February 1979

UNCERTAINTY AND INFORMATION IN ECONOMICS*

by

Jack Hirshleifer and John Riley**

Working Paper #94

*This is a slightly revised version of UCLA Working Paper #140.

**UCLA, and Boston College and UCLA.

CONTENTS

		Page
Part 1:	THE ECONOMICS OF UNCERTAINTY	5
1.1	Decision under Uncertainty	5
1.1.1	The Menu of Acts	5
1.1.2	The Probability Function	7
1.1.3	The Consequence Function	9
1.1.4	The Utility Function and the Expected Utility Rule	10
1.1.5	Risk-Aversion, and the Risk-Bearing Optimum of the Individual	14
1.2	Market Equilibrium under Uncertainty	22
1.2.1	Risk-Sharing	23
1.2.2	Insurance (including discussion of state dependent utility, adverse selection, and moral hazard)	24
1.2.3	Complete and Incomplete Market Regimes, the Stock Market Economy, and Optimal Production Decisions	36
1.2.4	Other Applications	41
Part 2:	THE ECONOMICS OF INFORMATION	42
2.1	Informational Decision Making	42
2.1.1	Acquisition of Information	42
2.1.2	Other Informational Activities	53
2.2	Emergent Information	55
2.2.1	The Value of Flexibility	55
2.2.2	Complete versus Incomplete Regimes of Markets	57
2.2.3	Speculation	64
2.3	The Economics of Research and Invention	67
2.4	Information Advantage and Market Revelation of Information	73
2.4.1	Signalling	73
2.4.2	Informational Inferences from Market Prices	78
2.4.3	Informational Efficiency	81

J. Hirshleifer
J. Riley

UNCERTAINTY AND INFORMATION IN ECONOMICS

Traditional economic analysis passes over, in more or less embarrassed silence, the problem of uncertainty. The central elements of economic reasoning have been shaped into models of ever-increasing precision, but models that presume a degree of knowledge on the part of economic decision-makers that is patently unreasonable -- for example, that the firm knows its demand function, not merely in the present but up to the economic horizon. Such an unrealisitc picture of the actual decision-making situation means that economic theory is of little help to a business man facing an actual marketing choice. At best the Law of Large Numbers operates to reduce the importance of uncertainty at the market level, thereby allowing the fiction of the average or "representative" individual.

Much more fundamentally, models postulating behavioral certainty are completely inconsistent with observable real-world activities of the first importance -- among them insurance, speculation, research, advertising, and even education.

In the past twenty years, however, an exciting new literature has addressed the problems of decision and market equilibrium under uncertainty. This literature has two main foundation-stones: (1) the Von Neumann-Morgenstern [1944] theory of preference for uncertain contingencies and in particular the "expected-utility theorem," and (2) Arrow's formulation of the ultimate goods or objects of choice in an uncertain universe as

<u>contingent</u> consumption claims: entitlements to particular commodities or commodity baskets valid only under specified "states of the world" (more briefly, "states") [Arrow 1953, 1964] Just as intertemporal analysis requires subscripting commodity claims <u>by date</u>, uncertainty analysis requires subscripting commodity claims <u>by state</u>. Among objects of choice so defined, as we shall see, production and exchange and consumption all take on recognizable forms as generalizations of the corresponding processes in the familiar theory of a world of certainty.

An alternative conceptualization of the objects of choice under uncertainty runs in terms of the <u>statistical parameters</u> of the probability distributions of commodity or income claims. In that formulation it is assumed that individuals prefer greater <u>mean</u> income but smaller <u>variance</u> of income; attention may or may not be paid to higher moments of the distribution [Markowitz, 1959]. It has been shown that the more general "state-preference" representation of the objects of choice under uncertainty can be reduced to such a "parameter-preference" representation by making a number of specializing assumptions [Tobin 1958, Borch 1968, Feldstein 1969]. In the particularly simple form of choice between mean return and variance of return on investment, the parameter-preference model has provided the basis for important modern developments in the theory of finance [Markowitz 1959, Sharpe 1964, Lintner 1965, Mossin 1966]. We will not be able to pursue the parameter-preference approach here; for recent surveys see Jensen [1972] and Merton [1978].

The modern literature on uncertainty and information divides into two rather distinct branches. The first deals with <u>market uncertainty</u>. Each individual is supposed to be fully certain about his own endowment and productive opportunities; what he is unsure about are the supply-demand offers of other economic agents. In consequence, search on the individual level, and disequilibrium and price dynamics at the market level, take the center stage -- replacing the traditional assumption of costless exchange at market-clearing prices [Stigler 1961, 1962; McCall 1965]. Explicit analysis of market uncertainty is leading toward a more realistic treatment of market "imperfections," with implications not only for microeconomics but for macroeconomics as well [Phelps <u>et. al.</u> 1970].

The second branch of literature has dealt with what is usually called <u>technological uncertainty</u>, though <u>event uncertainty</u> would be a better term. Depending upon the weather, for example, crops may be large or small. Here individuals are unsure about exogenous data such as resource endowments and/or productive opportunities. Analysts of event uncertainty generally have been content to employ, at least as an initial approach, the simpler traditional model of perfect markets in which all dealings (except possibly in the market for information itself) take place costlessly at equilibrium prices. However, as an important theme in the literature, while the markets that exist are supposed perfect the set of available markets may be assumed to be <u>incomplete</u> -- not every definable object of choice (commodity claim subscripted by state) may be separately exchangeable.

Ultimately, the two branches of investigation dealing respectively with market uncertainty and with event uncertainty must merge, but that day is not yet visible.

The present discussion is limited to the relatively more tractable topic of event uncertainty; reviews of the complex phenomena such as search and disequilibrium that emerge under market uncertainty may be found in Rothschild [1972] and Lippman and McCall [1976]. The purpose of the paper is mainly expository. We hope to provide a non-technical survey of the central underlying ideas, to introduce the novel tools of analysis that have proved fruitful in this area, and to suggest a few of the actual and potential applications of the economics of uncertainty and information.

Part 1

THE ECONOMICS OF UNCERTAINTY

1.1 Decision Under Uncertainty

In decision-making under uncertainty the individual chooses among acts while Nature may be metaphorically said to "choose" among states. In principle both acts and states may be defined over a continuum, but for simplicity here a discrete representation will ordinarily be employed. Table 1 pictures an especially simple 2x2 situation. The individual's alternative acts a = (1,2) are shown along the left margin, and Nature's alternative states s = (1,2) across the top. The body of the Table shows the consequences c resulting from the interaction of each possible act and state. Expressed in fuller detail the individual's decision problem requires him to specify: (1) a set of acts $a = \{1,\ldots,A\}$; (2) a probability function expressing his beliefs $\pi(s)$ as to Nature's choice of state $s = \{1,\ldots,S\}$; (3) a consequence function c(a,s) showing outcomes under all combinations of acts and states; and, finally (4) a preference-scaling or utility function v(c) defined over consequences. Using these as elements, the "expected-utility rule" (Sec. 1.1.4 below) enables the individual to order the available acts in terms of preferences, i.e., to assign a utility function over acts u(a) so as to determine the one most highly preferred.

1.1.1 The Menu of Acts

We shall consider here two main classes of acts: terminal and non-terminal or informational.

		States		
		s = 1	s = 2	Utility
Acts	a = 1	c_{11}	c_{12}	U_1
	a = 2	c_{21}	c_{22}	U_2
	Beliefs	π_1	π_2	

Table 1: Consequences of Alternative Acts.

Terminal actions represent making the best of one's existing combination of information and ignorance. For example, you might decide whether or not to take an umbrella on the basis of your past history of having been caught in the rain. In statistical theory, terminal action is exemplified by the balancing of Type I and Type II errors in coming to a decision (e.g., accepting or rejecting the null hypothesis) on the basis of the evidence or data now in hand. In contrast with the classical statistical problem, which may be likened to the decision situation of an isolated Robinson Crusoe, in the world of affairs studied by economics many interpersonal arrangements -- insurance contracts, futures markets, guarantees and collateral, the corporation and other forms of shared enterprise -- serve to widen the terminal-act options available to individuals. As we shall see, these market processes provide a variety of ways for sharing risks and returns among the decision-making agents in the economy.

Informational actions are non-terminal in that a final decision is deferred while awaiting or actively seeking new evidence that will, it is anticipated, reduce uncertainty. In statistics, informational actions involve decisions as to new data to be collected: choice of sampling technique, sample size, etc. Again, in the world of affairs interpersonal transactions open up ways of acquiring information apart from the sampling techniques studied in statistics: information may be purchased, or inferred by monitoring the behavior of others, or even stolen. To a degree, information acquisition and dissemination have become specialized functions (the "knowledge industry" [Machlup 1962] whose practitioners are rewarded by exchanges with other economic agents in the economy.

Part 1 of this paper will, apart from introductory discussions, cover only <u>terminal actions</u>, decisions made under fixed probability beliefs ("the economics of uncertainty"). The enlarged range of issues generated by admitting also non-terminal actions will be examined in Part 2 ("the economics of information").

1.1.2 <u>The Probability Function</u>

We will assume that each individual is able to represent his current <u>beliefs</u> as to the likelihood of the different states of the world (e.g., as to whether Nature will choose Rain or Shine) by a "subjective" probability distribution [Fisher 1912, Ch. 16, Savage 1954]. That is, an assignment to each state of a number between zero and one (end-points not excluded) whose sum equals unity. Subjective <u>certainty</u> would be represented by attaching the full probabilistic weight of unity to only one of the outcomes. The degree of subjective <u>uncertainty</u> is reflected in the dispersion of probability weights over the possible states.

Frank Knight [1921, p. 225] and others have denied the applicability of the probability concept to measure belief where there is no clear basis for objectively classifying instances: e.g., to such questions as whether or not Christchurch is the capital of New Zealand, or whether a cure for cancer will be discovered by the year 2000. It will not be possible to review here the philosophical and operational underpinnings of the probability concept; for our purposes, it is sufficient that the "subjective" or "degree of belief" interpretation has proved fruitful. But a reading of Knight's further discussion suggests that he was getting at a different and much more interesting point connected with the distinction between terminal and informational action -- the problem of "hard" versus "soft" probabilities.

Suppose an individual faces the terminal-action problem of betting Heads or Tails on a single toss of a coin. Imagine, in the one case, that he is absolutely certain that the coin is fair; i.e., that the probability of Heads is $\pi = .5$ (and of Tails is $1 - \pi = .5$). This is a "hard" probability, meaning that the individual has very high (indeed, perfect) confidence in the value he attaches to the parameter π of the probability distribution. Alternatively, imagine that the individual feels he knows nothing whatsoever about the parameter -- the coin could be so loaded that π, with equal likelihood so far as he knows and believes, might take on any value between 0 and 1 inclusive. Yet, for the terminal decision as to which way to bet on a single toss, the individual still has no basis for preferring Heads or Tails, and so must rationally still treat the two outcomes as equally likely. Thus in this second case also $\pi = .5$ for him, although here his probability estimate is very "soft", i.e., he has low confidence. Note that there is no operational difference between hard and soft probabilities if only terminal action is called for. The difference arises where there are opportunities for informational action; as we shall see in Part 2, the value of acquiring information varies inversely with one's prior confidence.

1.1.3 The Consequence Function

By consequence is meant a full definition of all relevant characteristics of the individual's environment resulting from the interaction of the specified act and state. A consequence can be regarded as a multi-commodity multi-date consumption basket; for simplicity, however, we will sometimes assume that it corresponds to the amount of a single summary variable like income.

In the case of a terminal action, the consequences contingent upon each state might either be certain or probabilistic -- depending upon the definition of "states of the world" for the problem at hand. If the states are defined deterministically, as in "Coin shows Heads" versus "Coin shows Tails," and supposing the act is "Bet on Heads," the contingent consequences are the simple certainties "Win" in the one state and "Lose" in the other. But states of the world might sometimes represent alternative probabilistic processes. For example, for some terminal decision problem the two alternative states might be "Coin is fair (has 50% chance of coming up Heads)" versus "Coin is biased to come up Heads with 75% chance." In this situation the act "Bet on Heads" will have probabilistic consequences: "50% chance of winning" for the first state, and "75% chance" in the second.

For an informational action, on the other hand, the consequences will in general be probabilistic even if the states of the world are defined deterministically, since acquisition of information does not ordinarily eliminate all uncertainty. If the states are "Rain" versus "Shine," and the informational action is "Look at barometer," the consequences will only be improved likelihoods of behaving appropriately in each state -- since the barometer reading is not a perfect predictor of Rain or Shine.

1.1.4 The Utility Function and the Expected-Utility rule

In the theory of decision under uncertainty, utility as an index of preference attaches both to consequences c and to acts a. We shall sometimes find it convenient to distinguish the two by the notations v(c) and u(a), the problem being to derive the u(a) for evaluating actions from the primitive utility valuations v(c) for consequences.

To choose an act is to choose a row of the consequence matrix, as in Table 1. Given the assignment of probabilities to states, this is also choice of a probability distribution or "prospect." A convenient notation for the prospect associated with an act a, whose consequences $c_a = (c_{a1},\dots,c_{aS})$ are to be received with respective probabilities $\pi = (\pi_1,\dots,\pi_S)$, is

$$a \equiv (c_{a1},\dots,c_{aS};\ \pi_1,\dots,\pi_S)$$

Or, more compactly:

$$a \equiv (c_a,\pi)$$

The connection between the preference scaling of acts and the preference scaling of consequences is provided by the Neumann-Morgenstern "expected-utility rule":

$$u(a) \equiv \pi_1\, v(c_{a1}) + \dots + \pi_S\, v(c_{aS}) \equiv \Sigma_{s=1}^{S}\, \pi_s\, v(c_{as}) \qquad (1.1)$$

That is, the utility of each act u(a) is the mathematical expectation or probability-weighted average of the utilities of the associated consequences $v(c_{as})$.

The expected-utility rule is of course a very specific and special procedure for inferring preferences u(a) over acts from the primitive preference scaling of consequences shown by the v(c) function. What is its justification? It turns out that the expected-utility rule is usable if and only if the v(c) function is determined <u>in a particular way that has been termed the assignment of "cardinal" utilities to consequences</u>. More specifically, the underlying theorem can be stated as follows:

> Given certain "postulates of rational choice," there is a way of assigning a cardinal preference-scaling function v(c) over consequences such that the preference ranking of any pair of prospects a', a" coincides with the ranking under the expected-utility rule (1) above.

The "postulates of rational choice" therefore justify the <u>joint</u> use of cardinal utilities and the expected-utility rule in dealing with choices among risky prospects -- a point worth emphasizing, since it would be quite invalid to infer that the theorem warrants or provides a cardinal utility measure for choices <u>not</u> involving risk (see the discussion in [Baumol 1951, Alchian 19 Strotz 1953]).

The postulate system serving as basis for the theorem has been set forth in a number of different ways in the literature [Friedman and Savage 1948, Luce and Raiffa 1957, Markowitz 1959, Marschak 1968], and involves technicalities that cannot be pursued here. Instead, what follows is an informal presentation (based mainly on Schlaifer [1959]) illustrating, by direct construction, the development of a personal cardinal preference-scaling function for use with the expected-utility rule (1.1).

For the purposes of this discussion, we will assume that the contingent consequences c are certainties, and also that c represents simply the quantity of generalized income. Let $\hat{c}$ represent the worst consequence (lowest level of income) contemplated by the individual, and $\hat{\hat{c}}$ the best consequence (highest level of income). As "cardinal" preference scales allow free choice of zero and unit interval, we can let $v(\hat{c}) = 0$ and $v(\hat{\hat{c}}) = 1$. Now consider the intermediate level of income c^*. We can suppose that the individual is indifferent between having c^* for certain and having some chance of success π^* in a prospect or "reference lottery" involving $\hat{c}$ and $\hat{\hat{c}}$. What numerical value can we attach to this common level of utility to allow use of the expected-utility rule? The answer is simply, the probability π^*. Thus

$$u(c^*) \equiv u(\hat{\hat{c}}, \hat{c}; \pi^*, 1-\pi^*) \equiv \pi^* \qquad (1.2)$$

Fig. 1 illustrates a situation in which $\hat{c} = 0$, $\hat{\hat{c}} = 1000$, $c^* = 250$, and $\pi^* = \frac{1}{2}$. That is, in the preferences of this individual a sure income of \$250 is indifferent to a 50% chance of winning in a lottery whose alternative outcomes are \$1000 or nothing. Hence, $v(250) = 1/2$. A similar introspective process generates the individual's entire $v(c)$ curve of Fig. 1, which is his preference-scaling function for consequences.

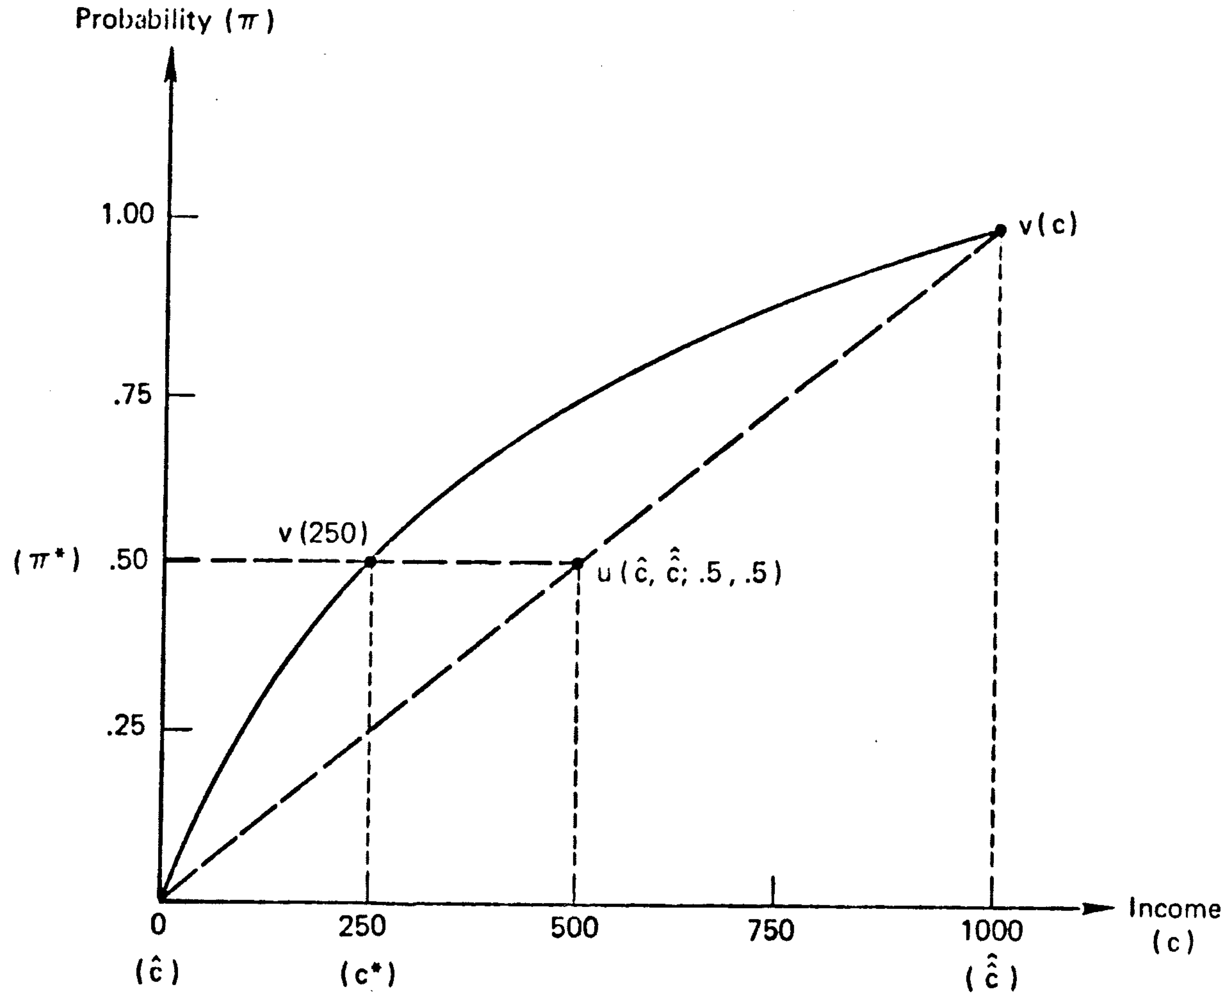

Fig. 1 — The preference-scaling function v(c) derived by the "reference-lottery technique"

We can immediately verify that the expected-utility rule (1.1) does give us the correct valuation at least of simple reference-lottery prospects like that in (1.2) above:

$$u(\hat{\hat{c}},\hat{c};\pi^*,1-\pi^*) = \pi^* v(\hat{\hat{c}}) + (1-\pi^*)v(\hat{c}) = \pi^*(1) + (1-\pi^*)0 = \pi^*$$

More generally, this argument can be extended to any prospect whatsoever -- for example, to the prospect $(c', c''; \pi, 1-\pi)$ where c' and c'' are any levels of income and π is any probability. This individual's $v(c)$ function tells us that the contingent income c' is indifferent (equivalent in utility terms) to some chance of success π' in the reference lottery -- $v(c') = \pi'$ -- and similarly c'' corresponds to some chance π'' -- $v(c'') = \pi''$. Then the prospect $(c',c''; \pi, 1-\pi)$ must be equivalent in preference terms to having some computable <u>overall</u> chance of success in the reference lottery. Specifically, the prospect gives us the chance π of an income c' equivalent to a probability of success π', and the chance $1-\pi$ of an income c'' equivalent to a probability of success π''. Using the laws of probability, the equivalent <u>overall</u> chance of success is $\pi(\pi') + (1-\pi)(\pi'')$. But this is just the expected-utility rule:

$$u(c',c'';\pi,1-\pi) = \pi(\pi') + (1-\pi)(\pi'') = \pi v(c') + (1-\pi)v(c'')$$

The expected-utility rule, combined with the constructed $v(c)$ function, works because the latter is <u>scaled as a probability</u>. The formula (1.1) for finding an overall $u(a)$ by weighting the utilities of contingent consequences $v(c)$ is exactly the formula for finding the overall probability associated with a set of contingent probabilities.

We have foregone presenting a formal statement of the "postulates of rational choice" that underly the expected-utility rule. But a few comments are in order here:

1. We have assumed that the v(c) scale is unique, applicable to every state of the world. This will be reconsidered below under the heading of "state-dependent utility."

2. We have implicitly ruled out complementarities in utility, whereby a higher income c_s in a state s might affect the v score attached to income c_t in another state t. The justification is that c_s and c_t are not received <u>in combination</u> but only <u>as alternatives</u>; no complementarity can exist because c_s and c_t can never be enjoyed simultaneously.

3. While we have emphasized that v(c) should be intuitively thought of as scaled in terms of probability, any fixed positive linear transformation of the v(c) scale would be equally satisfactory -- because cardinality permits free choice of zero and unit interval.

1.1.5. Risk-Aversion, and the Risk-Bearing Optimum of the Individual

The "concave" form of the cardinal preference-scaling v(c) function in Fig. 1 shows diminishing marginal utility of income-- $v''(c) < 0$ -- for this individual. Such a person is said to be <u>risk-averse</u>: he would always prefer a sure consequence (level of certain income) to any probabilistic mixture of consequences (lottery or prospect) having the same mathematical expectation. In Fig. 1 we have seen that the reference lottery with equal chances of $1000 or zero (and thus a mathematical expectation of $500) is the preference equivalent of a sure income of only $250. Thus, this person must prefer a sure income of $500 to a risky lottery with a mathematical expectation of $500. It is intuitively evident that this generalizes: any point P on a concave v(c) curve will lie <u>above</u> the corresponding (vertically aligned) point along the straight line connecting any pair of positions on v(c) that bracket P. The point <u>on the curve</u> represents the utility of a given sure income; the vertically aligned point <u>on the straight line</u> represents the utility of a lottery with a mathematical expectation equal to that given amount. The generalization of this result, often referred to as Jensen's inequality, can be expressed as:

$$v''(c) < 0 \rightarrow v(E(\tilde{c})) > Ev(\tilde{c})$$

Here E symbolizes the mathematical expectation operator, and the tilde indicates that c is a non-degenerate random variable.

It follows immediately that a risk-averse individual endowed with a given sure income would never accept a <u>fair gamble</u>, a lottery whose mathematical expectation of net return equals zero (since it would shift him from a position on the v(c) curve to a vertically aligned point below it). A gamble would have to be somewhat better than fair, offer some positive mean return (just how much depends upon his degree of risk-aversion) to be acceptable. On the other hand an individual whose $v(c)$ function had the opposite "convex" curvature, representing increasing marginal utility of income -- $v''(c) > 0$ -- would be happy to accept any fair gamble and even, up to a point, gambles worse than fair (offering a negative mean return). Such an individual is said to display risk-preference. An individual on the borderline, with a $v(c)$ function that is linear (constant marginal utility of income or $v''(c) = 0$) is said to be risk-neutral. A risk-neutral individual would accept, reject, or be indifferent to gambles that are respectively better than fair, worse than fair, or just fair.

It might be thought that the "concave" $v(c)$ function of Fig. 1 applies only to one psychological type of person, or perhaps only to people at particular times or stages in the life-cycle, so that the world would consist of a mixture of risk-averse, risk-neutral, and risk-preferring types. But the observed fact of diversification of assets suggests that risk-aversion is normal. An individual who is risk-neutral, for example, would plunge all of his wealth in that single asset which -- regardless of its riskiness -- offered the highest mathematical expectation of return. But we scarcely ever see this behavior pattern, and do observe more typically that individuals hold a variety of assets, thereby reducing their risk of ending up with an extremely low level of income.

What of the seemingly opposed evidence that fair gambles (and, indeed, gambles generally worse than fair) are accepted by throngs of bettors at Las Vegas and elsewhere? There have been some attempts to construct preference-scaling functions $v(c)$

that would be consistent with gambling over certain ranges of income and with avoiding gambles over other ranges [Friedman and Savage 1948, and see also Markowitz 1952]. These constructs run against the difficulty that individuals would never be at equilibrium in any risk-preferring range of their v(c) curves. To leave this range they would jump at the chance of accepting enormous riches-or-ruin gambles, provided only that these were available on a fair or nearly-fair basis. Such behavior is surely rare, and there is no indication of ranges of income that are thus depopulated. Except in more or less pathological cases, gambling at adverse odds is a recreational rather than income-status-determining activity for individuals. As evidence, we observe that actual gambling as in Las Vegas is mostly of a repetitive small-stakes nature, more or less guaranteed not to change one's overall income status in the long run [Hirshleifer 1966 p. 261].

That risk-aversion is the normal situation is indicated in a different way by Fig. 2. Here the familiar-looking indifference curves u^o, u', u'', ... represent the preference-scaling function v(c) of Fig. 1 in contingent-income or state-claim space. Specifically, assume for simplicity that there are only two states of the world s_1 and s_2, with corresponding fixed probabilities π_1 and $\pi_2 \equiv 1 - \pi_1$, and contingent consumption variables c_1 and c_2. Then (1.1) reduces to the special form (1.1'):

$$u(a) \equiv \pi_1 v(c_1) + \pi_2 v(c_2) \qquad (1.1')$$

If follows immediately, since du = 0 along any indifference curve,

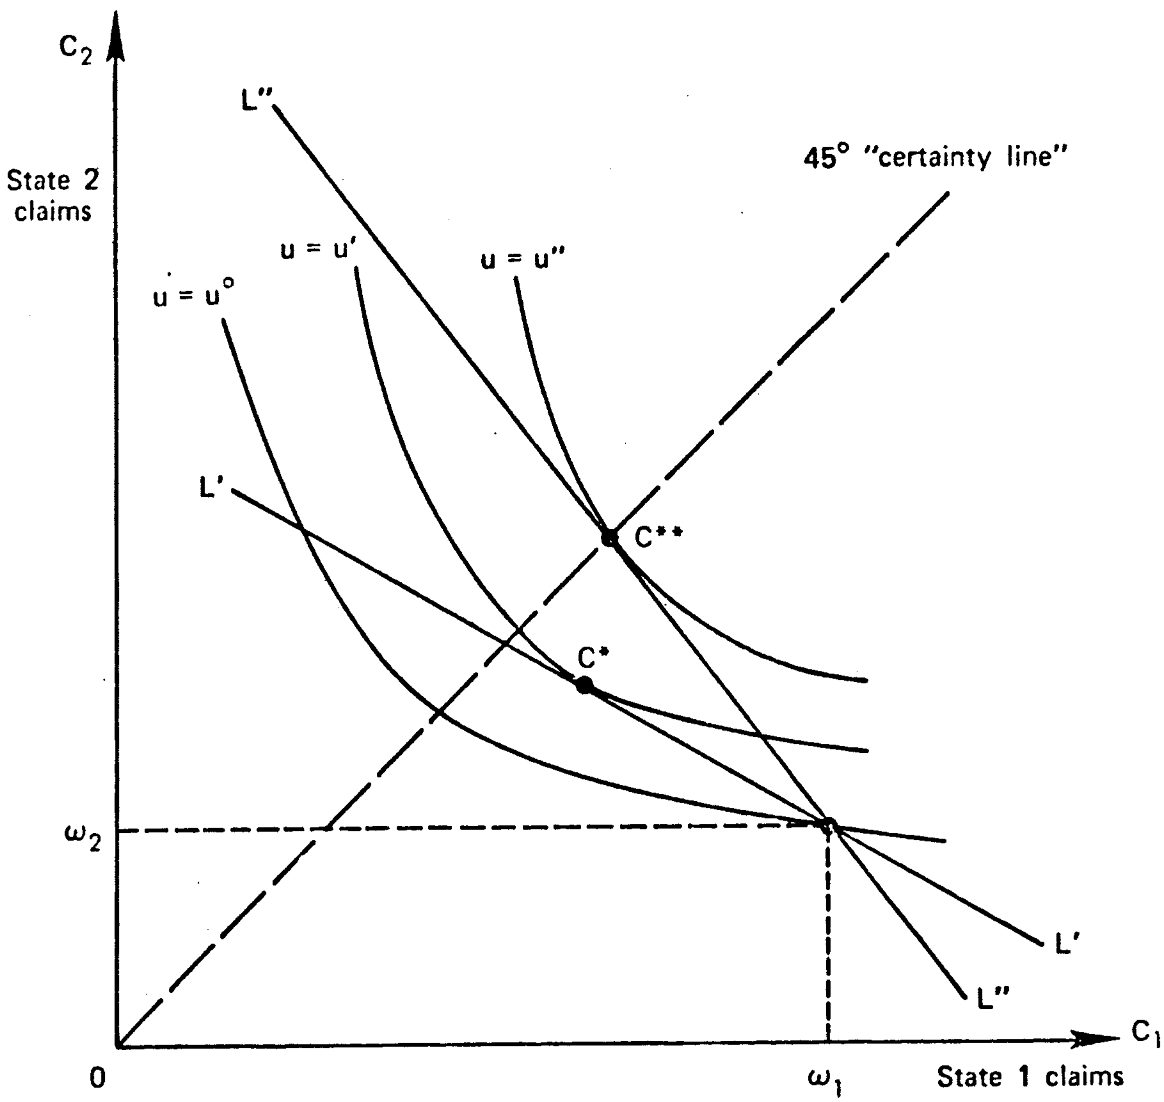

Fig. 2 — The preference map in contingent consumption or state-claim space (probability fixed)

that the indifference-curve slopes in Fig. 2 are related to the marginal utilities $v'(c)$ via:

$$\left.\frac{dc_2}{dc_1}\right|_{du=0} \equiv -\frac{\pi_1 v'(c_1)}{\pi_2 v'(c_2)} \tag{1.3}$$

It is then elementary to show that the state claims preferred to any given bundle C^* form a convex set if and only if $v''(c) < 0$ -- i.e., only if the preference-scaling function $v(c)$ is "concave."

Now let us suppose that the individual is a price taker in a market where contingent claims c_1 and c_2 can be exchanged in the ratio P_1/P_2. The price ratio, together with the individual's endowment position (ω_1,ω_2), determines his budget line L'L' in Fig. 2. It is then geometrically evident that, given the standard indifference-curve curvature that stems from risk aversion, C^* will not normally be at the intersection of the line L' with one of the axes -- i.e., the individual will want to "diversify" his holdings of state claims. Following standard techniques, the optimum position C^* along the budget line is the tangency determined by the condition:

$$-\left.\frac{dc_2}{dc_1}\right|_{du=0} \equiv \frac{\pi_1 v'(c_1)}{\pi_2 v'(c_2)} = \frac{P_1}{P_2} \tag{1.4}$$

We can arrive at a much stronger result for the special case where the price ratio P_1/P_2 equals the probability ratio π_1/π_2. Since the condition for "fair" gambles can be expressed as $\pi_1\,\Delta c_1 + \pi_2\,\Delta c_2 = 0$ -- the expected net return over consequences is zero -- and since in market exchange $\frac{\Delta c_2}{\Delta c_1} \equiv -\frac{P_1}{P_2}$, this equality of the price ratio and the probability ratio corresponds to the market offering fair gambles. Then the condition (1.4) simplifies to:

$$\frac{v'(c_1)}{v'(c_2)} = 1 \qquad (1.4')$$

Given the state-independent form of the v(c) curve as in Fig. 1, equation (1.4') corresponds to a solution where $c_1 = c_2$ -- i.e., to a C** tangency optimum at the intersection of the new budget line L"L" with the 45° "certainty line."

Thus, confirming our earlier result, starting from a certainty position the individual would never accept any gamble at fair odds. And, if endowed with a risky situation he would use the fair-odds condition to "insure" by moving to a certainty position. That is, he would accept just that risky contract -- offering income in one state in exchange for income in another -- which exactly offset his endowed gamble. (Correspondingly, of course, if the odds are not fair the individual would accept some risk so that C* would lie off the 45° line.) Note that mere acceptance of a risky contract does not tell us whether the individual is moving away or toward a certainty position (enlarging or reducing his risk exposure) -- the riskiness of the endowment position must also be taken into account.

A natural next step, from the theoretical point of view, would be to explore responses of the individual's optimum position C* to a variety of parametric changes. While space precludes a detailed presentation, two branches of this line of analysis are now briefly considered. First of all, consideration of the response of C* to increments of _income_ has led to a specification of interesting properties of the individual's utility function.

Specifically, an individual is said to have constant _relative risk-aversion_ (RRA) if an increase in income, price ratios held constant, leads to a new risk-bearing optimum C* which involves _proportionately_ more holdings of claims for each and every state (i.e., the new C* lies on a ray from the origin through the original C*). Geometrically, if this held everywhere the individual's preference map in Fig. 2 must be homothetic. It can be shown that the condition for constant RRA is:

$$-\frac{c_s v''(c_s)}{v'(c_s)} = R \qquad (1.5)$$

where R is a constant. If R is an increasing function of income (_increasing_ RRA), a rise in income leads to a new C* proportionately closer to the 45° line than the original C* -- and the reverse if R is a decreasing function of income (_decreasing_ RRA).

An individual is said to have constant _absolute risk-aversion_ (ARA) if a rise in income results in a final C* position representing equal absolute increases in state-claim holdings in comparison with the original C*, for each and every state. This corresponds geometrically to the C* position shifting parallel to the 45° line as income increases. The analytic condition for constant ARA is given by

$$-\frac{v''(c_s)}{v'(c_s)} = A$$

where A is a constant. If A is a rising function of income (the condition for increasing ARA), as income increases the new C* position will be _absolutely_ closer to the 45° line in comparison with original C* -- and absolutely farther from the 45° line in the case of decreasing ARA.

Assumptions about absolute and relative risk aversion are often critical in determining the optimal response to an increase in the _riskiness_ of a decision-maker's environment [Rothschild and Stiglitz 1971]. Since a general discussion would require a relatively long technical development we shall consider only a simple illustration.

Suppose an individual has a choice of holding a wealth endowment position ω on the 45° certainty line at a zero net return or of investing in a risky asset with a net return per dollar invested of α in state 1 and $-\beta$ in state 2. With \$x (≥ 0) invested in the risky asset, contingent final wealth is

$$(c_1,c_2) = (\omega+\alpha x,\ \omega-\beta x)$$

In state-claim space the opportunity locus is then a line with slope $-\beta/\alpha$ through the certainty point (ω,ω). This is depicted in Fig. 3. Also depicted is the indifference curve, $\pi_1 v(c_1) + \pi_2 v(c_2) = u^*$ through the optimum C*. Since the latter has a slope of $-\pi_1/\pi_2$ at the 45° line, a necessary and sufficient condition for investment in the risky asset is

$$\pi_1/\pi_2 > \beta/\alpha \qquad (1.6)$$

that is, the expected return, $\pi_1\alpha-\pi_2\beta$, is positive.

Intuitively, for any given investment x, the agent's environment becomes more risky if, without a change in its mean, the random return becomes more widely dispersed in the sense of

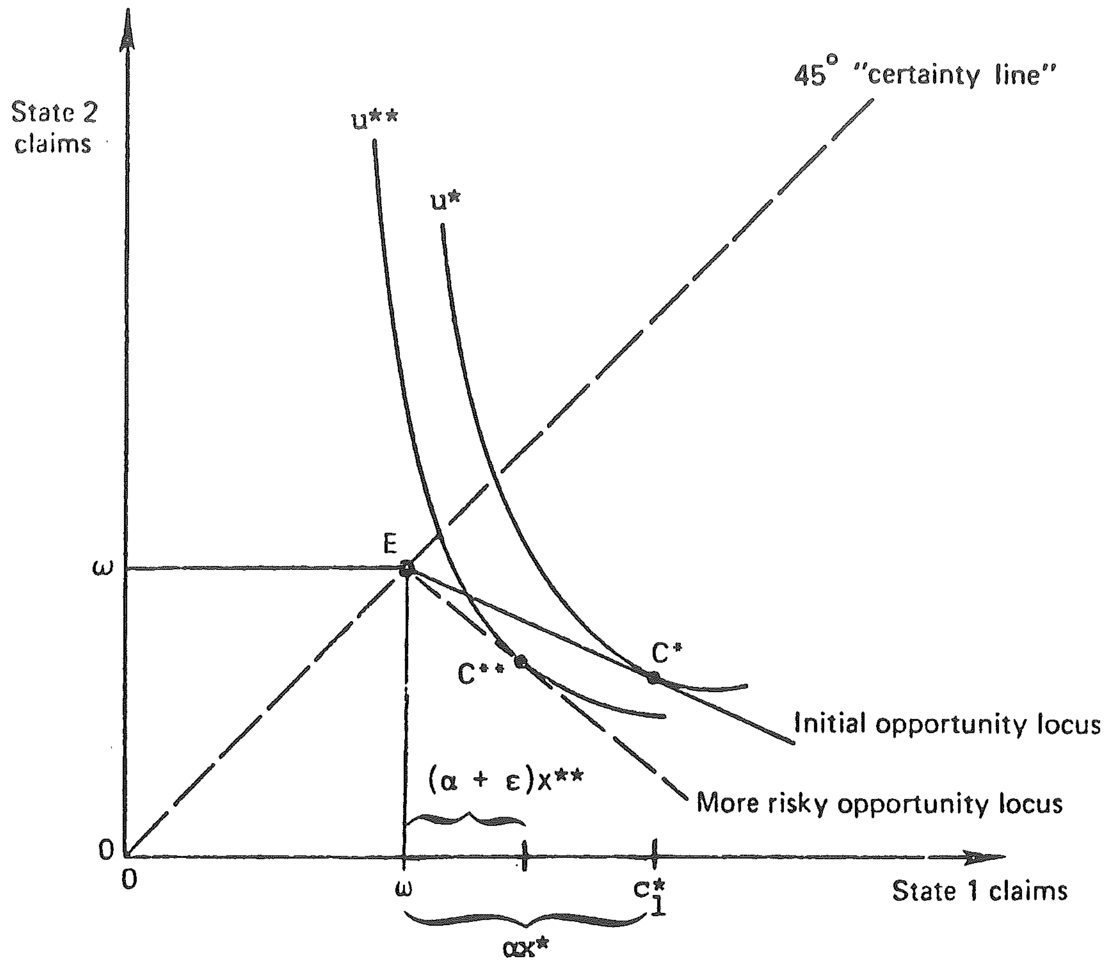

Fig. 3 — Investing in a risky asset

having more weight in the tails of the distribution [Blackwell and Girshick 1954, Rothschild and Stiglitz 1970]. For the two-state case a simple mean-preserving increase in the riskiness of the net return is achieved by increasing the state 1 return by ε and decreasing the negative state 2 return by $\varepsilon\pi_1/\pi_2$. The steepness of the new budget line is then

$$\frac{\beta + \varepsilon\pi_1/\pi_1}{\alpha + \varepsilon}$$

From (1.6) $\pi_1/\pi_2 > \beta/\alpha$ hence we have

$$\pi_1/\pi_2 > \frac{\beta + \varepsilon\pi_1/\pi_2}{\alpha + \varepsilon} > \beta/\alpha \tag{1.7}$$

The mean-preserving increase in risk thus has the effect of rotating the budget line in a clockwise direction. Applying the Hicks-Slutsky decomposition, the pure substitution effect is a movement around the indifference curve $u(c_1,c_2) = u^*$ towards the 45° line. Moreover, assuming that the individual exhibits decreasing absolute risk aversion, the income expansion path slopes away from the 45° line for higher incomes. The income effect therefore reinforces the substitution effect in the sense that it moves the optimal consumption bundle even closer to the 45° certainty line. It follows immediately that, given decreasing ARA, the optimal response to a mean-preserving increase in riskiness is to reduce investment in the risky asset.

1.2 Market Equilibrium Under Uncertainty

We now turn from individual decisions to examine market interactions under uncertainty. Recall however that we are not dealing with what is called market uncertainty (with its characteristic phenomena of search and of trading at non-clearing prices). Rather, we are dealing with event uncertainty. And we shall generally be assuming perfect but not necessarily complete markets: trading in consumption claims contingent upon alternative states of the world takes place at market-clearing prices, but not all definable claims may be separately tradable.

1.2.1. Risk-Sharing

If both parties in some transaction are risk-averse they will generally wish to enter into a contract in which the total risks and returns are shared. This can be illustrated by the Edgeworth box in Fig. 4 [Brainard and Dolbear 1971, Marshall 1976], which for concreteness may be thought of as illustrating a "share cropping" problem [Cheung 1969, Reid 1976]. The alternative states of the world are "good crop" or non-loss state N and "bad crop" or loss state L, with associated contingent claims c_N and c_L. Because of the difference in social totals of income in the two states, the box is vertically elongated. Given agreed-upon probabilities π_L, π_N $(=1-\pi_L)$, the indifference curves for each agent have the same absolute slope π_L/π_N along their respective 45° certainty lines. It follows that the contract curve TT must lie between the two certainty lines.

Suppose that the Worker W has an alternative opportunity which would yield the certain bundle E along his 45° line (a fixed wage independent of the state). If the Landowner O offers E and thereby bears all the risk, it can be seen that there are unexploited gains from trade. Landowner and Worker have an incentive to negotiate an alternative risk-sharing contract somewhere on the contract curve TT. For the limiting case of a perfectly elastic supply of workers all the gains go to the Landowner and the equilibrium contract is the point F.

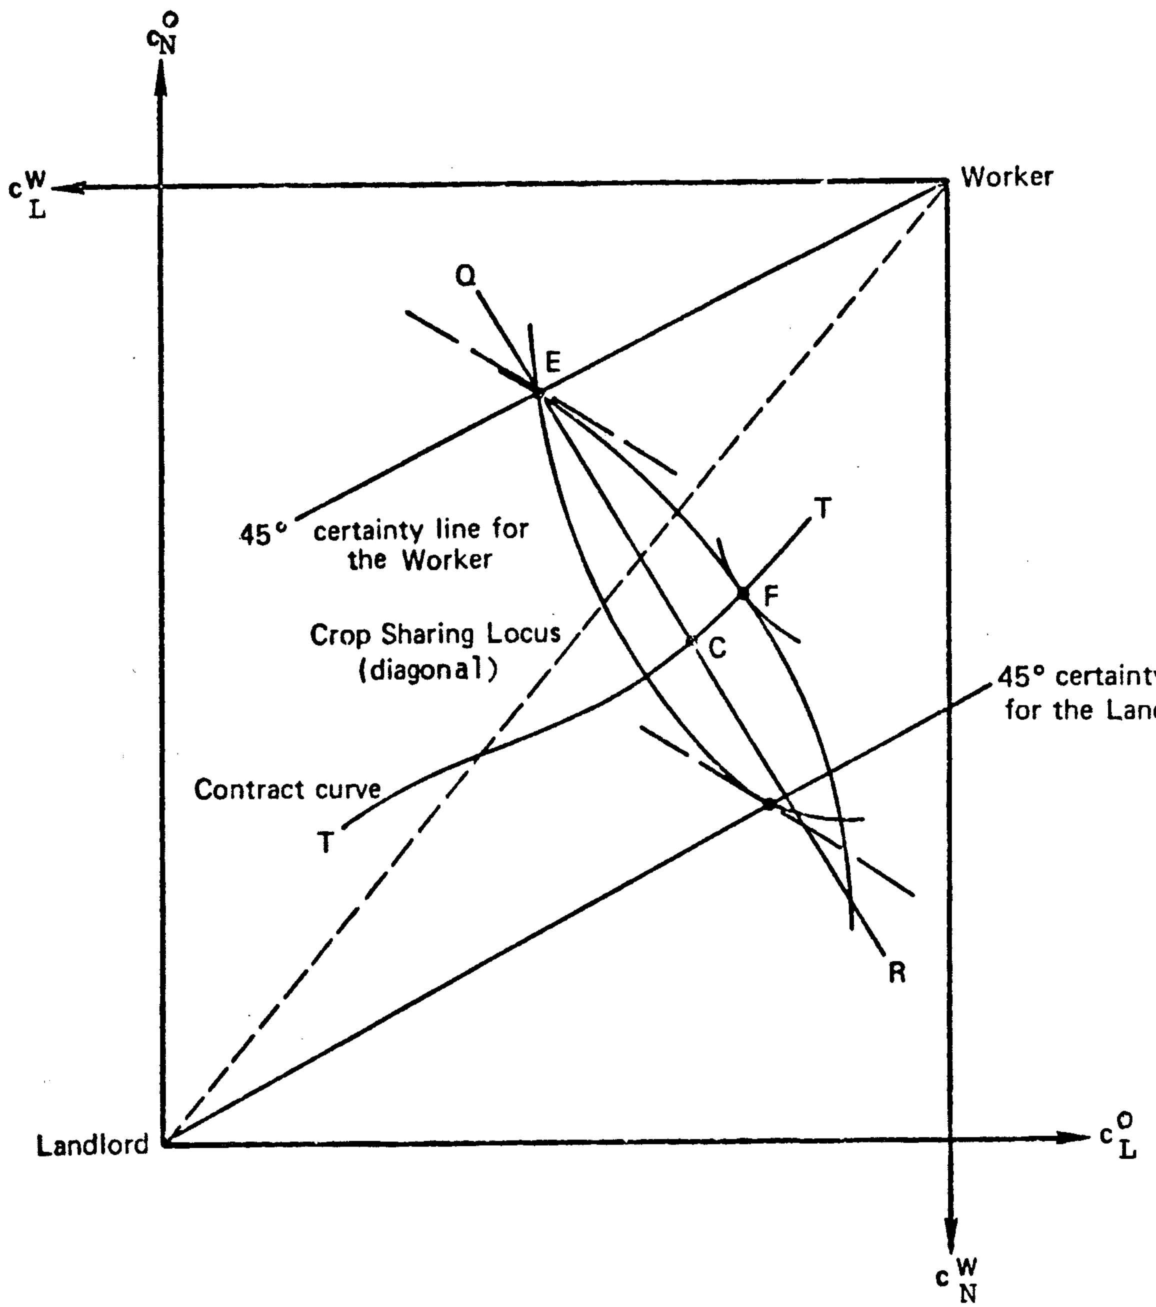

Fig. 4 — Risk Sharing

If the individuals were constrained to strict proportionate sharing of the income totals in the two states, the equilibrium would have to lie along the main diagonal of the Edgeworth box. (This would represent a kind of "incomplete market" for the trading of contingent claims.) Such a solution would not in general be Pareto-optimal, but it might be a rather close approximation of a point on the contract curve. Proportionate sharing in a world of unequal social totals of income would be strictly consistent with a Pareto optimal solution only if conjoined with side payments from one party to another. (With two states of the world a side payment in just one of the states would be required; with S states, a set of S-1 conditional side payments would be needed.)

1.2.2 Insurance

The Edgeworth box in Figure 4 can be given another interpretation: the risk sharing can be regarded as "mutual insurance." Indeed, all insurance is best thought of as mutual [Marshall 1974b]; insurance companies are only intermediaries in the risk-sharing process. Again L would correspond to a "loss" state of the world and N to a "non-loss" state of the world (recognizing that in general loss is a social phenomenon).

For the particular endowment point E all the initial risk is born by individual O since individual W is on his certainty line. Under complete contingent markets the two individuals trade contingent claims c_N and c_L at some price ratio P_L/P_N. If the price ratio were fair ($P_L/P_N = \pi_L/\pi_N$) we know from equation (1.4) above that a price-taking individual would move to a position on his certainty line. But as the 45° lines in Figure 4 do not coincide, both parties cannot simultaneously do this. So from the starting point E the two would move along a line like QR to a risk-sharing solution like C on the contract curve. Note that the absolute slope of the market line ($=P_L/P_N$) exceeds the absolute slope of the dashed lines ($=\pi_L/\pi_N$). That is, claims to income in the less affluent state L command a relatively high price. Since the ratio P_L/P_N corresponds to the "premium per dollar of insurance indemnity" as described above, insurance protection against loss is normally priced at higher than fair or "actuarial" prices.

Only if the Edgeworth box were square (the social totals of income are constant over states so that there is no "social risk"), would the competitive equilibrium prices be "fair." For example, two individuals 0 and W might have equal initial incomes I_0 and I_W, in a situation where a given hazard was however sure to impose a fixed loss ℓ on exactly one of them (with known, not necessarily equal, probabilities for each). In this limiting case the two 45° lines collapse down to a single diagonal 45° line, which would also be the contract curve. The two states are "loss strikes 0" and "loss strikes W." Each party will want to exchange income in his non-loss state (the "premium") for compensation to be received in his loss state (the "indemnity"). At the price ratio (premium/indemnity ratio) that corresponds to the respective probabilities, trading will take place to the contract curve (45° certainty line); at these fair prices, all private risk is also eliminated.

Social Risk

The case where a loss will strike exactly one of the two parties represents perfect negative correlation of risks. Apart from this extreme situation it would be necessary to take account of four distinct states of the world: loss suffered by (1) neither person (2) 0 only (3) W only, and (4) both persons. Let π_n be the probability that the number of losses is n. Continuing to assume equal initial incomes and constant loss amount, let us add the simplifying assumption of symmetry so that the probability of state 2 equals that of state 3 (and hence equals $\frac{1}{2}\pi_1$). The probabilities of each state can then be expressed in terms of the probability of loss by each person p ($\equiv\frac{1}{2}\pi_1 + \pi_2$), $q \equiv 1-p$, and the correlation coefficient r between individual outcomes as:

$$
\begin{aligned}
&\text{probability of state 1} &&\equiv \pi_0 &&= q(q+rp)\\
&\text{probabilities of state 2 and 3} &&= \tfrac{1}{2}\pi_1 &&= pq(1-r)\\
&\text{probability of state 4} &&\equiv \pi_2 &&= p(p+rq) \qquad (1.8)
\end{aligned}
$$

The main lesson to be derived from this development is that, in general in a world of uncertainty, <u>full insurance</u> (whereby each party attains his "certainty line") <u>is impossible</u> [Hirshleifer 1953, Brainard and Dolbear 1971, Marshall 1974]. There is a "social risk" in that the social total of losses may be 0, 1, or 2; there is no way of arranging affairs so that <u>everyone</u> can have the same personal income regardless of the social totals of income available. (Of course, <u>some</u> individuals can achieve certainty positions -- but only if others' risks are correspondingly greater). Note also that, as a result of diminishing marginal utility, claims to income in state 4 will be the most valuable and claims to income in state 1 the least valuable (relative to the respective probabilities).

In conventional insurance arrangements, protection would ordinarily be offered to an individual like 0 without distinction between states 2 and 4 -- more generally, <u>without</u> consideration of the social total of losses imposed upon the aggregate of persons in the insurance pool. This therefore is another kind of "incomplete-markets" situation. Since, as we have seen, claims to income in state 4 (more generally, to income in high-social-loss states) are particularly valuable, it follows that the price of conventional insurance must be higher than "fair." That is, the premium/indemnity ratio must exceed the "odds" p/(1-p) that any particular individual will suffer a loss. And, in consequence, the individual will (as we have seen) never "fully insure."

In the standard insurance discussions in the literature, it is usual to call upon the Law of Large Numbers to justify the assumption that the social totals of income are <u>approximately</u> constant over states. As the number of members N in the insurance pool grows, the variance of the proportion of members suffering losses falls -- so that the error committed by assuming away the "social-risk" problem tends to diminish. Nevertheless, this error will <u>not</u> tend toward zero unless the separate individual risks are on average uncorrelated [Markowitz 1959, p. 111]. In our special case above, assuming that the correlations between all pairs of risks equal some common r (which can only hold if $r \geq 0$), the variance σ^2 of the proportion of N exposures suffering losses is:

$$\sigma^2 = \frac{pq}{N} + \frac{rpq(N-1)}{N}$$

In the limit as N increases, σ^2 approaches the value rpq, which remains positive unless $r = 0$.

We see, therefore, that "social risk" is not exclusively due to small numbers; it persists even with large numbers if risks are on average correlated. In the language of portfolio theory, risks have a "diversifiable" element which can be eliminated by purchasing shares in many separate securities (equivalent to mutual insurance among a large number of individuals) and an "undiversifiable" element due to the average correlation between risks. It follows then that a particular asset will be more valuable the less is the correlation of <u>its</u> returns over states with the aggregate returns of all assets together -- the variability of which is the source of undiversifiable risk. As this concept is applied in modern investment theory, the correlation of returns on each particular security with the returns from the "market portfolio" consisting of all securities together is indicated by that security's "beta" parameter [Sharpe 1978, Ch. 6]. Securities with low or, even better, negative betas trade at relatively high prices (i.e., investors are satisfied with low expected rates of return on these assets) because they provide their holders with relatively large returns in just those states of the world where aggregate incomes are low (marginal utilities are high).

The "social risk" phenomenon therefore provides two reasons why insurance prices may not be fair or actuarial, so that purchase of coverage is ordinarily less than complete: (1) if the number of risks in the insurance pool is small, so that the Law of Large Numbers cannot fully work, or (2) even with large numbers, if risks are on average correlated.

Additional reasons for non-actuarial insurance exist. One obvious factor is transaction costs [Ehrlich and Becker, 1972], the consideration of which raises complex issues that cannot be considered here. We shall consider instead two other classes of phenomena: (1) state-dependent utilities, and (2) disparities of information between insurer and insured that go under the heading of "moral hazard" and "adverse selection."

State-dependent utilities:

Our discussion to this point has been based upon the unique state-independent preference-scaling function of Fig. 1. More generally, however, the utility we attach to income c may vary with the state of the world; in the insurance context, we may have a $v_N(c)$ curve for the non-loss state N and a separate (lower) $v_L(c)$ curve for the loss state L (Fig. 5a). This will be appropriate wherever the object insured cannot be regarded simply as an income-equivalent -- for example if it is an irreplaceable heirloom, or your own life, or your child's. There is no contradiction with the development in Sec. 1.1.4 above that led to the picture in Fig. 1, for there it was assumed (merely as a simplification) that utility was a function of the quantity of a single generalized "income" commodity c. Here we are saying that utility is a function of both c and an "heirloom" variable h, where $h = 0$ defines the loss state L and $h = 1$ the non-loss state N. Hence the two functions $v_N(c)$ and $v_L(c)$ do not represent distinct utility functions, but different sections through a single $v(c,h)$ function.

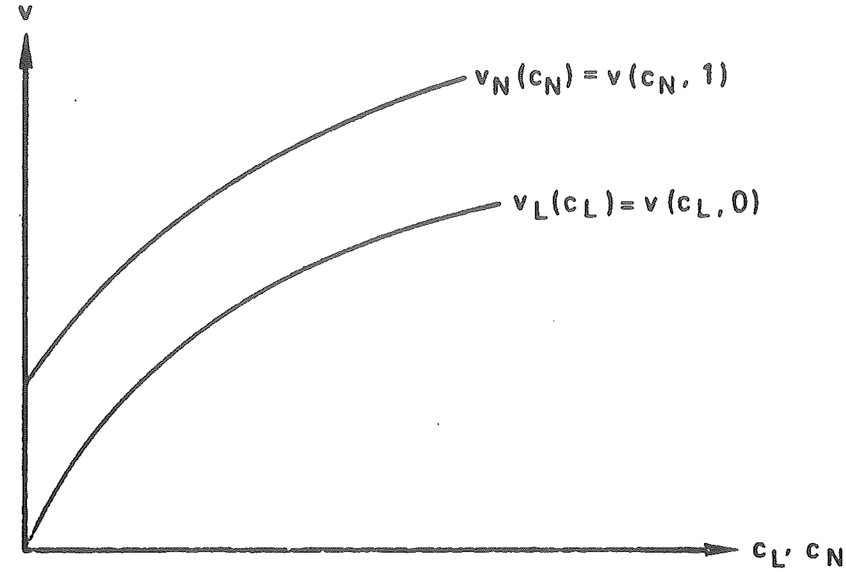

Fig. 5a — State-Dependent Utility

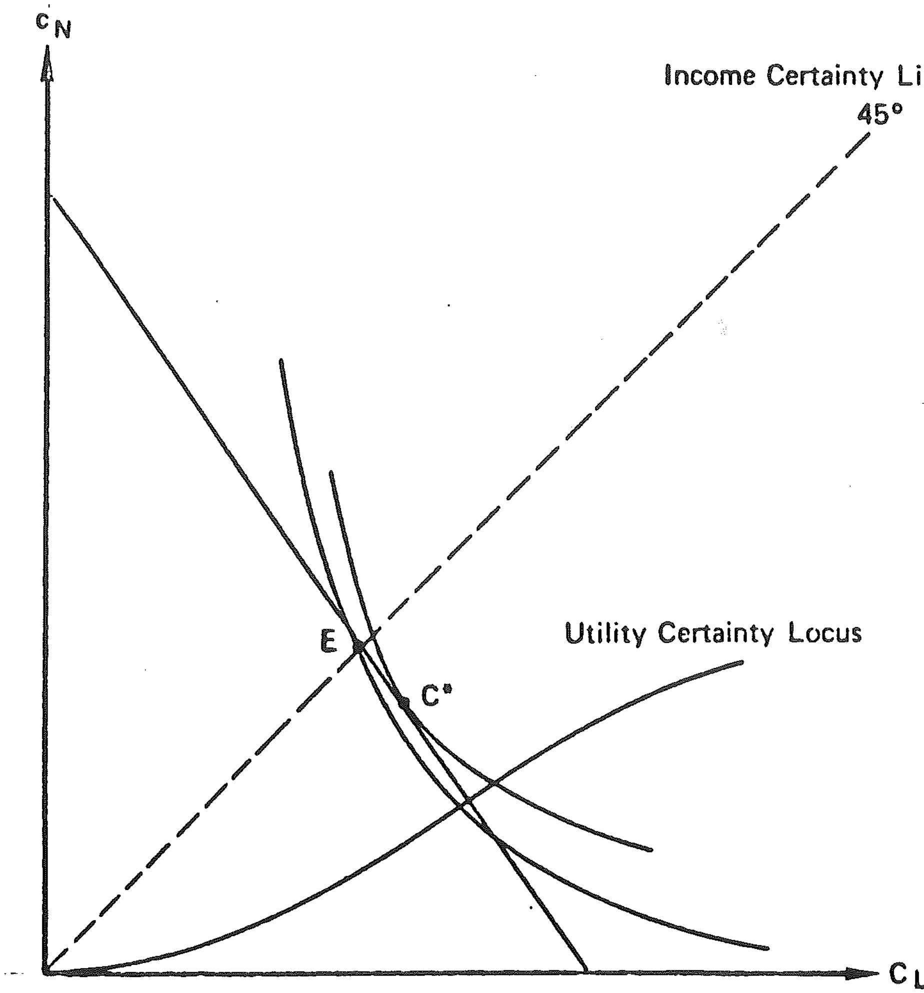

Fig. 5b — "Heirloom" Insurance

Indifference curves in state-claim space associated with the expected utility function $u(c_L,c_N) \equiv \pi_L v_L(c_L) + \pi_N v_N(c_N)$, are depicted in Figure 5b. Since utility at any given income level is lower in the loss state, the "utility certainty locus" [Cook and Graham 1977] is no longer the 45° line but instead lies everywhere below the latter. An individual purchasing enough insurance against the loss state to end up on the UCL has <u>fully insured utility</u>.

The individual optimization condition will, apart from the N or L subscripts attaching to the marginal utilities v', have the same form as equation (1.4):

$$\frac{\pi_L v_L'(c_L)}{\pi_N v_N'(c_N)} = \frac{P_L}{P_N} \tag{1.9}$$

Assuming, as a benchmark, that prices are fair, (1.9) reduces to:

$$v_L'(c_L) = v_N'(c_N) \tag{1.9'}$$

With actuarial insurance available, the individual thus equates his marginal utilities.

With uncertainty only over the possible loss of the "heirloom," the individual has an initial endowment point E on the 45° income certainty line. Therefore insurance against the loss state is optimal if and only if the indifference curve through E is steeper at this point than the budget line, that is if

$$v_L'(c_E) = \frac{\partial v}{\partial c}(c_E,0) > v_N'(c_E) = \frac{\partial v}{\partial c}(c_E,1)$$

The desirability of insuring against the loss state thus depends upon whether or not income c and the "heirloom" variable h are Edgeworth substitutes -- i.e., whether the cross-derivative of the cardinal utility function is negative. For an heirloom such as an ancestral painting with negligible cash value it is hard to establish an a priori case either way. We can thus expect to find that some people insure such objects while others, similarly situated, do not.

However if h = 0 represents an injury leaving the individual with major paralysis it seems reasonable that marginal utility will be higher in the loss state (when paralyzed one "needs" additional income to achieve a similar consumption bundle). In such cases the optimum must therefore lie to the southeast of the income certainty line -- but not necessarily southeast of the utility certainty locus. That is, the individual will buy insurance against loss, but not necessarily so much as to be "fully insured" in the sense of not caring whether or not the injury occurs.

The situation is very different if the variable h represents the life of one's child. It then seems plausible that h and c are complements; if your child dies (h=0), you have less need for income, since you planned to spend it mainly on him. In such a case it is optimal to transfer income from the loss state to the non-loss state. That is, such an individual would "reverse insure" -- would bet that the loss would not occur. (Contractually, instead of insuring his child's life he might buy a life annuity for him.)

We see that once allowance is made for state-dependent utility, it can no longer be presumed that individuals offered actuarial insurance terms will move to certainty positions -- either certainty with respect to income, or with respect to utility.

<u>Adverse selection and moral hazard</u>

We now turn back to the simple assumption of state-independent utility, and also assume away "social risk," in order to isolate another force operating upon individual decisions and market equilibrium: the inability of insurers to monitor perfectly the behavior or identify the risk-status of insureds. Note that this is an <u>informational</u> problem, and thus our analysis here has close ties with topics to be taken up in Part 2 below.

Consider first only two risk classes of individuals -- both with the same endowment point E defining initial income c_N^e in the non-loss state and $c_L^e \equiv c_N^e - \ell$ in the loss state. To simplify notation the probability of loss will be denoted here by π (rather than π_L). The lower loss probability for the low risk class, π', leads to the (dashed) indifference curve in Fig. 6a which is flatter than the (dotted) indifference curve of the high risk class with higher loss probability π''. For further simplicity assume that insurance is offered only on a full-coverage-or-none basis: each individual must choose between remaining at his endowment point or moving along his budget line to the specific point where it intersects the 45° certainty line. The slope of the budget line re-

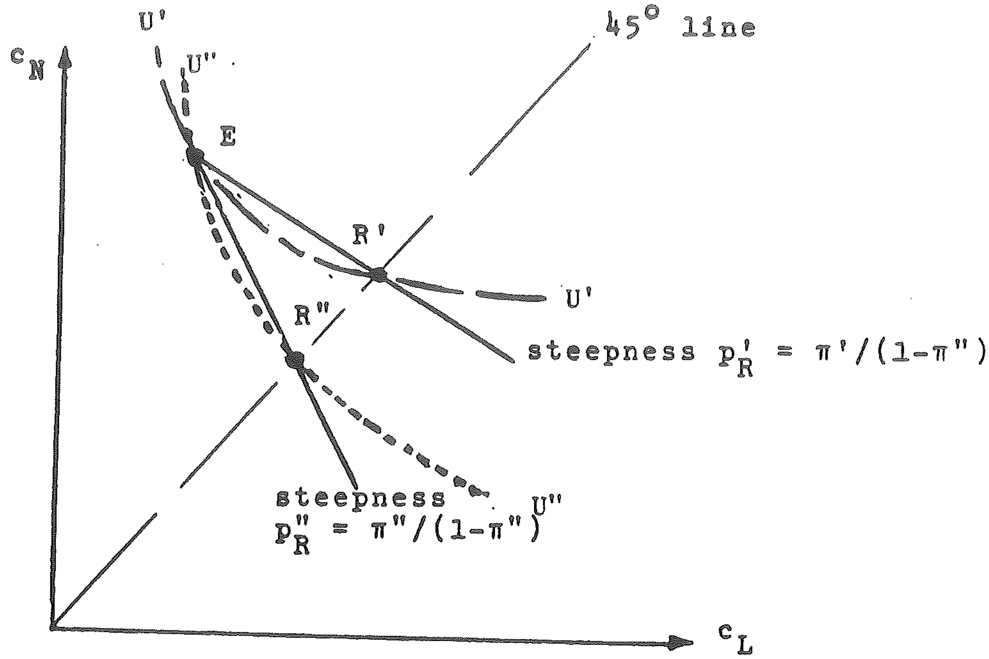

Fig. 6a: "Reservation price ratios" for two risk classes

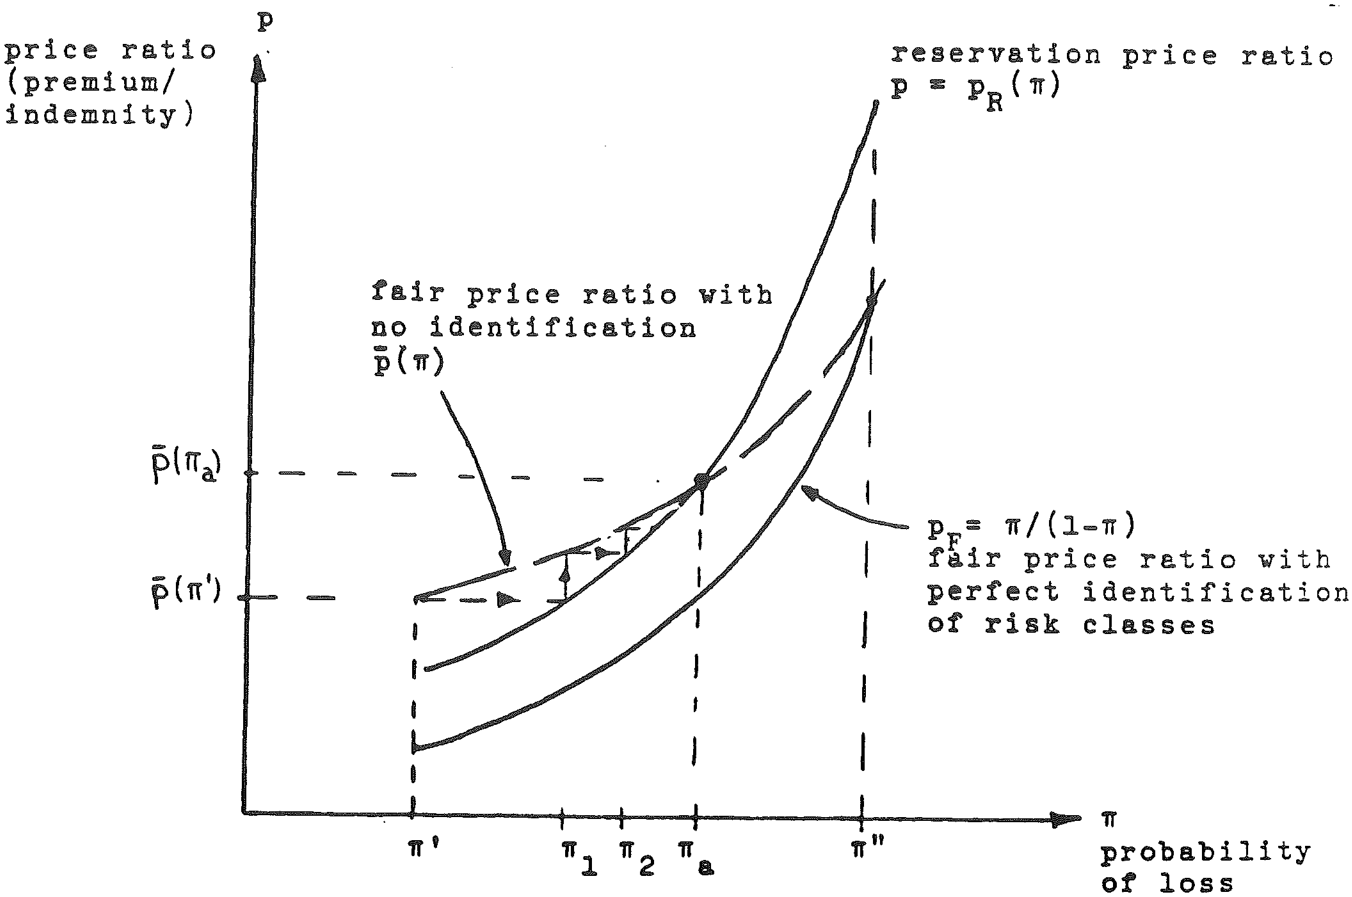

Fig. 6b: Adverse Selection

flects, of course, the premium/indemnity ratio p (= the price ratio p_L/p_N). If insurance were offered each group at fair terms, involving a higher price ratio for the high-risk (dotted) group, individuals in each class would find it optimal to insure to a certainty position (tangency solution along the 45° line) with of course a utility improvement over the endowment situation -- as illustrated by the C** solution in Figure 2. This familiar result is not shown in Figure 6a, which pictures not the "fair" budget lines for the two groups but rather the steeper "reservation price ratio" budget lines. At these reservation prices, each individual is just indifferent between his endowment position and the permitted "full-coverage" alternative on the 45° line. Of course, the reservation price also is higher for the high-risk group ($p_R'' > p_R'$). For either group, full-coverage insurance would not be purchased at any premium/indemnity ratio $p > p_R$.

Generalizing, Figure 6b illustrates a situation where risk classes are distributed continuously over the interval $[\pi', \pi'']$. Each risk class π has a characteristic reservation price ratio as shown by the curve $p_R(\pi)$. If insurance were offered at fair terms, on the other hand, the respective fair price ratios are shown by the curve $p_F = \pi/(1-\pi)$ lying below p_R throughout.

But now suppose that the insurers (other members of the mutual insurance pool) have no way of distinguishing individuals belonging to different risk classes. Instead they offer insurance based on the average probability of loss across risk classes. Let $\bar{p}(\pi)$ be the price ratio resulting in zero expected profits if full coverage were obtained only by those risk classes for whom the probability of loss exceeds π. As depicted in Figure 6b the curve $\bar{p}(\pi)$ must lie strictly below the fair price ratio with perfect identification, except at $\pi = \pi''$. If the difference between the highest and lowest risk classes is sufficiently great there is some probability of loss π_1 such that $p_R(\pi) < \bar{p}(\pi')$ for all risk classes $\pi < \pi_1$. That is, the lowest risk classes have a reservation price ratio which is lower than the fair price ratio when all risk classes are pooled. It follows that only those with loss probability greater than π_1 will purchase full-coverage insurance. To cover expected claims insurance companies must, therefore, raise the price ratio to $\bar{p}(\pi_1)$ which reflects the average probability of loss for all risk classes for whom $\pi \geq \pi_1$. But this in turn results in further exit by risk classes in the interval (π_1, π_2) and again insurance companies are forced to raise the premium/indemnity ratio $\bar{p}$. Only when all those risk classes with a loss probability π less than π_a have withdrawn is an equilibrium reached.

This is the problem of <u>adverse selection</u>. While we have described it in the insurance context, it is a much more general phenomenon. Wherever buyers are only able to observe average quality, there is a tendency for sellers not fully rewarded for high quality to withdraw from the market. In one extreme model of the used car market, Akerlof [1970] showed that only the lowest quality "lemons" would actually be traded. However, in general the pool of participants includes everyone from the lowest quality up to some cut off point [Zeckhauser 1972]. Certainly this must be the case in insurance markets.

So far, we have been assuming that the loss probability π_L is a fixed parameter for each individual. However, in many situations the actions taken by an individual effect the likelihood of different outcomes. For example by installing smoke detectors an individual can lower the probability of major residential fire damage. Similarly greater effort on the part of workers reduces the chance that the final product will be defective.

Continuing to focus on the insurance application, if it is possible to costlessly monitor the actions of each individual the solution is only slightly altered. Let $v(x,c)$ be the utility in a given state when the care level is x and income is c. For each care level, expected utility is then:

$$\pi_{L\cdot x}v(x,c_{L\cdot x}) + \pi_{N\cdot x}v(x,c_{N\cdot x})$$

where $\pi_{L\cdot x}$ is the probability of loss conditional upon the adoption of care level x. For the special case where actuarial rates are appropriate (no "social risk", and utilities not state-dependent) the solution is a rate of exchange between claims to income in the two states of:

$$\frac{P_{N\cdot x}}{P_{L\cdot x}} = \frac{\pi_{N\cdot x}}{\pi_{L\cdot x}}$$

That is, the insurance rate will reflect the level of care x. For any fixed level of x, fair insurance results in full coverage. Therefore, if the individual faces the prospect of a loss of ℓ and initial income is ω, expected utility after insurance is:

$$u(x) = v(x, \omega - \ell\pi_{L\cdot x}), \tag{1.11}$$

where $\ell\pi_{L\cdot x}$ is the premium so that $\omega - \ell\pi_{L\cdot x}$ is a constant income received in either state.

The individual then chooses that policy for which the marginal disutility of additional care is just high enough to offset the marginal utility of the lower premium resulting from the extra care.

More realistically, however, monitoring is at best imperfect. This leads to the problem known as <u>moral hazard</u>. In the extreme case, if insurance is offered at some fixed price ratio independent of any loss-prevention activity, it is evident that an individual will be motivated to entirely eliminate all such activity.

Insurers have two main ways of coping with the problem [Arrow 1963, Pauly 1968]. The first is to require the insured party to bear some portion of the risk, for example by a "deductible" provision (indemnity will be less than the loss by a fixed amount) or by "coinsurance" (indemnity will only be a proper fraction of the loss). Then insurance will continue to be provided [Shavell 1977], but moral hazard persists in the sense that insureds are motivated to engage in less preventive activity than would be efficient with costless information.

The second way of coping with the problem is to "rate" the insured individual according to his past record of claims (e.g., automobile insurance). The prospect of a higher premium after making a claim then provides further incentive to adopt some actions which reduces the probability of loss. However, such performance rating is really only a particular form of imperfect ex-post monitoring [Shavell 1977].

1.2.3 Complete and incomplete market regimes, the stock market economy, and optimal production decisions.

We have already alluded, in connection with both share cropping and insurance, to the fact that a complete set of markets will generally not be available to economic agents. If there are S distinct states of the world a market regime would be complete if the S elementary state-claims denoted $c_s (s=1,\ldots,S)$ were all separately tradable. Each individual i with endowment $\omega^i=(\omega_1^i,\ldots,\omega_S^i)$ and facing prices $P=(P_1,\ldots,P_S)$ chooses some vector of trades t^i satisfying

$$P\cdot t^i = 0$$

Equating marginal rates of substitution with the corresponding price ratios, the necessary condition for a utility-maximizing consumption choice of $c^i = \omega^i+t^i$ is then,

$$\frac{\pi_s v'(c_s^i)}{\pi_1 v'(c_1^i)} = \frac{P_s}{P_1} \quad \text{for all s.} \tag{1.12}$$

This is of course a generalization of equation (1.4) above. More generally, trading in any S distinct assets representing different combinations or packages of the c_s claims would also constitute a complete regime provided the assets were linearly independent (i.e., that none of them could be expressed as a linear combination of the others). We shall call this a regime of Complete Contingent Markets. From (1.12) each agent has the same marginal rate of substitution between every pair of state claims, hence in equilibrium the trades t^i result in a Pareto-efficient allocation.

Consider instead a "stock market economy". This will be defined as a situation where there are F distinct types of tradable assets, each consisting of some total vector of state claims $\omega^f = (\omega_1^f,\ldots,\omega_S^f)$. Individual i has an untradable endowment $\omega^i = (\omega_1^i,\ldots,\omega_S^i)$ plus endowed amounts of tradable <u>shares</u> $(\bar{\alpha}_1^i,\ldots,\bar{\alpha}_F^i)$ of the F "firms" in the economy.

If each firm's holding has market value V_f, the individual's decision problem is to choose a portfolio $(\alpha_1^i,\ldots,\alpha_F^i)$ subject to his marketable wealth constraint:

$$\Sigma_f \alpha_f^i V_f = \Sigma \bar{\alpha}_f^i V_f \tag{1.13}$$

His final consumption is $c^i = \omega^i + t^i$ where:

$$t^i = \Sigma_f (\alpha_f^i - \bar{\alpha}_f^i)\omega^f \tag{1.14}$$

The individual then chooses a portfolio to maximize :

$$u(\omega^i, \alpha_1^i,\ldots,\alpha_F^i) = \Sigma_s \pi_s \, v(c_s^i) \tag{1.15}$$

subject to (1.13) and (1.14). To achieve this he expands or contracts his holdings in the different firms until the expected marginal utility of a dollar invested in each asset is equated to his expected marginal utility of wealth, λ^i, that is:

$$\frac{\Sigma_s \Pi_s \, v'(c_s^i)\, \omega_s^f}{V_f} = \lambda^i \quad \text{for all f and all i,} \tag{1.16}$$

This directly implies:

$$\frac{\Sigma_s \pi_s \; v'(c_s^i) \; \omega_s^f}{\Sigma_s \pi_s \; v'(c_s^i) \; \omega_s^1} = \frac{V_f}{V_1} \quad \text{for all f and i.} \tag{1.17}$$

This relation, in comparison with (1.12), indicates an optimization constrained by the set of assets or claims packages (firms f = 1,...,F) through which trading may take place. Unless the set of tradable assets constitutes a Complete Contingent Market (that is, unless the F asset vectors span the full S-dimensional space) it will not in general be possible for individuals to achieve the Pareto-efficient vector of net trades by purchase and sale of shares. This is obvious for the extreme case in which all the state-claim vectors of the different firms are scalar multiples β^f of some constant vector $(\bar{\omega}_1, \bar{\omega}_2, \ldots, \bar{\omega}_S)$. For then it is evident that, for any two firms f' and f", $V_{f'}/V_{f''} = \beta^{f'}/\beta^{f''}$ -- that is, the asset price ratios will be equal to the ratios of the scaling factors. In such a world there can be no gains from trade; each individual must therefore end up consuming his own endowment.

For the one-commodity case it can be shown that the economy is efficient in the restricted sense of achieving Pareto preferred allocations of the tradable shares of different firms [Diamond 1967]. With more than one commodity the issue is complicated by the possibility of trading after the state of the world is known, that is, in a future spot market. Even with every agent correctly predicting future spot prices there may be multiple equilibria which are Pareto-ranked [Hart 1975]. Each such equilibrium can still be characterized as optimal in a weaker Nash sense [Grossman 1977a]. However, it can be shown that Pareto-dominated equilibria are sometimes unstable so the normative implications of the "Nash-optimum" are unclear.

The stock market model is particularly interesting when we pass from the realm of pure exchange to consider aggregate endowments and attendant uncertainty as being generated endogenously by <u>production</u> decisions. The vector ω^f for firm f now becomes the result of a production decision on the part of owners of the firm's shares, subject of course to constraints in the form of the production possibilities available. For expositional ease we shall characterize the firm's decision as a scalar x, generating a final output $\omega^f(x)$.

A question that has received considerable attention is whether maximization of a firm's market value V_f is in the interests of all its shareholders, and therefore would be unanimously chosen by them. In general the answer is in the negative. However, suppose that when firm f announces a new plan $x + \Delta x$, there is only a negligible effect on any shareholder's marginal utility of income in the different states. Then the expected marginal utility of wealth is unchanged and, from (1.16) the new market value of the firm must satisfy:

$$\frac{\sum_s \pi_s v_i'(c_s^i)\, \omega_s^f(x + \Delta x)}{V_f + \Delta V_f} = \lambda^i \qquad (1.18)$$

Multiplying both sides of (1.16) and (1.18) by the value of firm f and then subtracting we have:

$$\sum_s \pi_s v_i'(c_s^i)\, \{\omega_s^f(x + \Delta x) - \omega_s^f(x)\} = \lambda^i \Delta V_f, \text{ for all i.} \qquad (1.19)$$

The left hand side of this expression is individual i's expected marginal utility of the change in the firm's production plan. Since $\lambda^i > 0$ this is positive for all i if and only if ΔV_f is positive. That is, any plan will have the unanimous support of the stockholders if and only if it raises the value of the stock.

In general, however, the assumption that marginal utility is only negligibly affected by changes in the production decision of a firm is unsatisfactory. For example, suppose the only way for individuals to increase their holdings of state-s claims is to purchase a larger shareholding in firm f. (Firm f thus has a monopoly over state-s claims.) Then an increase in the size of this firm will have a significant proportional effect on the total supply of state-s claims and hence tend to lower the marginal utility of income in that state. The simple relationship given by (1.19) no longer holds and, as a result, utility maximization is inconsistent with value maximization. (The same would hold true, even under certainty, if owners of a firm with monopoly power were also consumers of the firm's product.) Moreover, in general a change in the production plan x will have different effects on the marginal utilities of different individuals. Then a change in x which is optimal for one shareholder will in general be non-optimal for another shareholder. That is, shareholders will not in general be able to agree on any jointly preferred production decision.

If, however, the elements of the state-claim distribution associated with firm f are all small, relative to the totals of state-claims generated by the entire pool of firms, an increase in the scale of production by firm f can have only a negligible effect on individuals' marginal valuations. The final term in (1.18) can then be ignored and value maximization is in the interest of every shareholder [Diamond 1967, Ekern and Wilson 1974].

Alternatively, suppose that all uncertainty is the result of randomness in production (rather than in individuals' untradable endowments). Then, even if a firm is a monopolist over certain classes of state-claims, its monopoly power (effect of its production decisions upon individuals' marginal valuations) shrinks toward zero as the number of consumer-shareholders becomes very large [Hart 1978]. Then shareholders again would unanimously support value-maximization as the goal of the firm.

1.2.4 Other Applications

In this Part 1 we have provided a relatively extensive treatment of insurance; under that heading we have been able to expound and illustrate, in rather simple format, most of the basic ideas of modern uncertainty theory. (Of course, we have scarcely been able to hint at the many exciting developments of a more advanced nature.) We have also referred, briefly, to two other applications of uncertainty theory: (1) investments and portfolios, and (2) share cropping. A number of other significant applications can only be mentioned here: (3) optimal contracts between agent and principal, for example to elicit ideal performance on the part of corporate managers [Marschak and Radner 1972, Harris and Raviv 1978, Shavell 1978, Cheung 1969, Groves 1973, Alchian and Demsetz 1972, Jensen and Meckling 1976, Zorn 1978]; (4) corporate finance, and in particular the balance between debt and equity funding [Modigliani and Miller 1958, Lintner 1962, Hirshleifer 1966, Fama and Miller 1972, Ch. 4]; (5) optimal behavior and equilibrium with respect to accidents [Vickrey 1968, Calabresi 1970, Baumol 1972, Diamond 1974]; the "value of life" appropriate for risk-taking decisions [Mishan 1971, Thaler and Rosen 1975, Conley 1976, Schelling 1968, Jones-Lee 1976, Bergstrom 1974]; and (6) choice of discount rate for public investment [Hirshleifer 1966, Arrow and Lind 1970, Sandmo 1972, Bailey and Jensen 1972, Mayshar 1977].

Part 2

THE ECONOMICS OF INFORMATION

In Part 2 we will examine the consequences of admitting non-terminal or informational actions into the repertory of economic agents. Paralleling the development in Part 1 we first consider the problem of optimization -- choice of action on the part of decision-making units (Section 2.1). The sections following cover aspects of the interaction of such decisions for determining market equilibrium.

2.1 Informational Decision-making

2.1.1 Acquisition of information

We continue to assume that the set of acts $a = 1,\ldots,A$, the set of states of the world $s = 1,\ldots,S$, and the associated consequences $c(a,s)$ are all known to the individual. He has, as before, a prior probability distribution of initial beliefs π_s as to the states of the world. The new element is that he can acquire information, receiving one of a known set of possible messages $m = 1,\ldots,M$ that in general will lead to a revision of probability beliefs. And thus in turn, to a possible revised choice of action. (In the extreme case, a message m might be conclusive as to be the occurrence of some particular state s^* -- in which case the revised or posterior distribution of beliefs $\pi_{s.m}$ will attach probability of unity to state s^* and zero to all other states.)

The warranted revision of beliefs, given any message m, is shown by Bayes' Theorem:

$$\pi_{s.m} = \frac{\pi_{sm}}{q_m} \tag{2.1}$$

That is, the revised or posterior probability $\pi_{s.m}$ assignable to state s after receiving message m equals the ratio of the joint probability π_{sm} of state s and message m both occurring, divided by the prior probability q_m of message m. Furthermore, using standard laws of probability, the numerator and the denominator on the right hand side of (2.1) can be expressed in terms of the prior probabilities π_s of the different states and the conditional probabilities or "likelihoods" $q_{m.s}$ of any message m given state s:

$$\pi_{sm} = \pi_s q_{m.s} \tag{2.2}$$

$$q_m = \sum_s \pi_{sm} = \sum_s \pi_s q_{m.s} \tag{2.3}$$

EXAMPLE: Suppose in a coin-tossing situation that an individual initially assigns equal prior probabilities of 1/3 to three states of the world: (1) coin is 2-headed, (2) coin is 2-tailed, and (3) coin is fair. The possible messages, on a single toss of the coin, are Heads (H) and Tails (T). Suppose Tails comes up. Then $\pi_{1.T}$, the posterior probability of state 1, must obviously be zero (since $q_{T.1}$, the likelihood of the message Tails given state 1, is zero). Using equations (2.2) and (2.3), the posterior probability of state 2 is the fraction with numerator $\frac{1}{3}(1)$ and denominator $0 + \frac{1}{3}(1) + \frac{1}{3}(\frac{1}{2}) = \frac{1}{2}$, whose value is 2/3. Similarly, the probability of state 3 can be found to be 1/3.

Fig. 7 is a suggestive illustration of Bayesian recalculation of probabilities on the basis of a given message m, where the possible states of the world are a continuum of values of s from zero to some upper limit S. The prior distribution shows that the bulk of the initial probability weight happens to lie toward the high end. But, the likelihood function indicates, the message (evidence) received is much more likely if s has a small rather than a large value. The posterior distribution is a compromise or average

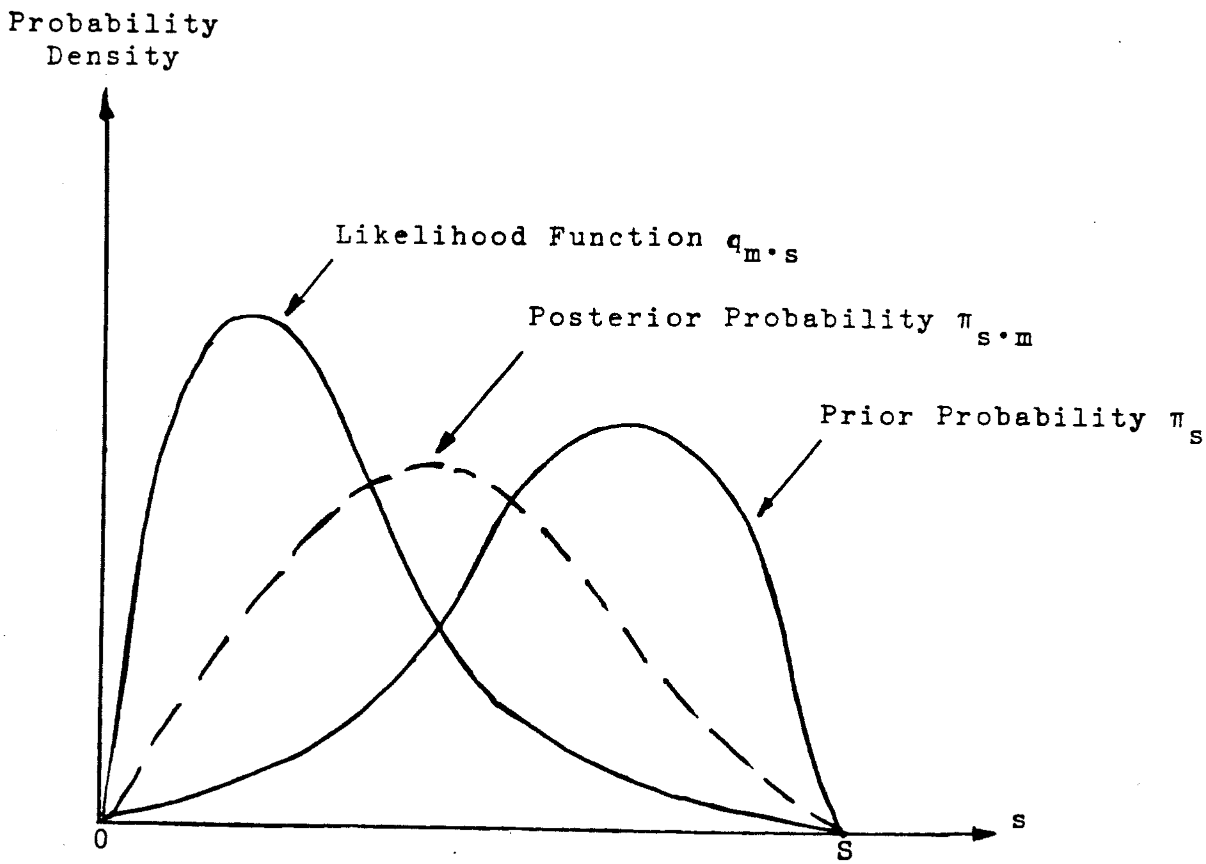

Figure 7: Bayesian Probability Recalculation

of the other two curves, derived by multiplying (for each s) the prior probabilities and likelihoods as in equation (2.2), and then re-scaling so that the integrated probability weight comes out to unity. We can conceptually picture the process as:

$$\boxed{\text{Prior beliefs } \pi_s} \quad \boxed{\text{Message m and Likelihoods } q_{m.s}}$$
$$\downarrow$$
$$\boxed{\text{Posterior beliefs } \pi_{s.m}}$$

The individual's confidence in his initial beliefs is indicated by the "tightness" of his prior probability distribution -- the degree to which he approaches assigning 100% prior probability to some single possible value for s. Evidently, the higher the prior confidence the more the posterior probability distribution will resemble the prior, for any given weight of evidence as summarized in the likelihood function. It follows quite directly, as we shall see again below, that greater confidence implies attaching lesser value to acquiring evidence.

While the prior beliefs will ordinarily be "personal" or "subjective" probabilities, in at least some cases the likelihoods might be "objective" in the sense of being calculable via the laws of probability. For example, if two tosses of a coin yield the message "Heads both times," then given the state of the world that the coin is fair probability theory says that the likelihood is $q_{m.s}=1/4$. A main theorem of Bayesian statistical theory is that as the sample size increases, the weight of the evidence tends to rise relative to the importance of the prior probabilities -- so that, in the limit, objective evidence tends to swamp out divergences in personal prior beliefs. (In some cases, however, there may be an irreducible subjective element in the likelihoods as well as in the initial probabilities.) One other point worth noting is that, other things equal, "more surprising" evidence (low

q_m in the denominator of (2.1)) has a bigger impact upon the calculation of posterior probabilities.

We now turn to the revision of optimal action. In general, the optimizing problem can be written:

$$\text{Max } u(a;\pi) = \sum_{s=1}^{S} \pi_s v(c_{as}) \tag{2.4}$$

For an informational action a, the consequences c_{as} will ordinarily be probabilistic. Then $v(c_{as})$ is the expected utility of the consequences of the ultimate deterministic outcomes possible in state s if the action taken is a. For example, if the state of the world is "Coin is fair" and the action is "Bet on Heads," the consequence is a probability distribution: 50% chance of winning, 50% chance of losing.

The values of <u>informational</u> actions are essentially based upon the expected utility gains from shifting to better choices among the set of <u>terminal</u> actions. In particular, denote as a_0 the optimal terminal action that would be chosen before receiving any message -- i.e., using the prior probabilities π_s in equation (2.4). If now a particular message m is received, the decision-maker would use (2.4) again, but with revised probabilities $\pi_{s \cdot m}$ possibly leading to a new choice of terminal action a_m. Then Δ_m, the "value of the message m" can be written:

$$\Delta_m = U(a_m, \pi_{s.m}) - U(a_0, \pi_{s.m}) \qquad (2.5)$$

Note that Δ_m, which is necessarily non-negative, is an ex post valuation. It represents the expected gain from revision of best action, estimated in terms of the revised probabilities.

However, the decision to seek information must necessarily be made ex ante. One is never in the position of choosing whether or not to receive the particular message m; the essence of the problem is that the information-seeker does not know in advance which of the set of possible messages m = 1,...,M he will obtain. What the agent can actually purchase is not a particular message but an information service μ -- generating a probability distribution of messages m.

An information service is best thought of as characterized by its matrix of likelihoods $Q = [q_{m.s}]$. In the example with three possible states of the world for a coin (2-headed, 2-tailed, and fair), the information services associated with sample sizes of one and two are represented by the matrices Q_1 and Q_2 below:

(Message) # of Heads

$$Q_1 = \begin{array}{c c} & \begin{array}{cc} 1 & 0 \end{array} \\ \text{(State)} & \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ .5 & .5 \end{bmatrix} \end{array} \qquad Q_2 = \begin{array}{c c} & \begin{array}{ccc} 2 & 1 & 0 \end{array} \\ \text{(State)} & \\ \begin{array}{c} 1 \\ 2 \\ 3 \end{array} & \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ .25 & .5 & .25 \end{bmatrix} \end{array}$$

Note that a purveyor of an information service μ could not "objectively" characterize it by its probabilities q_m of generating the different possible messages, since -- as equations (2.1) to (2.3) indicate -- not only the likelihoods $q_{m.s}$ but also the "subjective" prior probabilities π_s are involved in determining the message probabilities q_m. However, given his prior probability beliefs the individual can then assign "personal" q_m probabilities in order to calculate the value $\Delta(\mu)$ of an information service as an expectation of the values Δ_m of its associated messages:

$$\Delta(\mu) = E(\Delta_m) = \sum_m q_m[u(a_m,\pi_{s.m}) - u(a_o,\pi_{s.m})] \tag{2.6}$$

Since, as already indicated, each Δ_m represented by the square bracket in (2.6) is non-negative, an information service can never lower the agent's expected utility (before allowing for the cost of acquiring that service).

In Figure 8 the determination of $\Delta(\mu)$ is illustrated for a special case in which there are two states of the world, three available terminal actions, and an information service μ with two possible messages. In the diagram utility is measured vertically, while the probabilities of the two states are scaled along the horizontal axes. Each possible assignment of probabilities to states is represented by a point along AB, a 135° line in the base plane whose equation is simply $\pi_1 + \pi_2 = 1$.

The utilities of consequences attaching to the different actions if state 1 occurs, $v(c_{s1})$, are indicated by the points labelled $v(c_{11})$, $v(c_{21})$, and $v(c_{31})$ lying vertically above point A in the diagram. Similarly, the utilities of outcomes in state 1 -- $v(c_{a2})$ for a = 1,2,3 -- lie above point B. The expected utility $v(a;\pi)$ of any action a given any probability vector π is indicated by the vertical distance from the point $\pi = (\pi_1,\pi_2)$ along AB to the line joining $v(c_{a1})$ and $v(c_{a2})$ for that action. In the diagram, if π is the prior probability vector then the best terminal action is a=1 and the associated utility is indicated by the height of F above the base plane.

If the information service μ is costlessly acquired, each of the possible messages m = 1,2 will lead to a revised probability vector

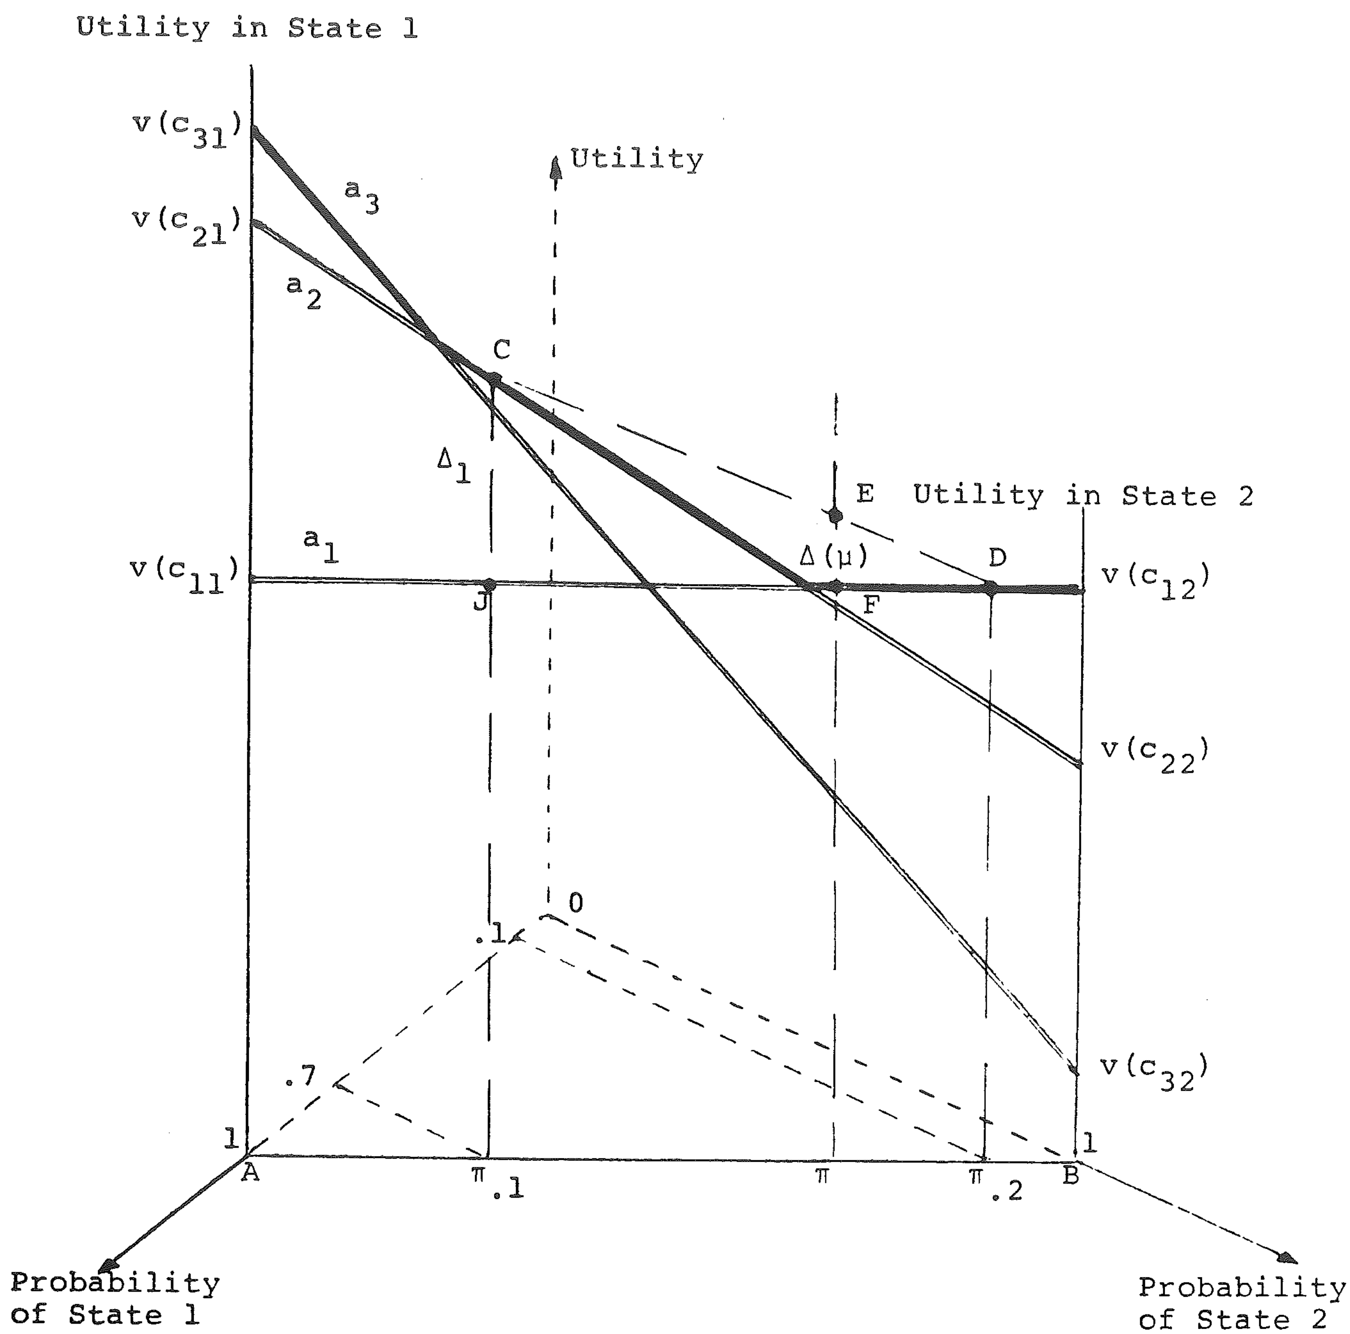

Figure 8: The Value of Information

$\pi_{.m} = (\pi_{1.m}, \pi_{2.m})$. If m=1 in the diagram, the revised choice of action is a=2 (point C) with Δ_1 (ex post utility gain over a=1) equal to the vertical distance CJ. If m=2, in the diagram the best action remains a=1 (point D), so $\Delta_2 = 0$. Then $\Delta(\mu)$, the value of the information service, is represented by the vertical distance EF above the point π along the line AB -- since $\Delta(\mu) = q_1\Delta_1 + q_2\Delta_2$ and:

$$q_1\pi_{.1} + q_2\pi_{.2} = q_1\begin{bmatrix}\pi_{1.1}\\ \pi_{2.1}\end{bmatrix} + q_2\begin{bmatrix}\pi_{1.2}\\ \pi_{2.2}\end{bmatrix} = \begin{bmatrix}\pi_{11} + \pi_{12}\\ \pi_{21} + \pi_{22}\end{bmatrix} = \begin{bmatrix}\pi_1\\ \pi_2\end{bmatrix} = \pi.$$

Fig. 8 also permits us to visualize why higher prior confidence implies lower value of information. Higher confidence -- a tighter prior probability distribution in Fig. 7 -- means that any given message or evidence will have a smaller impact upon the posterior probabilities. Then, in Fig. 8, the posterior probability vectors $\pi_{.1}$ and $\pi_{.2}$ would both lie closer to the original π. It is evident that the effect can only be to shrink the distance EF that represents the value of acquiring evidence.

An information service μ, we have seen, is characterized by its associated likelihood matrix $Q = [q_{m.s}]$, which may be objectively calculable. The matrix Q and the individual personal prior probability vector $\pi = [\pi_s]$ together imply a posterior probability matrix $\Pi = [\pi_{s.m}]$ and message probability vector $q = [q_m]$, so that (Π, q) provides a personal characterization of μ.

One information service $(\hat{\Pi}, \hat{q})$ is said to be "more informative" than another (Π, q), from the point of view of an economic agent, if it yields sometimes higher and never lower expected utility (regardless of the menu of actions). It would of course be very convenient if a measure of "informativeness" could be calculated from the characteristics of the message set, such that a more informative information service would be preferred (cost being equal) independently of the decision-maker's particular situation. There was some initial hope of using Shannon's "entropy" measure of communication theory for this purpose [Shannon 1948]. "Entropy" in the informational sense is the expected number of binary messages needed to communicate the output of an information service. This is indeed a measure of the amount of information, somewhat analogous to ton-miles as a measure of amount of transportation. However, Kenneth J. Arrow [1971] and Marschak [1973] have shown that only in a special logarithmic case does the entropy measure effectively scale the value of information.

Between some information services an informativeness ordering is clearly possible; a random sample of 2, we know, must be more informative than a sample of 1. But in general informativeness can only be partially ordered. The condition for $(\hat{\Pi},\hat{q})$ to be more informative than (Π,q) is that the posterior probability vector, associated with each message under the latter, is a "reduced version" in the sense of a convex combination of the posterior probabilities under the more informative service. That is, if:

$$\pi_{.m} = \sum_m \theta_m \hat{\pi}_{.m} \quad , \text{ where } \theta_m \geq 0 \text{ and } \sum_m \theta_m = 1 \qquad (2.7)$$

This condition is easily visualized in terms of Fig. 8, which pictures a particular information service (Π,q) leading to posterior probability vectors $\pi_{.1}$ and $\pi_{.2}$. Suppose an alternative information service $(\hat{\Pi},\hat{q})$ also had two possible messages 1 and 2, but $\pi_{.1}$ were to lie to the left of $\pi_{.1}$ and $\hat{\Pi}_{.2}$ to the right of $\pi_{.2}$. The alternative service must lead to higher utility so long as there is any change in best conditional action under either message. The fact that the original $\pi_{.1}$ lies <u>between</u> the new $\hat{\pi}_{.1}$ and $\hat{\pi}_{.2}$ means that it is a convex combination of them, and similarly if $\pi_{.2}$ lies between $\hat{\pi}_{.1}$ and $\hat{\pi}_{.2}$. If on the other hand the two posterior probability vectors of one service do not bracket the two posterior vectors of another, which one will be found "more informative" by an individual will depend upon the specifics of his personal situation.

<u>EXAMPLE</u>: For a coin with three possible states 1, 2, and 3 (2-headed, 2-tailed, and fair) we can compare the informativeness of random samples of sizes 1 and 2. Let information service μ_1 represent a single toss of the coin. The message "Heads" (m=H) implies a posterior probability vector $\pi_{.H} = (2/3,\ 0,\ 1/3)$; the message "Tails" (m=T) implies $\pi_{.T} = (0,\ 2/3,\ 1/3)$. Let μ_2 represent the alternative service, involving two tosses of the coin. The message "2 Heads" (m=2) implies $\pi_{.2} = (4/5,\ 0,\ 1/5)$; "1 Heads" (m=1) implies $\pi_{.1} = (0,\ 0,\ 1)$; and "no Heads" (m=0) implies $\pi_{.0} = (0, 4/5, 1/5)$. It may then be verified that $\pi_{.H}$ is a convex linear combination of $\pi_{.2}$, $\pi_{.1}$, and $\pi_{.0}$ with weights $\theta = (5/6,\ 1/6,\ 0)$; for $\pi_{.T}$ the weights are $\theta = (0, 1/6, 5/6)$. So we verify that μ_2 is indeed more informative than μ_1.

When an information service $\hat{\mu}$ is definitely more informative than μ, the fact that the posterior probability vectors $\pi_{.m}$ are convex linear combinations of the $\hat{\pi}_{.m}$ can be interpreted as saying that the less informative service provides outputs that are "garbled" versions of the other's messages. In particular, every other information service is a garbling of the <u>perfect</u> information service (μ^*) whose messages are conclusive as to which state will obtain. Thus, in the 2-state 2-message case, μ^* has posterior probability vectors $\pi^*_{.1} = (1,0)$ and $\pi^*_{.2} = (0,1)$; any other μ must have posterior vectors $\pi_{.1} = (\pi_{1.1}, \pi_{2.1})$ and $\pi_{.2} = (\pi_{1.2}, \pi_{2.2})$ that are convex combinations of

$\pi^*_{.1}$ and $\pi^*_{.2}$. An important special case of garbling is where distinctions are obliterated: two or more distinct messages of a more informative service are reduced to a single message of a less informative service [Radner 1968, Hart 1975]. This may lead to a situation where different agents have different partitionings of the states of the world, with consequent difficulty for the negotiation of contingent contracts.

Fig. 8 also indicates an important "non-concavity" (condition of <u>increasing</u> marginal returns) in the valuation of information services [Radner and Stiglitz 1975]. Starting with the null information service with posterior probabilities equal to the priors π, suppose a slightly informative $\hat{\mu}$ comes along with posterior probabilities $\hat{\pi}_{.1}$ just barely to the left of π and $\pi_{.2}$ barely to the right. If the probability changes are small, neither message changes the associated best action, and hence there can be no utility gain. So the <u>marginal</u> return of improved information will be zero over a certain range, before becoming positive at the point where the improvement begins to affect action taken after at least one of the possible messages.

Thus far we have interpreted messages essentially as sample evidence, for which likelihood functions that follow from the laws of probability can be calculated for use in Bayes Theorem. More generally, however, a message can take a form for which the likelihood function is not so easily assessed. One important example is "expert opinion." In general, it would seem that a decision-maker should take account of the opinions (if available) of all other parties who have some information not accessible to himself [Shavell 1976]. But such "expert" opinion would generally be a posterior probability vector that would in part depend upon objective evidence, but also in part upon the expert's subjective prior beliefs -- which the decision-maker might reject. Furthermore, even with regard to the objective evidence component, the decision-maker would not generally be in a position to know how much of it is independent of facts already available to himself and accounted for. So it seems inevitable that the decision-maker's likelihood function for the expert's opinion would be a "subjective" judgmental matter. More specifically, each distinct probability distribution that the expert might offer would represent a "message" to which the decision-maker would attribute a likelihood function as in Fig. 7 -- for use in a Bayesian calculation leading to a posterior probability distribution of his own.

A number of other complexities arise when the decision-making agent is not an individual but a group of people who must jointly arrive at some common choice (Raiffa 1970 , Ch. 8, Part 2). Of course this requires "constitutional" agreement as to the procedure for collective choice (e.g., unanimity). But whatever the rule followed, problems may stem from possible inter-individual conflicts of interest (differences in utilities attached to consequences) or conflicts of opinion (differences in probability beliefs as to states of the world).

Consider first differences of belief only. Since divergences of opinion commonly generate disagreement as to best terminal action, there is a tendency for the group to choose "excessive" informational action (from the vantage point of an outside impartial observer). In effect, each member of the group believes that the incoming evidence will likely support his own beliefs, and thus bring other members around to his own recommendation as to best terminal action. On the other hand, sometimes there may be agreement on terminal action among members that masks compensating disagreements; here members of the group may agree not to seek information where an impartial observer would advise them to do so.

If there are conflicts of interest as well as lack of consensus, the redistributive effect of information becomes important. The group tends to agree to acquire information, even at a collective loss, once each member has staked out a position whereby he expects to benefit thereby at the expense of the others -- as by a wager. Note that even if there were no initial conflicts of interest, wagering would convert mere differences of opinion into conflicts of interest.

2.1.2 Other informational activities

So far, under the heading of informational decision-making we have only considered the acquisition of evidence -- as by generation of sample data (the production of socially "new" information) or the receipt of expert advice (the interpersonal transfer of "old" information). But other types of informational activities can also be very important. The possibility of acquiring information from others, as discussed in connection with "expert opinion" above, immediately suggests the reverse activity -- the dissemination of information to other economic agents. This might be done for a price, as when one is hired as an expert, but (as we shall see below) sometimes it may pay to disseminate gratuitously, or even to incur cost to "push" information to others [Hirshleifer 1973]. Advertising is an obvious example. There is also a choice as to disseminating publicly ("publishing"), or else privately to a select audience. As a question of authenticity might arise in all such cases, the receiver of information may devote effort to the process of evaluation, possibly assisted by authentication activities (or hampered by deception activities) on the part of the disseminator. There is also the possibility of unintended dissemination, achieved by espionage or monitoring on the part of information-seekers -- possibly leading to countermeasures in the form of security (secrecy-maintaining) activities by the possessors of information.

Finally, there are classes of activities, apart from the direct acquisition or dissemination of information, that are explicable only as indirect consequences of informational patterns. Wagering, for example, follows from differences of opinion in a situation where conclusive information is anticipated (so as to determine who wins and who loses). Wagering is a special case of the more general activity of speculation, which turns on revision of prices anticipated in consequence of the arrival of information. Another form of activity, which will be considered in the next section below, involves the adoption of more or less "flexible" positions in anticipation of ability to make use of future information when it arrives.

2.2 Emergent Information

In the section preceding we thought of information as being newly generated by an informational action like a sampling experiment or, alternatively, as acquired from others via a transaction like the purchase of expert opinion. But in some cases information may autonomously emerge simply with the passage of time, without requiring any direct action by recipients. Tomorrow's weather is uncertain today, but the uncertainty will be reduced as more meteorological data flows in and will in due course be conclusively resolved when tomorrow arrives. Direct informational actions might still be useful, by providing knowledge earlier then it would autonomously arrive. But under conditions of emergent information a kind of "indirect" informational action becomes available -- simply waiting before taking terminal action.

2.2.1 The value of flexibility

Suppose a choice must be made now between immediate terminal action and awaiting emergent information. This choice can only be interesting, of course, where there is a trade-off between two costs: (1) a cost of waiting, versus (2) an "irreversible" element in the possible loss suffered from mistaken early commitment. Exactly these elements have been involved in analyzing the benefit of actions that irreversibly transform the environment [Arrow and Fisher 1974, Henry 1974] and in discussions of the value of "liquidity" [J. Marschak 1949, Hirshleifer 1972] or of "flexibility" [T. Marschak and Nelson 1967, Jones and Ostroy 1978].

The essential idea is pictured in Fig. 9 (which has the same structural framework as Fig.8 but suppresses the 3-dimensional background). The

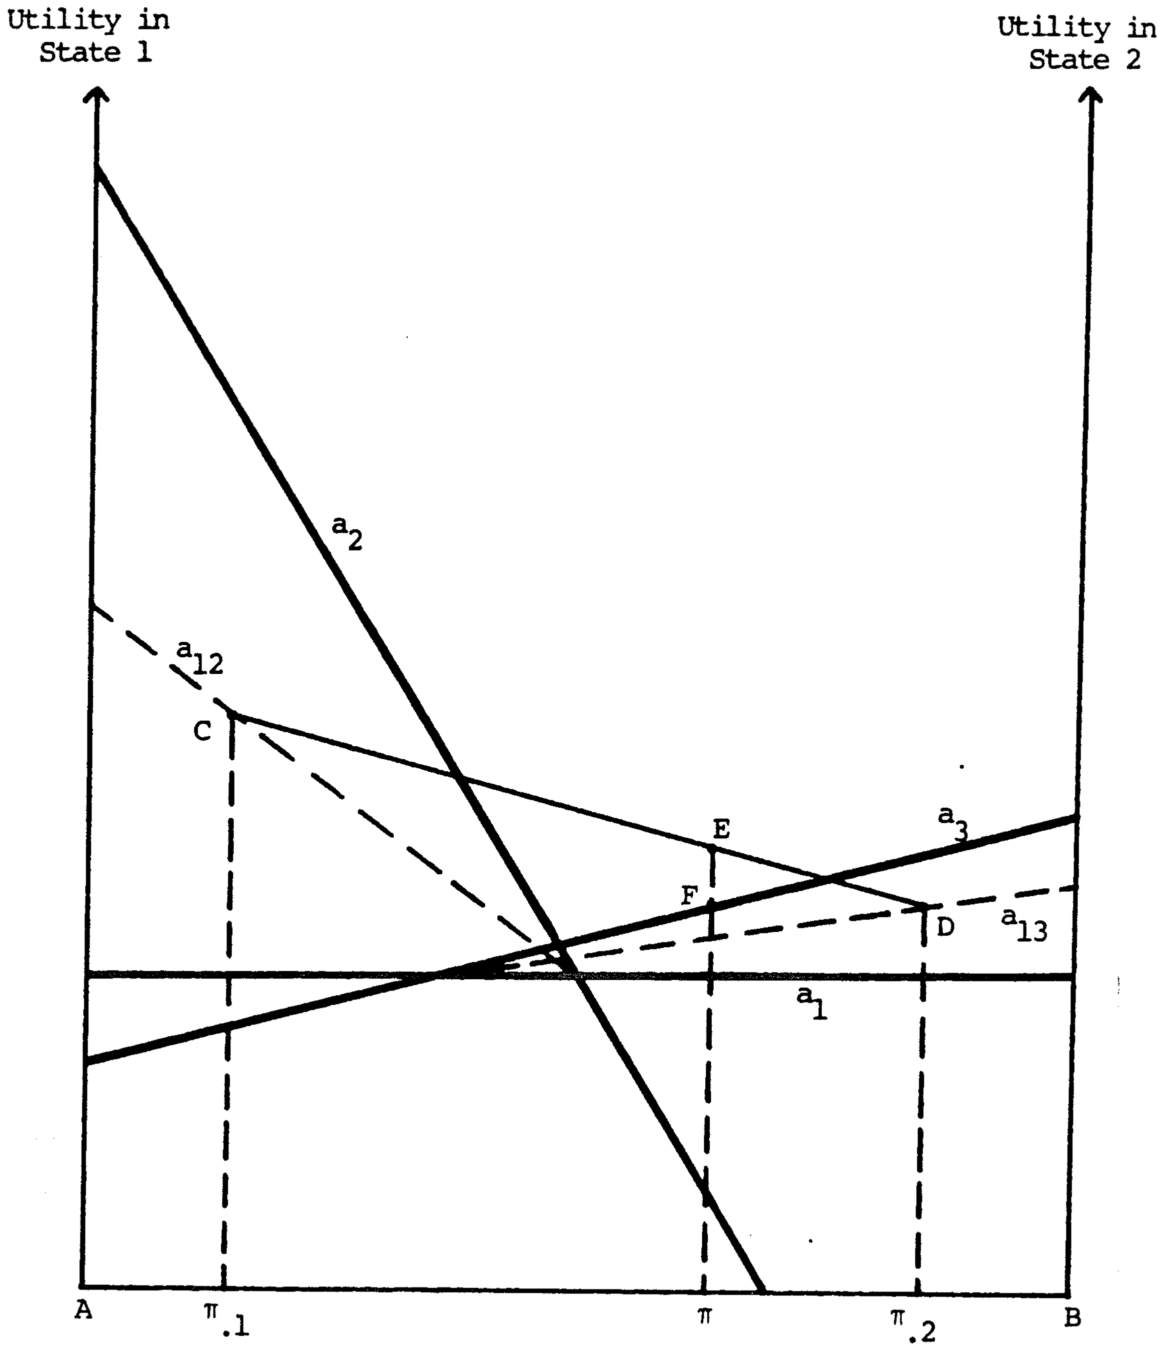

Fig. 9: The Value of Flexibility

individual, if he decides upon immediate terminal action, has a choice among a_1, a_2, or a_3. As shown here, he would choose either a_2 or a_3 depending upon his beliefs π (in the diagram, he prefers a_3 yielding utility F). As the diagram is drawn, he would never choose a_1 as a terminal action. But suppose that a_1 has a "flexibility" property. To wit, <u>after</u> receiving emergent information the individual can shift from a_1 to a_2, achieving the intermediate overall utility indicated by the line a_{12} or, should the information point the other way, he can shift from a_1 to a_3 with overall utility payoff indicated by line a_{13}. In the situation as drawn, had he initially chosen the "flexible" action a_1, then if message 1 is received (leading to the posterior probability vector $\pi_{.1}$) the individual would shift to a_2, thus attaining overall utility indicated by point C on line a_{12}. Similarly, message 2 would allow him to attain point D on line a_{13}. His expected utility is then E, superior by the amount EF to the result of the best immediate terminal action a_3.

The element of "irreversibility" appears here in the fact that line a_{12} lies below a_2 in the range where both of these are preferred to a_1, and similarly a_{13} lies below a_3 in the corresponding range. One has to pay a price to retain flexibility; the price is that you cannot do as well as if you had made the best choice among "irreversible" actions in the first place. Indeed, if the incoming information had somewhat lesser weight, so that the posterior belief vectors $\pi_{.1}$ and $\pi_{.2}$ were not so different from the original π, the price of flexibility can become too great; point E would then be somewhat lower in the diagram, and might well fall below the point F representing the utility of the best immediate terminal action (a_3).

2.2.2 Complete versus incomplete regimes of markets

Public emergent information will affect prices. It therefore divides market exchanges into two rounds: trading prior to, and trading posterior to the receipt of the information. The equilibria of these two rounds will generally be inter-related, but the relationship depends upon the completeness of markets in each round (see Sec. 1.2.3 above).

In Part 1 of this paper we mainly considered models with a single consumptive good c and S states of the world. Without anticipated emergence of information, a regime of Complete Contingent Markets (CCM) would exist if all the c_s claims (S in number) were separately tradable at prices P_s. (Trading in any set of S linearly independent combinations of the underlying c_s claims would also constitute a CCM regime, but we will generally ignore this complication.) Equilibrium would be characterized by the following optimality condition for each individual:

$$\frac{\pi_s \; v'(c_s)}{\pi_1 \; v'(c_1)} = \frac{P_s}{P_1} \tag{2.8}$$

If conclusive information as to which state of the world were anticipated, it would make no difference for this model. Equation (2.8) would have to hold in prior-trading equilibrium and there simply can be no posterior trading. Once it were known for certain that some state s* would obtain, c_{s*} would be the only claim retaining any validity; no-one has anything to trade against his holding of c_{s*} claims.

A more interesting model, continuing to assume that traders anticipate conclusive information, allows for multiple consumptive goods g=1,...,G. Here a regime of Complete Contingent Markets (CCM) in the prior round would allow trading in the GS different claims c_{gs} -- claims to any good g under state s -- at price P_{gs}. After state s* obtains, posterior trading becomes possible among the G remaining valid claims c_{gs*}. But suppose for the moment that individuals ignore the possibility of posterior trading in their prior-round dealings. Then the optimality conditions include ratios of the following form, where g' and g" are any two goods:

$$\frac{\pi s \dfrac{\partial v}{\partial c_{g's}}}{\pi s \dfrac{\partial v}{\partial c_{g''s}}} = \frac{P_{g's}}{P_{g''s}} \tag{2.9}$$

After the conclusive information arrives, and state s* is known to obtain, in the only condition of (2.9) that remains relevant π_{s*} now equals unity. But this makes no difference; in fact, π_s cancels out whatever its value. Therefore the price ratio on the RHS of (2.9) continues to sustain the solution attained in the prior round.

Thus we see that even though posterior trading is <u>possible</u> (there remain G tradable claims c_{gs*}), with CCM in the prior round no-one will find such trading <u>advantageous</u>. In effect, markets for GS prior claims plus G posterior claims are "more than are needed." Qualification: If prior-round traders failed to correctly forecast that the conditional price ratio on the RHS of (2.9) would remain unchanged in the posterior round, they would be led to make "erroneous" prior-round transactions, in turn requiring "corrective" posterior-round transactions.

So Complete Contingent Markets in the prior round suffice for Pareto-efficiency only subject to a proviso of "correct conditional price forecasting."

Consider now another special case. We return to the assumption of a single consumptive good, but we allow the anticipated information (taking the form of a message $m=1,\ldots,M$) to be <u>less than conclusive</u>. Here Complete Contingent Markets (CCM) would provide for trading in any of the SM claims c_{sm} -- entitlements valid if and only if message m and state s both occur -- at prices P_{sm}. As for the posterior round, some message m* will have been received so there can be just S tradable claims. The prior-round optimality conditions take the form:

$$\frac{\pi_{sm}v'(c_{sm})}{\pi_{1m}v'(c_{1m})} = \frac{P_{sm}}{P_{1m}} \qquad (2.10)$$

Here π_{sm} represents the individual's prior beliefs as to the <u>joint</u> probability of state s and message m. But this joint probability is the product of the <u>conditional</u> probability $\pi_{s.m}$ and the prior probability q_m of the message m:

$$\pi_{sm} = \pi_{s.m}q_m$$

Since q_m cancels out, and the conditional probability $\pi_{s.m}$ is nothing but the <u>posterior</u> probability attaching to state s after receiving message m, the price ratio on the RHS of (2.10) again continues to sustain the prior-round solution. So once more, under this CCM regime posterior trading is <u>available but not necessary</u> for Pareto-efficiency: markets in SM prior claims plus S posterior claims are more than are needed. (The same qualification, that the posterior price ratios must be correctly forecast, remains applicable.) Generalizing: if there are G goods, S states, and M messages, a prior-round CCM regime of trading in all the GSM distinct claims allows every agent to attain his optimum at one fell swoop -- provided that posterior-round prices are correctly forecast [Hirshleifer and Rubinstein 1976, Salant 1976, Feiger 1976].

Interesting questions arise, however, when we analyze market regimes that are incomplete (as, of course, they must actually be in the world). There are many different possible patterns of incompleteness. We will consider here only the particular case of emergent conclusive information; then the set of M messages collapses into the set of S states so that individuals are only concerned with c_{gs} claims, GS in number. Equation (2.9) represented the first-order optimality conditions holding under CCM for this case. Two types of incomplete-market regimes, Numeraire Contingent Markets (NCM) and Futures Markets (FM) will now be discussed.

Arrow [1953] has shown that the same allocation as indicated by conditions (2.9) under CCM (trading in GS claims) is achievable with prior contingent trading in only a single commodity. Let us think of this commodity as the numeraire good, g=1. Then under NCM only S claims of form c_{1s} would be tradable in the prior round. We can think of these as side-bets in numeraire units as to which state of the world is going to obtain. These bets determine the individual's posterior wealth under each possible state of the world, which he can then use to purchase a preferred consumption basket in the posterior round.

Under either regime of markets, the individuals is seeking to maximize expected utility $U = \sum_s \pi_s v(c_{1s},\ldots,c_{Gs})$. Under the CCM regime, the constraint is:

$$\sum_s\sum_g P_{gs} c_{gs} = \sum_s\sum_g P_{gs} \omega_{gs} \qquad (2.11)$$

where ω_{gs} indicates an endowment quantity. For the NCM regime, prices will be symbolized as $\emptyset_{1s}$ for the prior markets and $\emptyset_{g.s}$ for the posterior market (after it is known that state s obtains). In the prior round, the individual will transact at prices $\emptyset_{1s}$ to arrive at an intermediate "trading position" with elements denoted c^t_{gs}. All these elements are, apart from those involving $g = 1$, necessarily identical with the corresponding endowment elemts ω_{gs}. The elements c^t_{1s} that are the decision variables in the prior round may take on either negative or positive values, though of course none of the ultimate consumption elements c_{gs} can be negative.

The prior and the S contingent posterior constraints are:

$$\sum_s \emptyset_{1s} c^t_{1s} = \sum_s \emptyset_{1s} \omega_{1s} \tag{2.12}$$

$$\sum_{g=1}^{G} \emptyset_{g.s} c_{gs} = \sum_{g=2}^{G} \emptyset_{g.s} \omega_{gs} + \emptyset_{1.s} c^t_{1s} \tag{2.13}$$

$$(s = 1, \ldots, S)$$

It can be shown by direct substitution that (2.12) and (2.13) together are equivalent to (2.11) if:

$$\emptyset_{1s} = P_{1s} \text{ and } \frac{\emptyset_{g.s}}{\emptyset_{1.s}} = \frac{P_{gs}}{P_{1s}} \tag{2.14}$$

Thus, NCM markets will achieve the same efficiency conditions as under CCM, provided once again that individuals correctly forecast the posterior price ratios $\emptyset_{g.s}/\emptyset_{1.s}$, in particular, that each such ratio will equal the P_{gs}/P_{1s} ratio under CCM. Note however that this proviso as to "correct conditional price forecasting" becomes more stringent in comparison with CCM, where the correct forecast was simply "no change" from the prior to posterior ratio (for the state that has actually obtained). Here under NCM the correct forecast is not in general computable from data available to traders in the prior round [Radner 1968]. Conditions under which it might become, at least approximately, computable have been discussed by Hirshleifer and Rubenstein [1976], Rubinstein [1975], Feiger [1976], and Hirshleifer [1976].

A numerical example is provided in Table 2. Agents J and K are a "representative pair" of traders who diverge both in beliefs and in endowments, but not in their preferences. Under CCM each can use the equilibrium price matrix P_{gs} to trade in a single round to his preferred consumptive position.

Table 2

COMPLETE CONTINGENT MARKETS (CCM) VERSUS NUMERAIRE CONTINGENT MARKETS (NCM)

Equilibrium for Representative Trader-Pair J, K

DATA		J	K
Preference-scaling function v		$\ln c_1 + \ln c_2$	$\ln c_1 + \ln c_2$
Beliefs:	(π_1, π_2)	(.7,.3)	(.5,.5)
Endowment:	$\begin{bmatrix} \omega_{11} & \omega_{12} \\ \omega_{21} & \omega_{22} \end{bmatrix}$	$\begin{bmatrix} 200 & 200 \\ 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 \\ 400 & 160 \end{bmatrix}$
COMPLETE CONTINGENT MARKETS (CCM)			
Equilibrium prices:	$P_{gs} = \begin{bmatrix} .6 & .4 \\ .3 & .5 \end{bmatrix}$		
Consumption:	$\begin{bmatrix} c_{11} & c_{12} \\ c_{22} & c_{22} \end{bmatrix}$	$\begin{bmatrix} 116\frac{2}{3} & 75 \\ 233\frac{1}{3} & 60 \end{bmatrix}$	$\begin{bmatrix} 83\frac{1}{3} & 125 \\ 166\frac{2}{3} & 100 \end{bmatrix}$
NUMERAIRE CONTINGENT MARKETS (NCM)			
Equilibrium prices:			
Prior:	$\phi_{1s} = (.6, .4)$		
Posterior:	$\phi_{g.s} = \begin{bmatrix} 1 & 1 \\ .5 & 1.25 \end{bmatrix}$		
Trading position:	$\begin{bmatrix} c^t_{11} & c^t_{12} \\ c^t_{21} & c^t_{22} \end{bmatrix}$	$\begin{bmatrix} 233\frac{1}{3} & 150 \\ 0 & 0 \end{bmatrix}$	$\begin{bmatrix} -33\frac{1}{3} & 50 \\ 400 & 160 \end{bmatrix}$
Consumption	$\begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{bmatrix}$	$\begin{bmatrix} 116\frac{2}{3} & 75 \\ 233\frac{1}{3} & 60 \end{bmatrix}$	$\begin{bmatrix} 83\frac{1}{3} & 125 \\ 166\frac{2}{3} & 100 \end{bmatrix}$

Note that J, attaching higher probability belief than K to state-1 (.7 versus .5), has arranged for relatively higher contingent consumptions than K if state-1 occurs, ([116 2/3,75] versus [83 1/3,125]). Under NCM, on the other hand, the less complete market regime requires trading in both rounds to attain the same ultimate solution. The prior-round price vector $\emptyset_{1s}$ and posterior-round price matrix $\emptyset_{g \cdot s}$ are related to P_{gs} as in equation (2.14). Here agent J, who attaches relatively higher belief to state-1, expresses that belief by a trading position involving a heavier commitment in state-1 claims to the numeraire good, and the reverse for agent K. (In fact, the c_{11}^t element of K's trading position is actually negative.)

The CCM and NCM regimes allow trading in state-contingent claims. Such trading does take place to some extent in the actual world, directly as in some insurance transactions or indirectly via trading in assets like corporate shares that can be regarded as packages of state-claims. But, most of the time, the trading observed in the world represents exchange of unconditional claims to goods. In a market regime allowing only the exchange of unconditional claims to G consumptive goods, under conclusive emergent information the prior and posterior rounds can be respectively identified with "futures" and later "spot" markets. The question to be considered is whether such a regime of unconditional or, as we shall say, Futures Markets (FM) can attain the same final outcome as the more ample market regimes previously considered.

For the model of conclusive emergent information, under the Complete Contingent Markets CCM regime the GS tradable claims in the prior round made it unnecessary to re-trade at all among the G claims in the posterior round.

Under NCM the S tradable claims in the prior round just sufficed, with re-trading among the G claims in the posterior round, to reproduce the same result. (The proviso as to "correct conditional price forecasting" being assumed to hold in each case.) It will be evident that the G+G tradable claims in two rounds under FM cannot do as well as the S+G under NCM, if G < S. However, this negative conclusion tends to be overcome once we allow for the fact that emergent information in the world is only rarely conclusive. If improved but not yet conclusive information emerges repeatedly, the multiple opportunities for trading among the G goods that are re-created after each informational input increase the effectiveness of FM relative to NCM and CCM. (Again, the proviso as to correct conditional price forecasting is required under FM as well.)

We shall not provide here a formal statement of the optimality or equilibrium conditions for FM (see Hirshleifer [1977], Townsend [1978]). The crucial point is that, in the prior round, the difference between the elements c^t_{gs} of the trading position for any good g and the corresponding elements ω_{gs} of the endowment position is a constant g^t:

$$c^t_{gs} - \omega_{gs} = g^t \quad , s=1,\ldots,S \qquad (2.15)$$

Here g^t is simply the unconditional purchase (or sale, if negative) of good g in the prior round. Table 3 illustrates that, *if* G=S (=2 in this example), then with conclusive emergent information the same final result for the representative trader-pair J,K can be attained under FM as shown in Table 2 for CCM and NCM. Table 3 also illustrates an alternative

Table 3

FUTURES MARKETS (FM)

Equilibria for Alternative Representative Trader-Pairs

J, K Versus L, M

($v = \ln c_1 + \ln c_2$ for all traders)

DATA		J	K	L	M
Beliefs:	(π_1, π_2)	(.7,.3)	(.5,.5)	(.6,.4)	(.6,.4)
Endowment:	$\begin{bmatrix} c^e_{11} & c^e_{12} \\ c^e_{21} & c^e_{22} \end{bmatrix}$	$\begin{bmatrix} 200 & 200 \\ 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 \\ 400 & 160 \end{bmatrix}$	$\begin{bmatrix} 200 & 200 \\ 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 \\ 400 & 160 \end{bmatrix}$
FUTURES MARKETS (FM)					
Equilibrium prices:					
Prior:	$\phi_g = (1, .8)$				
Posterior:	$\phi_{g.s} = \begin{bmatrix} 1 & 1 \\ .5 & 1.25 \end{bmatrix}$				
Trading position:	$\begin{bmatrix} c^t_{11} & c^t_{12} \\ c^t_{21} & c^t_{12} \end{bmatrix}$	$\begin{bmatrix} 288\frac{8}{9} & 288\frac{8}{9} \\ -111\frac{1}{9} & -111\frac{1}{9} \end{bmatrix}$	$\begin{bmatrix} -88\frac{8}{9} & -88\frac{8}{9} \\ 511\frac{1}{9} & 271\frac{1}{9} \end{bmatrix}$	$\begin{bmatrix} 200 & 200 \\ 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & 0 \\ 400 & 160 \end{bmatrix}$
Consumption:	$\begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{bmatrix}$	$\begin{bmatrix} 116\frac{2}{3} & 75 \\ 233\frac{1}{3} & 60 \end{bmatrix}$	$\begin{bmatrix} 83\frac{1}{3} & 125 \\ 166\frac{2}{3} & 100 \end{bmatrix}$	$\begin{bmatrix} 100 & 100 \\ 200 & 80 \end{bmatrix}$	$\begin{bmatrix} 100 & 100 \\ 200 & 80 \end{bmatrix}$

representative trader-pair L,M -- differing from J,K only in that L,M share identical beliefs. The L,M world generates the same market equilibrium prices as the J,K world though the individual outcomes are different. Note that L and M choose not to trade in the prior round, a point that will take on significance in the discussion of speculation that follows.

2.2.3 Speculation

The term "speculation" has caused a good deal of confusion. Some authors loosely apply the word to arbitrage between markets, or to storage of goods over time or carriage over space -- activities which do not involve uncertainty in any essential way. For our purposes, speculation is purchase with the intention of re-sale, or sale with the intent of re-purchase, where the uncertainty of the future spot price is the source of both risk and gain. The probabilistic variability of price is in turn due to anticipated emergence of information. Each possible message (in the conclusive-information case that we shall be assuming here, this is equivalent to the advent of a single possible state) leads to an associated equilibrium posterior price vector, benefiting agents who adopted trading positions generating relatively high conditional wealths for that state.

From the discussion in the preceding section showing that re-trading possibilities are not needed in a regime of Complete Conditional Markets (CCM), we see that speculation is a response to incompleteness of prior-round markets /Feiger 1976/. (However, if the proviso as to correct conditional price forecasting fails to hold, "corrective" re-trading would be needed even under CCM.) Also involved as determinants of speculative activity are initial risk-exposure in the form of unbalanced state-endowments, degree of risk-aversion, and probability beliefs as to future states of the world.

Among these determinants Keynes [1930] and Hicks [1946], followed by many others, have emphasized differential risk-aversion. In their view, in the prior round of an FM regime relatively risk-tolerant speculators accept the risks of price change from relatively risk-intolerant "hedgers." Thus, the former buy commodity "futures," achieving a small average gain (excess of later mean spot price over futures price) that represents the price paid by hedgers for the bearing of the price risk. Later developments along this line [Houthakker 1957, 1968, Cootner 1968] have brought out that hedgers can be on either side of the futures market; speculators need bear only the imbalance between "long" and "short" hedgers' commitments, and risk-compensating average price movement could go either way. In contrast with these views Holbrook Working [1953, 1962] has denied that there is any systematic difference as to risk-tolerance between those conventionally called speculators and hedgers. Working emphasizes, instead, differences of belief (optimism or pessimism) as motivating futures trading.

The Keynes-Hicks concentration upon aversion to "price risk" is seriously misleading. Individuals' prior-round trading decisions are affected by "quantity risk" (unbalanced endowments over states) as well as by price risk [McKinnon, 1967]. Indeed, from the social point of view price uncertainty is the (inverse) reflection of an underlying uncertainty as to the future aggregate commodity totals. The price risks and quantity risks faced by individuals can be offsetting, so that no prior-round hedging activity would be engaged in. This is illustrated in the equilibrium of traders L and M in Table 3 The two have the opportunity of mutually hedging in the prior round, and of doing so "without cost". That is, in terms of their common beliefs the prior-round price of the risky good $g = 2$, in terms of the other as numeraire, is

$\phi_2 = .8$ -- the mathematical expectation of the conditional posterior prices $\phi_{2.1} = .5$ (with probability .6) and $\phi_{2.2} = 1.25$ (with probability .4) -- so that neither need pay anything, on average, for divesting himself of price risk via prior-round trading. Nevertheless, in the situation of Table 3 any such trading would make it impossible for the individuals to attain a Pareto-optimal position.

Strong results that are fully general are hard to come by [Salant 1976, Feiger 1976, Hirshleifer 1976, 1977]. But under reasonable simplifying assumptions Working's contention tends to be borne out: differences of belief, and not differences in risk-tolerance, are the sine qua non of speculative activity. On the other hand, given any difference in belief the degree of risk-aversion affects the extent of traders' speculative commitments. And, our above discussion suggests, unbalanced state-endowments may also play a role under incomplete market regimes. Even without differences of belief or of risk-aversion, the restricted trading opportunities under a FM regime may require both prior and posterior trading to attain a Pareto-efficient allocation.

2.3 The Economics of Research and Invention

Research and invention activities are prime instances of the informational actions studied, on the level of the individual economic agent, in Sec. 2.1 above. We are not dealing here with situations like those examined in Sec. 2.2, where information simply emerges with the passage of time. Since Nature will not autonomously reveal her secret, it must be sought out by costly direct search for the still-unknown "message". The topic of concern here however is not the informational decisions of an isolated Robinson Crusoe, but rather of individuals in a market environment facing rivalrous competition from some agents, but also having opportunities for mutually advantageous exchanges with others.

The central problem considered by modern analysts [for example Machlup 1968, Arrow 1972, and Cheung] has been the conflict between the social goals of achieving efficient use of information once produced versus providing ideal motivation for production of information. With regard to optimal use, already-produced information is a "public good" in the sense that its use by any member of society does not reduce the amount that could be made available to others. Then any barrier to use, as may stem from legal enforcement of patents or copyrights or property in trade secrets, is inefficient. On the other hand, as in the standard public-good situation, there will be inadequate motivation to invest in generation of information if the product cannot be reduced to legally protected property.

Under ideal conditions the efficient-use problem could be solved by charging perfectly-discriminating fees to license (non-exclusively) all uses of a given idea. If the discoverer were granted full property rights in the idea, as by a perpetual copyright or patent, he in turn would have the optimal incentive to produce (search for) ideas. But in practice owners of copyrights or patents cannot impose perfectly-discriminating royalty fee structures on licensees.

A patentee might instead maximize returns by granting exclusive licenses (in which case the social value of the excluded uses is of course lost) or by imposing fee structures that distort the marginal production decisions of licensees. On the other side of the picture, because of the elusiveness of property in ideas there is a good deal of uncertainty and unreliability in the legal protection of patents and copyrights, and of course relatively little protection for trade secrets not made subject to patent or copyright. The result is that a good deal of unlicensed uses escape control. Short of ideal conditions there will be losses from both underutilization and underinvestment, and in practice something of a trade-off: provision of greater legal protection to inventors tends to ameliorate the underinvestment problem, but to worsen the underutilization problem.

More recent investigations have indicated, however, that not all the important elements of the picture have been captured by this analysis. These newer results turn upon the possibility of overinvestment in the production of ideas ("a rush to invent").

The first such factor is the fugitive resource (or common-property resource) nature of undiscovered ideas [Barzel 1968]. For concreteness, we can use as metaphor the "over-fishing" model of Gordon [1954]. Suppose there are perfect property rights in fish caught, but complete free entry into fishing (i.e., there are no property rights that exclude others from engaging in fishing as an activity). Then in competitive equilibrium there will be over-fishing; private marginal cost will equal price, but the true social marginal cost in fishing exceeds the private marginal cost. The reason is that a certain fraction of each fisherman's catch consists of fish that would have been caught anyway, by other fishermen -- so the true social product of fishing effort is less than appears in private calculations. The upshot is that too many fish are caught, too soon. Among the remedies discussed in the fugitive-resource literature are the imposition of taxes or production quotas to reduce the amount of fishing activity, or alternatively the assignment of exclusive property rights to engage in such activity.

This last point may be clarified by explicitly distinguishing "rights in fish" (the right to exclude others from a fish you have caught) from "rights to fish" (the right to exclude others from competing with you in fishing activity). Assuming fully protected rights in fish, the fugitive-resource problem can be solved by also vesting rights to fish -- for example by auctioning the right to exclusively exploit a fishery. Indeed, once rights to fish are defined such an auction would tend to occur of its own, via Coase-Theorem negotiations. In the context of research, the equivalent distinction is between rights in an idea and rights to (engage in search for) an idea. Again, one might imagine auctioning off the right to search for an idea -- for example, the right to invent an alloy with specified properties. As the lowest-cost inventor would bid highest for this right, the "rush to invent" problem would be solved [Cheung].

In research, the difficulty of defining the nature of an "uncaught" idea seems t make the assignment of rights to (search for) them unfeasible. (Even in fishing, it is often impractical to define exclusive rights to such a wandering resource.) In contrast with fishing, however, for inventors property rights in ideas when caught are still very far from perfect. But, in the circumstances, this is not necessarily bad; being somewhat like a tax on inventive activity, defective rights in ideas reduce what otherwise might be an excessive "rush to invent."

Recapitulating at this point, we have seen two distinct possible justifications for limiting property rights in ideas -- for example, by granting patents only for a term of years. The first is that some protection to inventors is traded off against protection to users of invention; e.g., after the patent expires use of the idea is unconstrained. The second is that the "rush to invent" tendency is moderated by reducing the capturable value of the invention itself.

There is still another motivation that may lead to excessive devotion of resources to invention. Ideas of course vary enormously in their significance, and some among them will have far-reaching consequences. This opens up a new channel of reward for inventors. Instead of, or possibly in addition to selling the information via patent license or otherwise, an inventor might be able to speculate by taking long or short positions in assets whose values will be affected by the invention [Hirshleifer 1971]. The cotton gin, for example, had speculative implications for the prices of cotton and cotton land, the business prospects of firms engaged in cotton warehousing and shipping, the site values of key points in the cotton transportation network -- in addition to more indirect implications for competitor industries like wool and complementary ones like textile machinery. Nor does an idea have to be as momentous as the cotton gin to have profitable speculative implications. An oil firm that has developed a new method of deep recovery might, for example, reap a speculative payoff by buying up options on tracts whose petroleum now lies too deep to be recovered. One important implication of the speculative reward for invention is that it motivates the possessor of information to disseminate it widely and even gratuitously -- after having made his speculative commitment.

Looking at this more generally, individual ideas will -- unlike individual fish, or whole boatloads of fish -- often have important pecuniary externalities. The "ideal conditions" referred to above, that would have reserved for the inventor the entirety of the technological benefit flowing from his idea, would then generally lead to overcompensation if the inventor can also capture some fraction of the pecuniary externalities as well.

There are classes of research activity for which the reward element stemming from the technological benefit is negligible, where the potential pecuniary return is almost the whole picture. Stockmarket research ("security analysis"), whether engaged in by full-time professionals or by ordinary investors, is essentially of this nature [Fama and Laffer 1971]. There may be a technological benefit (increase in society's productive opportunities) due to accurate security analysis, to the extent that the analysis is generally believed and therefore affects market prices. For one thing, the earlier movement of stock prices more quickly rewards effective corporate decision-making and punishes inefficient decision-making. Also, those firms whose better prospects are now more accurately known to the public will find it easier on that account to acquire new capital to exploit those prospects. But whether the technological benefit of security-analysis information is substantial or minor, the reward reaped by security analysts stems almost entirely from the pecuniary revaluations -- the correctly interpreted rises or falls in the stock prices themselves.

Recognition of the "rush to invent" problem, while undermining the traditional argument for patents (or other forms of protection for discoverers of ideas) that was based upon a presumption of underinvestment in research, does not warrant going to the other extreme. It would not be in order to conclude that patent protection is <u>not</u> justified, but only that the arguments pro and con are more complex than had previously been realized.

2.4 Informational Advantage and Market Revelation of Information

In the preceding sections we have considered situations in which an individual could profit by timely publication of knowledge in his private possession. For example, an agent obtaining new information might transact at existing prices (take a speculative position), planning to sell out at the revised prices that ensue once he publicly discloses his knowledge. But there are two problems here. First, the agent must be able to move to a trading position (make a speculative commitment) without thereby revealing his secret -- this is the problem of *information leakage*. Second, at the disclosure stage he must be able to authenticate the information he is trying to publicize -- this is the problem of *signalling*. We shall consider these two problems, in reverse order, in the sections following.

2.4.1 Signalling

The particular signalling problem that has aroused greatest interest arises when sellers of a higher-quality product or service are attempting to convey that information (that their product *is* higher-quality) to buyers. Of course, any seller is motivated to *claim* that his is a high-quality product. Signalling as a solution to this difficulty takes place when sellers of truly higher-quality products engage in some activity *which would not be rational for those selling lower-quality products*. Any activity is a potential signal if sellers of higher-quality products can engage in it at lower marginal cost than producers of lower-quality products. For example, it has been argued [Nelson 1974, 1975] that advertising tends to be especially advantageous for producers of higher-quality goods, in contexts where *repeat purchases* are a significant consideration. Since the high-quality firm will be acquiring a pool of satisfied customers, its marginal advertising cost per unit of sales will be lower. Even if there is zero information content to the promotion itself, a message is being conveyed: *that the product is worth promoting*.

For the labor market Spence [1973], Stiglitz [1975], and Riley [1976, 1979] have argued that educational credentials may constitute signals with regard to jobs in which productivity is difficult to determine. As long as there is a negative correlation between productivity and the (money and time) costs involved in achieving any education level, the marginal cost of education is lower for the higher-quality workers. The latter are then able to signal by attaining higher educational credentials. On the other side of the market, the complementary process to signalling is called screening; employers are able to use education signals to screen for quality differentials.

Rothschild and Stiglitz [1976] and Wilson [1977] make parallel arguments for the insurance market. Section 1.2.3 above illustrated how, in the absence of ability to distinguish between better and poorer risks, insurance premiums reflect the average risk quality. Hence adverse selection occurs, with lower-quality risk classes tending to insure more than others. However, *ceteris paribus*, the higher the probability of loss the higher is the marginal loss in utility associated with accepting less than full coverage. Thus, the marginal cost of accepting a high deductible is greater for those with low-quality risks (high loss probabilities). Insurance companies can then screen for differences in risk by offering a menu of policies. Some policies can be offered at the low premiums appropriate for high-quality risks but with high deductibles to deter acceptance on the part of poor risks. Other policies, intended to appeal to poorer risks, offer smaller deductibles but steeper premiums.

In contrast to the situations examined in section 2.2, in signalling models the flow of information from seller to buyer is generated endogenously. This has important consequences for the stability of informational equilibria.

It has been established that unless the gap between the quality of different products is sufficiently large, there is no Cournot-Nash equilibrium [Rothschild and Stiglitz 1976, Riley, 1975, 1977, Spence 1976]. That is, starting from a situation in which all traders adopt some complementary signalling/screening pattern, there is always an alternative which yields someone greater profit.

Figure 10 illustrates a simple case where informational equilibrium can be achieved. There are only two quality levels (of workers, if we think of a labor-market interpretation). The dashed indifference curves $U_1(s,p)$ represent , for a low-quality seller (worker), equivalent combinations of price p (wage) and level s of signalling activity (e.g., education). For a high-quality seller, the solid indifference curves $U_2(s,p)$ will be applicable. The lesser slope of the latter indicates that a more qualified worker can more easily or cheaply acquire the educational attainment that serves as signal, i.e., he requires a smaller wage increment to warrant his investing in education.

It is supposed here that, with full information, buyers (employers) would be willing to pay θ_1 for the low-quality product and θ_2 for the high-quality product. Knowing only average quality, however, their maximum offer (to begin with, with both classes of suppliers in the market) is $\bar{\theta}$. But, as depicted here, the reservation indifference curve U_2^R of the high-quality sellers lies always above $\bar{\theta}$. In the absence of signalling, high-quality sellers would therefore withdraw from the market, changing the average quality $\bar{\theta}$. Ultimately, the equilibrium price would be $p = \theta_1 = \bar{\theta}$, with trade only in the low-quality product taking place. This is of course the extreme adverse-selection outcome (compare Section 1.2.2 above).

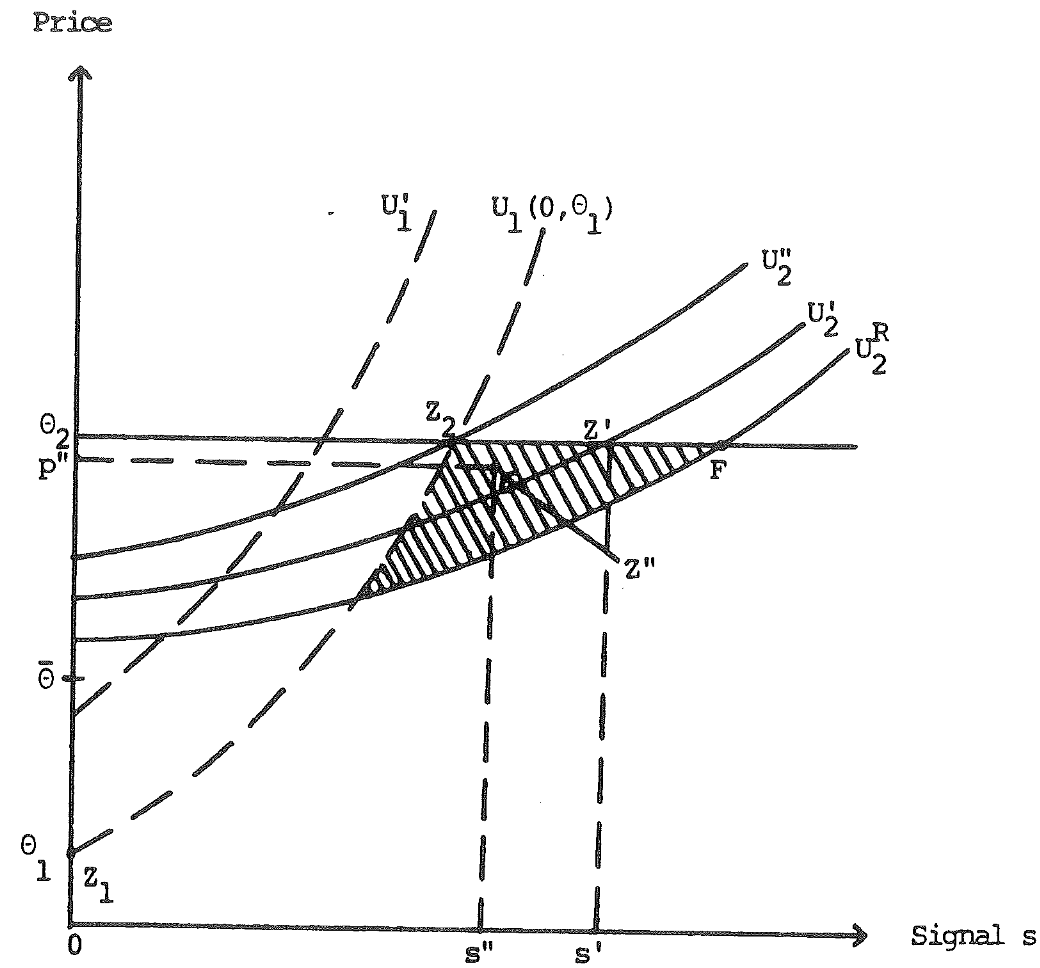

Fig. 10: Cournot-Nash Equilibrium

Now let signalling opportunities be introduced. Starting from the above adverse-selection outcome, suppose a single employer made some offer $\langle s,p_2\rangle$ -- a wage of p_2 to individuals with qualifications s -- lying within the shaded region in the diagram. High-quality workers would accept the offer, since it lies above their reservation indifference curve U_2^R; low-quality workers would not, since it lies below their indifference curve $U_1(0,\theta_1)$. In addition, the employer would be making a profit (since $p_2 < \theta_2$). So we can conclude that such offers would be made. But then other buyers would follow, putting upward pressure on price, so that ultimately $p_2 = \theta_2$.

Consider then the tentative equilibrium consisting of the pair of offers, or offered wage schedule, $Z_1 = \langle 0,\theta_1\rangle$ and $Z' = \langle s',\theta_2\rangle$ where s' is some level of s lying between the points Z_2 and F in the diagram. Acting as price-takers, the lower-quality workers will accept the offer Z_1 and the higher-quality workers the offer Z'. On the other side of the market the buyers, also acting as price-takers, find that the products purchased have the anticipated characteristics. The pair of offers $\{Z_1,Z'\}$ is thus an equilibrium in the Walrasian sense (and in the sense of Spence [1973]). Three points are of interest about this equilibrium: (1) it is not fully determinate, because of the range of values permitted for s'; (2) it is obviously inefficient, since resources are surely being wasted in signalling so long as s' lies to the right of Z_2; (3) nor is this equilibrium even approximately stable in the Cournot-Nash sense. As illustrated in Fig. 10, there exists an alternative offer pair or wage schedule $\{Z_1,Z''\}$ which would strictly benefit both employer and high-quality supplier, without injuring the low-quality workers. However, within the set of $\{Z_1,Z'\}$ Walrasian equilibria there is one that is a stable Nash-Cournot equilibrium: the pair $\{Z_1,Z_2\}$. The

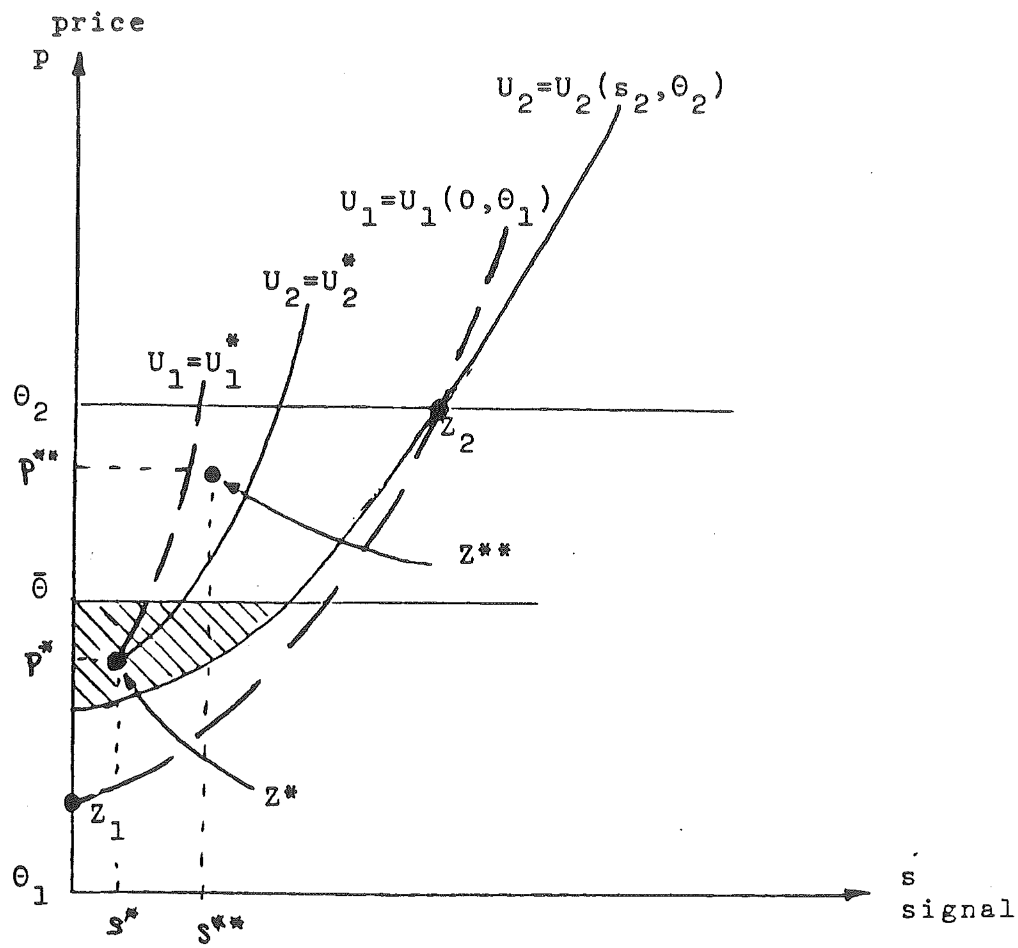

Fig. 11: Cournot-Nash Instability

high-quality workers accept Z_2; as for the low-quality workers, being indifferent between Z_2 and Z_1 they have no motive to shift from the latter.

Unfortunately, this result is not general. Suppose instead the indifference curves of the high-quality sellers differ less from those of the low-quality sellers, as in Figure 11. Starting from the $\{Z_1,Z_2\}$ solution, consider the alternative offer $Z^* = \langle s^*,p^* \rangle$ in the shaded region. It is strictly preferred by sellers of both quality levels; moreover, since $p^* < \bar{\theta}$, it yields expected profits to the buyer. Therefore the set of offers $\langle Z_1,Z_2 \rangle$ is no longer a Cournot-Nash equilibrium. However, no "pooling" offer like Z^* can be a Cournot-Nash equilibrium either since it is always possible to skim the high-quality sellers off the top of the pool. In Fig. 11, if all workers were receiving Z^*, a firm could enter and offer Z^{**} to high-quality sellers only. Since $p^{**} < \theta_2$ such an offer is profitable.

How then would such a market behave? Plausibly, in the absence of collusion, each buyer would eventually expect some <u>reaction</u> by other agents to changes in his own list of offers. Suppose that a new offer would be profitable in the absence of any reaction, but yields losses once another buyer reacts with a strictly profitable counter-offer. Suppose furthermore that the latter's response is riskless, in the sense that <u>further</u> response by any other buyers would not impose losses on the first reactor. Then it seems reasonable that the potential initial "defector" would eventually recognize that his new offer would bring on such a reaction, and hence would be deterred from making it. This suggests the following strategic equilibrium concept [Riley 1977], which builds on the development by Wilson [1977].

REACTIVE EQUILIBRIUM: <u>A set of offers is a reactive equilibrium if, for any additional offer which yields an expected gain to the agent making the offer</u>

there is another which yields a gain to a second agent and losses to the first. Moreover, no further addition to or withdrawal from the set of offers generates losses to the second agent.

The general derivation of the existence and uniqueness of the reactive equilibrium is somewhat delicate. However, it is relatively easy to check that in Fig. 11 $\{Z_1,Z_2\}$ is a reactive equilibrium. The initial "defector" must make an offer like Z* to generate an expected profit. But then another buyer can counter with Z**, thereby attracting away some high-quality products. As this process continues, $\bar{\theta}$ will fall until Z* generates losses, while Z** remains strictly profitable since $p^{**} < \theta_2$.

To conclude, the endogenous revelation of information via markets is, after all, explainable as a non-cooperative equilibrium phenomenon. While in general there is no Cournot-Nash equilibrium, recognition of reasonable reactions by other agents always result in a stable equilibrium.

2.4.2 Informational Inferences From Market Prices

We now consider the problem of information leakage. In Section 2.2.3 the process of speculation was interpreted as largely due to differences of information and belief. Nevertheless, the problem of leakage did not arise there, because no trader regarded any other individual's knowledge or beliefs as intrinsically superior to his own. Here, we will suppose, everyone recognizes that some traders do and others do not possess an informational advantage. (Though, it may or may not be the case that traders with an informational advantage are publicly identified as such.) In Section 2.4.1 above, better-informed individuals were seeking to overcome the informational disparity by signalling to potential trading partners. In this section, in contrast, the better-informed individuals are trying to capitalize on the disparity, by adopting a speculative position before their informational advantage disappears.

For concreteness, we can imagine that an information service μ is available which, at a certain price K, will (non-exclusively) provide any purchaser with perfect information as to which state of the world will obtain. Initially, all the potential traders (speculators) may be assumed to have the same beliefs. But anyone can become better informed, and everyone knows that this is the case. The first, rather obvious point is that the speculative profit to those who become better-informed will decrease the larger the number purchasing the information. Figure 12 illustrates a 2-state situation. An individual has an endowment $E = (\omega_1, \omega_2)$ and beliefs (π_1, π_2). With initial state-claim prices (P_1, P_2), his optimal consumption point is C_0. If he purchases information at a price K (to be paid regardless of state), his endowment shifts to $E' = E - K$. He then anticipates that with probability π_1 he will learn that the true state is $s = 1$. In this case he will exchange all his state-2 claims for additional consumption in state 1. Similarly, with probability π_2 he anticipates learning that the true state is $s = 2$, in which case he sells all his state-1 claims. The final consumption vector C_μ yields him an expected utility gain of $U_\mu - U_o$.

But, of course, the higher expected utility associated with the information service μ attracts other individuals. If the message is that the true state is $s = 1$, all the informed individuals will be in the market purchasing state-1 claims. This pushes up the relative price of these claims (steeper dashed budget line through E'). Similarly if the true state is $s = 2$ the price of state-2 claims is bid up (flatter dashed budget line). Final consumption is thus lower in each state, reducing the expected utility of informed agents from U_μ to U'_μ. As there is still a gain over U_0, entry continues -- until the utility gain due to informed trading is exactly offset by the cost K of obtaining and processing the information.

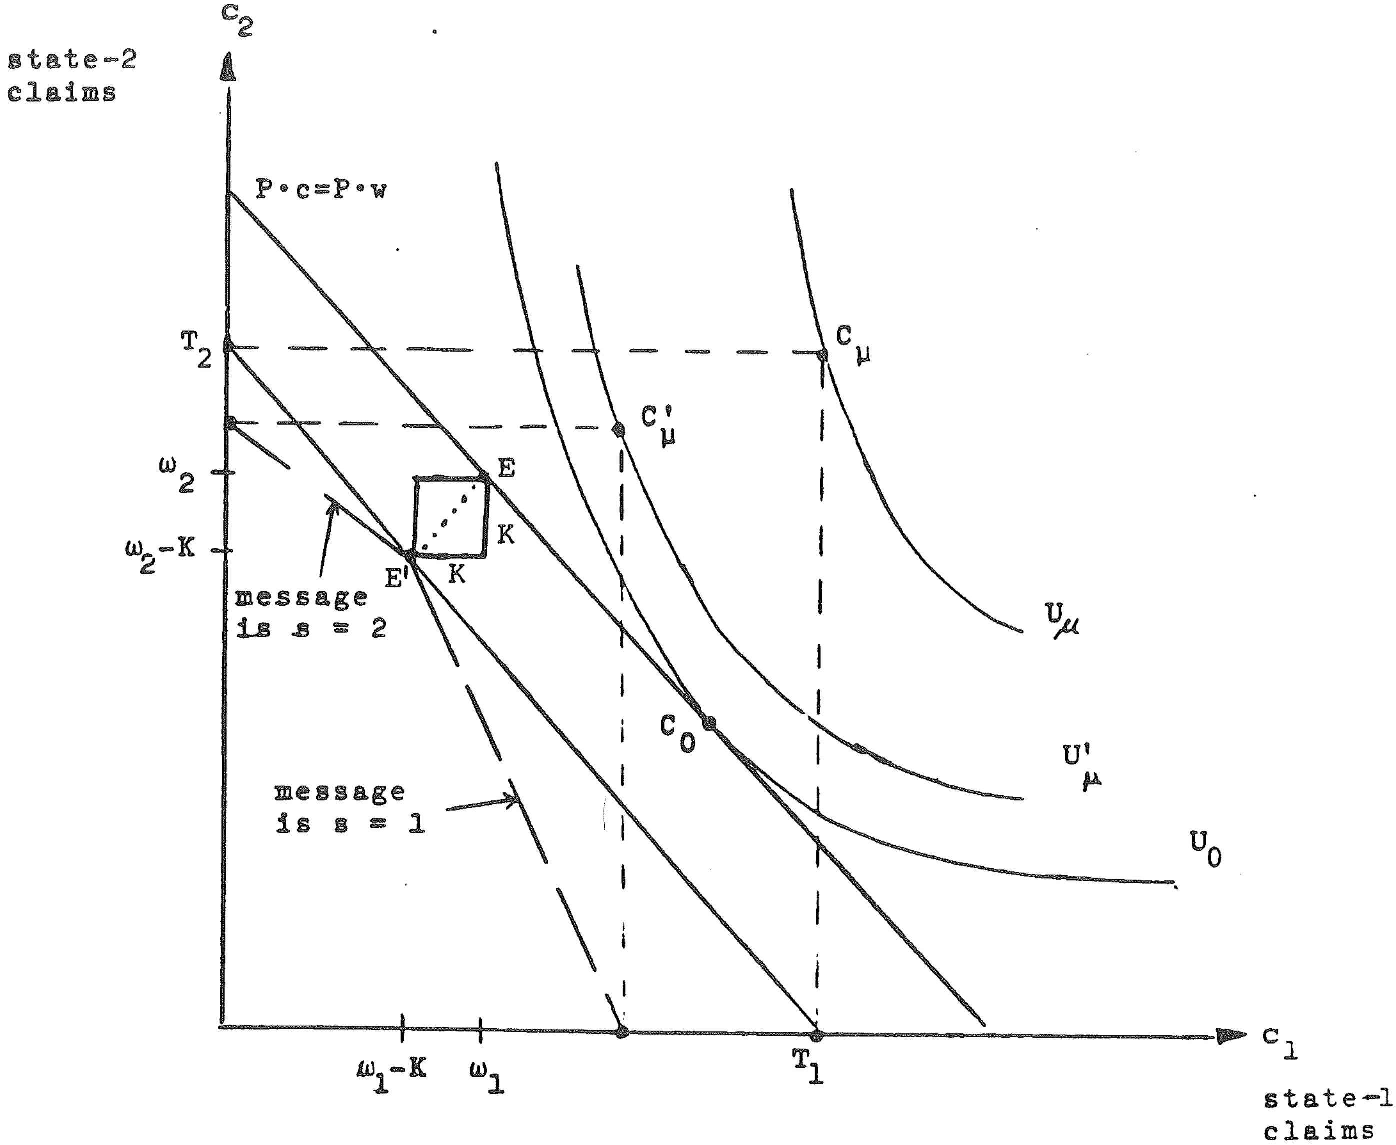

Figure 12: Trading by informed agents

So far, it looks as if there may be an equilibrium in terms of a fraction of traders that choose to become informed. But there is a further complication. Should the true state be s = 1, as long as any traders at all are buying the information the price of state-1 claims will tend to rise in comparison with the initial uninformed situation -- and, of course, the reverse if s = 2 is going to obtain. Those not purchasing the information can therefore <u>infer</u> it, simply by observing the market prices! [Green 1973, 1977, Grossman and Stiglitz 1976]. They will therefore speculate in the same direction as those who are informed. With sufficiently hair-trigger reaction functions on the part of the uninformed, there will not even be a gross profit to those choosing to buy the information -- so that, on net, these latter must lose.

In general, of course, prices may depend upon a great number of unknown or partially known determinants or parameters, apart from the uncertain element here that defines the state (e.g., the weather). But the same general result will continue to hold, so long as the price vector p(I) which would prevail if <u>all</u> agents had all the available information I differs for each different I. Then the function p(I) is invertible, and $I = f^{-1}(a)$. That is, the information can be computed from the prices; or, the price vector p is a sufficient statistic for I [Kihlstrom and Mirman 1975, Grossman 1977b].

With a finite number of information states I, it is almost certainly the case that even in very incomplete markets, the function $p(I)$ is invertible [Radner 1977]. Thus there is almost certainly a "fulfilled-expectations" equilibrium in which each agent correctly infers aggregate information from the price vector p.

As in the case of signalling models there is a market externality -- tending to break down any equilibrium in which information is obtained only at a cost: if none are informed there is potential profit in becoming informed, yet if anyone invests in information and trades accordingly he loses relative to those not having invested. The analog of the "reactive equilibrium" concept here would evidently be the corner solution with no informational investments. However, in contrast with the signalling case, an interior solution can be obtained by introducing noise or lags. If only imperfect information about the state of nature can be inferred by observing prices (as will generally be the case with a continuum of information states [Grossman 1977b, Jordan 1976]), or if the informed individual can make his commitments before the uninformed can fully react, there will tend to be an equilibrium <u>fraction</u> of traders who choose to become informed.

2.4.3 Informational Efficiency

In the Arrow-Debreu world of complete information about traded commodities and competitive markets for all contingent claims, market equilibrium is efficient in the sense that an omniscient observer would not be able to reallocate resources to the advantage of all agents. However, if information is costly, and as a result incomplete, it is hardly surprising that markets no longer have this same efficiency property. This observation does not, however, rule out the possibility that markets may be "efficient" in some narrower sense. For example, in the stock market model developed by Diamond [1967], markets are efficient with respect to a restricted set of omnisciently planned allocations.

Another restricted notion of efficiency has been the subject of considerable attention in the finance literature. A market is said to be "informationally efficient" if it is one "in which prices always <u>fully reflect</u> available information" [Fama 1970 , p. 383].

Unfortunately, the meaning of the term "fully reflect" has proved elusive. Rubinstein [1975] has argued that new public information is fully reflected in prices only if the latter adjust in such a way that, upon arrival of the information, traders have no incentive to change portfolios. But if individuals begin with heterogeneous beliefs they will almost certainly wish to adjust their portfolios. That is to say, prices almost never fully reflect public information!

Even if individuals begin with homogeneous beliefs they will, in general, again end up with heterogeneous beliefs. For example, in the previous section equilibrium was reached when the gains to informed trading were just offset by the cost of processing public information. Those individuals who do not purchase the information in such an equilibrium will again wish to revise their portfolios once the information becomes generally available. Thus again prices do not fully reflect available information in the Rubinstein sense.

In contrast, Grossman and Stiglitz [1976] have interpreted the "efficient market hypothesis" to mean that prices fully _convey_ information. In terms of our illustration, efficiency then requires that the uninformed are able to infer (from the price changes) the content of the evidence processed by the informed group. As we have seen in the previous section, under these conditions the market for private information is not viable.

The underlying problem here, in our opinion, is confusion between "efficiency" as used to characterize _equilibrium configurations_ and as applied to _disequilibrium processes_. Consider an analogous problem with regard to "profit." In an efficient competitive equilibrium, there are no competitive profit opportunities -- and yet, this outcome is the end-result of a process in which profit provides the motivating force. Whether competitive profit-seeking takes place at the ideal _rate_, and is therefore dynamically efficient, is quite another question, to answer which we as yet lack adequate tools. In its recent usage, "informational efficiency" means essentially that there is no way of making money from informational activities. Since information-involved activities (information seeking, processing, disseminating, etc.) by their very success tend, like profit-seeking activities, to eliminate their own _raison d'etre_, it is certainly true that _in equilibrium_ there will be no way of making money via such activities. But this says nothing whatsoever about the dynamically optimal level of informational activities -- which are, of their very nature, disequilibrium processes.

REFERENCES

Akerlof, G. A., "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism," Quart. J. Econ. Aug. 1970, 84, pp. 488-500.

Alchian, A. A., "The Meaning of Utility Measurement," Amer. Econ. Rev. March 1953, 42, pp. 26-50.

_________ and H. Demsetz, "Production, Information Costs, and Economic Organization," Amer. Econ. Rev. Dec. 1972, 62, pp. 777-795.

Arrow, K. J., "Alternative Approaches to the Theory of Choice in Risk Taking Situations," Econometrica Oct. 1951, 19, pp. 404-437.

_________, "Le Rôle des Valeurs Boursières pour la Répartition la Meilleure des Risques," International Colloquium on Econometrics, 1952, Centre National de la Recherche Scientifique, Paris, 1953.

_________, "Uncertainty and the Welfare Economics of Medical Care," Amer. Econ. Rev. Dec. 1963, 53, pp. 941-973.

_________, "The Role of Securities in the Optimal Allocation of Risk-Bearing," Rev. Econ. Stud. April 1964, 31 pp. 91-96.

_________, "The Value of and Demand for Information," Ch. 12 in Essays in the Theory of Risk Bearing, Markham, Chicago, 1971.

_________, "Economic Welfare and the Allocation of Resources to Invention," in The Rate and Direction of Inventive Activity: Economic and Social Factors, Universities - N.B.E.R. Conference Series, Princeton, 1972.

_________ and A. C. Fisher, "Environmental Preservation, Uncertainty and Irreversibility," Quart. J. Econ. May 1974, 88, pp. 312-319.

_________ and R. C. Lind, "Uncertainty and the Evaluation of Public Investments," Amer. Econ. Rev. June 1970, 60, pp. 364-378.

Bailey, M. J. and M. C. Jensen, "Risk and the Discount Rate for Public Investment," in Michael C. Jensen, ed., Studies in the Theory of Capital Markets, New York, 1972.

Barzel, Y., "Optimal Timing of Innovations," Rev. Econ. Stat. Aug. 1968, 50, pp. 348-55.

W. J. Baumol, "The Neumann-Morgenstern Utility Index - An Ordinalist View," J. Polit. Econ. Feb. 1951, 59, pp. 61-66.

_________, "On Taxation and the Control of Externalities," Amer. Econ. Rev. June 1972, 62, pp. 307-322.

Bergstrom, T., "Preference and Choice in Matters of Life and Death," Appendix 1 of J. Hirshleifer, T. Bergstrom, and C. Rappaport, Applying Cost-Benefit Concepts to Projects which Alter Human Mortality, UCLA School of Engineering and Applied Science, ENG-7478, Nov. 1974.

Blackwell, D. and M. A. Girshick, Theory of Games and Statistical Decisions, John Wiley, New York, 1954.

Borch, K., "Indifference Curves and Uncertainty," Swed. J. of Economics March 1968, 70, pp. 19-24.

Brainard, W. and F. T. Dolbear, "Social Risk and Financial Markets," Amer. Econ. Rev. Proc. May 1971, 61, pp. 360-370.

Calabresi, G., The Costs of Accidents: A Legal and Economic Analysis, New Haven: Yale Univ. Press, 1970.

Cheung, S., The Theory of Share Tenancy, Chicago, 1969.

________, Transaction Costs, Risk Aversion, and the Choice of Contractual Arrangements," J. Law and Econ., April 1969, 12, pp. 23-42.

____

Conley, B. C., "The Value of Human Life in the Demand for Safety," Amer. Econ. Rev. March 1976, 66, pp. 45-55.

Cook, P. J. and D. A. Graham, "The Demand for Insurance and Protection: The Case of Irreplaceable Commodities," Quart. J. Econ. Feb. 1977, 91, pp. 143-156.

Cootner, P., "Speculation, Hedging and Arbitrage," Intnl. Encyclopedia of the Soc. Sci., 1968, 15.

Debreu, G., Theory of Value, New York, John Wiley and Sons, 1959.

Diamond, P. A., "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty," Amer. Econ. Rev. Sept. 1967, 57, pp. 759-776.

________, "Accident Law and Resource Allocation," Bell J. of Econ. and Mgt. Sci., Autumn 1974, 5, pp. 366-405.

Drèze, J. H., "Market Allocation under Uncertainty," Eur. Econ. Rev. Winter 1970-71, 2, pp.

Ehrlich, I. and G. S. Becker, "Market Insurance, Self-Insurance, and Self-Protection," J. Polit. Econ. July/Aug. 1972, 80, pp. 623-648.

Ekern, S. and R. Wilson, "On the Theory of the Firm in an Economy with Incomplete Markets," Bell J. of Econ. Spring 1974, 5, pp. 171-180.

Fama, E. F.,"Efficient Capital Markets: A Review of Theory and Empirical Work," J. Finance May 1970, 25, pp.

________ and A. B. Laffer, "Information and Capital Markets," J. Bus. July 1971, 44, pp. 289-98.

________ and M. H. Miller, The Theory of Finance, New York, 1972.

Feiger, G., "What Is Speculation?", Quart. J. Econ. Nov. 1976, 90, pp. 677-88.

Feldstein, M. S., "Mean-Variance Analysis in the Theory of Liquidity Preference and Portfolio Selection," Rev. Econ. Stud. Jan. 1969, 36, pp. 5-12.

Fisher, Irving, The Nature of Capital and Income, New York: Macmillan, 1912.

Friedman, M. and L. J. Savage, "The Utility Analysis of Choices Involving Risks," J. Polit. Econ. 1948, 56, pp. 279-304.

Gordon, H. S., "The Economic Theory of a Common Property Resource: The Fishery," J. Polit. Econ. 1954, 62, pp. 124-42.

Green, J., "Information, Efficiency and Equilibrium," Harvard Inst. of Econ. Res., Discussion Paper #284, March 1973.

________, "The Non-Existence of Informational Equilibria," Rev. Econ. Stud. Oct. 1977, 44, pp. 451-63.

Grossman, S., "A Characterization of the Optimality of Equilibrium in Incomplete Markets," J. Econ. Theory June 1977, pp. 1-15.

________, "The Existence of Futures Markets, Noisy Rational Expectations and Informational Externalities," Rev. Econ. Stud. October 1977, 44, pp. 431-450.

________ and J. E. Stiglitz, "Information and Competitive Price Systems," Amer. Econ. Rev., May 1976, 66, pp. 246-253.

Groves, T., "Incentives in Teams," Econometrica, July 1973, 41, pp. 617-631.

Harris, M. and A. Raviv, "Some Results on Incentive Contracts with Applications to Education and Employment, Health Insurance, and Law Enforcement," Amer. Econ. Rev., March 1978, 68, pp. 20-30.

Hart, O., "On the Optimality of Equilibrium When the Market Structure Is Incomplete," J. Econ. Theory Dec. 1975, 11, pp. 418-443.

________, "On Shareholder Unanimity in Large Stock Market Economies," mimeo, Cambridge Univ. April 1978.

Henry, C., "Investment Decisions under Uncertainty: The 'Irreversibility Effect'," Amer. Econ. Rev. Dec. 1974, 64, pp. 1006-1012.

Hicks, J. R., Value and Capital, 2nd ed., London, 1946.

Hirshleifer, J., "War Damage Insurance," Rev. Econ. and Stat., May 1953, 25, pp. 144-153.

________, "Investment Decisions under Uncertainty: Applications of the State-Preference Approach," Quart. J. Econ., May 1966, 60, pp. 252-277.

________, Investment, Interest, and Capital, Englewood Cliffs, N.J., 1970.

________, "The Private and Social Value of Information and the Reward to Inventive Activity," Amer. Econ. Rev., Sept. 1971, 61, pp. 561-574.

________, "Liquidity, Uncertainty, and the Accumulation of Information," in C. F. Carter and J. L. Ford., eds., Uncertainty and Expectations in Economics: Essays in Honour of G. L. S. Shackle, Oxford 1972.

________, "Where Are We in the Theory of Information?," Amer. Econ. Rev. Proc. May 1973, 63, pp. 31-39.

________, "Speculation and Equilibrium: Information, Risk, and Markets," Quart. J. Econ. Nov. 1975, 89, pp. 519-42.

________, "Reply to Comments on 'Speculation and Equilibrium: Information, Risk, and Markets,'" Quart. J. Econ., Nov. 1976, 90, pp. 689-96.

________, "The Theory of Speculation under Alternative Regimes of Markets," J. Fin. Sept. 1977, 32, pp. 975-998.

________ and M. E. Rubinstein, "Speculation and Information in Securities Markets," Proceedings -- XX Meeting, The Institute of Management Science 1973, v. 2, Jerusalem 1976.

Houthakker, H., "Can Speculators Forecast Prices?", Rev. Econ. Stat., May 1957, 39, pp. 143-51.

________, "Normal Backwardation," in J. N. Wolfe, ed., Value, Capital, and Growth, Edinburgh, 1968.

Jensen, M. C., and W. H. Meckling, "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure," J. Fin. Econ., Oct. 1976, 3, pp. 305-360.

Jones, R. A. and J. Ostroy, "Uncertainty and Flexibility," UCLA Discussion Paper #73, July 1976.

Jones-Lee, M. W., The Value of Life: An Economic Analysis, Chicago, 1976.

Jordan, J. S., "Expectations Equilibrium and Informational Efficiency for Stochastic Environments," Univ. of Minnesota Disc. paper #71, August 1976.

Keynes, J. M., A Treatise on Money, New York, 1930.

Kihlstrom, R. E. and L. J. Mirman, "Information and Market Equilibrium," Bell J. of Econ., Spring 1975, 6, pp. 357-76.

Knight, F. H., Risk, Uncertainty and Profit, Houghton Mifflin, New York, 1921.

Lintner, J., "Dividends, Earnings, Leverage, Stock Prices and the Supply of Capital to Corporations," Rev. Econ. & Stat. Aug. 1962, 44, pp. 243-269.

________, "The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets," Rev. Econ. Stat., Feb. 1965 47, pp. 13-37.

Lippman, S. A. and J. J. McCall, "The Economics of Job Search: Survey," Econ. Inquiry June 1976, 14, pp. 155-189.

Luce, R. D. and H. Raiffa, Games and Decisions, New York, John Wiley, 1957.

Machlup, F., The Production and Distribution of Knowledge in the United States, Princeton, 1962.

________, "Patents," Intnl. Encyclopedia of the Social Sciences, New York, 1968, 11.

Markowitz, H., Portfolio Selection, Cowles Foundation Monograph 16, New York, John Wiley, 1959.

________, "The Utility of Wealth," J. Polit. Econ., April 1952, 60, pp. 151-58.

Marschak, J., "Role of Liquidity under Complete and Incomplete Information," Amer. Econ. Rev. Proc. May 1949, 39, pp. 182-195.

________, "Decision-Making: Economic Aspects," Intnl. Encyclopedia of the Soc. Sci., New York, 1968, 4.

________, "Limited Role of Entropy in Information Economics," West. Mgt. Science Inst., UCLA, Working Paper no. 204A, Oct. 1973.

________ and R. Radner, The Economic Theory of Teams, North Holland, Amsterdam, 1972.

Marschak, T. and R. Nelson, "Flexibility, Uncertainty and Economic Theory," Metroeconomica, April-Dec. 1962, 14, pp. 42-58.

Marshall, J., "Private Incentives and Public Information," Amer. Econ. Rev. June 1974, 64, pp. 373-390.

________, "Insurance Theory: Reserves versus Mutuality," Econ. Inquiry, Dec. 1974, 12, pp. 476-492.

________, "Moral Hazard," Amer. Econ. Rev., Dec. 1976, 66, pp. 880-890.

Mayshar, J., "Should Government Subsidize Risky Private Projects?", Amer. Econ. Rev. March 1977, 67, pp. 20-28.

McCall, J. J., "The Economics of Information and Optimal Stopping Rules," J. Bus. July 1965, 38, 300-317.

McKinnon, R., "Futures Markets, Buffer Stocks, and Income Stability for Primary Producers, J. Polit. Econ., Dec. 1967, 75, pp. 844-861.

Merton, R. C., "Investment Theory," in Handbook of Mathematical Economics, K. J. Arrow and M. D. Intriligator (eds.) North Holland, Amsterdam, 1978.

Mishan, E. J., "Evaluation of Life and Limb," J. Polit. Econ., July/Aug. 1971, 79, pp. 687-705.

Modigliani, F. and M. Miller, "The Cost of Capital, Corporation Finance and the Theory of Investment," Amer. Econ. Rev., June 1958, 48, pp. 261-297.

Mossin, J., "Equilibrium in a Capital Asset Market," Econometrica, Oct. 1966, 34, pp. 768-783.

Nelson, P., "Advertising as Information," J. Polit. Econ., July/Aug. 1974, 82, pp. 729-754.

________, "The Economic Consequences of Advertising," J. Bus., April 1975, 48, pp. 213-41.

Pauly, M., "The Economics of Moral Hazard: Comment," Amer. Econ. Rev. June 1968, 58, pp. 531-536.

Phelps, E. S. (ed.), Microeconomic Foundations of Employment and Inflation Theory, W. W. Norton, New York, 1970.

Radner, R., "Competitive Equilibrium under Uncertainty," Econometrica, Jan. 1968, 36, pp. 31-58.

________, "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," U. C. Berkeley Discussion Paper #IP-245, July 1977.

__________ and J. E. Stiglitz, "A Nonconcavity in the Value of Information," July 1976.

Raiffa, H., Decision Analysis: Introductory Lectures on Choices under Uncertainty, Reading, Mass., 1970.

Reid, J. D., Jr., "Sharecropping and Agricultural Uncertainty," Econ. Develop. and Cultural Change April 1976, 24, pp. 549-576.

Riley, J. G., "Competitive Signalling," J. Econ. Theory, April 1975, 10, pp. 174-186.

__________, "Information, Screening, and Human Capital," Amer. Econ. Rev. Proc., May 1976, 66, pp. 254-260.

__________, "Informational Equilibrium," Rand Report, R-2059-NSF, April 1977, forthcoming in Econometrica.

__________, "Testing the Educational Screening Hypothesis," J. Polit. Econ. 1979, 87,

Rothschild, M., "Models of Market Organization with Imperfect Information: A Survey," J. Polit. Econ., Nov./Dec. 1973, 81, pp. 1283-1308.

__________ and J. E. Stiglitz, "Increasing Risk I: A Definition," J. Econ. Theory, Sept. 1970, 2, pp. 225-243.

__________, "Increasing Risk II: Economic Consequences," J. Econ. Theory, March 1971, pp. 66-84.

__________, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," Quart. J. Econ., Nov. 1976, 90, pp. 629-649.

Rubinstein, M. E., "Securities Market Efficiency in an Arrow-Debreu Economy," Amer. Econ. Rev. Dec. 1975, 65, pp. 812-24.

Salant, S. W., "Hirshleifer on Speculation," Quart. J. Econ. Nov. 1976, 90, pp. 667-76.

Sandmo, A., "Discount Rates for Public Investment under Uncertainty," Int. Econ. Rev. June 1972, 13, pp. 287-302.

Savage, L. J., The Foundations of Statistics, John Wiley and Sons, New York, 1954.

Schelling, T. C., "The Life You Save May Be Your Own," in S. Chase, ed., Problems in Public Expenditures Analysis, Brookings 1968.

Schlaifer, R., Probability and Statistics for Business Decisions, New York, McGraw-Hill, 1959.

Shannon, C., "The Mathematical Theory of Communication," Bell Syst. Tech. J. July 1948, 27, pp.

Sharpe, W. F., "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," J. Finance, Feb. 1964, 25, pp. 65-86.

________, Investments, Englewood Cliffs, N.J., 1978.

Shavell, S., "On Valuable Opinion and the Efficiency of the Price System under Uncertainty," Harvard Inst. of Econ. Research, Discussion Paper 458, March 1976.

________, "On Moral Hazard and Insurance," Quart. J. Econ. June 1977.

________, "Risk Sharing and Incentives in the Principal and Agent Relationship," Harvard Inst. of Econ. Research, Discussion Paper #   , 1977.

________, "Do Managers Use their Information Efficiently?" Amer. Econ. Rev., Dec. 1978, 68, pp. 935-937.

Spence, A. M., Market Signalling: Information Transfer in Hiring and Related Processes, Harvard Univ. Press, Cambridge, Mass., 1973.

________, "Competition in Salaries and Signalling Prerequisites for Jobs," Quart. J. Econ., Feb. 1976, 90, pp. 51-75.

Stigler, G., "The Economics of Information," J. Polit. Econ., June 1961, 69, pp. 213-225.

________, "Information in the Labor Market," J. Polit. Econ., Oct. 1962, 70, pp. 213-225.

Stiglitz, J. E., "The Theory of Screening, Education and the Distribution of Income," Amer. Econ. Rev., June 1975, 65, pp. 283-300.

Strotz, R. H., "Cardinal Utility," Amer. Econ. Rev. May 1953, 43, pp. 384-397.

Thaler, R. and S. Rosen, "The Value of Saving a Life: Evidence from the Labor Market," in Nestor Terlecky, ed., Household Production and Consumption, Nat. Bur. Econ. Res. Studies in Income and Wealth, 40, New York, 1975.

Tobin, J., "Liquidity Preference as Behavior Towards Risk," Rev. Econ. Stud. Feb. 1958, 25, pp. 65-86.

Townsend, R. M., "On the Optimality of Forward Markets," March 1978, 68, pp. 54-66.

Vickrey, W., "Automobile Accidents, Tort Law, Externalities, and Insurance: An Economist's Critique," Law and Contemporary Problems, Summer 1968, 33, pp. 464-487.

Von Neumann, J. and O. Morgenstern, Theory of Games and Economic Behavior, Princeton, 1944.

Wilson, C. A., "A Model of Insurance Markets with Incomplete Information," J. Econ. Theory, Dec. 1977, 16, pp. 167-209.

Working, H., "Futures Trading and Hedging," Amer. Econ. Rev. June 1953, 43, pp. 314-343.

________, "New Concepts Concerning Futures Markets and Prices," Amer. Econ. Rev. June 1962, 52, pp. 431-59.

R. Zeckhauser, "Risk Spreading and Distribution," Kennedy School of Govt., Harvard Univ., Discussion Paper #10, Aug. 1972.

Zorn, T., "Information Differences, the Stockmarket, and Management Incentives," UCLA Ph.D. dissertation, 1978.