See, for instance, Leland and Pyle (1977), Ross (1977), Stiglitz (1982), and Stiglitz and Weiss (1980, 1981).

Between the beginning of the period and the end of the period there are, say, T trading periods, during which individuals can trade their ownership claims with one another. There is little loss of generality in assuming only two trading periods and it simplifies the discussion considerably. In the intervals between the beginning of the period and the trading points, individuals can expend resources to acquire information about the output of the different factories, or about the probabilities of different states of nature.

We assume everyone has the same tastes and the same initial endowments (i.e. their ownership shares in the different firms are identical). Assume that the output in the different states of nature of each of the firms is known¹ but the relative probabilities of the different states are unknown. Assume, again for simplicity, that there are only two states of nature. Initially, all individuals assume the two states equally likely. Consider the following possible scenarios:

(1) No one does research. Then at the first trading date, the relative price of the different factories will reflect the individual's priors that the two states of nature are equally likely. Since all individuals have the same priors, the same tastes, and the same endowments, no trade will actually take place. The prices established on the market at the two trading dates are identical.

(2) One person does research. Could the situation described above be an equilibrium? Hirschleifer argues that it is not. For an

¹Indeed conventionally, we define a state of nature as a complete specification of all the outcomes of the different firms.

individual can come along, invest some resources in obtaining information about which state of nature will occur. Assume he obtains the information before the first trading point. He then knows that the market value of some firms is too high, some too low. He sells short the overpriced firms and buys the underpriced firms with the proceeds. In the interval between the first and second trading period he makes the information public; this raises the price of the underpriced firms and lowers that of the overpriced firms, enabling our investor-in-information to make a capital gain on both parts of the transaction and to pay for the cost of research. Thus, the expenditure on information turns out to be profitable. Since the amount available for consumption is unchanged by the expenditure on information, Hirschleifer has termed this expenditure on information socially wasteful. One person is made better off, at the expense of others who are worse off, and because it requires resources to obtain information the gains of the former are smaller than the losses to the latter.

There are three objections to this argument:

a) The uninformed individuals in this argument are particularly naive. At the minimum, they should discover that, were they to have held on to their original portfolios, they would have done better than to have engaged in speculative activity. They should have, in other words, pursued what I call later, a non-speculative strategy. Their expected utility from pursuing such a strategy is higher

than that from pursuing the naive policy assumed by Hirschleifer-- and it is equally unsophisticated.

b) It might be argued, on the other hand, that the uninformed individuals are more sophisticated than suggested just now; that they observe that the market price conveys information--when the price of a security for a particular state is bid up, it is because the informed know that the state is going to occur. In this particular case, the implication of this is that again there will be no trade, but whereas in the previous case, they refused to trade, here they have a demand function for securities, but the only market clearing price entails no trade.

c) Finally, there is no reason to believe that only one individual will do research. Indeed, consider what happens if only two individuals do research, and they both come up with the same result about which state of nature is to occur before the first trading period. Assume that they behave competitively (rather than collusively). Then they will bid against each other and drive up the price of the securities which otherwise would have been undervalued, and drive down the price of the securities which otherwise would have been overvalued. Thus, when the next trading period occurs, the prices will be the same as they were at the first trading point, and neither individual makes any gains (or losses); but they are both worse off than those who did no research, since they have invested in resources to obtain information.

Thus, Hirschleifer's argument for excessive expenditure on information implicitly requires both irrationality on the part of the uninformed and a monopoly on the part of the informed. It is

hardly a model of a rational competitive securities market.¹

This leaves open the question of the possibility of a rational, competitive, informed securities market. We shall argue that

(i) A rational pure exchange market with trades cannot exist if individuals have identical endowments, tastes, and information.

(ii) Moreover, if the level of information is endogenous, a rational, informed market cannot exist.

III. Pure Gambling Markets

There are two central properties of what we shall refer to as pure gambling markets. (i) The gains of one individual are completely at the expense of some other individual, i.e. it is a pure exchange market; and (ii) the outcomes (events on which the bet is placed) are uncorrelated with all participants' other sources of income. The latter condition is what distinguishes gambling markets from insurance markets. The reason that trade (gambling) occurs is that individuals have different opinions about the probability distribution of the outcomes. The prototype of a pure gambling market is a horse race. One of the questions to which we shall turn later is whether the stock market ought to be viewed as a pure gambling market.

There is a widespread view that rational individuals who are risk averse should not (or would not) gamble. This, it is believed, would be true even if individuals could acquire information about

¹It is not an answer to this to suggest that it is unlikely that two individuals will come up with exactly the same information at exactly the same moment. For if one individual comes up always first, then he is in effect a monopolist in information and the market is not competitive; if different individuals come up with the information first in different periods, then the individual is a temporary monopolist--the market is still not competitive in the conventional sense. In either case, for Hirschleifer to be correct requires that the uninformed also be irrational, in the sense to be defined below.

the various events. The fact that some information about the relative merits of different horses can be acquired does not make it any more rational to gamble on horses. Equilibrium in a pure gambling market with rational risk averse individuals thus entails no trades (no gambling) and no information. These contentions, if true, clearly have profound implications for understanding the nature of the stock market, if the stock market can be described as a pure gambling market.

In this section, we examine the nature of equilibrium in pure gambling markets with rational, risk averse, individuals. Critical to the analysis is an understanding of the concept of rationality. The central question is, what is reasonable for the individual to know? What does he observe, and how does he make inferences on the basis of these observations?

Our concern here is with repeated markets, like horse races. The reason that there is a widespread presumption that betting on horse races is irrational is that there is enough experience with gambling in such markets that rational "individuals should have learned that they can't win." (This is not a completely persuasive argument, because some individuals do win, and it is this observation which keeps others in the market.) The intuitive belief that gambling is irrational is, however, I think, correct. Our objective is to clarify the precise sense in which this is so.

In the subsequent discussion we shall explore several alternative notions of rationality. We begin our analysis with a concept we refer to as weak rationality. We shall say that an individual's

expectations are weakly rational if the expected value of his gains are equal to his average gains.

Note that rationality in this sense is much weaker than that which usually goes under the rubric "rational expectations" where it is the expectations on the whole distribution of returns that conform to ex post realizations.

We now show:

Proposition 1. If all individuals have weakly rational expectations, and are risk averse, there will be no trading in a pure gambling market.

Proof. For simplicity, assume individuals' income apart from the gamble is given. Let the gamble be defined as follows: if a particular even E -occurs, the individual receives 1; if it does not occur, he pays $p/1-p$. The amount he bets is chosen to

$$(3.1) \max U(Y^j + B^j)\pi^j + U(Y^j - \frac{B^j p}{1-p})(1-\pi^j)$$

where π^j is the j th individual's estimate of the probability; hence

$$(3.2) \frac{U'(Y^j + B^j)\pi^j}{U'(Y^j - \frac{pB^j}{1-p})} = \frac{p}{1-p}$$

Clearly

$$(3.3) B^j > 0 \text{ as } \frac{\pi^j}{1-\pi^j} > \frac{p}{1-p}$$

The expected gain is

$$(3.4) |B^j [\pi^j - \frac{p}{1-p} (1-\pi^j)]|$$

Thus if betting occurs ($|B^j| > 0$),

some individuals must have expectations of gains which are not realized.

q.e.d.

On the other hand, the following converse of Proposition 1 can also be easily established:

Proposition 2: If there are some individuals who do not display rationality, and some who do, then equilibrium in a pure betting model entails some gambling.

Proof. We established earlier that

$$(3.5) \quad B^j \begin{matrix} > \\ < \end{matrix} 0 \quad \text{as} \quad \frac{\pi^j}{1-\pi^j} \begin{matrix} > \\ < \end{matrix} \frac{p}{1-p}.$$

Let π^* be the "true" probability of the event. By assumption, for some individuals

$$(3.6) \quad \pi^j = \pi^*$$

while for others

$$(3.7) \quad \pi^k \neq \pi^*.$$

Hence, for at least two individuals, j and k

$$(3.8) \quad \pi^j \neq \pi^k$$

Assume there were no trade. Then for all individuals

$$(3.9) \quad \frac{\pi^j}{1-\pi^j} = \frac{p}{1-p}$$

contradicting (3.8).

In this discussion so far, we have said nothing either about the consequences of differential information or about the incentives to acquire information. It seems plausible that in most interesting situations (including the stock market) individuals have (or can

obtain) information, that different individuals obtain different information, and these differences are what make such markets function. We shall show that even when individuals have different information, provided they are rational in a particular sense to be explained below, there will be no trade in a pure gambling market.

First, however, we examine a somewhat different notion of rationality, in which there will be trades in the pure gambling market. We shall say that the market exhibits group rationality if the average value of individuals' estimates of the probability of the states is equal to the true probability. Assume, for instance, that individuals are betting on the proportion of red balls in a very large urn containing red and black balls. Each individual is allowed to draw 100 balls from the urn. The proportion of red balls in their sample provides an unbiased statistic on the proportion of red balls in the urn. Thus the average value of individuals' estimates of the probability of red balls is equal to the true probability. If individuals bet on the basis of their estimate, and if the distribution of estimates is symmetric we can show

Proposition 3: The market odds provide an estimate of the probabilities which are biased towards .5, for events with probability near .5 or for concentrated distributions. More generally, market odds will not equal the true probability, even with a large population.

Proof:

Let Π^* denote the true probability. Consider the demand and supply of bets of individuals who over and underestimate Π by an amount x . By the assumption of a symmetric distribution, the

number of such individuals is the same. Assume $p = \Pi^*$. Consider first the case of $\Pi^* = 1/2$. It is clear that $B(-x) = -B(x)$ where $B(x)$ is the demand for bets of a person whose estimate of Π is $1/2 + x$. For if

$$(3.10a) \frac{U'(Y+B)(\frac{1}{2}+x)}{U'(Y-B)(\frac{1}{2}-x)} = 1 = \frac{p}{1-p},$$

then

$$(3.10b) \frac{U'(Y-B)(\frac{1}{2}-x)}{U'(Y+B)(\frac{1}{2}+x)} = 1 = \frac{p}{1-p}.$$

We now show that if $\Pi^* > \frac{1}{2}$, $B(x) < -B(-x)$, $x > 0$:

$$(3.11) \frac{d \ln \left\{ \frac{U'(Y+B)(p+x)}{U'(Y-\frac{pB}{1-p})(1-p-x)} / \frac{U'(Y-B)(p-x)}{U'(Y+\frac{pB}{1-p})(1-p+x)} \right\}}{dp} = \frac{2x(2p-1)}{(p^2-x^2)((1-p)^2-x^2)} \\ + \frac{B}{(1-p)^2} \left[\frac{U''(Y-\frac{pB}{1-p})}{U'(Y-\frac{pB}{1-p})} + \frac{U''(Y+\frac{pB}{1-p})}{U'(Y+\frac{pB}{1-p})} \right].$$

At $\Pi^* = \frac{1}{2}$ this is negative for all $B > 0$. Hence for Π^* sufficiently near $\frac{1}{2}$, at $p = \Pi^*$, when $\frac{U'(Y-B)(p-x)}{U'(Y+\frac{pB}{1-p})(1-p+x)} = \frac{p}{1-p}$,

$\frac{U'(Y+B)(p+x)}{U'(Y-\frac{pB}{1-p})(1-p-x)} < \frac{p}{1-p}$. Hence $B(x) < -B(-x)$. Thus, for

$$\int B(x) dx = 0, p < \Pi^*.$$

Similarly, for $x = 0$, $B = 0$; hence the R.H.S. of (3.11) also equals zero.

Taking the derivative with respect to x , we obtain at $x = 0$ (since

$$\frac{dB^j}{dx} = \frac{dB^j}{d\Pi^j} = 1/2\Pi^j(1-\Pi^j)A(Y) \quad \text{where } A = -U''/U'$$

$$\frac{2(2p-1)}{p^2(1-p)^2} - \frac{1}{p(1-p)^3} < 0, \quad \text{for } 0 < p \leq 1,$$

so again $B(x) < -B(-x)$.

For $|\Pi^* - \frac{1}{2}|$ large and x large, the market odds may provide either an under or overestimate.¹

Of course, in any market in which the market odds differ from the true probability there is an incentive to obtain accurate information concerning the true probability, provided such information is not too expensive.

Still, in the situation just described, with the market displaying group rationality, it seems foolish for individuals not to realize that others are betting on the basis of their sample information; their betting behavior ought to convey information, and the rational individual ought to take this into account in forming his expectations.

We shall say that an individual is rational if the expected return, conditional on whatever is observable to the individual, is equal to the average realized return on an (arbitrarily) large number of repetitions of the gamble. This concept differs from our earlier notion of weak rationality in that we require the individual to form his expectations conditional on whatever is observable;

¹Let $U'(Y) = e^{-Y}$. Let $\Pi^* = 3/4$, $x = \pm 1/8$. Then for $x = \frac{1}{8}$, if $p/(1-p) = 3$, from (3.2),

$$\left(\exp \frac{-B}{1-p}\right) \cdot 7 = 3 \quad \text{or} \quad B = \frac{1}{4} \ln \frac{7}{3}$$

while for $x = -1/8$, $\exp \frac{|B|}{1-p} \frac{5}{3} = 3$ or $|B| = \frac{1}{4} \ln \frac{9}{5} < \frac{1}{4} \ln \frac{7}{3}$

Hence $p > \frac{3}{4}$.

but like the concept of weak rationality, it entails a far less stringent notion of rationality than that associated with conventional rational expectations, since it requires only that the mean of the distribution coincide with the mean of the true distribution. Because, in general, different individuals may observe different signals, they will have different information and therefore different beliefs. One might have thought that therefore, in this situation, there could be trade on a gambling market. We shall now show that this is not the case.

Proposition 4. If all individuals have rational expectations, there can be no trade in a pure gambling market.

Proof. For simplicity, let us assume that we have two groups in the population, one of which only observes the market odds of the gamble, the other of which observes something else, which is correlated with the event on which the gamble is being made. The former group we shall refer to as the uninformed, the latter as the informed. Since the uninformed observe only the market odds, if they trade it must be the case that (using (3.2))

$$E(\pi^u | p) \neq p$$

where the superscript u indicates that it is the uninformed individuals' expectations.

The informed individuals form their expectations on some signal Ω , and possibly on p too. Clearly, their expected utility, if they ignore the information Ω , must be less than or equal to their expected utility if they do not ignore their information

$$\max_{\{B\}} E(U^I | \Omega, p) > \max_{\{B\}} E(U^I | p) . \quad (3.13)$$

Assume $B^*(\Omega, p)$ is the optimum trade function of the informed and $|B^*(\Omega, p)| > 0$ for some Ω .

But if, say,

$$E\{\Pi^u | p\} > p \text{ and } B^u > 0 , \quad (3.14)$$

clearly $B^I < 0$, and the expected return to gambling, for that p , is negative. Moreover, by hypothesis, both the signal Ω and the event are uncorrelated with the individuals' income. Hence, the gamble represents a mean reducing increase in the spread of the distribution of income, and hence expected utility is lowered:

$$E(E\{U^I | \Omega, p\} | p) < E\{U^I | p\} . \quad (3.15)$$

This contradicts (3.13). Hence

$$B^*(\Omega, p) \equiv 0 .$$

The result can easily be generalized to more than two groups. The basic intuition is simple: gambling cannot increase the mean income of all groups; at least one group must have a lower mean income. At the same time, gambling increases the dispersion of the distribution of income. Hence, there must be at least one group for whom expected utility is lowered as a result of gambling. Hence, there cannot exist a gambling market with all rational individuals.

The argument we have just presented can also be easily generalized to cases where Y is not constant and endowments are not identical. The critical property of a pure gambling market is that it serves no

insurance function, i.e., if the bet is about the event S,

$$\frac{EU^i[Y^i|S]}{EU^i[Y^i|\sim S]} = \frac{EU^j[Y^j|S]}{EU^j[Y^j|\sim S]}, \quad \text{all } i,j \quad (3.16)$$

(where $\sim S$ denote all the states in which S does not occur). With perfect information, a market on the gamble on S would have no trade. (The initial resource allocation, relative to the event S, is ex ante pareto optimal.) Assume some individuals have information Ω , such that

$$\frac{E\{U^i[Y^i|\Omega, S]\}}{E\{U^i[Y^i|\Omega, \sim S]\}} = \frac{E\{U^j[Y^j|\Omega, S]\}}{E\{U^j[Y^j|\Omega, \sim S]\}} \quad (3.17)$$

This asserts that if everybody had the information, no further trade will occur.

Assume now only a subset of individuals have the information. Our "no gambling market" theorem asserts that there still can be no trade. For $B(p) > 0$, for the uninformed

$$\frac{EU^i[Y|p, S]}{EU^i[Y|p, \sim S]} = \frac{\Pi(p, S)}{1-\Pi(p, S)} > \frac{p}{1-p} \quad (3.18)$$

where $\Pi(p, S)$ is the probability of S occurring conditional on observing the market odds p.

p can only depend on Ω . Assume p does not depend on Ω . Then B must be independent of Ω .

The informed individual will only "sell" bets if

$$\frac{EU^j[Y^j|S]\Pi(p, S)}{EU^j[Y^j|\sim S](1-\Pi(p, S))} < \frac{p}{1-p}, \quad (3.19)$$

which, using (3.16) and (3.18), is clearly impossible.

Assume p does depend on Ω . Then for the informed to be willing to sell bets

$$\frac{EU^j[Y^j | S, p, \Omega] \Pi(p, S, \Omega)}{EU^j[Y^j | \sim S, p, \Omega] (1 - \Pi(p, S, \Omega))} < \frac{p}{1-p}, \quad (3.20)$$

for each Ω , and the given p . Since p is a function of Ω

$$\frac{EU^j[Y^j | S, \Omega] \Pi(\Omega, S)}{EU^j[Y^j | \sim S, \Omega] (1 - \Pi(\Omega, S))} < \frac{p}{1-p}. \quad (3.21)$$

This must be true for all Ω for which

$$p(\Omega) = p. \quad (3.22)$$

But for each Ω , for which (3.22) is true, (3.17) holds.

Hence,

$$EU^i[Y^i | S, \Omega] \Pi(\Omega, S) < \frac{p}{1-p} EU^i[Y^i | \sim S, \Omega] (1 - \Pi(\Omega, S)).$$

Taking expectations over all Ω for which (3.22) holds, we obtain

$$EU^i[Y^i | S, p] \Pi(p, S) < \frac{p}{1-p} EU^i[Y^i | \sim S, p] (1 - \Pi(p, S))$$

which contradicts (3.18).

The intuition behind this result may be put another way: Any individual who is uninformed, and knows that he is uninformed, will not bet, since he knows that his opponent will only bet when the odds are in his favor. Assume that the individual does not know whether he is more or less informed than his rival, but believes that his rival knows whether he is

more or less informed. The rival will only bet when he is the more informed and the odds are in his favor. Since the individual knows this, he will not be willing to engage in the bet. Thus, the only possible case in which betting can occur is when both bettors do not know whether they are the more or less informed. It is as likely that the individual is the less informed as that he is the more informed. But then betting simply represents a mean preserving spread in their incomes, and if both are risk averse (no matter what their utility functions look like), neither will be willing to engage in a bet.^{1,2}

IV. Incentives to Acquire Information in a Pure Exchange Market

Proposition 4 has one important corollary. Since in a pure gambling market, with rational expectations, there will never exist trades, regardless of informational differences, there will never be an incentive to acquire information.

If we observe an exchange market in which there is trade, either we can infer that some participants in the market are not rational or that the market is not a pure gambling market. Trade may be desirable if individual's marginal rates of substitution across states of nature differ in the absence of trade. Insurance markets clearly serve this function. Although some trading in the stock market clearly serves this function, it is questionable whether this provides a primary motivation for most of the trading which occurs.

In an exchange market with trade, the question again arises, whether there will be an incentive to acquire information. Since there is trade, there is an argument that by becoming better informed about the true probabilities, say of the two alternative states, expected utility

¹Recent work by Milgrom and Stokey (1979) has generalized this result using Aumann's concept of common knowledge, (1976).

²For the definition and interpretation of a "mean preserving spread" in a distribution. see Rothschild-Stiglitz (1977)

would be increased (for sufficiently small costs of acquiring information). But that argument is not necessarily correct, for it is possible that if some individuals were informed, then the market price would convey all the information; but if the market price conveyed all the information, it would not pay any individual to pay to become informed.

competitive equilibrium with costly information in which prices fully reflect the information.

A simple example may make this clearer.¹

Assume we had an insurance market for rain. Whether it rains or not depends on certain factors which can be observed at a cost; some of the factors are observed, some are not, so that from the observed factors, one cannot predict precisely whether it will or will not rain. Let η represent the information, and let Π be the estimate of the probability of rain when the information is η

$$\Pi = f(\eta) .$$

For simplicity we assume $f' > 0$.

We assume that everyone who observes the weather observes exactly the same information (later, we shall consider the case where different individuals obtain different information).

The net demand for "bets" on the weather (insurance for the event rain) will be a function of the individual's estimate of the

¹This kind of example has subsequently been developed at greater length in a series of papers by Grossman (1977) and Grossman and Stiglitz (1976, 1980a).

probability (and of the "price" of the bet, i.e. the market odds).

Let p represent the price. Then

$$B = B_I(p, \Pi)$$

where B is the (per capita) demand by informed individuals.

$$\frac{\partial B}{\partial p} < 0, \quad \frac{\partial B}{\partial \Pi} > 0.$$

An increase in the price lowers the demand but an increase in the estimated probability of the event increases the demand.

The uninformed (those who do not observe η) must base their betting on a priori information (which we assume is given, and unaltered throughout the analysis) and on market price (odds); for the individual knows that if there are informed individuals, the market odds will reflect their information. If we write

$$B_u = B_u(p), \text{ with } \partial B_u / \partial p < 0$$

to represent his demand function, then market clearing requires

$$B_I(p, \Pi) + B_u(p) = 0.$$

Under our assumptions, for every value Π there is a unique value of p

$$p = p(\Pi) \quad p' > 0$$

i.e. the uninformed can infer precisely from the market odds the information of the informed (even though the market odds need not equal Π). But then there is no incentive for anyone to obtain information. /

Note that if no one obtains information, the market odds will be invariant period to period and cannot reflect any information.

It would pay then someone to obtain the information if the

information were not too costly.

Thus, there is no non-stochastic competitive equilibrium to the exchange model with information which is costly (but not too costly). If the demand functions for bets depend on some other set of variables which were stochastic, and not observable (and the values of which were not inferable indirectly, say, by observing some other set of markets), then the market price (odds) will depend on this variable as well, and the uninformed will not be able to infer precisely the value from the market odds. Thus, in exchange markets which are not pure gambling markets there may exist market equilibrium with trade in which prices do not always perfectly reflect the information of the informed, although they do on average.

There is an incentive in this system for the uninformed to study the demand functions of the informed and uninformed, rather than η directly; for by knowing (the demand function and the market price they can infer the information about η , and this may be a cheaper (or indeed the only) way for the uninformed to obtain the information which we have assumed is available to the informed.

V. The Stock Market as a Gambling Market

The stock market is clearly not a pure gambling market. Yet, the arguments of the preceding section suggest both that much of the trading on the securities markets is based on some kind of irrationality, and that there are only limited incentives to acquire information about the returns of various securities. In

this section, we attempt to show this.

For simplicity, we assume that all individuals have a constant wage income, are endowed initially with an equal fraction of the shares of each of the firms, and are risk averse. Initially, we shall assume that they have the same degree of risk aversion. Let us define a non-speculator as an individual who purchases (holds) a representative market share, paying no attention to price; we define a speculator as an individual who attempts to form expectations of the returns of various securities and allocates his portfolio to maximize his expected utility, given those expectations and market prices.

Individuals may form expectations about all securities or only about a subset of securities.

The expectations of the j^{th} individual with respect to the i^{th} security are said to be unbiased if (where tildes denote the true values, carets the estimated values),

$$E\tilde{\mu}_i^j = \hat{\mu}_i, \quad E\tilde{\sigma}_{ik}^j = \hat{\sigma}_{ik}$$

where $\hat{\mu}_i^j$ is the j^{th} individual's expectations about the mean of the i^{th} security and $\hat{\sigma}_{ik}^j$ is his expectations of its covariance with the returns of the k^{th} security.

Proposition 5. If individuals have a choice between obtaining unbiased information at a cost or following the non-speculation strategy, all individuals will follow the non-speculation strategy.

Assume only one individual speculated. Clearly, if all others are pursuing a non-speculative strategy, he must end up with his representative market share of each of the firms, and hence his

income pattern is the same, but he has spent resources to obtain the information, and so is worse off. If several individuals speculate, and they obtain identical information the argument is the same. If they obtain different (but on average unbiased) information on average they must obtain the market return, but sometimes they do better than the market, sometimes worse. If the "quality of the information" is uncorrelated with the output of the economy (so that the individual does not consistently do better than the market average when the market does poorly, and conversely) the variance of his income is larger than with the non-speculative strategy and, as we argued above, his net mean income is lower: again it does not pay to speculate.

The result we have obtained is valid even if some individuals have a comparative advantage in obtaining information. If, for instance, one individual could find out costlessly for certain the returns to different securities, others would be unwilling to trade with him. He would simply determine the prices of the different securities. If two individuals believed that they have a comparative advantage--and they agreed on the returns (for certain) of the different securities--again there would be negative net returns to acquiring information, since in fact they would end up with exactly the same portfolio that they would have had had they followed the alternative non-speculation strategy.

The result can also be extended to cases where individuals have different utility functions.

Assume all individuals evaluate alternative strategies in terms of the mean and variance of income which they generate. Assume they have choice between (a) obtaining unbiased information about individual securities, (b) following the non-speculation strategy, or (c) following the non-speculation strategy with respect to the risky assets, but obtaining information about the mean and variance of the risky securities together, at a lower cost than obtaining information about individual securities. We now show they will follow one of the latter two strategies.

All individuals who obtain each period unbiased sure information would purchase the same bundle of risky securities, i.e. the market bundle. An individual who purchased only information about the mean and variance of the market bundle would, however, have just as good a knowledge of the "market line"¹ as the individual who knew the returns to individual securities. But he would have paid less for information, so he would be better off.

¹In the terminology of mean variance analysis, the "market line" is defined by

$$\sigma_y = b(\mu_y - r)$$

where μ_y and σ_y are the mean and standard deviation of the income from a portfolio consisting of all the risky securities on the market, and r is the (safe) rate of interest.

Thus, in a mean variance model, individuals would never acquire information about the properties of individual securities.^{1,2}

Note that screening information (discussed below in Section VII), that some stock is above average, some other stock below average, does not, in general, affect individual's beliefs about the return of the market as a whole. Much of the information provided by stockbrokers appears to be of this type.

¹ In subsequent work, S. Grossman has formalized and considerably extended these arguments. Grossman (1977).

² This result can be seen as an immediate corollary of Proposition 4. In a mean variance model, information which leaves unchanged expectations about the mean and variance of the market portfolio leaves unchanged individual's portfolio allocations. Thus (3.16) and (3.17) are satisfied.

VI. Two Alternative Explanations of Trade and Information Acquisition on the Stock Market

Trades do occur on the stock market, and there is considerable expenditure on information about particular securities. How do we reconcile these observations with the results of the preceding two sections. There are two basic explanations for trading: not all individuals are rational, in the sense in which we have defined the term, and individuals' endowments differ so that trade is desirable, even with identical information (and, perhaps more to the point, their endowments of assets, including non-traded assets, change in such a way as to warrant continuous trade in the stock market). We discuss these alternative explanations in the next two subsections.

VI.1 Irrationality

In this section we shall see that as long as individuals are finitely lived, and there is a continual stream of new individuals being born (entering maturity) it may be optimal for there to be speculation (from the private point of view). The argument requires not only that a fool be born every moment but that each of us believes he is not that fool.¹

An individual entering the market observes a distribution of returns from speculating. If he assumed he were simply average, he would not speculate. But if he believes that the reason that

¹This corresponds to the observation of most teachers that more than 1/2 of their students believe that they are in the upper half of their class.

individuals who have done above average is that they have a comparative advantage in obtaining information (and are not just "lucky"-- alternative hypotheses between which the data may not discriminate) and he believes that he is one of those individuals, he will go ahead and speculate. If there are many such individuals, there will be a "competitive" market for securities. As some individuals "win," their estimate of their comparative advantage may actually increase; as other individuals "lose" their estimate decreases, and if they lose long enough they stop being speculators and adopt the alternative non-speculative strategy.

The "thinness" of the market will depend then on the flow of new entrants into the market, the speed with which individuals revise their expectations, and the variance of the returns. If they are stubborn, and revise their expectations only slowly (the gambler whose luck is about to turn), then only if individuals have a long string of bad luck will they drop out of the market.

IV.2 Differences in endowments and tastes

Some trading on the stock market is life cycle trading: young individuals purchasing securities which they will sell when they are old. But if our earlier analysis is correct, rational individuals would simply purchase a mutual fund; there would not be trading in individual securities.

When entrepreneurs are lucky, and the firms which they have started do particularly well, they will have a disproportionate share of their wealth in their own firm. If they are risk averse, they will thus wish to trade their own firm for a mutual fund consisting

of the market as a whole. This gives rise to some trade; but again, if the analysis of the preceding sections is correct, it cannot give rise to sustained trading. Once portfolios are "balanced" so that each individual has a proportionate share in all firms, there will be no further trading.

Not all assets are tradeable; in particular, markets for human capital are notoriously imperfect. To the extent that (a) the human capital of different individuals yield returns which are imperfectly correlated; and (b) different securities on the market have different correlations with the returns to different individuals' human capital, if individuals initially had identical endowments of securities, there would be an incentive for them to trade, to obtain portfolios that are appropriately matched to their human capital. Moreover, changes in their human capital will, in general, give rise to changes in the optimal portfolios. Thus coal miners are likely to sell coal short in their portfolio (assuming the skills of being a coal miner are specific to the industry). I suspect this kind of "insurance" or "matching" function of trading in securities is relatively unimportant. In any case, if this is the primary motivation, it has interesting and important implications for the nature of the incentives for information acquisition, which we shall discuss in the next section.

The assumption that individuals evaluate portfolios simply in terms of their means and variances was critical to the result that only information about the market portfolio had value. (Information about the mean and variance of the market portfolio has social value, when individuals' attitudes towards risk differ, since it allows a more efficient distribution of the burden of risk; formally, for such information, assumption (3.16) is not valid.) When individuals' attitudes towards risk are not described by means and variances (or by one of the other utility functions for which a generalized mutual fund theorem (Cass-Stiglitz, 1970) is valid), then, again, information about a particular security may have value. (Again, assumption (3.16) is not valid.)

But note that if the returns to all securities can be written as a linear combination of a set of market factors,

$$Q_i = \sum_j M_j \mu_j + \varepsilon_i$$

where $E\varepsilon_i = E\varepsilon_i \varepsilon_j = 0$, and where there are enough securities so that individuals can diversify out of the idiosyncratic risk (ε_i) (as is commonly asserted), then again, information about individual securities is again of no value. Individuals' expected utility will be a function of the implicit prices associated with each of the market factors (assumed to be unaffected by information about a small firm) and the probability distribution of these market factors. This will determine their demand for each of the market factors. Thus ex post realized utility level will be a function only of the realization of the vector \underline{M} and the implicit market prices associated with each of the market factors (and, of course, the subject

probability distribution of these market factors). Again, if assumption (3.16) is valid, (3.16) will be; information about particular securities will not generate any trades and such information will have no value.

VII. The Returns to Information

In this section, we shall examine in more detail the nature of the incentives for information acquisition in the stock market. We shall show that there are marked discrepancies between the social and private returns to information.

VII.1 The taxonomy of information and the capital market

First, however, we must distinguish among several different kinds of information.

In another paper, on information in the labor market, (Stiglitz 1975) I distinguished between two kinds of information: general and specific. The former was information about a characteristic of an individual which affected his productivity in a variety of jobs, the latter was information relevant for a specific job.

Here we need to distinguish four kinds of information, depending on how the information relates to both buyers and sellers.

(1) General-general: information which affects all securities and all purchasers, for instance, that pertaining to the relative probability of different states of nature which are of importance to all individuals (e.g. the probability of a recession).

(2) General-specific: information which is of importance to specific individuals, but which affects their attitudes towards a whole class of securities, e.g. the probability of a recession in the coal mining industry is of importance for workers in that industry, but if the decline of the coal mining industry is uncorrelated with, say, the business cycle, it may be of relatively little importance to individuals who work in other industries.

(3) Specific-general: information which is specific to a firm (e.g. the mental stability of the manager) but which is of interest to all individuals who own the security.

(4) Specific-specific: information which is of value both to a specific individual (or class of individuals) and to a specific firm. If I am risk neutral, the only information I want is the mean return of the security; if I am risk averse, I will want information about other risk properties of the security.

Most information obviously falls in scope between the very general and the very specific: it affects not all firms but more than one; it is of value not to all individuals but to more than one. The dichotomy is important, however, both because the mechanisms for obtaining the returns to information and the relationship between social and private returns differ in the different categories.

There is another important distinction which must be made: some information may be of "value" to different individuals, but the information may increase the value of the security in the eyes of some, decrease it in the eyes of others.

Information that the firm will do very well if a certain contingency arises, but will do very badly otherwise, may raise the valuation of the security in the eyes of those who think the event likely, lower it in the eyes of those who think the event unlikely.¹ For instance, an individual who works in coal mining and whose wage is, as a consequence, correlated with the prosperity of the coal industry, will value firms whose profits are negatively correlated with the prosperity of the coal industry more highly than those

¹ Obviously, if there is a complete set of Arrow Debreu securities markets, then the judgments concerning the relative probabilities of different events affects the relative prices of different contingent commodities, but given the prices of Arrow Debreu securities, the firm need not concern itself with the probability of different events (see Stiglitz (1970), Grossman-Stiglitz (1977, 1980b)).

whose profits are positively correlated; for someone in another industry, a negative correlation with the coal industry may correspond to a positive correlation with his own wage, and thus he will find such a security unattractive.

Consider a bit of information which affects two securities which are initially indistinguishable. If, as a result of the information, all individuals now agree that security A is more valuable than security B, then we call that information hierarchical; that is to say, at least with respect to the characteristic being identified, all individuals agree that, say, more of the characteristic is better than less. For instance, information related only to the mean of the securities is hierarchical. On the other hand, when one subset of the population values A more highly as a result of the information, while another values B more highly, then we say that the information is matching, i.e. it matches specific individuals to specific securities.¹

VII.2 Social return to information in an exchange economy

From this discussion, we can see the nature of the social returns in a pure exchange economy:² to information/ given that individuals are different, it is in general not optimal for individuals to have the same portfolio. Information (both general and specific) allows a better "matching" of securities with individuals.

¹This distinction played an important part in my analysis of information in labor markets. Information about the wage a specific firm pays is hierarchical, information, about certain non-pecuniary characteristics of the firm is likely to be "matching" information. See Stiglitz (1975c).

²With production, there are, of course, further returns. See below.

VII.3 Private returns to information

What is the private return to providing information?

The private return to providing information takes two forms:

(a) a direct consumption effect, particularly of specific-specific information, allowing the individual who acquires the information to obtain a portfolio better suited to his needs; (b) a market valuation effect; if announcement of the information results in a change in the market value of the securities, then by buying the security if it is underpriced or selling the security short if it is overpriced, before releasing the information, and then making the information public, the individual is able to reap a capital gain.

In this sub-section, we are concerned primarily with the latter effect, particularly with the incentives for individuals to obtain, and disseminate, hierarchical information.¹ For individuals to be able to appropriate the returns from this kind of information

- (i) they must own, or be able to acquire, the asset before the information is disseminated;
- (ii) there must be an incentive for someone to disseminate the information, so that the market price can adjust to reflect the information; and
- (iii) the information, when disseminated, must be believed; it must be credible.

¹Hierarchical information, it will be recalled, is information which affects all individuals' evaluation of the asset identically; it does not result in any better allocation of the asset among the population.

There are problems at each of the three stages of the analysis. Our previous discussion (sections II and IV) argued that in a pure betting market there would be no trade. The fact that one individual (who is believed to be informed) is willing to buy some shares from me conveys the information that he believes the shares are under-priced; I will not, as a result, be willing to sell to him.

This has one interesting implication: it is sometimes suggested that firms that are undervalued will be subjected to take over bids. If all takeover bids were so motivated then they would never be successful (if those making the takeover bids on average were correct in identifying undervalued firms).¹

Information dissemination poses a standard public good (free rider) problem: all individuals who own shares in the firm would like someone else to pay the costs of information dissemination. All who own shares in the firm gain from the increase in the price.²

¹There may, of course, be other motivations for take-overs; a wealthy individual might, for instance, wish to buy an asset which is negatively correlated with the return to his other assets. As we argued earlier, trade can occur to match assets with individuals, but it seems that this provides little of the motivation for most take-overs.

²A similar free rider problem arises in take-overs aimed at improving the quality of the management. If on average, those who engage in takeovers do improve the quality of management, so that the returns on the asset are increased, it will result in an increase in the value. But then no small shareholder has an incentive to sell his shares to the firm attempting the take-over. He shares in the gains from the improved management. The only situation in which there is an incentive for a take-over to occur is when the new management can appropriate some of the increase in returns for itself. This point has been developed by Grossman and Hart (1980).

Only if a single individual owns all the shares, or if the owners act collusively, can the full benefits of the change in price from the information transmission be captured by the provider of the information.¹

¹Note that if there is a cost of transmitting the information, and these costs are large, and are a function of the number of individuals to whom the information is transmitted, the only individuals to whom it pays to transmit information are other potentially large shareholders (the wealthy). Thus the wealthy are supplied information by other wealthy sellers of securities.

There is, however, a natural method by which the owners of the firm can act collusively: to have the information provided by the firm itself. Assume that the owners of the firm have perfect knowledge about the returns to their firm (and hence are in agreement). Then clearly, provided the transmission costs are not too high, it will pay the firm to make the information public.

It would of course pay firms with low returns not to have information disclosed. But if some firms disclose information, and others do not, the market will assume that the firms which have not disclosed information must have a low return. Since there is "no information" about any of these firms, they will all be treated identically. But firms who are above the average of this "below average group" can increase their value by disclosing information, or having information disclosed about them (at the expense of course of those below average in the below average group). The process repeats itself until information about all the firms is revealed.

The process is closely parallel to that which I have used elsewhere to describe the education system as a screening device. It is in the self-interest of the bright to become so identified, and although this identification need have no social returns (if it does not result in a change in production), it will yield private returns to the bright at the expense of the stupid.

There are, however, some important differences which make the analysis of screening in the capital market considerably more complex than in the education market. The most important differences arise from the fact that human capital is (at least currently) not

normally bought or sold short; that is, even if I know that you are overvalued, I cannot sign a contract to sell your services next period, announce the true value of your productivity, and then, next period having arrived, buy your services from you at the now much lower price. Nor can I buy a fraction of your human capital for later resale.¹ Thus, in our discussions of screening of individuals, we assumed that it was the individual who made the decision about whether to have himself screened. Since in the capital market, through speculation, every individual is a potential owner of every firm, there is the possibility, at least, that it might be in the interests of someone other than the original owner to certify the characteristics of the firm. Our analysis has suggested that so long as the screening industry is competitive, the returns to screening will be captured by the original owners of the security, and thus it is they who--as in the education market--have the incentive to provide the screening.

In the education market, it is natural to assume (although not necessarily the case) that individuals know more about their own ability than anyone else. Similarly, in the context of the capital market it is natural to assume that the individuals who do have more information about the security are the original owners of the firm.²

¹These statements are not quite correct. There are firms which attempt to find "undervalued" individuals, screen them, and then sell their labor services at a higher price, and some individuals do incorporate themselves, in effect selling a portion of their human capital. But these instances are more the exception than the rule.

²It is, of course, possible that others have more information about the prospects of the firm than the original owners. In that case, these individuals may take over the firm, and have the firm provide the information. But if the original owners know that this is the motive for the takeover, they will be unwilling to sell.

VII.4 Relationship between the social and private returns to information

In the previous subsection, we showed how certain aspects of the problem of information on the capital market can be recast as **problems** in the economics of screening. We can thus borrow some results from the general theory of screening. We obtain the following two important conclusions:

(a) The returns to the provision of information which changes the market value of firms is captured primarily by the original owners of the shares; it is accordingly they who have an economic incentive to provide the information.

(b) The private returns to the provision of information do not, in general, correspond to the social returns, just as in the case of screening of individuals. There is likely to be excessive spending on "hierarchical screening," i.e. in screening for characteristics which individuals value in the same way, and too little spending on "matching screening", i.e. on screening for characteristics which individuals value in different ways. The latter can be seen most easily by considering a case where the provision of information leaves the price of the asset unchanged, the increase in the demand by one group being offset, say, by a decrease in the demand by another group. Even when there is essentially no cost to providing this information, and even though there may be a high social return to doing so, there will be no incentive for the firm to provide this information.

In the next two sub-sections, we provide both examples showing the relationship between social and private returns for hierarchical and matching information.

The third problem we noted above is that of credibility: although firms have an economic incentive to provide information, they also have an economic incentive to provide misinformation, just as in the education market individuals have an incentive to have themselves overrated. In the education market, an individual who has misrepresented himself is usually found out, and thus is not able to enjoy the benefits of his misrepresentation (a higher wage) for long. In the capital market it may be more difficult to ascertain misrepresentation, and by the time it is ascertained, the original owner of the security has already absconded with his gains. This makes it even more necessary than in the case of the education market for the information to be certified by "public" outside institutions. These institutions in the capital market are the investment consultants, the stock brokerage firms, etc.¹

¹The division between direct production of information and certification depends, presumably, on comparative costs of producing different kinds of information and the fact that some information may be relevant to several different firms. That is, GM may have an advantage in producing information about the production characteristics of GM, but a stock brokerage firm will produce information about the demand for cars in general.

VII.5 An Example of Hierarchical Screening

We assume that everyone is risk averse and evaluates the portfolio in terms of the mean and standard deviation of the return.

We assume, moreover, that there are enough securities of each type so that the individual, if he completely diversifies, faces essentially no risk. There are two kinds of securities, type 1 with mean θ_1 , and type 2 with mean θ_2 . $\theta_1 > \theta_2$. All securities are known to have the same variance, σ^2 . There are equal numbers of the two types of securities. For simplicity, we assume all securities are independently distributed, and the distribution of returns of each is normal. All individuals have the same utility function, which we assume is of the form $-e^{-\alpha Y}$ where Y is income. (constant absolute risk aversion). Hence, the individual seeks to maximize

$$(7.1) \quad -\bar{Y} - \frac{\alpha \sigma^2 \bar{Y}}{2}$$

where \bar{Y} is his mean income

σ^2 is the variance of his income.

Each individual is endowed initially with one firm. He knows what kind of firm it is, but does not know what type any other firm is. He can supply to the market information about the quality of his firm.

The cost of providing this information is c ; we assume c is neither very small nor very large.¹

$$(7.2) \quad \left(1 - \frac{\theta_1 - \theta_2}{2\alpha\sigma^2}\right) (\theta_1 - \theta_2) > c > (\theta_1 - \bar{\theta}) \left(1 - \frac{\theta_1 - \bar{\theta}}{2\alpha\sigma^2}\right).$$

¹The role of these inequalities will be apparent shortly.

There are then two equilibria:

- (a) The no-information equilibrium. The relative price of all securities is unity. The individuals who own type one firms retain a fraction λ of the wealth in their own security and divide the remainder equally among all other shares. λ is chosen to

$$(7.3) \quad \max \bar{\theta} + \lambda(\theta_1 - \bar{\theta}) - \frac{\alpha \sigma^2 \lambda^2}{2}$$

when $\bar{\theta}$ is the mean return on all securities sold on the market. For the moment, we shall assume that the owners of low productivity firms retain the same fraction of their shares as do the owners of high productivity firms.¹ This implies that the average return to a randomly selected share on the market is $\frac{\theta_1 + \theta_2}{2}$. If the number of firms is large, we can ignore the effect of the i th firm's action on the average return, so

$$(7.4) \quad \lambda^* = \frac{\theta_1 - \bar{\theta}}{\alpha \sigma^2}$$

Using (7.4) it is clear that for the entrepreneur owning a type 1 firm

$$(7.5) \quad \bar{Y} - \frac{\alpha}{2} \sigma_Y^2 = \bar{\theta} + \frac{1}{2} \frac{(\theta_1 - \bar{\theta})^2}{\alpha \sigma^2}$$

He is slightly better off than the average person; how much depends on his aversion to risk (α) and the variance of his own security.

If he screened, the value of his wealth would go up by an amount $\theta_1 - \bar{\theta}$, and he could then completely diversify, so

$$(7.8) \quad \bar{Y} - \frac{\alpha}{2} \sigma_Y^2 \approx \theta_1.$$

¹This assumption will be justified in the next section. The basic argument is simple: if the fraction of shares retained by the original owners is an observable variable, if the low productivity firms retained a smaller fraction than λ , it would signal to the market that they were low productivity, and this would have a marked effect on their market value.

The results would not be significantly affected if we assumed that the original owners of poor quality firms retain a smaller fraction of their shares.

Under the assumption of (7.2), it does not pay him to undertake the screening.

(b) The full information equilibrium. The relative price of the two securities is θ_1/θ_2 . For the upper group, after paying for screening

$$(7.7) \quad \bar{Y} - \frac{\alpha_\sigma^2}{2} \bar{Y} \approx \theta_1 - c$$

and for the lower

$$(7.10) \quad \bar{Y} - \frac{\alpha_\sigma^2 Y}{2} \approx \theta_2.$$

If an individual who owns a type 1 firm does not screen, his expected utility is

$$(7.11) \quad -\exp - \left[\theta_2 + \frac{1}{2} \frac{(\theta_1 - \theta_2)^2}{\alpha_\sigma^2} \right].$$

Thus, it pays the individuals of type 1 to screen. This example illustrates several of the important aspects of markets with imperfect information:

- (1) There may be multiple equilibria.
- (2) Some of these equilibria may be Pareto inferior to others.

In this example, it can be shown that the no screening equilibrium is pareto superior to the screening equilibrium.

- (3) The type 2 securities owners exert an externality on the type 1 security owners. If all securities were good (type "1") the expected utility of the type 1 security holder would be $-\exp(-\theta_1 \alpha)$. In both the full information and no information equilibrium, the type 1 individuals are worse off. On the other hand, the losses to the type 1 individuals exceed the gains to the type 2 individuals, relative to what they could have had if all securities were of type 2. This is obviously the case in the screening (full information) equilibrium where in fact the type 2 security holders are no better off than they would have been in the absence of type 1 security holders, but the latter are unambiguously worse off: the externality is purely dissipative. But even in the no screening equilibrium, the "income equivalent loss" to type 1 security holders is

$$(\theta_1 - \bar{\theta}) \left(1 - \frac{1}{2} \frac{(\theta_1 - \bar{\theta})}{\alpha \sigma^2} \right)$$

while the gain to type 2 security holders is only

$$(\bar{\theta} - \theta_2) - \frac{(\theta_1 - \bar{\theta})}{\alpha \sigma^2} (\bar{\theta} - \theta_2 + \frac{1}{2} (\theta_1 - \bar{\theta}))$$

for a net loss of

$$\frac{(\theta_1 - \theta_2)^2}{\alpha \sigma^2}$$

Matching Screening.

The example developed in some detail above involved one group of securities being unambiguously better than another group. The example which we ^{now} present involves differences in opinions among different individuals about what states of nature are most likely to occur; this in turn affects which securities maximize ex ante expected utility.

Consider a particular firm. It is known that it either produces using technique A or technique B. For simplicity we assume individuals are risk neutral, but we assume individuals are not allowed to sell securities short.¹

¹ These assumptions may easily be modified.

There are three groups in the population; type A individuals believe that, if the firm produces using A, it will have a very high profit, but if it produces using B, its profits will be low; conversely, type B individuals believe expected profits with B are high but with A are low. Type C individuals believe that the two techniques have the same expected profits. In the absence of information about which technique is being employed, all individuals assume that there is a fifty-fifty chance that a particular technique is used. More precisely, let $\Pi^A(A)$ be expected profits if technique A is used in the judgment of person of type A, and ω^A be the aggregate wealth of individuals of type A. There is an alternative investment opportunity (security C) whose expected return r^* is the same in everybody's judgment. Then we assume

$$\Pi^A(A) > \Pi^C(A) > \Pi^B(A), \quad \omega^A < \frac{\Pi^C(A)}{r^*}$$

$$\Pi^B(B) > \Pi^C(B) > \Pi^A(B), \quad \omega^B < \frac{\Pi^C(B)}{r^*}$$

$$\frac{\Pi^A(A) + \Pi^A(B)}{2} < \Pi^C(A) = \Pi^C(B), \quad \omega^C > \frac{\Pi^C(A)}{r^*}$$

$$\frac{\Pi^B(A) + \Pi^B(B)}{2} < \Pi^C(A) = \Pi^C(B).$$

Thus, in the absence of information, individuals of type A or type B prefer security C to the given firm. The firm will be entirely owned by individuals of type C. By assumption, their wealth exceeds the value of the firm. (The value of the firm will be

$$\frac{\Pi^C(A)}{r^*}$$

which is less than the aggregate wealth of type C individuals).

The remainder of the wealth of type C individuals as well as the wealth of type A and B individuals is invested in type C securities. The expected income of individuals of each type is

$$EY^A = r^* \omega^A$$

$$EY^B = r^* \omega^B$$

$$EY^C = r^* \omega^C$$

Now assume we have perfect information, e.g. everybody knows that A was the technique used. Then type A individuals would all purchase securities of the given firm. But by assumption, their aggregate wealth is insufficient to buy all the securities, so C is the "marginal purchaser". Hence the value of the firm is unchanged at $\frac{\Pi^C(A)}{r^*}$. Expected incomes are now given by

$$EY^A = \frac{\Pi^A(A)}{\Pi^C(A)} r^* \omega^A > r^* \omega^A$$

$$EY^B = r^* \omega^B$$

$$EY^C = r^* \omega^C$$

Thus, there is a Pareto improvement: A is better off and no one is worse off (in an expected utility sense). If the cost of the information is less than

$$\left(\frac{\Pi^A(A)}{\Pi^C(A)} - 1 \right) r^* \omega^A$$

it is clearly worth procuring. Note, however, that although this information has value, there is absolutely no change in the market value of the firm. Thus the firm would have no incentive to provide the information -- even if it could do so almost costlessly. Individuals of type A have an

incentive to get together, to form a "consumer cooperative" to procure the information; but this is likely to be impossible, since it is difficult for individuals of type A to identify each other. (Thus, although there is a natural "producer" cooperative for providing information by sellers of securities, i.e. the firm itself, there is no corresponding natural consumer cooperative.) And even were it possible to identify who was in fact an individual of type A, there are all the classical public good (free rider) problems which argue that there is likely to be underinvestment in this type of information. /

VIII. Self-Selection Equilibria

We argued in the previous section that the original owners of the more productive firms have an incentive to provide information

to the market establishing their productivity. There is, however, a fundamental problem of credibility: information supplied by (or paid for by) the firm is not likely to be believed. It is well known that actions speak louder than words; in the recent literature on screening, actions which convey information (e.g. about the productivity of the firm) are referred to as self-selection devices. (Stiglitz (1982)). An entrepreneur who is willing to hold on to a large fraction of the shares of his company may be conveying information that he believes the market is undervaluing his firm.

Of course, if investors come to believe this, then firms which are not productive may attempt to imitate the more productive firms; the share of the firm retained by the original owners would not then convey information. This discussion should make clear that analysis of self-selection equilibria is a fairly complex question. Indeed, as we shall see, it may not even be clear what the appropriate notion of equilibrium should be.

The essential property of a self-selection device is that the relevant cost curves (indifference curves) for one group differ from those of the other. An increase in the share of the firm retained by the original owners has a cost in the reduced diversification (increased risk) which the entrepreneur must bear. If, however, the market undervalues the security, there is a gain: mean income will increase as the fraction retained increases. Thus, the net cost to the more productive firms is lower than to the less productive firms, and it is this difference which in some circumstances may enable the fraction of shares retained by the original owners to serve as a self-selection device.

To see this more formally, we return to the simple model of section 7.5. Let p_i denote the price at which an owner of a firm of type i can sell his shares. The mean return of shares purchased on the market is normalized at unity. In figure 7.1 we have depicted the individuals' indifference curves in (λ, p) space (where λ is the fraction of the shares retained). Clearly,

$$\left(\frac{dp}{d\lambda}\right)_u \geq 0 \text{ as } \lambda > \lambda^*(p)$$

where $\lambda^*(p)$ is the optimal value of λ . (See, e.g. equation 7.4.)

Moreover,

$$\left(\frac{dp}{d\lambda}\right)_2 > \left(\frac{dp}{d\lambda}\right)_1 .$$

The owners of the less productivity (type 2 firms) require a larger increase in the price to compensate them for an increase in the fraction of the firm they retain.

Pooling equilibrium

Assume first that λ is observable. Consider the no information equilibrium, when $\lambda_1 = \lambda_2$ (Rothschild and Stiglitz (1976) refer to this as a pooling equilibrium). Clearly, in equilibrium, $p=1$, and the optimal value of λ (from the point of view of the more productive firms) is just λ_1 .

Is this, however, an equilibrium? Clearly, the good firms would like to signal that they are good. What would happen if one such firm were to announce that it was willing to sell a fraction $\hat{\lambda}$ at a price p greater than 1, with the point $(\hat{\lambda}, \hat{p})$, located in the shaded area in figure 7.1, i.e. below type 2 individuals' indifference curve, but above type 1 individuals' indifference curve. An investor might be tempted to infer that only a type 1 entrepreneur,