THE ANALYTICS OF UNCERTAINTY AND INFORMATION --
AN EXPOSITORY SURVEY*

by

J. Hirshleifer

and

John G. Riley

University of California,
Los Angeles

*A slightly condensed version of this paper is forthcoming in the Journal of Economic Literature.

UCLA
Department of Economics
Working Paper #159
July 1979

The Analytics of Uncertainty and Information --

An Expository Survey

by

J. HIRSHLEIFER

and

JOHN G. RILEY

University of California,

Los Angeles

Among the very large number of people who provided comments and suggestions, we would like to thank most especially Fritz Machlup, Robert A. Jones, Mark Perlman, and Richard J. Zeckhauser. J. Hirshleifer's work on this paper was supported in part by National Science Foundation grant No. SOC75-15697 and by a grant from the Foundation for Research in Economics and Education. John G. Riley's work was supported in part by National Science Foundation grant No. SOC79-07573.

CONTENTS

J. Hirshleifer
J. Riley

July 16, 1979

THE ANALYTICS OF UNCERTAINTY AND INFORMATION -- AN EXPOSITORY SURVEY

All human endeavors are constrained by our limited and uncertain knowledge of the world, as leading economic thinkers have always recognized. Limitations of knowledge play a notable role (to cite but a few instances) in: (1) Joseph A. Schumpeter's theory of economic development [1936(1911)], where the key figure is the entrepreneur as creator of previously unknown factor combinations; (2) Frank H. Knight's theory of profit [1921] as essentially the return (often negative) received by individuals willing to bear unquantifiable risks; (3) Friedrich A. Hayek's view [1945] of the price system as a device for communicating and integrating the divided information possessed by the separate individuals of a society. And behavior under uncertainty has been central to important theories of such fundamental economic phenomena as money, the firm, and the business cycle. An overview of the economic role of knowledge in its many aspects is forthcoming in a multi-volume treatise by Fritz Machlup under the general title _Knowledge: Its Creation, Distribution, and Economic Significance_; this work contains an extended historical bibliography [1979].

But despite this longstanding recognition, until relatively recently (until the post-World War II years) there was no rigorous microeconomic foundation for the _analysis_ of individual decisions and market equilibrium under uncertainty. (A classic instance is Irving Fisher's work on interest, where uncertainty entered only as the "third approximation"; significantly, Fisher titled his relevant chapter "The Third Approximation Unadapted to Mathematical Formulation" [1930, p. 316].) This foundation lacking, standard textbooks presented analytical models, typified by the familiar apparatus of supply and demand, that allowed no role for uncertainty. Given the overwhelming practical importance of uncertainty, it is not surprising that decision-makers in the world of affairs have often found academic economics to be of little operational value.

Recent explosive progress in the economics of uncertainty has changed this picture. The subject is now an essential topic for students not only in economics departments, but in professional schools and programs oriented toward business, government and administration, and public policy. In the world of affairs, commercial stockmarket analysts now regularly report measures of shareprice uncertainty devised by economic theorists. Even government and the law are beginning to appreciate the need for formal analysis of uncertainty in dealing with such problems as malpractice, safety and health, allowable return on investment, and income distribution. And academic economists, armed with the new developments in the economics of uncertainty, are much more successfully analyzing previously intractable phenomena such as insurance, research and invention, advertising, speculation, and the functioning of financial markets.

It will be impossible to provide any adequate review here of all the important developments under the headings of uncertainty and information. What we hope to do is to expound the central underlying ideas in non-technical fashion; to introduce the novel tools of analysis that have proved fruitful in this area; and to go somewhat more deeply into selected areas of application in order to convey some impression of the potential richness and power of the theory. Wherever possible, we will provide citations to major branches of the literature that we have been unable to survey here.

The theoretical developments that have brought about this intellectual revolution have two main foundation-stones: (1) the theory of preference for uncertain contingencies and in particular the "expected-utility theorem" of John von Neumann and Oskar Morgenstern [1944], and (2) Kenneth J. Arrow's formulation of the ultimate goods or objects of choice in an uncertain universe as contingent consumption claims: entitlements to particular commodities or commodity baskets valid only under specified "states of the world" (more briefly, "states") [1953, 1964]. Just as intertemporal analysis requires subscripting commodity claims by date, uncertainty analysis require subscripting commodity claims by state. 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 world of certainty.

An alternative conceptualization of the objects of choice under uncertainty runs in terms of the statistical parameters of the probability distributions of commodity or income claims. In that formulation it is assumed that individuals prefer greater mean income but smaller variance of income; attention may or may not be paid to higher moments of the distribution [Harry M. 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 [James Tobin, 1958; Karl Borch, 1968; Martin S. 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; William F. Sharpe, 1964; John

Lintner, 1965; Jan Mossin, 1966]. We will not be able to pursue the parameter-preference approach here; for recent surveys see Michael C. Jensen [1972] and R. C. Merton [1978].

The modern analytical literature on uncertainty and information divides into two rather distinct branches. The first branch deals with market uncertainty. 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, on the individual level the search for trading partners, and at the market level disequilibrium and price dynamics, take the center stage -- replacing the traditional assumption of costless exchange at market-clearing prices [George J. Stigler, 1961, 1962; John J. 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 [Edmund S. Phelps, 1970]. The second branch of literature deals with technological uncertainty or (a preferable designation) event uncertainty. Here individuals are uncertain not about the terms on which they might make market exchanges but rather about exogenous data -- such as resource endowments (will the wheat crop be large or small?) or productive opportunities (will fusion power be available?) or public policy (will taxes be cut?). Put another way, market uncertainty concerns the endogenous variables of the economic system, event uncertainty the exogenous data.

The present survey is limited to the relatively more tractable topic of event uncertainty. This limitation permits us to employ the simpler traditional model of perfect markets in which all dealings take place costlessly at equilibrium prices. Recent studies of the complex search and disequilibrium phenomena that emerge under market uncertainty are reviewed in Michael Rothschild [1973] and Steven A. Lippman and McCall [1976].

The paper is divided into two main parts, the first covering the economics of <u>uncertainty</u> and the second the economics of <u>information</u>. The two categories correspond to what might be called passive versus active responses to our limitations of knowledge. In Part 1 individuals may be said to <u>adapt</u> to the fact of uncertainty; in Part 2 they are allowed also to <u>overcome</u> uncertainty by engaging in informational activities.

Part 1

THE ECONOMICS OF UNCERTAINTY

1.1 Decision Under Uncertainty

In decision-making under uncertainty the individual chooses among acts while Nature may metaphorically be 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.

More generally, 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 him 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 of acts
Acts	a = 1	c_{11}	c_{12}	u_1
	a = 2	c_{21}	c_{22}	u_2
Beliefs as to states		π_1	π_2	

Table 1: Consequences of Alternative Acts and States

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 (whether to accept or reject 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 there are interpersonal arrangements -- insurance contracts, futures markets, guarantees and collateral, the corporation and other forms of combined enterprise -- which 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 which 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 terminal actions -- 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 The Probability Function

We will assume that each individual is able to represent his beliefs 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 [Irving Fisher, 1912, Ch. 16; Leonard J. 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 certainty would be represented by attaching the full probabilistic weight of unity to only one of the outcomes. The degree of subjective uncertainty is reflected in the dispersion of probability weights over the possible states.

Frank Knight [1921] attempted to distinguish between "risk" and "uncertainty," depending upon whether probability estimates are or are not calculable on the basis of an objective classification of instances. At times he suggested (pp. 20, 226) that the probability concept is inapplicable under true uncertainty, for example, to such questions as whether or not a cure for cancer will be discovered in the next decade. 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 even for Knightian uncertainty situations. But elsewhere Knight's discussion is much more in line with modern developments, as when he suggests (p. 227) that a man's actions may depend upon his estimate of the chances that his beliefs are correct -- or, we shall say, upon his confidence in his beliefs.

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; 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. There is no operational difference between hard and soft probabilities (high and low confidence) if only terminal action is called for. The difference arises where there are opportunities for informational action. We will show explicitly in Part 2 how 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. However, we will sometimes assume that it corresponds simply 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, 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" in one state of the world, "75% chance" in the other.

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 -- 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 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 preference scaling 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 $c_a = (c_{a1},\ldots,c_{aS})$ are to be received with respective probabilities $\pi = (\pi_1,\ldots,\pi_S)$, is:

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

Or, more compactly:

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

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

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

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 v(c). What is its justification? It turns out that the expected-utility rule is usable <u>if and only if the v(c) function is determined 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.

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 discussions in William Baumol, 1951; Armen A. Alchian, 1953; Robert H. Strotz, 1953).

The "postulates of rational choice" serving as basis for the theorem have been set forth in a number of different ways in the literature [Milton Friedman and Leonard J. Savage, 1948; R. Duncan Luce and Howard Raiffa, 1957; Markowitz, 1959; Jacob Marschak, 1968], and involves technicalities that cannot be pursued here. Instead, what follows is an informal presentation (based mainly on Robert 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 some intermediate level of income c^*. We can suppose that the individual is indifferent between (assigns equal utility to) having c^* for certain versus having some particular chance of success π^* in a "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^* of success in the reference lottery. Making use of "prospect" notation:

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

Figure 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. Then the utility assigned to the sure consequence $250 is just 1/2 -- $v(250) = .5$. Repeating this process, the reference-lottery technique generates the individual's entire $v(c)$ curve of Fig. 1, which is his preference-scaling function for consequences.

We can immediately verify that, with the $v(c)$ function determined as just shown, the expected-utility rule (1.1) does give us the correct valuation for any prospect whatsoever -- for example, 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 specific chance of success π' in the reference lottery -- $v(c') = \pi'$ -- and similarly c'' corresponds to some other 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 him the chance π of an income c' (equivalent to the 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:

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

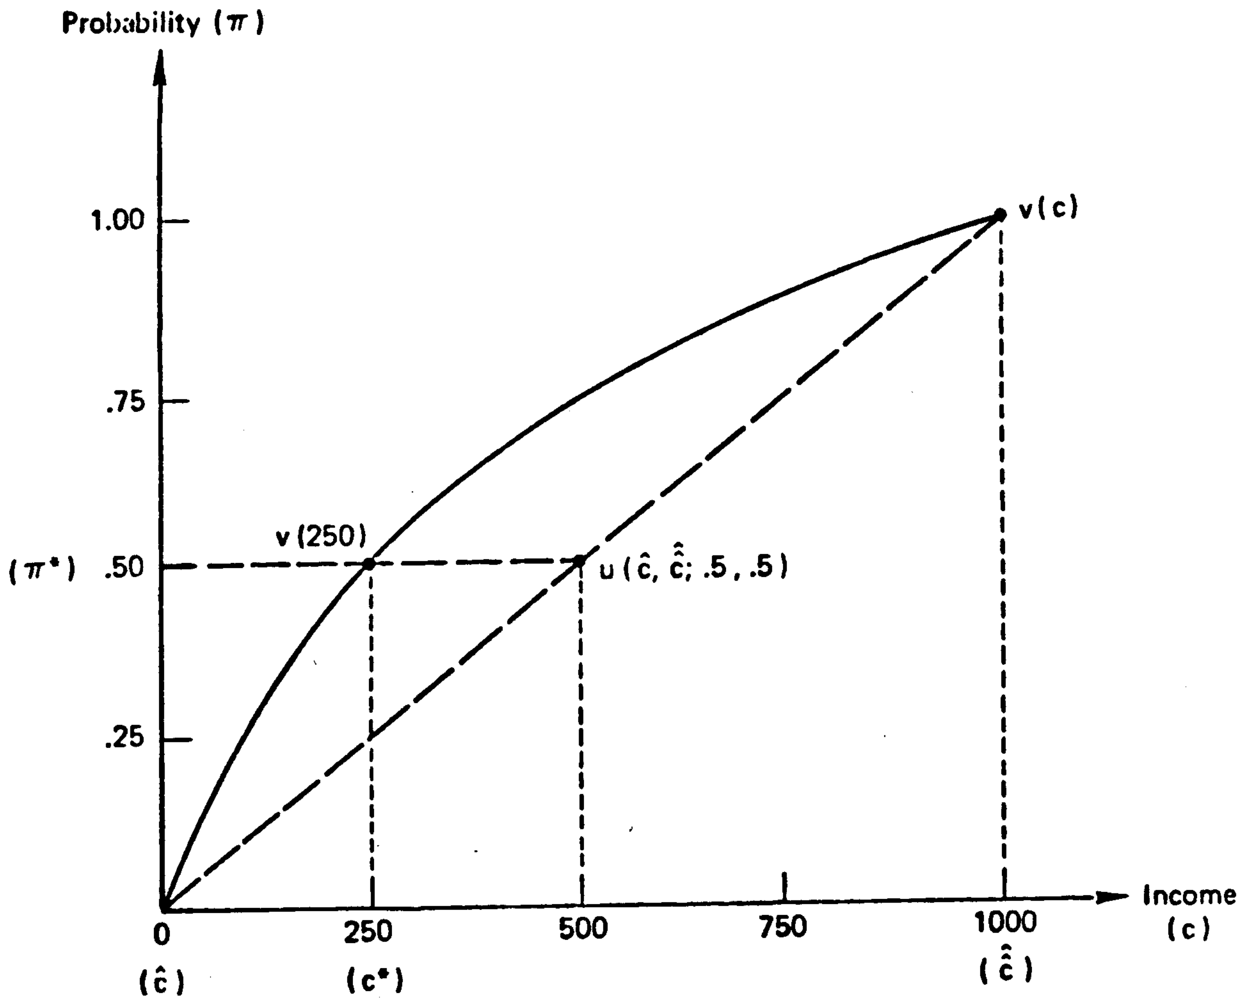

Fig. 1 — The preference-scaling function v(c) derived by the "reference-lottery technique"

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 to be 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. Fig. 1 illustrated a situation where the reference lottery with equal chances of \$1000 or zero (and thus with a mathematical expectation of \$500) is the utility equivalent of a sure income of only \$250. Such a person must then prefer a sure income of \$500 to this risky lottery whose mathematical expectation is \$500. It is intuitively evident that this generalizes: any point P on a concave v(c) curve will lie above the corresponding (vertically aligned) point along the straight line connecting any pair of positions on v(c) that bracket P. The point on the curve represents the utility of a given sure income; the vertically aligned point on the straight line 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 fair gamble, 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 on a straight line 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. One 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 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; Markowitz, 1952]. These constructs run against the difficulty that if gambles are available on a fair or nearly-fair basis, no one could ever be at an optimum in any risk-preferring range of his v(c) curve. To leave such a range, individuals would accept even enormous riches-or-ruin gambles. 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, therefore, gambling at fair or adverse odds appears to be 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.

That risk-aversion is the normal situation is indicated in a different way by Fig. 2. Here the familiar-looking indifference curves $u^o, u', u'', \ldots$ show the expected utility of gambles, based on 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'):

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

This family of equations corresponds to the indifference curves of the diagram. It follows immediately, since $du = 0$ along any indifference curve, that the indifference-curve slopes in Fig. 2 are related to the marginal utilities $v'(c)$ via:

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

It is elementary though tedious to show that the indifference curves have the normal "convex to the origin" curvature 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)

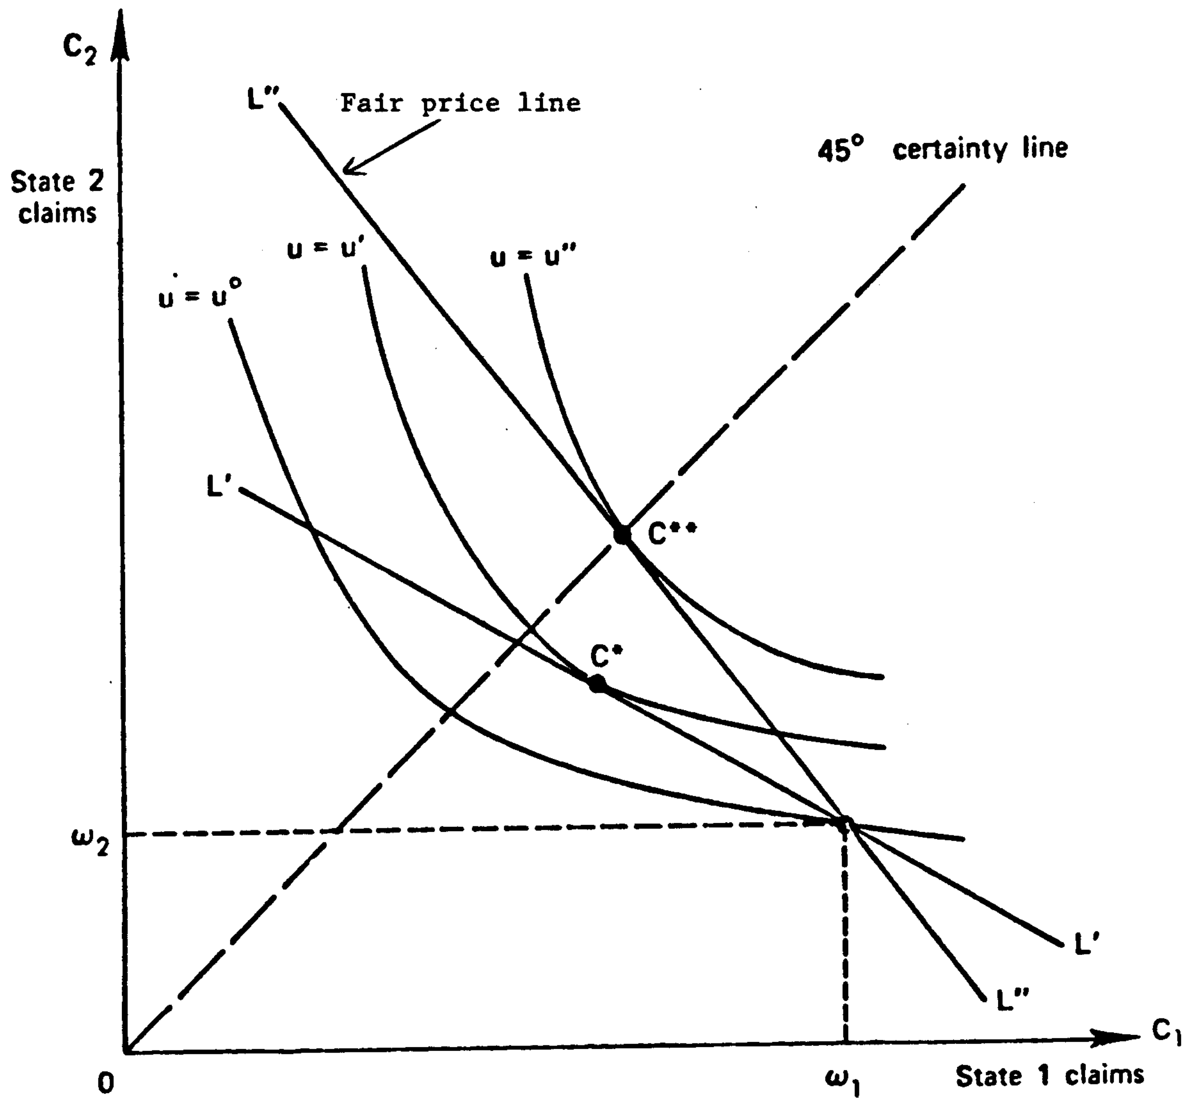

Fig. 2 — The preference map in contingent consumption or state-claim space (probability fixed)

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, the risk-bearing optimum position C* will normally be in the interior -- i.e., the individual will want to "diversify" his holdings of state-claims. Following standard techniques, 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 mathematical expectation of gain 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 coresponds to the market offering fair gambles. Then the condition (1.4) simplifies to:

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

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 tangency optimum like C** at the intersection of the 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 offsets his endowed gamble. (Correspondingly, if the market 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 his endowment position must also be taken into account.

A natural next step would be to explore the responses of the individual's risk-bearing optimum C* (and thus of his implied state-claim transactions) to a variety of parametric shifts: to changes in prices, in probability beliefs, in the size of endowed income and its state-distribution, etc. Limited space precludes a detailed presentation here. However, it will be useful to describe certain measures of risk-aversion which help indicate the directions of these parametric responses.

Starting from some initial C* position in Fig. 2, suppose that income increases while prices remain unchanged. If the Income Expansion Path so generated is a straight line out of the origin, the individual is said to have constant relative risk-aversion (R). The implication is that as he becomes richer he holds proportionately more of each and every state-claim. If the utility function $u(c_1,c_2)$ of Fig. 2 is homothetic, R will be everywhere constant. If, on the other hand, R is increasing with income (increasing relative risk-aversion), the Income Expansion Path will curve relatively closer to the 45° line at higher incomes, and the reverse if R decreases with income (decreasing relative risk-aversion). It can be shown [John W. Pratt, 1964; Arrow, 1965] that R is related to the v(c) function by the condition:

$$- \frac{cv''(c)}{v'(c)} = R \tag{1.5}$$

As an example, it can be verified that the logarithmic preference-scaling function $v(c) = \log_e c$ implies constant $R=1$.

An individual is said to have constant <u>absolute risk-aversion</u> (A) if his final C* position, after an increase in income, represents equal absolute increases in each and every state-claim holding. This corresponds geometrically to an Income Expansion Path parallel to the 45° line. The associated analytic condition on the v(c) function is:

$$- \frac{v''(c)}{v'(c)} = A \tag{1.6}$$

Increasing absolute risk-aversion A implies an Income Expansion Path converging absolutely toward the 45° line at higher incomes, and the reverse if A is decreasing with income. As an example, it can be verified that the negative exponential preference-scaling function $v(c) = 1 - e^{-c}$ implies constant $A=1$.

Pratt [1964] has argued that aversion to bearing a given absolute risk (to holding a portfolio with a given gap between c_1 and c_2) normally decreases as income rises. This is plausible; a rich individual, other things equal, should be more willing to hold any kind of asset, including assets with a given absolute risk. Note that diminishing <u>absolute</u> risk-aversion A is not inconsistent with constant <u>relative</u> risk-aversion R.

Figure 3 illustrates how risk-aversion enters into the determination of one kind of parametric response -- to a "mean-preserving spread" [Michael Rothschild and Joseph E. Stiglitz, 1970, 1971] of the returns (distribution of state-claims) associated with a particular asset. Let the endowment position lie on the 45° certainty line (the individual is initially holding

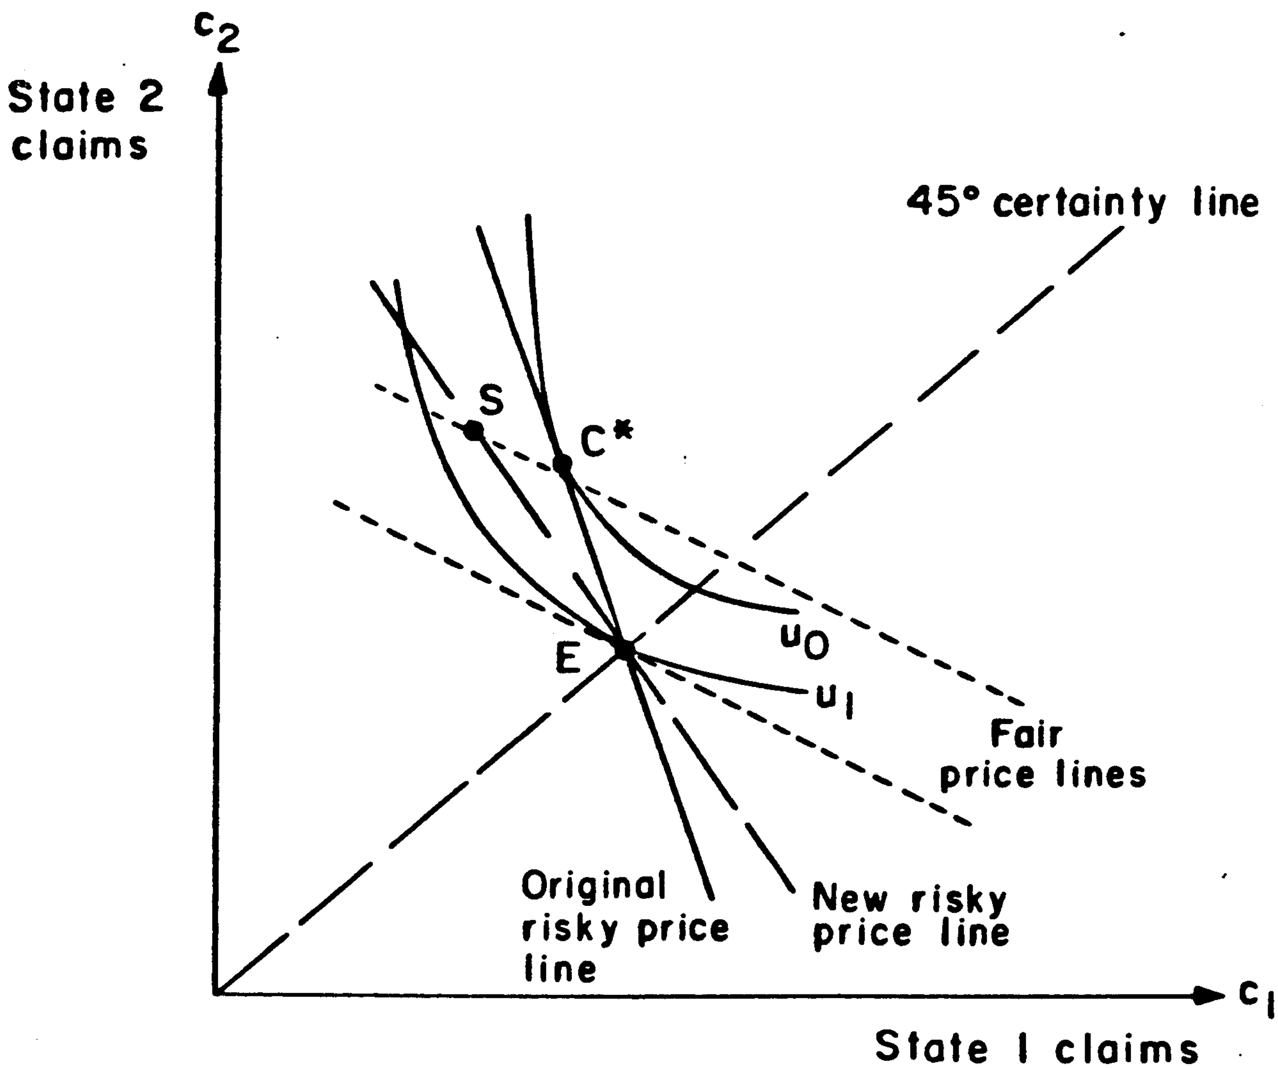

Fig. 3: Effect of a "mean preserving spread" of asset returns.

only riskless assets). If a risky asset were purchasable but only at "fair" prices, the individual would of course prefer to remain at his riskless position E on the (dotted) fair price line. If the risky asset represents a favorable gamble, however, the price line through E would be steeper and the individual would purchase some amount of it (would exchange some of his riskless for risky assets) in moving to his optimum C* position.

Now suppose the nominal asset price is unchanged, but there is a "mean-preserving spread" of asset returns. Then each point like C* on the old risky price line is shifted to a new point S involving a higher risk (gap between c_1 and c_2). Note that S and C* lie on the same dotted fair price line, since the mean return has been preserved. The new effective risky price line through points E and S (bold dashed line) is flatter than the original risky price line, since exchanging the riskless and risky asset now involves giving up more c_1 per unit of c_2 acquired. As the terms of the gamble are less favorable than before, the pure substitution effect implies that a smaller absolute amount of the risky asset will be acquired. There is also an adverse income effect; if the individual is characterized by "normal" diminishing absolute risk-aversion A, his impoverishment will increase his aversion to risk and thus reinforce the substitution effect. In these circumstances the new optimum along the price line through E and S would surely lie closer to the 45° certainty line than the old optimum point C*. This concludes our illustrative discussion of the individual's response to particular changes in the parameters of his opportunity set.

1.2 Market Equilibrium Under Uncertainty

We now shift the level of analysis, from the decisions of the individual to market interactions and the conditions of equilibrium. 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 contract to share the total risks and returns. This can be illustrated by the Edgeworth box in Fig. 4 [William C. Brainard and F. Trenery Dolbear, Jr., 1971; John M. Marshall, 1976], which for concreteness may be thought of as illustrating a "share cropping" problem [Steven N.S. Cheung, 1969; Joe D. Reid, Jr., 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 ($\pi_N = 1 - \pi_L$), the indifference curves for each agent have the same absolute slope π_L/π_N along their respective 45° certainty lines. It is obvious that the two traders (Landlord I and Worker II) cannot both attain certainty positions. At a position like E the Landlord is bearing all the risk (the Worker is receiving a fixed wage independent of which state obtains). At a position like V the opposite

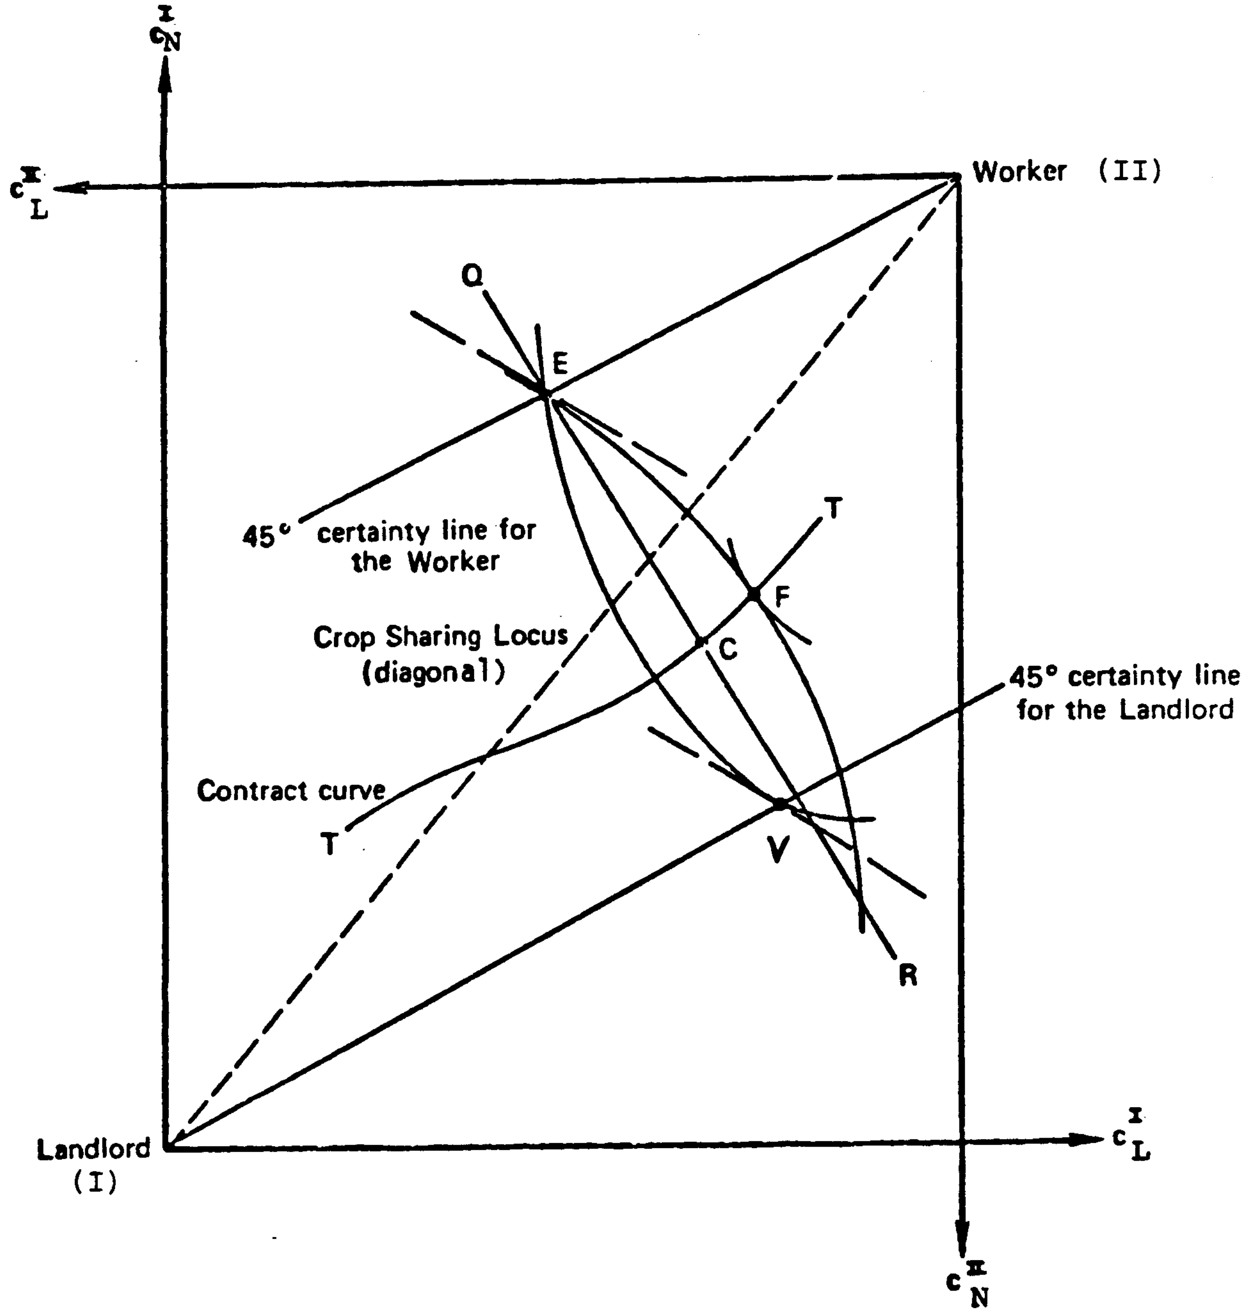

Fig. 4 — Risk Sharing

holds; the Landlord is receiving a fixed rent regardless of state, while the Worker bears the risk. Starting from a position like E, state-claim trading will lead to an equilibrium at a point like C on the contract curve TT within the region of mutual advantage. It is geometrically evident that the contract curve necessarily lies between the two 45° certainty lines, so some of the risk will be borne by each party.

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 interaction in Fig. 4 can be given another interpretation: the risk-sharing that takes place there 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. We will be providing here a relatively extended discussion of insurance markets in order to illustrate in somewhat greater depth a number of the salient issues of uncertainty theory.

In the insurance context, once again the Edgeworth box will in general be elongated; we can imagine a social "loss" state of the world L (e.g., an earthquake occurs) versus a "non-loss" state N. From any given endowment point like E, price-taking traders arrive at a risk-sharing equilibrium like C on the contract curve. The absolute slope of the equilibrium market line ($=P_L/P_N$) exceeds the absolute slope of the dashed lines representing the fair prices or probability ratio ($=\pi_L/\pi_N$). That is, claims to income in the less affluent state L command a relatively high price, the marginal utility of income in that state being higher.

In economic analyses of insurance there has been a tendency to assume that fair or "actuarial" insurance terms would be normal were it not for transaction costs ("loading") [Isaac Ehrlich and Gary S. Becker, 1972]. In what follows we shall survey some of the major elements, transaction costs aside, that generally lead to non-fair equilibrium prices.

Social Risk:

Suppose two individuals I and II have equal initial incomes, but there is a hazard that will surely impose a fixed loss ℓ on exactly one of them (with fixed, but not necessarily equal probabilities for each). The Edgeworth box would be square. Then the two 45° lines collapse into the single main 45° diagonal, which also becomes the contract curve. Here there is private risk without social risk. The two states of the world are "loss strikes I" versus "loss strikes II." 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"). The equilibrium price ratio (premium/indemnity ratio) corresponds to the respective probabilities; at these fair prices, all private risk is eliminated by mutual insurance.

Apart from this extreme special case of perfect negative correlation of risks, four distinct states of the world can be defined in a two-party situation -- according as the loss is suffered by neither person, I alone, II alone, or both. And the social total of losses can be 0, 1, or 2. Evidently, there is no way of arranging affairs so that everyone can have the same income regardless of state; universal full insurance (whereby everyone attains his "certainty line") is generally impossible. It follows that equilibrium prices cannot be "fair"; each person's premium/indemnity ratio must exceed the odds that he will suffer a loss.

For larger insurance pools with M members, the Law of Large Numbers is sometimes thought to justify treating the per-capita loss $\gamma \equiv \frac{1}{M} \Sigma_{i=1}^{M} \ell_i$ as approximately constant over states. As M increases, the variance of γ declines and thus the error committed by assuming away social risk diminishes. Nevertheless this error does not tend toward zero unless the separate risks are on average uncorrelated [Markowitz, 1959, p. 111]. If for each individual i the variance of loss ℓ_i has the same value σ^2, and if the correlations between all pairs of risks equal some common r (which can only hold if $r \geqq 0$), the variance of the per-capita loss γ is:

$$\sigma_{\gamma}^2 = \frac{1}{M^2} (M\sigma^2 + M(M-1)r\sigma^2) \tag{1.7}$$

In the limit as M increases, the variance of per-capita loss approaches the value $r\sigma^2$, 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 its 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 (where 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.

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

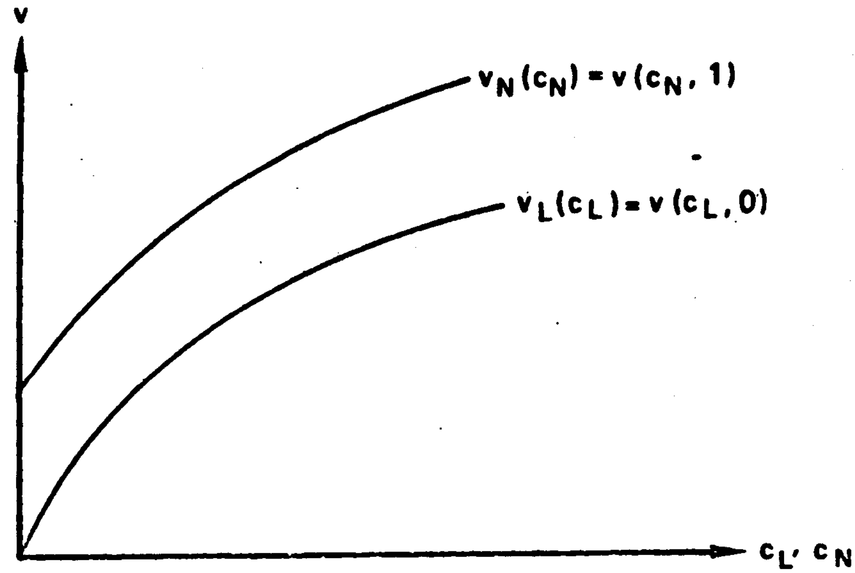

Fig. 5a — State-Dependent Utility

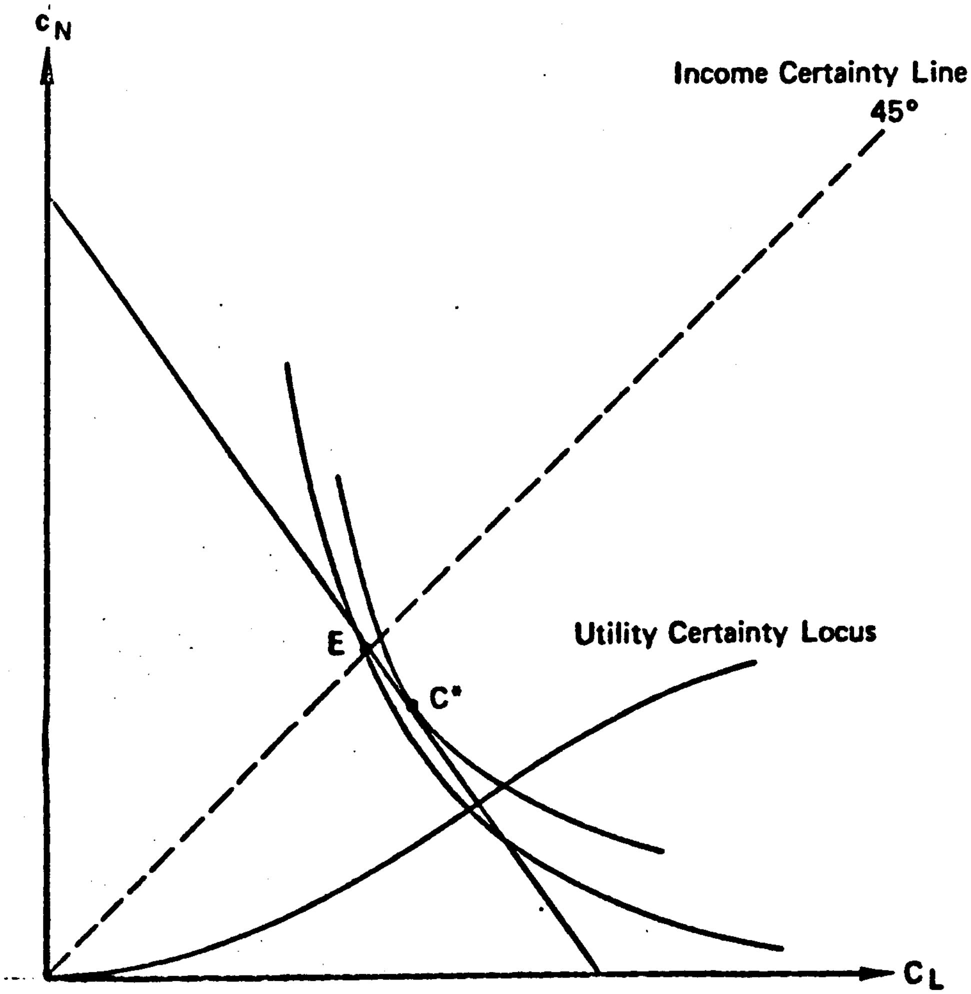

Fig. 5b — "Heirloom" Insurance

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 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 curves $v_N(c)$ and $v_L(c)$ do not represent distinct utility functions, but different sections through a single v(c,h) function.

With state-dependent utility, the 45° income certainty line (ICL) is no longer the individual's utility certainty locus (UCL). Someone on the 45° line would not be indifferent as to whether state N or L occurs. In Fig. 5b, the UCL lies toward the c_L axis, since the individual requires more income in state L than in state N if he is to "fully insure utility" [Philip J. Cook and Daniel A. Graham, 1977].

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):

$$\frac{\pi_L v_L'(c_L)}{\pi_N v_N'(c_N)} = \frac{P_L}{P_N} \tag{1.8}$$

Assuming, as a benchmark, that prices are fair, (1.8) reduces to:

$$v_L'(c_L) = v_N'(c_N) \tag{1.8'}$$

With actuarial insurance available, the individual thus equates his marginal utilities in the two states as before, but these marginal utilities are now slopes along differing v(c) curves.

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, at actuarial prices, 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 two-dimensional cardinal preference-scaling function v(c,h) 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.

But suppose h=0 represents a major injury. Then the marginal utility of income will probably be higher in the loss state (one "needs" income c more than before). In such cases the optimum C* lies to the southeast of the income certainty line, though not necessarily southeast of the utility certainty locus. The individual will buy insurance against injury, 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 largely on him. In such a case it is optimal to transfer income from the loss state to the non-loss state. That is, to "reverse insure" -- to bet that the loss would not occur. (Contractually, instead of insuring your child's life you 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 perfectly monitor the behavior or identify the risk-status of insureds. (Since this is a kind of <u>informational</u> problem, our analysis here has close ties with topics to be taken up in Part 2 below.)

To stick to essentials, we need only consider two risk classes with the same initial wealth ω and facing the same potential loss ℓ. In the absence of insurance the high-risk class, with loss probability π', has an expected utility of $\pi' v(\omega-\ell) + (1-\pi')v(\omega)$. This is represented by the distance A'B' in Fig. 6. If members of this risk class are identifiable by the insurers (by the other members of the mutual insurance pool), fair insurance will result in their being offered full coverage for a premium of $\pi'\ell$ -- equal to the mathematical expectation of loss. These individuals then choose certainty position C' in Fig. 6 (equivalent to being on the 45° line in Fig. 2). Expected utility is $v(\omega-\pi'\ell)$, equal to the distance A'C' in the diagram.

Similarly, if members of a low-risk class with loss probability π'' can be identified, insurance at actuarial terms will raise their expected utility from A"B" to A"C".

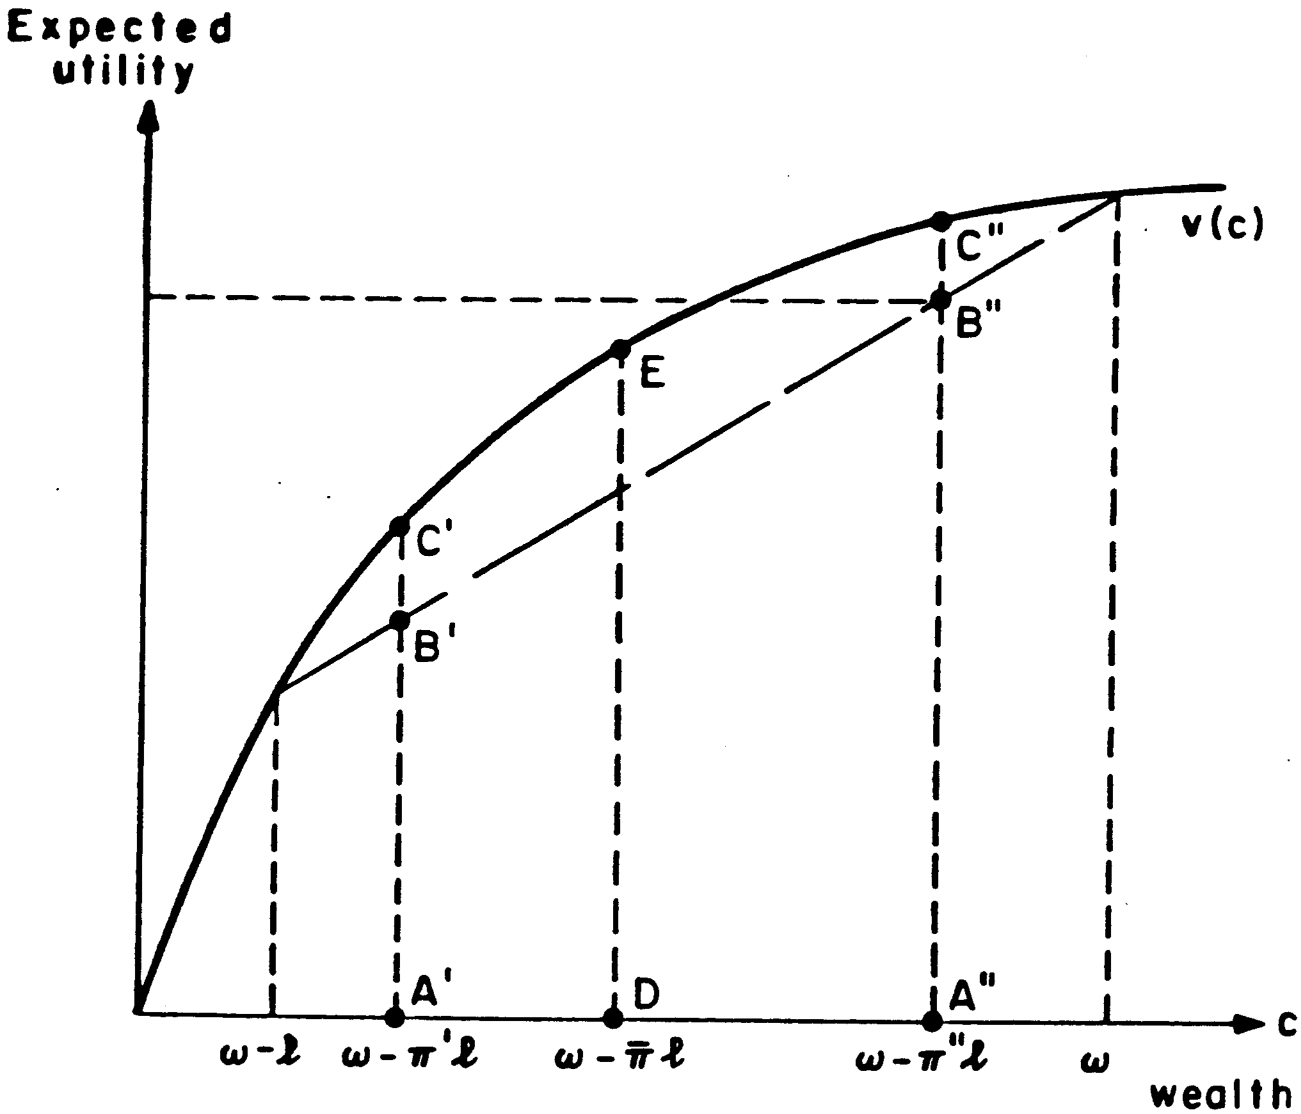

Fig. 6: Adverse Selection

But what if the insurers have no way of distinguishing individuals belonging to different risk classes? Suppose they offer insurance on a full-coverage-or-none basis, initially using the average probability of loss $\bar{\pi}$. The resulting expected utility of those purchasing the insurance, $v(\omega-\bar{\pi}\ell)$, is given by the distance DE in Fig. 6. The high-risk class would be getting a bargain, but the low-risk class may (as shown here) be better off without any coverage (A"B" > DE). If so, the latter drop out of the insurance pool; only the high risks insure. The premium in equilibrium would of course be $\pi'\ell$, reflecting the loss probability of the high-risk class.

This is the problem of adverse selection. While we have described it in the insurance context, it is a much more general phenomenon. Whenever buyers can only observe average quality, there is a tendency for sellers not fully rewarded for high quality to withdraw from the market. In one extreme model George A. Akerlof [1970] showed that even if the used cars in existence represent a merit continuum, only the lowest-quality "lemons" would actually be traded.

More generally, however, the equilibrium may not be quite so extreme. With somewhat greater risk-aversion, in Fig. 6 the v(c) curve might have warranted participation even of the low-risk class in the insurance pool, at prices based on the average odds $\bar{\pi}$. With a risk continuum also, the pool of participants may include everyone from the lowest quality up to some cutoff point [Richard Zeckhauser, 1974]. Returning to Fig. 6, suppose now that the risks are distributed over a range of loss probabilities, the interval $[\pi'',\pi']$. Since insurance yields an expected-utility gain of B'C' to the highest-risk class, those with not too dissimilar loss probabilities are also better off purchasing the same type of policy. Such participation lowers the average probability of loss and improves the actuarial premium, thereby drawing still better risks into the insurance pool. The process continues until the marginal risk class is just indifferent between no coverage and full coverage at a premium reflecting the average probability of loss for all those in the pool.

So far we have been assuming, in effect, a world of pure exchange: we have allowed people to trade risks (to engage in mutual insurance) but not to modify risks by productive activities. Such modifications might take the form of committing resources to loss reduction (reducing the gap between c_N and c_L) or to loss prevention (reducing the loss-probability π_L) [Ehrlich and Becker, 1972; Marshall, 1976] -- apart from or in addition to purchase of market insurance. More commonly the problem is viewed the other way. Might individuals who purchase insurance be inclined not to undertake protective measures that would reduce the scale or chance of loss? This is what is called moral hazard.

Since efficient prices are proportioned only to the probability of loss, if π_L is known and fixed there is no need for insurers to guard against inadequate loss-reduction activity by price-taking risk-owners; the decisions of the insureds lead automatically to a Pareto-efficient production and distribution of risks [Marshall, 1976]. And if π_L is variable but subject to costless monitoring by insurers, prices would respond appropriately and thus would continue to induce efficient loss-prevention activity as well as market risk-sharing transactions [Michael Spence and Richard Zeckhauser, 1971]. Another way of looking at this is to note that in principle, variation of probability of loss is equivalent to variations in the amounts of loss under a suitably extensive specification of states of the world [J. Hirshleifer, 1970, p. 217]. Thus, with perfect monitoring, insurance terms can always be based upon the fixed probabilities of the true underlying states. If insurers could offer suitably different premium/

indemnity ratios for losses under perfectly observable contingencies like a 10-foot flood, a 20-foot flood, etc., they need not be concerned with how high a dike the insured chooses to build.

Realistically speaking, however, monitoring of states will be imperfect. Insurers cannot be certain about the true flood hazards, fire hazards, medical hazards experienced by insureds. Consequently, very often contracts have to be written in terms of "result-states" [Marshall, 1976] -- the actual occurrence of loss -- rather than the true underlying states. Then, if in the extreme case price did not respond at all to loss-prevention activity, no such activity would be undertaken.

Insurers have two main ways of coping with the problem [Arrow, 1963; Mark Pauly, 1968; Richard Zeckhauser, 1970]. 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 be a proper fraction of the loss). Then insurance will be provided, but moral hazard persists in that insureds will engage in less preventive activity than would be efficient with costless monitoring. In addition, risk-spreading through insurance is less than ideal. The second way of coping is to price insurance in accordance with the actual loss-prevention behavior of insureds (the height of the dike built), the idea being that to monitor behavior may be easier than monitoring the underlying states. Again, as this process is subject to slippage and uncertainty, there is less preventive activity and less risk-spreading through insurance than would be optimal.

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

We have briefly alluded above to the possibility that, under uncertainty, a complete set of markets may not be available to economic agents. More formally, a regime of Complete Contingent Markets (CCM) will exist if, with S distinct states of the world, the S elementary state-claims c_s ($s=1,\ldots,S$) are all separately tradable. In such a regime each individual i with endowment $\omega^i \equiv (\omega_1^i,\ldots,\omega_S^i)$ and facing prices $P \equiv (P_1,\ldots,P_S)$ 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:

$$\frac{\pi_s v'(c_s^i)}{\pi_1 v'(c_1^i)} = \frac{P_s}{P_1} \quad \text{for all s.} \tag{1.9}$$

This is of course a generalization of equation (1.4) above. More generally, trading in any S distinct <u>assets</u> representing combinations or packages of the c_s-claims will also constitute a CCM regime provided the assets are linearly independent (i.e., that none of them can be expressed as a linear combination of the others). For, any desired vector c^i can then be attained by holding an appropriate combination of the S assets.

There is however a rather large gap between the CCM model and reality. Given the infinite variety of conceivable contingencies of economic interest (possible inventions, disasters, political developments, taste changes, etc.), in practice market regimes will necessarily be severely incomplete. Economic agents cannot in fact trade, directly or indirectly, in every distinct contingent claim. There are a number of different incomplete market regimes, some of which will be studied in more detail in Part 2. For example, it might be the case that for some or all commodities only <u>certainty</u> claims rather than <u>contingent</u> claims are tradable: one might be able to contract to deliver wheat, but there might be no effective market in wheat contingent upon the Republicans winning the next election.

We will consider here one interesting regime of incomplete markets: a "stockmarket economy." Here each individual i has an untradable endowment $\omega^i = (\omega^i_1,\ldots,\omega^i_S)$ plus endowed amounts of tradable <u>shares</u> $(\bar{\alpha}^i_1,\ldots,\bar{\alpha}^i_F)$ of the F "firms" in the economy. To each firm there corresponds a total vector of state claims $\omega^f = (\omega^f_1,\ldots,\omega^f_S)$.

If each firm's holding has market value V_f, the individual's decision problem is to choose a portfolio $(\alpha^i_1,\ldots,\alpha^i_F)$ subject to his marketable wealth constraint:

$$\Sigma_f \alpha^i_f V_f = \Sigma \bar{\alpha}^i_f V_f \tag{1.10}$$

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

$$t^i = \Sigma_f (\alpha^i_f - \bar{\alpha}^i_f)\omega^f \tag{1.11}$$

The individual then chooses a portfolio to maximize:

$$u(\alpha^i_1,\ldots,\alpha^i_F) = \Sigma_s \pi_s \; v(c^i_s) \tag{1.12}$$

subject to (1.10) and (1.12). 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:

$$\frac{\Sigma_s \pi_s \; v'(c^i_s) \; \omega^f_s}{V_f} = \lambda^i \quad \text{for all f and all i.} \tag{1.13}$$

This directly implies:

$$\frac{\Sigma_s \pi_s \; v'(c^i_s) \; \omega^f_s}{\Sigma_s \pi_s \; v'(c^i_s) \; \omega^1_s} = \frac{V_f}{V_1} \quad \text{for all f and i} \tag{1.14}$$

This relation, in comparison with (1.9), indicates an optimization constrained by the set of assets or claims packages (firms $f=1,\ldots,F$) through which trading may take place. Unless the set of tradable assets constitutes a Complete Contingent Market (which cannot be the case if $F < S$, or more generally if the F asset vectors fail to 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 all tradable securities are in effect perfect substitutes. There can be no gains from trade; each individual must therefore end up consuming his own endowment. However, it can be shown that the stockmarket economy is efficient in the restricted sense of achieving Pareto-preferred allocations of the tradable shares of different firms [Peter A. Diamond, 1967].

Two related questions have received considerable attention: (1) Will shareholders in general be <u>unanimous</u> in support of the firm's production decision, and (2) if so, will the optimal decision be such as to maximize the firm's <u>market value</u> V_f? In the simpler model of certainty choices, it is well-known that unanimous support for maximization of market value follows when a "separation theorem" holds. If there are perfect competitive markets and no technological externalities among firms, maximization of firm value implies that every shareowner's wealth and thus his consumption opportunities will be maximized. In the absence of the stated conditions, the separation theorem does not in general hold. For example, if the firm has significant monopoly power, the shareowner must balance increases in wealth against the loss he suffers as a consumer having to pay higher prices. And if differing shareowners have different tastes or endowments, the failure of separation will imply non-unanimity as well.

Very much the same holds for the firm's decisions in a world of uncertainty. In a stockmarket economy, in particular, shareowners will unanimously support value maximization if the firm's decision can have only a negligible perceived effect upon their marginal utilities in the different states. This condition will be violated if the firm can have a significant effect upon the aggregate supply of claims to any particular state s, akin to its having a degree of monopoly over c_s-claims. As an important special case, the condition will fail if the productive options before the firm enable it to create otherwise unavailable patterns (ratios) of state-claims. And again, unanimity fails if there is technological interdependence between this firm and any other firm, since there generally will be overlapping ownership between the two [Diamond, 1967; Steinar Ekern and Robert Wilson, 1974; Harry DeAngelo, 1977].

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 other applications of uncertainty theory such as share cropping and portfolio selection. A number of other significant applications can only be mentioned here: (1) optimal contracts between agent and principal, for example to elicit ideal performance on the part of corporate managers [Jacob Marschak and Roy Radner, 1972; Milton Harris and Artur Raviv, 1978; Steven Shavell, 1978; Steven N.S. Cheung, 1969; Theodore Groves, 1973; Armen A. Alchian and Harold Demsetz, 1972; Michael C. Jensen and William A. Meckling, 1976; Thomas S. Zorn, 1978]; (2) corporate finance, and in particular the balance between debt and equity funding [Franco Modigliani and Merton H. Miller, 1958; John Lintner, 1962; Jack Hirshleifer, 1966; Eugene F. Fama and Merton H. Miller, 1972, Ch. 4]; (3) optimal behavior and equilibrium with respect to accidents [William Vickrey, 1968; Guido Calabresi, 1970; William Baumol, 1972; Peter A. Diamond, 1974]; (4) the "value of life" appropriate for risk-taking decisions [Ezra J. Mishan, 1971; Richard H. Thaler and Sherwin Rosen, 1975; Bryan C. Conley, 1976; Thomas C. Schelling, 1968; Michael Jones-Lee, 1976; Joanne Linnerooth, 1979; Theodore Bergstrom, 1974]; and (5) choice of discount rate for public investment [J. Hirshleifer, 1966; Kenneth J. Arrow and Robert C. Lind, 1970; Agnar Sandmo, 1972; Martin J. Bailey and Michael C. Jensen, 1972; Joram Mayshar, 1977].

Part 2

THE ECONOMICS OF INFORMATION

In Part 1 individuals were limited to terminal actions, permitting them only to adapt to uncertainty. In Part 2 we examine the consequences of informational actions, which allow them to overcome uncertainty. Paralleling the sequence of topics in Part 1, Section 2.1 first analyzes the optimizing choices of the decision-making unit. Then Section 2.2 covers market equilibrium, and in particular the inter-related prior and posterior equilibria associated with the receipt of public information. This is followed in Section 2.3 by a discussion of the incentives to seek out private information, as in inventive effort. In Section 2.4 we examine the process by which information is revealed in market prices. Finally, a brief discussion of rational expectations and informational efficiency appears in Section 2.5.

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 <u>conclusive</u> as to 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 <u>posterior</u> probability $\pi_{s.m}$ assignable to state s after receiving message m equals the ratio of the <u>joint</u> probability π_{sm} of state s and message m both occurring, divided by the <u>prior</u> probability q_m of message m. Furthermore, using standard laws of probability, the numerator and 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 an individual tossing a coin 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). Let Tails come 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. In the prior distribution pictured, the bulk of the initial probability weight happens to lie toward the high end. But, the likelihood function indicates, some message (evidence) has been received that is much more likely if s has a small rather than a large value. The posterior distribution is a compromise or average 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:

$$\boxed{\text{Prior beliefs } \pi_s} \quad \longrightarrow \quad \boxed{\text{Message m and associated likelihoods } q_{m.s}}$$

$$\downarrow$$

$$\boxed{\text{Posterior beliefs } \pi_{s.m}}$$

The individual's <u>confidence</u> 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 we shall see in detail below, <u>greater confidence implies attaching lesser value to acquiring evidence</u>.

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,

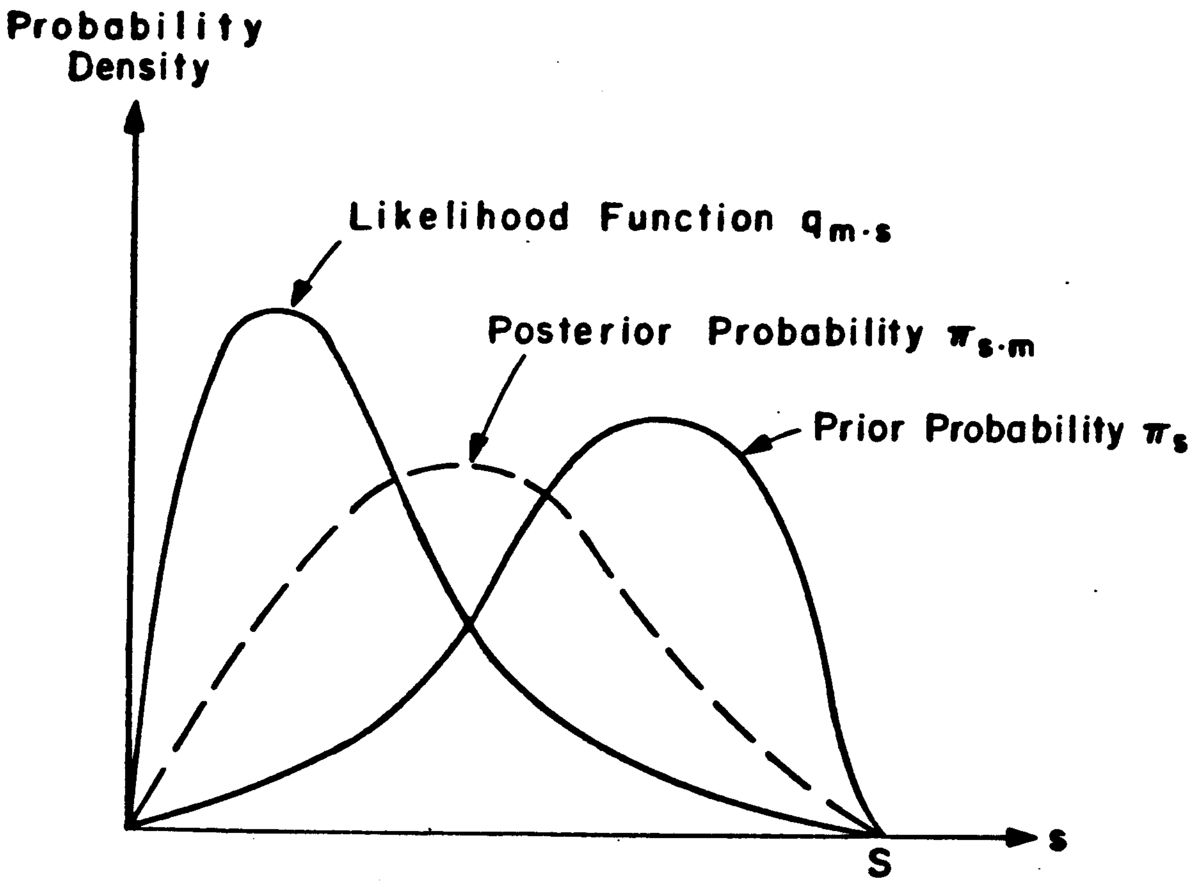

Fig. 7: Bayesian Probability Recalculation

if two tosses of a coin yield the message "Heads both times," then <u>given</u> the state of the world that the coin is fair the likelihood is $q_{m.s}=1/4$. The 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. 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 terminal actions consequent upon acquisition of information. The terminal-action decision problem of Section 1.1 above, for any given set of probability beliefs π, can be written as the instruction to maximize expected utility:

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

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 $\pi_{s.m}$ possibly leading to a new choice of terminal action a_m. Then Δ_m, the "value of the message m" can be written:

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

Note that Δ_m, which is necessarily non-negative, is an <u>ex post</u> 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 <u>ex ante</u>. 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 $m=1,\dots,M$ he will obtain. What the agent can actually purchase is not a particular message but an <u>information service</u> μ -- generating a probability distribution of messages m [Jacob Marschak and Koichi Miyasawa, 1968].

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 "objective" matrices Q_1 and Q_2 below representing the calculable likelihoods of the different messages given each state:

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

Note that a purveyor of an information service μ could not "objectively" characterize it simply by its probabilities q_m of generating the different possible messages, since -- as equation (2.3) indicates -- not only the likelihoods $q_{m.s}$ but also the "subjective" prior probabilities π_s are involved in determining the message probabilities q_m.

Given his prior probability beliefs the individual can then assign "personal" q_m probabilities in order to calculate <u>the value $\Delta(\mu)$ of an information service</u>, as an expectation of the values Δ_m of its associated messages:

$$\Delta(\mu) = E(\Delta_m) = \sum_m q_m[u(a_m,\pi_{s.m}) - u(a_0,\pi_{s.m})] \tag{2.6}$$

Since, as already indicated, each Δ_m represented by the 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 Fig. 8 the determination of $\Delta(\mu)$ is illustrated for a special case in which there are two states of the world (s=1,2), three available terminal actions (a=1,2,3) and an information service μ with two possible messages (m=1,2). 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 $v(c_{a1})$ attaching to the different actions if state 1 occurs are indicated by the points labelled $v(c_{11})$,$v(c_{21})$, and $v(c_{31})$ lying vertically above A in the diagram. Similarly, the utilities of outcomes in state 2 -- $v(c_{a2})$ for a=1,2,3 -- lie above point B. The expected utility $u(a;\pi)$ 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 (point F), 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 $\pi_{.m} = (\pi_{1.m},\pi_{2.m})$. If m=1, the revised choice of action in Fig. 8 is a=2 (point C) with Δ_1 (ex post utility gain over a=1) equal to the vertical distance CJ. If m=2, the best action in the diagram remains a=1 (point D),

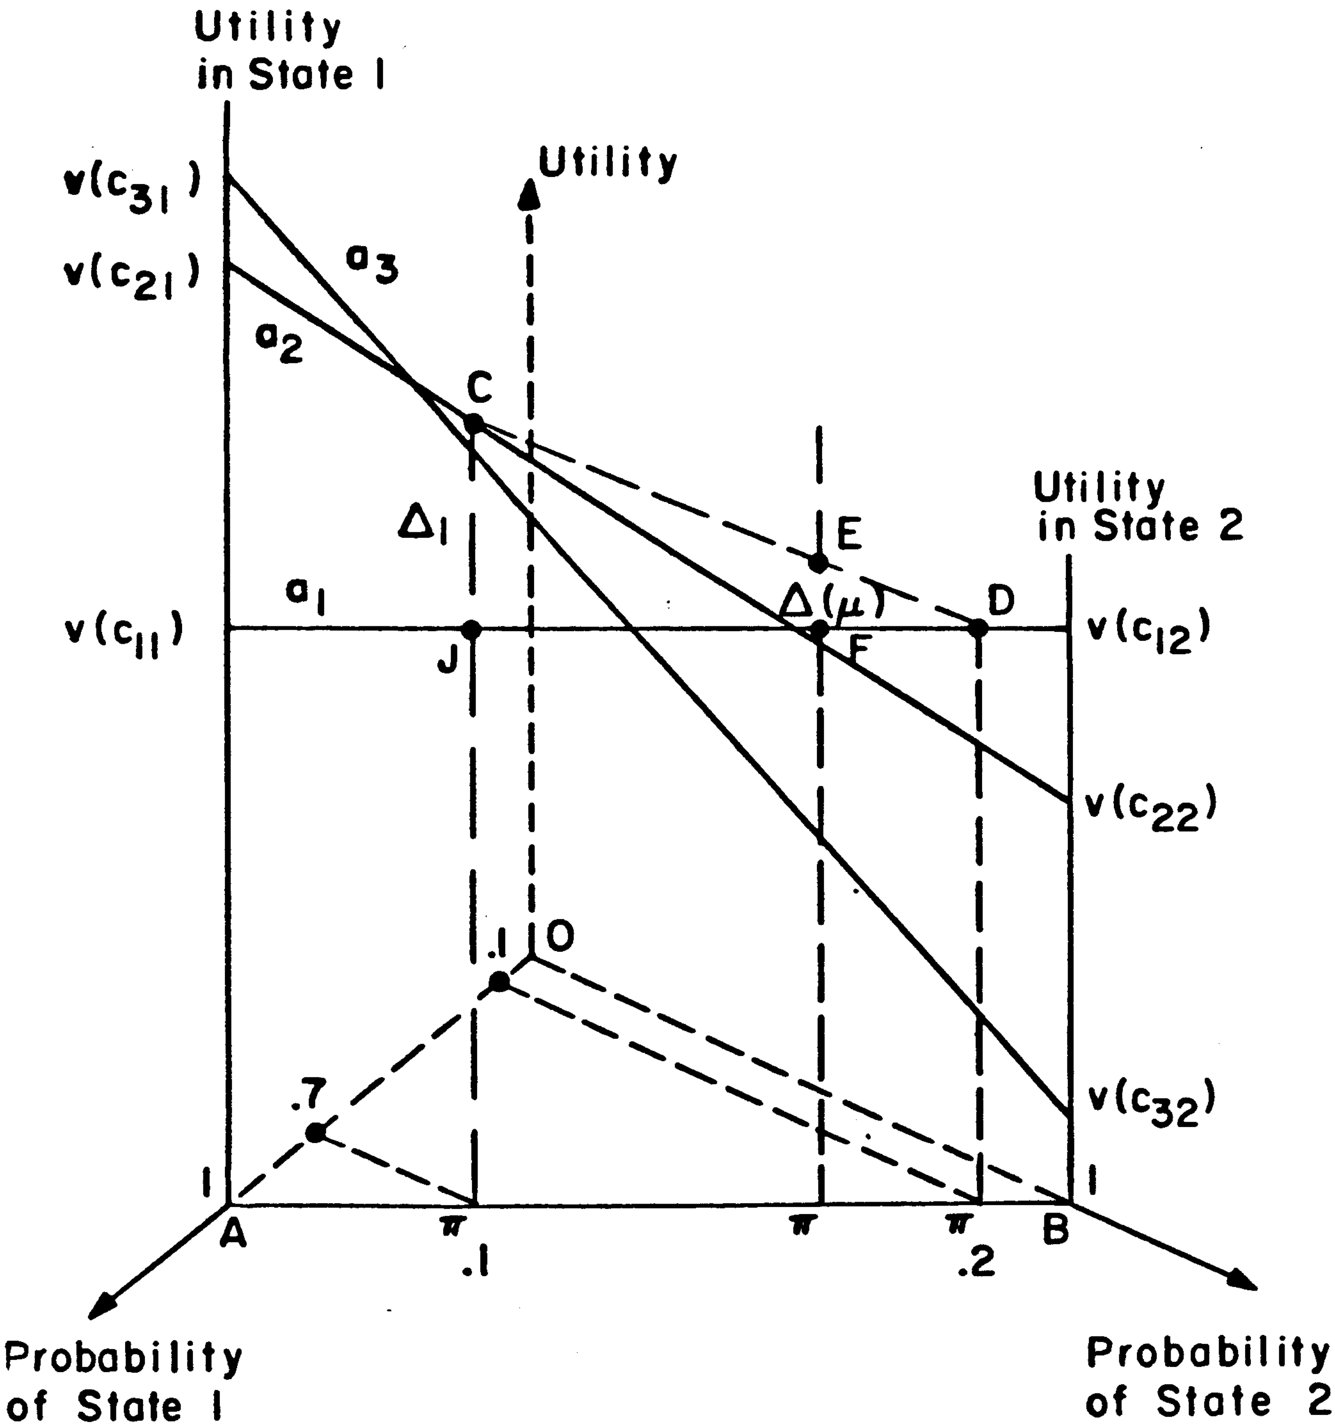

Fig. 8: The Value of Information

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 helps us see 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 (if any) 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 [Jacob Marschak, 1971]. 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 must be a convex combination of the posterior probabilities under the more informative service. For example, with two messages in each case, $(\hat{\Pi},\hat{q})$ is more informative than (Π,q) if for some α,β between zero and unity:

$$\pi_{.1} = \alpha\hat{\pi}_{.1} + (1-\alpha)\hat{\pi}_{.2} \quad \text{and} \quad \pi_{.2} = \beta\hat{\pi}_{.1} + (1-\beta)\hat{\pi}_{.2} \qquad (2.7)$$

One interpretation of these conditions is that the recipient of the information service (Π,q) knows that the true messages 1 and 2 would imply revised probabilities $\hat{\pi}_{.1}$ and $\hat{\pi}_{.2}$. However the messages have become garbled in transmission, so that he is not sure which message he has actually received. For example, his received message 1 has chance $(1-\alpha)$ of really being the true message 2.

These conditions are 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 $\hat{\pi}_{.1}$ 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. 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.

Fig. 8 also illustrates a general "non-concavity" (condition of <u>increasing</u> marginal returns) in the valuation of information services [Roy Radner and Joseph E. Stiglitz, 1976]. Starting the with 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 $\hat{\pi}_{.2}$ barely to the right. If the probability changes are small, neither message changes the associated best action and there can be no utility gain. So the marginal return of improved information will be zero over a certain range, before becoming positive at the point where the improvement begins to affect the actions taken. (If the action set were a continuum rather than discrete, the marginal return would never be zero but would nevertheless be increasing over an initial range.)

We now turn to some complications that arise when there are multipersonal aspects of the informational decision process.

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." It would seem that a decision-maker should take account of the opinions, whenever available, of all other parties who have some information not accessible to himself [Steven Shavell, 1976]. But an expert's opinion would generally represent his own posterior probability vector, which 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 might not 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. An additional complicating factor is that the expert may himself be affected by the action taken as a result of his provision of information. Unless his interest perfectly coincides with that of the decision-maker, the latter must take into account the expert's incentive to distort the information provided [Jerry Green, forthcoming].

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 [Howard Raiffa, 1970, Ch. 8, Part 2]. Any such arrangement 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 <u>conflicts of interest</u> (differences in utilities attached to consequences) or <u>conflicts of opinion</u> (differences in probability beliefs as to states of the world).

Consider first differences of belief only. Divergences of opinion commonly generate disagreement as to best <u>terminal</u> action; in consequence, there is a tendency for the group to choose "excessive" informational investment. 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, an agreement to opt for immediate terminal action might mask <u>compensating disagreements</u> as to the facts; here members of the group may agree <u>not</u> 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 <u>redistributive effect</u> 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 action 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 [J. Hirshleifer, 1973]. Advertising is an obvious example. There is also a choice between 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 those involved with acquisition or dissemination of information, that are indirect consequences of informational patterns. Wagering, for example, typically follows from differences of opinion in a situation where conclusive information is anticipated (so as to determine who wins and who loses). Speculation is a somewhat analogous activity, which turns on a prospective revision of prices 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.

2.1.3 Emergent Information and the Value of Flexibility

In the sections 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 automously 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 than it would autonomously arrive. But under conditions of emergent information one might choose a kind of indirect informational action -- adopting a flexible position and waiting before taking terminal action.

Suppose a choice must be made now between immediate terminal action and awaiting emergent information. This choice can only be interesting 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 [Kenneth J. Arrow and Anthony C. Fisher, 1974; Claude Henry, 1974] and in discussions of the value of "liquidity" [Jacob Marschak, 1949; Jack Hirshleifer, 1972] or of "flexibility" [Thomas Marschak and Richard Nelson, 1967; Robert A. Jones and Joseph 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 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}. Having initially chosen the "flexible" action a_1, 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

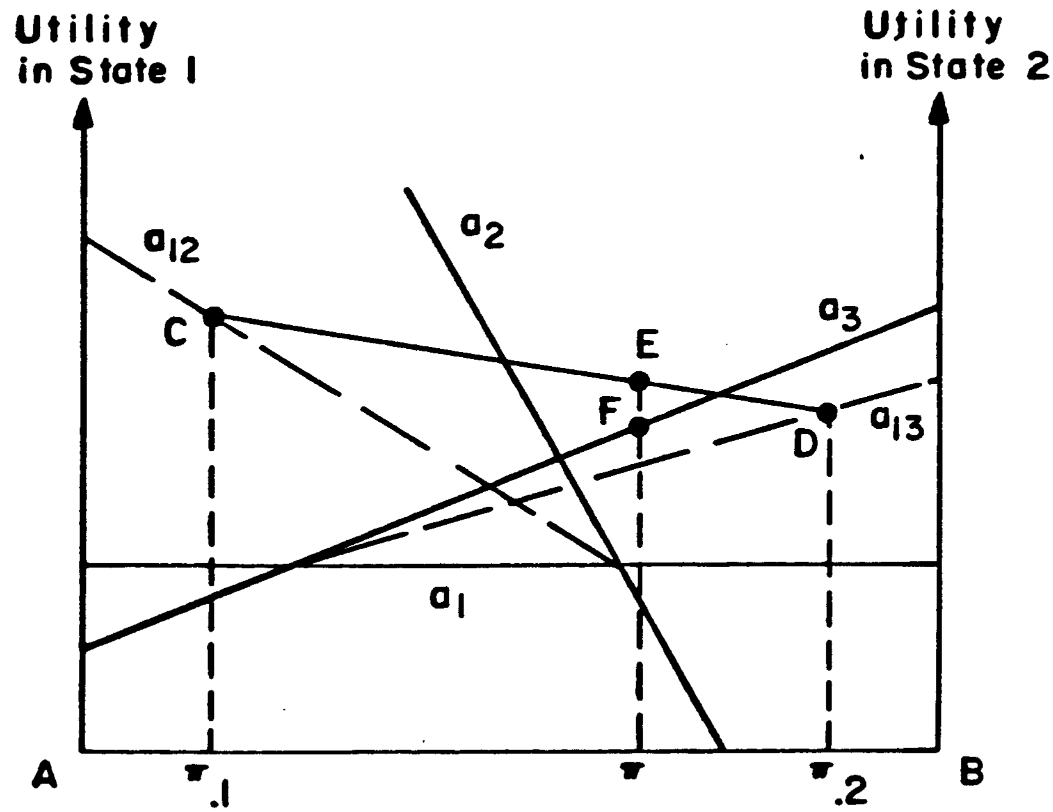

Fig. 9: The Value of Flexibility

pay a price to retain flexibility; the price is that you cannot do as well as if you had made the best choice among the "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 Public Information and Market Equilibrium

Emergence of new public information will affect prices. In particular, relative market values will rise for those assets paying off more handsomely in states of the world now regarded as more likely. The anticipated arrival of public information requires economic agents to contemplate market exchanges in two distinct rounds -- trading prior to, and posterior to receipt of the message. The equilibria of the two trading rounds will generally be inter-related, but the form of the relationship depends upon the completeness of markets (see Section 1.2.3 above) in each round.

2.2.1 Equilibrium in complete versus incomplete market regimes

In Part 1 we mainly considered models with S states of the world and a single consumptive good c. In the realm of the economics of uncertainty, where arrival of public information is not anticipated, a regime of Complete Contingent Markets (CCM) was said to exist if all the distinct c_s-claims (S in number) are separately tradable at prices P_s. (Trading in any set of S linearly independent asset combinations of the underlying c_s-claims would also constitute a CCM regime, but we will generally ignore this complication.) In equilibrium, the optimality condition (1.9) holding for each individual can conveniently be repeated here:

$$\frac{\pi_s \; v'(c_s)}{\pi_1 \; v'(c_1)} = \frac{P_s}{P_1} \; (s = 2,\ldots,S) \tag{2.8}$$

To achieve this condition, individuals will generally undertake productive transformations (e.g., loss-reduction and loss-prevention activities) as well as market exchanges (e.g., purchase of insurance).

Now consider that public information is expected to arrive before the close of trading. We will generally be assuming, however, that the message is not timely enough to permit productive adaptations to the changed probabilities. For example, a message as to increased flood danger comes in time to affect the market terms of flood-insurance transactions, but not in time to permit construction of dikes.

As an initial special case, assume it is known that the message will be conclusive as to which state of the world will obtain. Then to each and every state s corresponds exactly one message m. In these circumstances Complete Contingent Markets (CCM) in the prior round will, just as before, provide for separate trading in the S distinct c_s-claims at prices P_s. Under the timing assumption of the previous paragraph, posterior trading is in general possible. But such trading would be meaningless in this special case; once it becomes known that some single state s* is the true one, c_{s*} is the only state-claim retaining any market value, and there is nothing available for exchanging against it. So equations (2.8) in the prior round are the only relevant conditions of equilibrium.

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 G·S different claims c_{gs} -- claims to any good g under any state s -- at prices P_{gs}. After the conclusive message arrives that some state s* will obtain, complete <u>posterior</u> markets would permit exchanges among the G commodity claims c_{gs*}. But suppose for the moment that individuals were not aware, in their prior-round dealings, of this possibility of posterior trading. Then the prior-round optimality conditions would have included ratios of the following form, where g' and g" are any two goods:

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

Note that π_s, whatever its value may be, is not involved in the optimality condition between different goods <u>contingent upon</u> state s. It follows that receipt of the incoming message, revising the probabilities π_s (in this particular case, making $\pi_{s*}=1$) does not affect this condition. Therefore the price ratio on the right-hand-side of (2.9) continues to sustain the solution arrived at in the prior round.

Thus we see that even though posterior exchanges are <u>possible</u> among the G remaining tradable claims c_{gs*}, with CCM in the prior round no-one will find such exchanges <u>advantageous</u>. Trading in the G posterior claims is "not needed" if there have been markets for G·S prior claims. We must however emphasize a very important qualification to this result: prior-round traders must <u>correctly forecast</u> that the price ratio on the right-hand-side of (2.9) will remain unchanged in the posterior round. If they mistakenly thought that it would change, they would be led to make "erroneous" prior-round transactions, affecting the market equilibrium and thus requiring corrective posterior-round transactions. The result would be a loss of efficiency. So CCM in the prior round (without posterior trading) suffices for Pareto-optimality, but subject to a proviso of "correct conditional price forecasting" [J. Hirshleifer, 1977]. (This proviso corresponds to one of the meanings of that delphic phrase, "rational expectations," to be discussed in Section 2.5 below.)

We can now generalize still further to the case where messages are <u>not conclusive</u> as to the advent of any particular state. With Complete Contingent Markets in the prior round, there would be G·S·M distinct tradable claims c_{gsm}. And there remain G·S valid claims c_{gsm*} in the posterior round after receipt of message m*. But here also it is not difficult to verify that CCM in the prior round permits every agent to attain his optimum at one fell swoop -- provided once again that everyone correctly predicts that the relevant posterior-round price ratios are unchanged [George Feiger, 1976]. So quite generally, given "correct conditional price forecasting" and a CCM regime for prior-round trading, posterior-round markets are available but not necessary for Pareto-efficiency.

Very interesting and important issues arise, however, when we analyze prior-round market regimes that are <u>incomplete</u> (as of course they actually must be in the world). To maintain simplicity, we will however return to the particular case of <u>conclusive</u> information: then the set of M messages collapses into the set of S states so that individuals are concerned only with c_{gs} claims, G·S in number.

Among the many possible patterns of market incompleteness, three will be briefly discussed here.

Absence of prior-round markets:

Total absence of prior-round markets is of course the most extreme form of incompleteness. In effect, what has happened is that the information arrives before any exchanges have taken place, while individuals are still at their endowment positions.

This situation has aroused considerable interest, as it implies the surprising result that incoming public information may be socially disadvantageous in the sense that everyone in the economy might be willing to pay not to have it! [Jack Hirshleifer, 1971; John M. Marshall, 1974; Richard Zeckhauser, 1974; Nils H. Hakansson, J. Gregory Kunkel, and James A. Ohlson, 1979.] As among a group of traders who would otherwise have mutually insured against fire, a conclusive message (as to whose houses would actually burn down) would negate the possibility of insurance. The prospective arrival of such information prior to the opening of markets imposes an undiversifiable wealth-redistribution risk on the economy; no-one can hedge against the price impact of the message to be received. (On the other hand, if earlier arrival of information permits more effective productive adaptations, as in loss-reduction measures against fire, this socially valuable feature must be weighed against elimination of the ability to spread risks.)

Numeraire Contingent Markets (NCM):

Suppose now instead that there is prior-round trading, but only in contingent claims to a single commodity -- which might as well be taken as the numeraire good g=1. In the prior market, individuals cannot purchase claims to any good contingent upon state s, but can purchase claims to (say) corn contingent upon state s. Under this Numeraire Contingent Markets (NCM) regime, these purchases are in effect side-bets as to which state of the world is going to obtain, whose outcome will determine the individual's posterior wealth. After receipt of the message, of course, the individual will use his enhanced or reduced wealth to purchase a preferred consumption basket in the posterior round of trading.

Kenneth J. Arrow [1964] has shown that the same equilibrium allocation as indicated by conditions (2.9) under CCM (with prior trading in G·S claims) is achievable under NCM with prior contingent trading only in the S claims to a single commodity. But given the prior-round incompleteness of this NCM regime, the availability of G markets in the posterior round becomes now quite essential. And once again, we must also specify the important proviso of "correct conditional price forecasting". Furthermore, this proviso here becomes more stringent than under CCM, where the correct forecast was simply "no change" from the prior price ratio P_{gs*}/P_{1s*} to the posterior ratio $P_{g.s*}/P_{1.s*}$ (for the particular state s* pointed to by the incoming message). Here, under NCM, the correct price forecast is not in general computable from data available to traders in the prior round [Roy Radner, 1968; Jacques H. Drèze, 1970-1971]. The conditional relative supplies and demands determining the posterior ratios $P_{g.s*}/P_{1.s*}$ are not publicly "visible" in the prior round, where only claims to g=1 are being traded. Only the absence of utility complementarities between the numeraire and other goods could make the posterior prices computable.

Futures Markets (FM):

The CCM and NCM regimes both 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 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 current "futures" markets versus later "spot" markets.

Under such a regime of unconditional or, as we shall say, Futures Markets (FM) there will be just G tradable claims in the prior round followed by possible re-trading in the same claims in the posterior round. Since it is reasonable to assume that $G < S$ (there are many more conceivable contingencies than goods), it is evident that the G+G markets in two rounds under FM cannot in general achieve the same efficiency as the S+G markets under NCM (or, *a fortiori*, as the $G \cdot S$ plus G markets under CCM) [Robert M. Townsend, 1978].

This negative conclusion is somewhat mitigated, however, once we allow for the fact that emergent information in the world is only rarely conclusive. If improved though not yet conclusive public information emerges repeatedly, the multiplication of rounds of trading recreated after each informational input increases the effectiveness of FM relative to NCM and CCM. (Once again, the proviso as to correct conditional price forecasting retains its relevance, and is indeed increasingly difficult of achievement.) Also, in general multiple rounds can only partially offset market incompleteness, so that full Pareto efficiency is not achieved. More troubling, there may be different self-fulfilling predictions about prices in future trading rounds, leading to different equilibrium allocations -- and these allocations may be Pareto-rankable [Oliver D. Hart, 1975].

2.2.2 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 a 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.

Among the possible determinants of speculative activity, John Maynard Keynes [1930] and J.R. Hicks [1946] followed by many others have emphasized differential risk-aversion. In their view, in the prior round of a Futures Markets (FM) regime the relatively risk-tolerant speculators accept risks of price variability from relatively risk-intolerant "hedgers." In the prior trading round, speculators buy commodity futures, achieving on average a small gain (excess of later mean spot price over futures price) which represents the return they receive from suppliers unwilling to bear the price risk. For example, a wheat-grower hedges by accepting a firm price now from a speculator, both of them anticipating that the unknown spot price will on average be a little higher. Later developments along this line [Hendrik S. Houthakker, 1957, 1968; Paul H. 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, so that the 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 [Ronald I. 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 [J. Hirshleifer, 1977]. And these risks tend to be offsetting, reducing the need to engage in prior-round hedging activity. For example, when the crop of a representative wheat-grower is big (good news) he will find that the wheat price tends to be low (bad news) and vice versa. Since in a Futures Markets (FM) regime it is in general not possible to divest oneself of quantity risk (since only certainty claims can be traded), it can easily follow that individuals might find it preferable not to hedge against the offsetting price risk.

In a world where people have a spectrum of beliefs as to probabilities of future states, an individual's speculative activity proper (his adoption of a trading position in anticipation of arriving public information that will change market prices) depends in a rather complex way upon the degree of deviation of his beliefs from average opinion, upon his willingness to tolerate risks, and upon his endowment position conjoined with the trading limitations imposed by a regime of incomplete markets [J. Hirshleifer, 1975, 1977]. With regard to the first of these determinants, a speculator with strongly deviant beliefs thinks that others will be surprised by the incoming message (and thus will be forced to make unanticipated posterior transactions). He will in consequence have adopted a trading position enabling him to benefit from these transactions. With regard to the second determinant, degree of risk-aversion affects the scale of preferred speculative exposure. As to the regime of markets, we saw in the preceding section that re-trading possibilities are not needed under Complete Conditional Markets (CCM). Thus, incompleteness of prior-round markets is also a necessary condition for speculation [Stephen W. Salant, 1976; George Feiger, 1976].

2.3 The Economics of Research and Invention

Research and invention activities are prime instances of the informational actions studies, on the level of the individual economic agent, in Section 2.1 above. We are not dealing here with situations like those examined in Section 2.2, where public information simply emerges with the passage of time. Since Nature will not autonomously reveal her secret, it must be sought out by costly (generally private) 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.

In particular, as we shall see, successful private search generally leads to more or less universal dissemination of the discovery, with price impacts akin to that of the public information studied in Section 2.2.

The central problem considered by modern analysts [Fritz Machlup, 1968; Kenneth J. Arrow, 1972] has been the conflict between the social goals of <u>achieving efficient use of information once produced</u> versus <u>providing ideal motivation for production of information</u>. With regard to optimal use, already-produced information is a "public good" in the sense that its availability to 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 production 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 <u>exclusive</u> 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 uncertainty and unreliability in the legal protection of patents and copyrights, and even less protection for trade secrets not covered by patent or copyright. The result is that unlicensed uses often escape control. Short of ideal conditions there will be losses from both underproduction and underutilization, and in practice something of a trade-off: provision of greater legal protection to inventors tends to ameliorate the underproduction 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 [Yoram Barzel, 1968]. For concreteness, we can use as metaphor the "over-fishing" model of H. Scott 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 of fishing (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 of fishing -- for example by auctioning the right to exclusively exploit a fishery. Indeed, once rights of fishing are defined such as auction would tend to occur of its own, via Coase-Theorem negotiations. In the context of research, the equivalent distinction is between a right in an idea and a right of engaging in search for an idea. Again, one might imagine auctioning off the right of searching for an idea -- for example, the right of inventing an alloy with specified properties. As the lowest-cost inventor would bid highest for this right, the "rush to invent" problem would be solved [Steven N.S. Cheung, 1979a,b].

In research, the difficulty of defining the nature of an "uncaught" idea seems to make the assignment of rights of searching for them unfeasible. (Even in fishing, it is often impractical to define exclusive rights of hunting for such a wandering resource.) In contrast with fishing, however, property rights in ideas when caught are also 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. 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 <u>speculate</u> by taking long or short positions in assets whose values will be affected by the invention [J. 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 of 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 <u>pecuniary</u> 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 <u>overcompensation</u> 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 [Eugene F. Fama and Arthur B. Laffer, 1971]. There may be a technological benefit (improvement 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 to acquire new capital so as 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 not 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 <u>information leakage</u>. Second, at the disclosure stage he must be able to authenticate the information he is trying to publicize -- that is the problem of <u>signalling</u>. 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 message (that their product <u>is</u> high-quality) to buyers. Of course, any seller is motivated to <u>claim</u> 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 <u>which would not be rational for those selling lower-quality products</u>. Any activity is a potential signal if sellers of higher-quality products can engage in it at lower marginal cost (or higher marginal return) than producers of lower-quality products. For example, it has been argued [Phillip Nelson, 1974, 1975] that advertising tends to be especially advantageous for producers of higher-quality goods, in contexts where <u>repeat purchases</u> 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 in the advertising itself, a message is thus being conveyed: <u>that the product is worth promoting</u>.

For the labor market A. Michael Spence [1973], Joseph E. Stiglitz [1975], and John G. Riley [1976, 1979b] have argued that educational credentials 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.

Michael Rothschild and Stiglitz [1976] and Charles A. Wilson [1977] make parallel arguments for the insurance market. Section 1.2.2 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 large deductible is greater for low-quality risks (those with high loss probabilities). In effect, then, high-quality risks can signal by willingness to accept a big deductible. Insurance companies can thus screen for differences in risk by a menu of policies: some policies can be offered at the low premiums appropriate for high-quality risks but with large deductibles, and others with smaller deductibles but the steeper premiums appropriate for low-quality risks.

In contrast to the autonomously emergent information 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, 1979a]. 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.

Fig. 10 illustrates this for the simple case in which there are only two quality levels of the item for sale (insurance risk, consumer good, labor service, etc.) For concreteness we shall use the labor market interpretation. The dashed indifference curves $u_1(e,w)$ represent, for a low-quality worker, equivalent combinations of the price of his labor services w and the level e of signalling activity (education). For a high-quality worker the solid indifference curves $u_2(e,w)$ 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 is $\bar{\theta}$. Suppose that initially all workers are offered the same signal-wage pair $Z^* = \langle e^*,w^* \rangle$ as depicted in Fig. 10, where w^* is no greater than $\bar{\theta}$. Given the divergence in signalling costs, some firm would be motivated

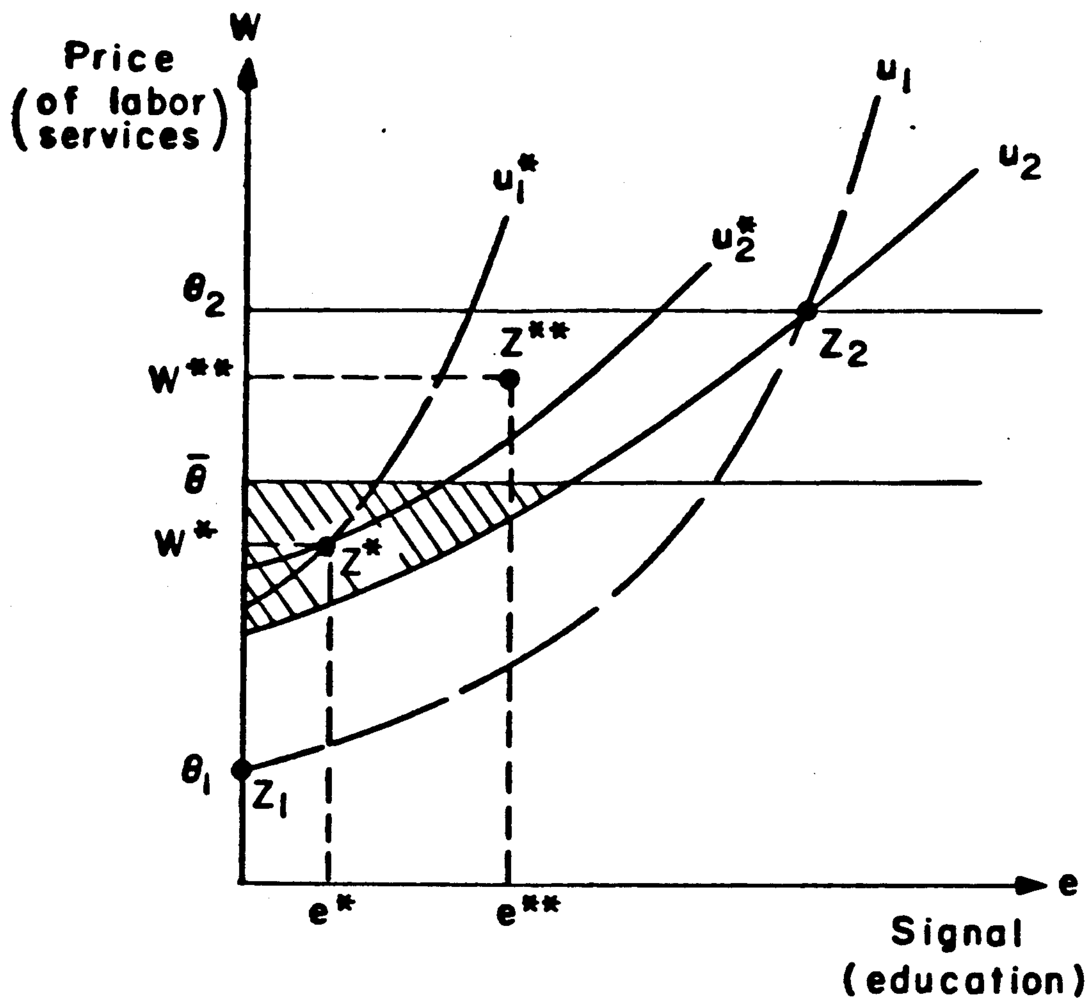

Fig. 10: Reactive Signalling Equilibrium

to make the new offer Z^{**}. This would attract only the high-quality workers and, since $w^{**} < \theta_2$, such an offer would generate a positive profit. Here pooling of different quality levels is impossible if the equilibrium is required to have the Cournot-Nash property that no ceteris paribus opportunities for gain exist.

Alternatively, suppose the two classes of workers are successfully separated with the pair of offers $Z_1 = \langle 0,\theta_1\rangle$ and $Z_2 = \langle e_2,\theta_2\rangle$ depicted in Fig.10. Acting as price takers, the lower-quality workers will accept the offer Z_1 and the higher-quality workers the offer Z_2. 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 $\{Z_1,Z_2\}$ is thus an equilibrium in the Walrasian sense (and in the sense of Spence [1973]). It is also efficient in the sense that all other pairs of offers (more generally all other wage schedules) which yield zero profit and which separate workers yield lower utility to the high-quality workers.

However, this equilibrium does not have the Cournot-Nash stability property. A firm can now enter offering the signal-wage pair Z^* which is strictly preferred by both classes of workers and yields an expected profit to the entering firm. As we have already seen that an offer like Z^* is not itself an equilibrium, there is no Cournot-Nash equilibrium.

How then would such a market behave? Plausibly, in the absence of collusion, each buyer would eventually expect some reaction 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 further 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: A set of offers is a reactive equilibrium if, for any additional offer which yields an expected gain to the agent making the offer 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. 10 $\{Z_1, Z_2\}$ 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 $w^{**} < \theta_2$.

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 results in a stable equilibrium.

2.4.2 Informational Inferences from Market Prices

We now consider the problem of information leakage. In Section 2.2.2 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 instead, everyone recognizes that some traders do and others do not possess an _informational advantage_. (Though traders with an informational advantage may not be 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 conclusive 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. Fig. 11 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 $E' = (\omega_1-k,\omega_2-k)$. 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 (point T_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

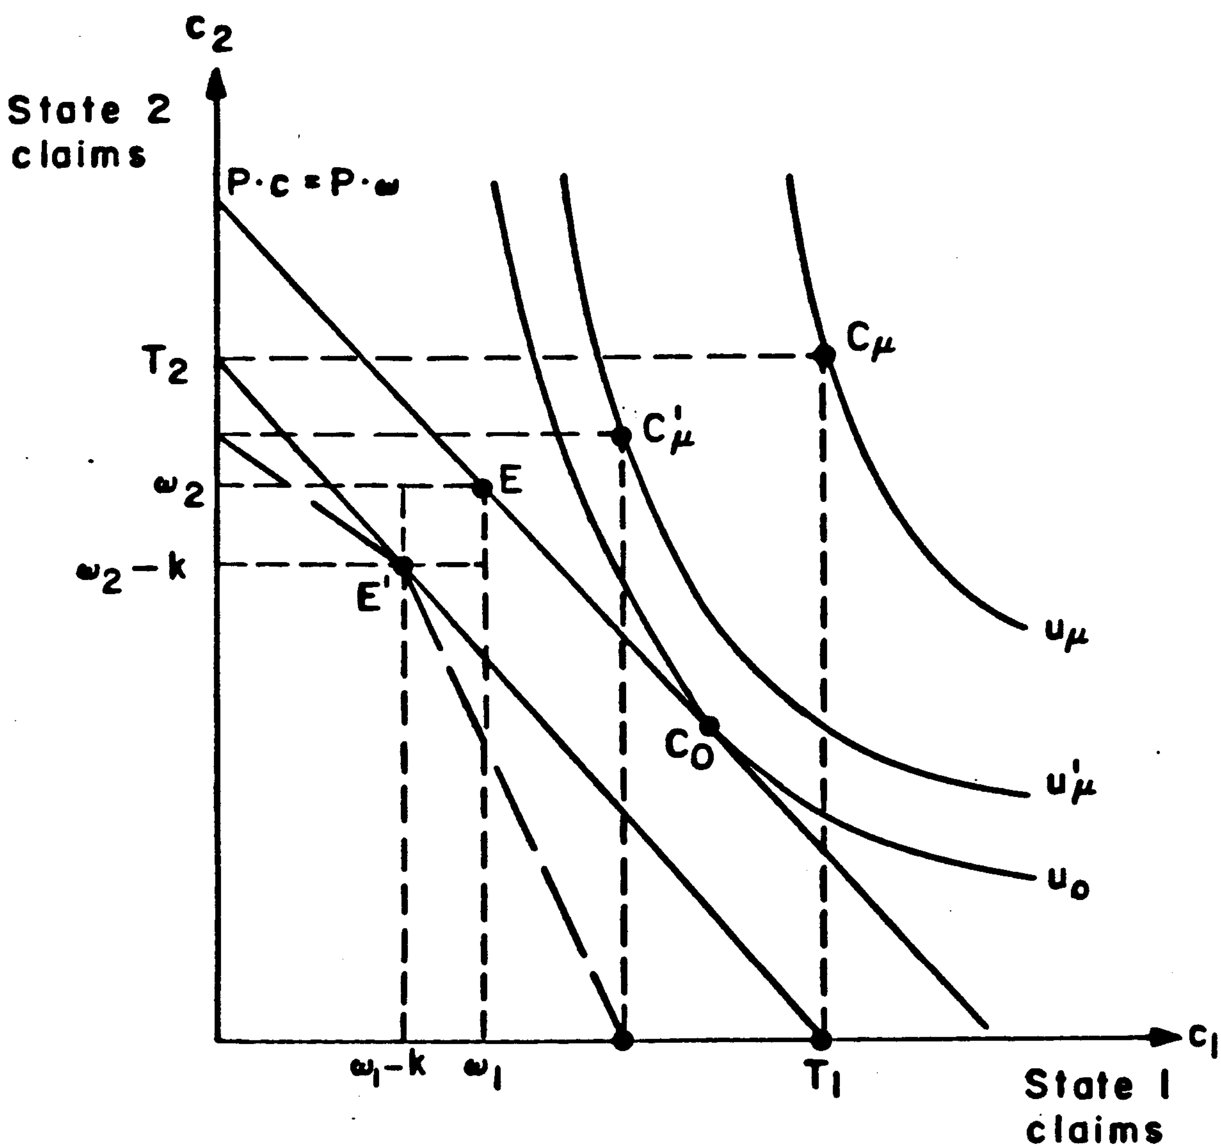

Fig. 11: Trading by informed agents

(point T_2). The final consumption vector C_μ yields him an expected utility gain of $u_\mu - u_0$.

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, purchase of information continues -- until the utility gain due to informed trading is exactly offset by the cost k of obtaining and processing the information.

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. So individuals not purchasing the information can infer it, simply by observing the movement of market prices! [Jerry Green, 1973, 1977; Sanford J. Grossman and Joseph E. Stiglitz, 1976]. They will therefore speculate in the same direction as those who have paid to become 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 all agents had all the available information I differs for each different I. Then the function p(I) is invertible, and $I = f^{-1}(p)$. That is, the information can be computed from the prices; or, the price vector p is a sufficient statistic for I [Richard E. Kihlstrom and Leonard J. Mirman, 1975; Grossman, 1977]. With a finite number of states, it is almost certainly the case that even in very incomplete markets the function p(I) is invertible [Roy Radner, 1979]. Thus there is almost certainly a "fulfilled-expectations" equilibrium in which each agent correctly infers aggregate information from the price vector p. (As in Section 2.2, however, this conclusion relies on a rather extreme "correct conditional price forecasting" assumption, in which each trader is able to compute the equilibrium price vector associated with each state of the world.)

As in the case of the signalling models considered above, there is a market externality here that tends 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 states [Grossman, 1977; James S. Jordan, 1976]), or if the informed individual can make his commitments before the uninformed can fully react, there will tend to be an equilibrium fraction of traders who choose to become informed.

2.5 Rational Expectations and Informational Efficiency

There is much confusion about both the logical meaning and the descriptive realism of the inter-related concepts of "rational expectations" and "informational efficiency."

The original idea of rational expectations is that anticipations "are essentially the same as the predictions of the relevant economic theory" [John F. Muth, 1961, p. 316]. This can be visualized in a very simple temporal model without exogenous uncertainty -- that is, a world where present and future endowed supplies, productive opportunities, utility functions, etc. are all perfectly determinate. Even though there can be no information input in such a world, there can be trading at different dates. And since the relative supply-demand situations for the various goods may change as the economy moves on its world-line, there is no reason to expect the spot price ratios to remain constant over time.

The role of traders' anticipations about the outcomes of future-dated markets (about future spot prices) is connected with the problem of incompleteness of market regimes discussed in Section 2.2.1 above. The relevant incompleteness here is the possible lack of futures markets -- of markets today for tomorrow's goods [Kenneth J. Arrow, 1978].

For concreteness, think of a two-period world. Complete markets at the current date (the analog of Complete Conditional Markets as discussed earlier, but in the absence of exogenous uncertainty there is only one state) would provide for trading in the 2G claims c_{tg} at prices P^0_{tg} -- where the first subscript represents the effective date of the claim (t=0,1) and the second designates the good (g=1,...,G), while the superscript indicates the <u>trading date</u>. In an incomplete market regime with <u>no</u> futures markets, the future-dated claims c_{1g} would not be tradable at t=0 so that the prices P^0_{1g} would not exist. (In either case, however, there could be later spot trading at t=1, at prices P^1_{1g}.)

In either market regime, individuals' trading decisions today will depend upon their anticipations as to the later spot prices P^1_{1g}, since there will in general be both productive interdependencies (e.g., storage possibilities) and utility complementarities between the two dates. And, in consequence, the equilibrium prices in today's and tomorrow's markets will generally both depend upon the anticipations that the individuals hold today. But, following a familiar theme of the previous discussion, a <u>complete</u> regime at t=0 makes "correct price forecasting" for t=1 easy: "no change" will be the correct prediction. That is, if all traders forecast that the <u>later spot</u> price ratios will equal the <u>current futures</u> price ratios -- $P^0_{1g'}/P^0_{1g''} = P^1_{1g'}/P^1_{1g''}$ -- and make their current trading decisions accordingly, their anticipations will be borne out.

It is thus the absence of futures markets that creates the forecasting problem. We have seen that, with incomplete markets today, individuals cannot actually compute (on the basis of information privately available to them) tomorrow's spot prices necessary to guide today's productive and consumptive decisions. The model of rational expectations nevertheless assumes that, at least on average, they <u>can</u> do so. Each person in effect makes a guess on the basis of his private bit of information (as well as his general knowledge of relationships like the law of supply and demand), and the errors of the various independent guesses balance out in the aggregate:

> ...allowing for cross-sectional differences in expectations is a simple matter, because their aggregate effect is negligible as long as the deviation from the rational forecast...is not strongly correlated with those of the others. [Muth, 1961, p. 319].

The central idea of rational expectations can of course be applied more generally than in our bare-bones illustrative example: there can be more than two dates, not all futures markets need be lacking, individuals can have subjective probability distributions rather than simple point estimates for the unknown future spot prices, and finally some agreed patterns of exogenous uncertainty might be introduced -- as in a random shift factor for future supply and/or demand. These generalizations lead to complications which cannot be pursued here. But even in the simplest version the assumption remains a strong one, whose virtue is in enabling us to close our intertemporal models and force out solutions. We shall not attempt to comment here on the descriptive validity of these models; just how validity might be tested is not at all evident, and has been the subject of controversy. One interesting point brought out by Arrow [1978] is that the rational-expectations assumption in effect stands on its head the famous argument by Friedrich A. Hayek [1945] about the informational function of the market system. Hayek's view was that market prices convey to traders all they need to know about the vastly detailed particular circumstances of other economic agents. Without a price system, a central planner would require an impossibly elaborate data-gathering and data-analyzing scheme to reproduce its results. But rational expectations implies that the price signals from the missing markets are not needed after all; individual traders can, at least on average, reproduce the missing signals on their own!

In these microeconomic applications, rational expectations essentially means correct prediction (at least on average) of the prices that will reign given different exogenous contingencies -- "correct conditional price forecasting." No restrictions are placed upon the beliefs of individuals as to the probabilities of the different states. In contrast, a central feature of macroeconomic applications [Robert E. Lucas, Jr., 1972; Thomas J. Sargent and Neil Wallace, 1975] has been that individuals are not only superior econometricians but clairvoyant about events as well. Frank H. Knight seems to have anticipated this view: "We are so built that what seems to us reasonable is likely to be confirmed by experience, or we could not live in the world at all" [1921, p. 227]. In the present context, it is supposed that in addition to being able to predict the effect of money supply upon prices, individuals can also decipher the actual monetary rule being followed.

This Lincolnesque idea that "the people can't be fooled" may be based upon viewing the underlying world process as stochastically stationary, so that individuals can gradually learn both about the effects of events upon prices and about the probability distribution of events. However, this learning evolution does not imply that beliefs would be on average correct except in the limit.

Prices in an economy will reflect the knowledge and beliefs of all participating traders, weighted by factors like endowed wealths, degree of risk-aversion, etc. An issue which is almost the obverse of rational expectations concerns what is called the "informational efficiency" of prices. Where rational expectations asserts that people can know not-yet-observed prices, informational efficiency asks if prices "fully reflect" people's current information [Eugene F. Fama, 1970].

Unfortunately, the meaning of the term "fully reflect" has proved elusive. Mark E. Rubinstein [1975] has proposed that information be said to be already reflected in prices if, upon arrival of the message, traders have no incentive to revise portfolios. On this definition, he shows that prices can almost never "fully reflect information." As follows from our Section 2.2 above, there are two main reasons: (1) even with agreed beliefs, incomplete regimes of <u>prior</u> markets generally make posterior trading (portfolio revision) unavoidable, and (2) with diverging beliefs and consequent speculative prior trading, it will be necessary to close out speculative positions in the posterior round.

In contrast, Sanford J. Grossman and Joseph E. Stiglitz [1976] have interpreted the "efficient market hypothesis" to mean that prices fully <u>convey</u> information. Informational efficiency then implies 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 Section 2.4, under these conditions the market for private information is not viable.

Perhaps the most widely held interpretation is that markets are "informationally efficient" if, as a pragmatic matter, there is no way to make a profit (more precisely, to achieve an expected utility gain) from information already in the public domain. In particular, no way to outsmart the stockmarket by detecting patterns of price movement in the publicly recorded data.

More specifically, this idea has been taken to imply that prices should follow a martingale process. That is, except perhaps for time and risk adjustment factors, the price today should be the mathematical expectation of the price tomorrow [Paul A. Samuelson, 1965]. However, it has been shown that even supposing known or agreed probabilities (without which it would be impossible to calculate mathematical expectations), only under very special conditions does a martingale in prices result [Susan E. Woodward, 1979]. First of all, prices are ratios; if a given ratio followed the martingale property, in general its reciprocal (the ratio taken the other way) would not. And even if expressed in terms of some standard numeraire commodity, prices would not follow a martingale unless there were no utility complementarities between the numeraire and other goods [Stephen W. Salant, 1976]. It has also sometimes been argued that failure of prices to follow a martingale would by definition create an arbitrage-like profit opportunity in the sense of a positive expectation of gain from holding an asset, over and above the normal interest yield. But, as we have seen from our early discussion of "private and social risk" in Section 1.2, a higher mathematical expectation of income does not in general represent higher expected utility, so this argument is erroneous.

Informational activities, finally, have an unusual relation to economic equilibrium. Information generation is in large part a disequilibrium creating process [Joseph A. Schumpeter, 1936 (1911)], and information dissemination a disequilibrium-repairing process. The two are intertwined, as we have seen, in very complex ways. It does not yet seem that we are very close to having an efficiency concept that can usefully be employed to measure the dynamically optimal level of such activities.

REFERENCES

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

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

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

Arrow, Kenneth J., "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. Apr. 1964, 31, pp. 91-96.

__________, "Aspects of the Theory of Risk Bearing, Helsinki, Yrgo Jahnssonin Saatio, 1965.

__________, "Economic Welfare and the Allocation of Resources to Invention," in The Rate and Direction of Inventive Activity: Economic and Social Factors, Universities - NBER Conference Series, Princeton, 1972.

__________, "The Future and the Present in Economic Life," Economic Inquiry, Apr. 1978, 16, pp. 157-169.

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

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

Bailey, Martin J. and Michael 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, Yoram, "Optimal Timing of Innovations," Rev. Econ. Stat., Aug. 1968, 50, pp. 348-355.

Baumol, William, "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, Theodore, "Preference and Choice in Matters of Life and Death," Appendix 1 of Jack Hirshleifer, Theodore Bergstrom, and Edward Rappaport, Applying Cost-Benefit Concepts to Projects Which Alter Human Mortality, UCLA School of Engineering and Applied Science, ENG-7478, Nov. 1974.

Borch, Karl, "Indifference Curves and Uncertainty," Swed. J. of Econ., Mar. 1968, 70, pp. 19-24.

Brainard, William C. and F. Trenery Dolbear, Jr., "Social Risk and Financial Markets," Amer. Econ. Rev., Proc. May 1971, 61, pp. 360-370.

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

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

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

__________, "Property Rights and Inventions," University of Washington Institute of Economic Research, Report No. 79-11 (1979).

__________, "The Right To Invest and the Right To an Invention," University of Washington Institute of Economic Research, Report No. 79-13 (1979).

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

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

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

DeAngelo, Harry, "Three Essays in Financial Economics," UCLA Ph.D. dissertation (1979).

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

Diamond, Peter 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, Jacques H., "Market Allocation Under Uncertainty," Eur. Econ. Rev. Winter 1970-71, 2, pp.

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

Ekern, Steinar and Robert Wilson, "On the Theory of the Firm in an Economy With Incomplete Markets," Bell J. of Econ and Mgt. Sci. Spring 1974, 5, pp. 171-180.

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

__________ and Arthur B. Laffer, "Information and Capital Markets," J. Bus. July 1971, 44, pp. 289-298.

__________ and Merton H. Miller, The Theory of Finance, New York, 1972.

Feiger, George, "What Is Speculation?" Quart. J. Econ. Nov. 1976, 90, pp. 677-688.

Feldstein, Martin 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.

__________, The Theory of Interest, New York, Macmillan, 1930.

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

Gordon, H. Scott, "The Economic Theory of a Common Property Resource: The Fishery," J. Polit. Econ. Apr. 1954, 62, pp. 124-142.

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

__________, "The Non-Existence of Informational Equilibria," Rev. Econ. Stud. Oct. 1977, 44, pp. 451-463.

__________, "Statistical Decision Theory Requiring Incentives for Information Transfer," in John J. McCall (ed.), The Economics of Information and Uncertainty, Universities - NBER Conference Series Volume, Chicago, forthcoming.

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

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

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

Hakannson, Nils H., J. Gregory Kunkel, and James A. Ohlson, "Sufficient and Necessary Conditions for Information to Have Social Value in Pure Exchange," School of Business Administration, University of California, Berkeley, June 1979.

Harris, Milton and Artur Raviv, "Some Results on Incentive Contracts With Applications to Education and Employment, Health Insurance, and Law Enforcement," Amer. Econ. Rev. Mar. 1978, 68, pp. 20-30.

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

Hayek, Friedrich A., "The Use of Knowledge in Society," Amer. Econ. Rev. Sept. 1945, 35, pp. 519-530.

Henry, Claude, "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, Jack, "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., Prentice-Hall, 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. 519-542.

__________, "The Theory of Speculation Under Alternative Regimes of Markets," J. Finance Sept. 1977, 32, pp. 975-998.

Houthakker, Hendrik S., "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, Michael C., "Capital Markets: Theory and Evidence," Bell J. of Econ. and Mgt. Science Autumn 1972, 3, pp. 357-398.

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

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

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

Jordan, James S., "Expectations Equilibrium and Informational Efficiency for Stochastic Environments," Univ. of Minnesota Disc. Paper #71, Aug. 1976.

Keynes, John Maynard, A Treatise on Money, New York, 1930.

Kihlstrom, Richard E. and Leonard J. Mirman, "Information and Market Equilibrium," Bell J. of Econ. and Mgt. Sci. Spring 1975, 6, pp. 357-376.

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

Linnerooth, Joanne, "The Value of Human Life: A Review of the Models," Economic Inquiry Jan. 1979, 17, pp. 52-74.

Lintner, John, "Dividends, Earnings, Leverage, Stock Prices and the Supply of Capital to Corporations," Rev. Econ. and 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. and Stat. Feb. 1965, 47, pp. 13-37.

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

Lucas, Robert E., Jr., "Expectations and the Neutrality of Money," J. Econ. Theory Apr. 1972, 4, pp. 103-124.

Luce, R. Duncan and Howard Raiffa, Games and Decisions, New York, John Wiley, 1957.

Machlup, Fritz, The Production and Distribution of KNowledge in the United States, Princeton, 1962.

__________, "Patents," Intnl. Encyclopedia of the Soc. Sci. New York, 1968, 11.

__________, "The Economics of Knowledge and Information," Part Seven of Knowledge: Its Creation, Distribution, and Economic Significance (Manuscript, 1979).

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

__________, "The Utility of Wealth," J. Polit. Econ. Apr. 1952, 60, pp. 151-158.

Marschak, Jacob, "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.

__________, "Economics of Information Systems," in M.D. Intriligator (ed.), Frontiers of Quantitative Economics, Amsterdam, 1971.

__________ and Koichi Miyasawa, "Economic Comparability of Information Systems," Int. Econ. Rev. June 1968, 9, pp. 137-174.

__________ and Roy Radner, The Economic Theory of Teams, North Holland, Amsterdam, 1972.

Marschak, Thomas and Richard Nelson, "Flexibility, Uncertainty and Economic Theory," Metroeconomica Apr.-Dec. 1962, 14, pp. 42-58.

Marshall, John M., "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, Joram, "Should Government Subsidize Risky Private Projects?" Amer. Econ. Rev. Mar. 1977, 67, pp. 20-28.

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

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

Merton, Robert C., "Investment Theory," in Handbook of Mathematical Economics, K.J. Arrow and M.D. Intriligator (eds.) North Holland, Amsterdam, forthcoming, 1980.

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

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

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

Muth, John F., "Rational Expectations and the Theory of Price Movements," Econometrica July 1961, 29, pp. 315-335.

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

__________, "The Economic Consequences of Advertising," J. Bus. Apr. 1975, 48, pp. 213-241.

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

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

Pratt, John W., "Risk Aversion in the Small and in the Large," Econometrica Jan-Apr. 1964, 32, pp. 122-136.

Radner, Roy, "Competitive Equilibrium Under Uncertainty," Econometrica Jan. 1968, 36, pp. 31-58.

__________, "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," Econometrica 1979.

__________ and Joseph E. Stiglitz, "A Nonconcavity in the Value of Information" (unpublished), July 1976.

Raiffa, Howard, Decision Analysis: Introductory Lectures on Choices Under Uncertainty, Reading, Mass., 1970.

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

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

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

__________, "Informational Equilibrium," Econometrica Apr. 1979.

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

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

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

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

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

Rubinstein, Mark E., "Securities Market Efficiency in an Arrow-Debreu Economy," Amer. Econ. Rev. Dec. 1975, 65, pp. 812-824.

Salant, Stephen W., "Hirshleifer on Speculation," Quart. J. Econ. Nov. 1976, 90, pp. 667-676.

Samuelson, Paul A., "Proof That Properly Anticipated Prices Fluctuate Randomly," Indust. Mgt. Rev. Spring 1965, , pp. 41-49.

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

Sargent, Thomas J. and Neil Wallace, "'Rational' Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule," J. Polit. Econ. Apr. 1975, 83, pp. 241-254.

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

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

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

Schumpeter, Joseph A., The Theory of Economic Development, Cambridge, Mass., Harvard University Press, 1936. (Original German publication, 1911).

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

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

Shavell, Steven, "On Valuable Opinion and the Efficiency of the Price System Under Uncertainty," Harvard Inst. of Econ. Resch., Discussion Paper #458, March 1976.

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

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

__________ and Richard Zeckhauser, "Insurance, Information, and Individual Action," Amer. Econ. Rev. May 1971, 61, pp. 380-387.

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

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

Thaler, Richard H. and Sherwin Rosen, "The Value of Saving a Life: Evidence From the Labor Market," in Nestor Terlecky, (ed.) Household Production and Consumption, Nat. Bur. Econ. Resch. Studies in Income and Wealth, 40, New York, 1975.

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

Townsend, Robert M., "On the Optimality of Forward Markets," Amer. Econ. Rev. Mar. 1978, 68, pp. 54-66.

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

Von Neumann, John and Oskar Morgenstern, Theory of Games and Economic Behavior, Princeton, 1944.

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

Woodward, Susan E., "Two Essays in the Theory of Competitive Markets for Contingent Claims," UCLA Ph.D. dissertation (1979).

Working, Holbrook, "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. 431-459.

Zeckhauser, Richard, "Medical Insurance: A Case Study of the Tradeoff Between Risk Spreading and Appropriate Incentives," J. Econ. Theory Mar. 1970, , pp. 10-26.

__________, "Risk Spreading and Distribution," in Harold M. Hochman and George E. Peterson (eds.) Redistribution Through Public Choice, New York, 1974.

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