Limited Knowledge and Economic Analysis

Kenneth J. Arrow

The American Economic Review, Vol. 64, No. 1 (Mar., 1974), 1-10.

Stable URL:
http://links.jstor.org/sici?sici=0002-8282%28197403%2964%3A1%3C1%3ALKAEA%3E2.0.CO%3B2-X

The American Economic Review is currently published by American Economic Association.

Limited Knowledge and Economic Analysis

By KENNETH J. ARROW*

The content of presidential addresses to this Association provides a fine example of a random variable with a high variance. It might even be a good subject for econometric analysis; the variation might be explained in terms of the economic conditions of the moment, previous intellectual investments, or even, for boldly interdisciplinary analysis, the psychological states or class origins of the speakers or the audience. But no doubt captious theorists like myself will object that the endogenous variable is not cardinally measurable and probably not even ordinally measurable; tougher-minded econometricians will worry about collinearity in the predetermined variables; and practical-minded policy analysts will see no discernible effect on the gross national product, the price level, or the balance of payments through effects on either fiscal policy or the stock of money. The last group, the policy-oriented, are perhaps the least accurate; at least according to Keynes, the effect of ideas on policy is dominant, though the lag may be as variable as and a good deal longer than that of the stock of money on money gross national product.

* Presidential address delivered at the eighty-sixth meeting of the American Economic Association, New York, New York, December 29, 1973. The views expressed above are a personal synthesis of a widespread viewpoint which has appeared in many different contexts in the published work of and in some cases personal discussions with G. Akerlof, A. Alchian, G. Calabresi, R. H. Coase, H. Demsetz, P. A. Diamond, J. Green, F. H. Hahn, A. G. Hart, F. A. von Hayek, J. Hirshleifer, L. Hurwicz, C. Kaysen, F. H. Knight, A. Leijonhufvud, F. Machlup, J. Marschak, T. Marschak, J. Mirrlees, M. Nerlove, R. Radner, M. Rothschild, T. Schelling, H. A. Simon, A. M. Spence, and G. J. Stigler. For some additional aspects of the role of information problems in economics, see my lecture, *Information and Economic Behavior*, Stockholm 1974.

It is now more fashionable than it used to be for statisticians to be told to take a good look at their data before fitting models. Taking presidential addresses as our data, we find most frequently a review of the speaker's main research concerns but also expressions of methodological or ethical concerns, historical surveys of varying degrees of erudition and humor, and, least frequently, new points of view on significant problems of economics.

I am taking a somewhat different tack today; it will be an expression of discontents and expectations. As I shall try to argue, the uncertainties about economics are rooted in our need for a better understanding of the economics of uncertainty; our lack of economic knowledge is, in good part, our difficulty in modelling the ignorance of the economic agent.

Critical aspects of this need for reorientation of theory have been recognized by many scholars in the last quarter-century and particularly in the last decade. I view my remarks today as a summary and perspective on a widely shared development of thinking.

The starting point of discussion must still be the much-abused neoclassical theory. No really cohesive alternative which aspires to the same level of completeness exists. The neoclassical model is founded on two concepts, which are considerably different in nature. One is the notion of the individual economic agent, whose behavior is governed by a criterion of optimization under constraints which are partly peculiar to the agent, such as production functions, and partly terms of trade with the economic system as a whole.

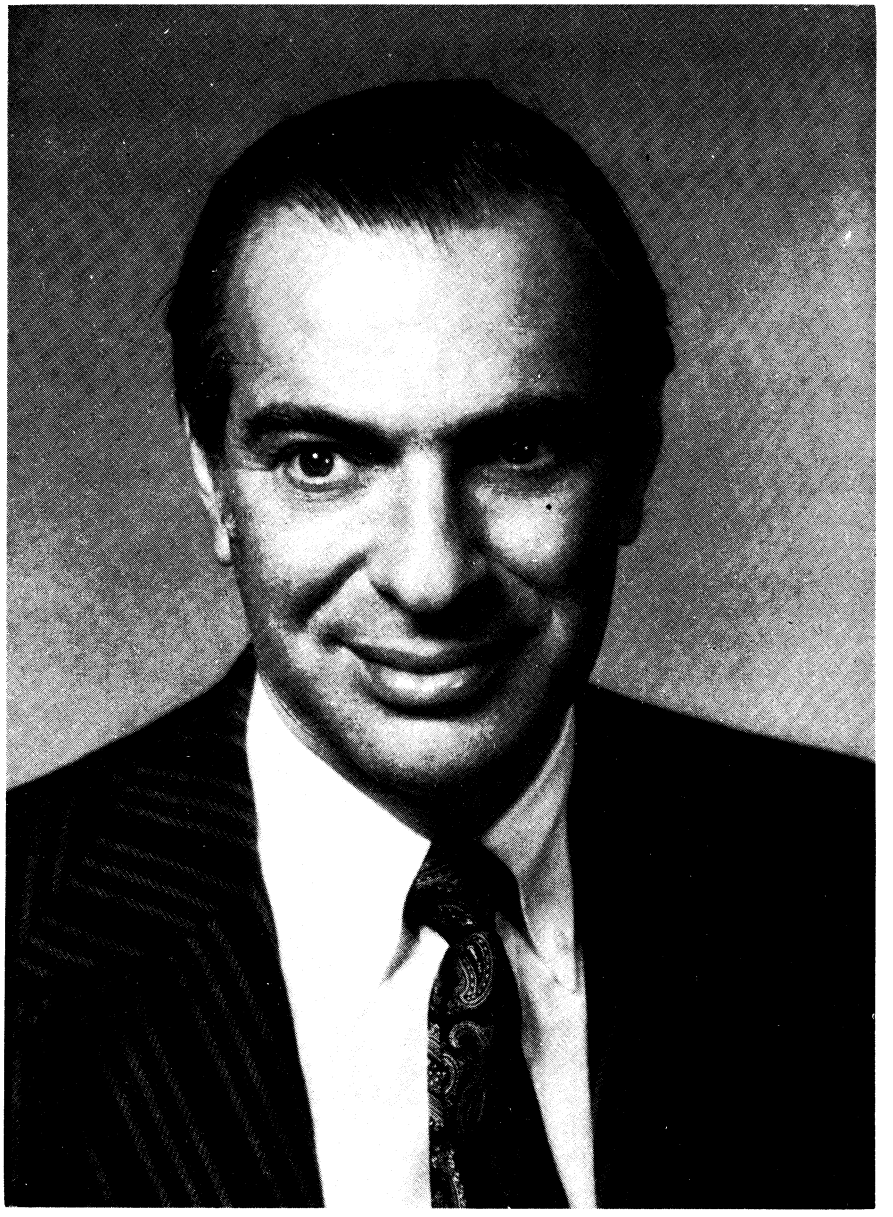

Kenneth J. Arrow

Number 75 of a series of photographs of past presidents of the Association

The other is the market; here, the aggregate of individual decisions is acknowledged, and the terms of trade adjusted until the decisions of the individuals are mutually consistent in the aggregate, i.e., supply equals demand.

The neoclassical theory, especially in its competitive form, can be and has been given a rich formal development. Parenthetically, one cause for the persistence of neoclassical theory in the face of its long line of critics is precisely that for some reason of mathematical structure, the neoclassical theory is highly manipulable and flexible; when faced with a specific issue, it can yield meaningful implications relatively easily. Although I intend to air complaints and desires for change today, I must express my unabashed admiration for the accomplishments of the neoclassical viewpoint. In its most formal statement, we simply use for analysis the equilibrium conditions of the individual agent and of the market, without inquiry as to how they come to hold. Yet even these statements turn out to yield revealing insights in the workings of resource allocation. Why have medical costs risen so rapidly relative to other prices since 1967? The upward shift in demand due to Medicare and Medicaid with a price-inelastic supply of physicians and hospitals provides a simple straightforward answer; I cannot really imagine how a Marxian or a neo-Ricardian would even approach the question, though I suppose they might dismiss it as unimportant. The explanation of environmental problems as due to the nonexistence of markets is similarly an insight of purely neoclassical origin. The now-demonstrated fact that flexible exchange rates are a feasible way of conducting international finance is a triumph of theoretical insights over practical men's convictions. More broadly, the shifts in long-run resource allocation as motivated by returns and, in particular, the absence of a secular trend in technological unemployment to the perpetual surprise of the layman fit in well with the neoclassical formulation but have no ready explanation in alternate models.

Of course, the implications of neoclassical theory have also been conspicuously falsified in important ways. Most notably, the recurrent periods of unemployment which have characterized the history of capitalism are scarcely compatible with a neoclassical model of market equilibrium. A post-Keynesian world in which unemployment is avoided or kept at tolerable levels by recurrent alterations in fiscal or monetary policy is no more explicable by neoclassical axioms, though the falsification is not as conspicuous.

Inequality in economic development among countries and among groups and regions within a country provides a second, somewhat complicated difficulty for neoclassical theory. A purely neoclassical answer would explain differences in per capita income by differences in physical and human assets per capita. This of course raises the further question, how this came to be, a question which would require a fully dynamic model to answer; but I think the more compelling problem is that the differences in income seem much too vast to be explained by factor differences. Indeed, in the presence of international trade and especially international capital movements, wage differences should be very strongly reduced compared with what would occur in autarchic states where domestic capital is the limiting factor. Hence, we come immediately to the explanation that there are differences in the production possibility sets of the different countries. This conclusion is a legitimate and important use of neoclassical analysis; but obviously it raises new questions, to which we will return.

I pass by the whole tangle of questions relating to the holding of money and the general level of prices. In its pure form,

neoclassical theory is a theory of relative prices. Monetary theories vaguely related to it in spirit can be grafted on to it, but none have succeeded in achieving a genuine synthesis.

The two failures of the neoclassical explanatory mechanism reflect on its foundations in quite different ways. The existence of unemployment is clearly a direct contradiction to the notion of the smoothly clearing market. One must of course be aware that the official measure of unemployment is by no means a simple inequality between supply and demand; it aggregates a whole range of distinct markets, it does not separate out voluntary and involuntary unemployment according to the tests of economic theory, and it does not take account of unfilled jobs. I do not subscribe, I hasten to say, to the sometimes expressed view that all unemployment is essentially voluntary, an unwillingness to search or whatnot; indeed, the official measure may underestimate the degree of disequilibrium in the labor market, particularly with regard both to underutilization of advanced skills and discouraged job seekers. With all these qualifications, it is clear that statistical unemployment does correspond to a disequilibrium as that term is used in the basic neoclassical model; there are two individuals, identical in productive capacity and both willing to work at a given wage, but one is working at that wage and the other is not.

Differential levels of economic development, on the other hand, point to a difficulty with the other fundamental concept; the conditions of optimization. If countries differ in their production possibility sets, then firms, occupying similar economic positions, are facing different constraints on their optimization. This does not contradict the fundamental assumption of optimizing behavior, but it does raise severe questions about its interpretation. The simplest hypothesis is to take the technological conditions as data, possibly varying over time due to exogenous changes in scientific knowledge. But here we are asserting that two contemporaries have different access to productive knowledge. Clearly, we are saying something about the conditions of transmission of knowledge across national boundaries, and of course the same questions arise among firms or workers within a single economy. The constraints upon the firm's optimization begin to seem more like variables to be explained than like constants exogenously given.

Let me look now at the two basic concepts from the inside, from the point of view of our direct perceptions which motivate the modelling. The two are far from parallel. The optimization by individual agents has a sense of concreteness about it, for all the sophisticated mathematical ability with which we theorists endow the agents. They behave in ways whose logic we understand. They seek to achieve goals which are reasonable to postulate, and we can specify constraints which clearly are real. It can be and has been correctly objected that our models are too simple; we ignore other arguments in the utility function, power, status, social approval, or whatnot that also motivate individuals, and we ignore some constraints, capacity for calculation and social controls. But the model is comprehensible, and the motives and constraints we deal with are real and important.

The market, on the other hand, is a much more ethereal construct. Who exactly is it that is achieving the balancing of supply and demand? Where in fact is the information on bids and offers needed for equilibration actually collected and stored? Right from the beginning of neoclassical theory, the difficulty of explaining markets in terms of individual self-seeking behavior was perceived. Parenthetically, this is one example of the superiority of neoclassical analysis to its predecessors, de-

spite the current fashion for exalting Ricardo over his successors; Ricardo implicitly equated supply and demand on all his markets without ever realizing the problematic nature of this process. Jevons felt obliged to enunciate explicitly a Law of Indifference, enforced by arbitrage; but this does not really meet the problem when the market is out of equilibrium, for the arbitrage might well not be feasible. Menger, at least according to Hayek and Streissler, concentrated on individual trades and ignored the market completely. It is Walras's auctioneer which has proved to have had the most enduring effect on subsequent theoretical development, and the stability theory which flows from that concept is still the subject of vigorous theoretical development, though very little empirical application. What is envisioned is a feedback mechanism in which errors in the price are successively corrected by reference to the disequilibria they generate. This view specifies and makes feasible the operations of the market. But on one hand the stability models are far from adequate representations even of the dynamics of the neoclassical models and, what may be connected, the results are by no means necessarily favorable to the stability of the adjustment process; and on the other hand, the motivations for the feedback to operate are obscure.

Let me be clear on one methodological point. The fact that our intuitive understanding, our *verstehen* as the German social methodologists call it, of the market as an institution is not entirely satisfactory does not mean that we should not use the perfect market as a model, at least pending further development. Certainly, as Popper and Friedman hold, the acceptability of a theory is to be judged by its ability to predict and understand phenomena. The theory of the perfect market is in an interesting way complementary to Keynesian theory. We have never been able to integrate Keynesian viewpoints into standard neoclassical theory, in terms of individual motivation, yet this theory, with its various modifications, has been a most serviceable tool of prediction and control. In fact, it is useful in domains where competitive theory fails and *vice versa*. Neither theory is good, however, at predicting dynamic processes, the short-run changes which are responses to disequilibria, and it is here that the pressure for a more satisfactory model arises.

Hold in abeyance for a moment our considerations about the market. Let us return to the optimization problem of the individual. One aspect on which we put a good deal of weight, particularly in our less formal discussions, is that a market system is informationally economical. That is, we tend to regard it as a virtue of the system that the individual agent need not know very much. Specifically, he is supposed to know the motivation and production conditions which define him, i.e., his utility function and production possibility set, together with the prices of the commodities he buys and sells. The economic system, taken as a whole, has vastly more in it than any one individual knows; it contains the utility functions and production possibilities of all individual agents. Indeed, the apparent modesty of the information needed is one of the most appealing aspects of the neoclassical model, both in the descriptive sense that the individual's decision problems appear manageable for him and for the economist studying him, and in the normative sense that the system permits its members to spend their time and effort at producing goods rather than in unnecessary duplication of information.

But clearly this simplification of the individual's decision making is made possible only because the markets have supplied the information economized on, in the form of prices. In equilibrium, at least, the system as a whole gives the impression of

great economy in the handling of information, presumably because transmission of prices is in some significant sense much cheaper than transmission of the whole set of production possibilities and utility functions. It is this point which emerged in the great debate over the feasibility of socialism begun by Ludwig von Mises's attack and usually thought of as concluding with the work of Oskar Lange and Abba Lerner in the 1930's; though it should be added many of the essential points had already been made earlier by Vilfredo Pareto and Enrico Barone. What was argued, in effect, was that a socialist system could use the price system and therefore achieve whatever economies in information it does achieve; and if the equilibrium conditions are written out they do give the appearance of relative simplicity. But what was left obscure is a more definite measure of information and its costs, in terms of which it would be possible to assert the superiority of the price system over a centralized alternative. Though I feel that current work has brought about a considerable clarification, we still have no definite measure. Indeed, in some respects, more recent developments have made the answers less clear. Several writers, in both Western and socialist countries, have noted that alternative decentralized schemes exist where quantity messages rather than price messages are transmitted in the successive stages of approximation and that such schemes also have efficient equilibrium points. Indeed, with the development of mathematical programming and high-speed computers, the centralized alternative no longer appears preposterous. After all, it would appear that one could mimic the workings of a decentralized system by an appropriately chosen centralized algorithm. While there is more to the story than these few remarks, they do make the point that if we are going to take informational economy seriously, we have to add to our usual economic calculations an appropriate measure of the costs of information gathering and transmission.

But actually the comparisons between socialist and capitalist resource allocation systems have tended to overlook some of the most obvious facts while examining finer points closely. As we all know, both production and consumption decisions are in fact made with reference to the future as well as to the present. A rational production plan includes very importantly decisions or at least plans about the future; and similarly with consumption plans. Investment and savings are not only integral parts of our current decisions but in the long run shape the possibilities for further development. As we know, the formal neoclassical model can be extended to decisions over time by dating commodities and regarding the same commodity at different dates as different commodities. All previous conclusions follow; allocative efficiency, for example, is achieved with the same appearance of informational efficiency.

But of course there is a slight problem with this reasoning. The information about future commodities needed includes their prices. These prices must be those found on a suitable market, one in which future supply and future demand are equated. Unfortunately, no such markets exist. Even the futures markets in certain commodities, limited in extent as they are, do not in fact lead to balancing *all* future decisions. Rather they balance present commitments to the future; but it is understood by all parties that when the future becomes the present, there will be a spot market on which the futures commitments may be undone; and indeed those making no futures commitments at all can enter and know now that they will be able to enter.

Even as a graduate student, I was somewhat surprised at the emphasis on static allocative efficiency by market socialists,

when the nonexistence of markets for future goods under capitalism seemed to me a much more obvious target.

However that may be, the nonexistence of these markets must be faced. Now in general equilibrium any part of the system affects every other part in at least two different ways. Thus, we may ask two questions about the nonexistence of futures goods markets: what are its implications for the rest of the system and what are the reasons for its nonexistence.

The implication first of all is that the information needed by the optimizer is not provided by an existing market. It will be provided by a market which will exist in the future, but that is a bit too late to help in decisions made today. Hence, the optimizer must replace the market commitment to buy or sell at given terms by expectations: expectations of prices and expectations of quantities to be bought or sold. But he cannot know the future. Hence, unless he deludes himself, he must know that both sets of expectations may be wrong. In short, the absence of the market implies that the optimizer faces a world of uncertainty.

The exact modelling of behavior under uncertainty is probably not crucial to the subsequent discussion; let us use the conventional expected-utility hypothesis. When there is uncertainty, risk aversion implies that steps will be taken to reduce risks. This partly affects decisions within the firm, such as the holding of inventories and preference for flexible capital equipment, and partly leads to new markets which will shift risks to those most able and willing to bear them, particularly through the equity market. The rich development of inventory theory and portfolio theory in the last twenty years or so reflects growing understanding of these matters.

But when we speak of expected utilities, we need to have some probabilities. Where do these come from? We may in the first instance regard them as subjective. But the economic agent observes his world and has the opportunity to learn from his experience, for there is a considerable degree of continuity. By Bayes' Theorem or perhaps psychological learning theory, the probabilities, say of future prices, will gradually adjust so as to conform to the facts. If indeed the economic world exhibited the same structure in some sense from period to period, and if everybody observed everything relevant, then the probabilities ascribed by different individuals to the same events might be expected gradually to converge to the correct values and therefore be the same for all. In fact, of course, the basic economic facts are changing, partly endogenously because of capital accumulation in its most general sense, partly exogenously with predictable and unpredictable changes in technology and tastes; equally if not more important, though, is the fact that the dispersion of information which is so economical implies that different economic agents do not have access to the same observations. Hence, it is reasonable to infer that they will never come into agreement as to probabilities of future prices.

A further implication is that the past influences the future. Jevons's well-known slogan, "bygones are forever bygones," ceases to be fully accurate. The past is relevant because it contains information which changes the image of the future; the probabilities which govern future actions are modified by observations on the past. It follows that present decisions with implications for the future are functions of past values of variables as well as present values.

This point of view has been exploited in the econometric models which have used distributed lags in explaining investment decisions. What still needs to be exploited more, however, is that the inference to the

future is necessarily uncertain, and the decisions made still exhibit risk aversion.

Expectations for the future are related to quantities as well as prices. The importance of quantity expectations has been stressed in macroeconomic models, even in such pre-Keynesian concepts as the acceleration principle, and most especially in relation to inventories. It sometimes is held that in a neoclassical world only prices matter; in the absence of prices, presumably they are replaced by price expectations. But that is not strictly true. Under constant returns, at least, quantity information for the individual firm is needed even when neoclassical assumptions are strictly fulfilled. Neoclassically founded investment theories usually predict capital-output ratios or capital-labor ratios; they still need output forecasts explicitly or implicitly. This gives considerable, perhaps major weight to past quantity information in predicting the future and therefore in guiding current investment decisions. It is perhaps along these lines that Keynesian theory, with its overwhelming emphasis on quantity changes as equilibrating variables, can be founded firmly on individual optimizing behavior.

I have referred to the fact that information is dispersed throughout the economy but have not suggested how. In the pure neoclassical model, each agent knows only his own production possibilities and his tastes, together with market information on the rest of the economy. In the world I have just sketched, however, any variables which improve his ability to predict the future have a very meaningful economic value to him. He will seek to acquire additional information. Such information is presumably costly; that is the basis for such great emphasis on the value of informational economy. But there is clearly a great incentive to acquire information of predictive value, and, as neoclassical theory would predict, there will be an incentive to produce such information. We have then an economic information industry: data assembly and analysis, business journalism, economic forecasting, with a longer-run perspective business education. Since information as a commodity does not satisfy all the neoclassical norms, it is not surprising that the government plays a large role in this process. Information-acquisition activities and information markets now appear on the economic landscape. Efficiency in the operation of firms ceases to be purely productive efficiency; it involves efficiency in prediction as well.

I would conjecture that the incomplete diffusion of information along the lines just sketched has a good deal to do with the operations of the securities markets and the decisions on corporate financing. The predominant role of internal financing and indeed the whole special importance of the managerial factor in corporate decision making are clearly connected with differential access to information about the firm.

You may have forgotten by now, but I earlier promised to consider not only the implications of but also the causes for the absence of markets for future goods. One might wonder why one should explain the absence of a phenomenon. Sherlock Holmes once maintained to the dimwitted local police inspector so typical of English detective stories that the significant question in the case at hand was the dog's barking at night. "But," said the inspector, "the dog didn't bark." "That," said Holmes, "is what is significant." So too is the absence of these markets significant for a full neoclassical theory. A truncated theory of temporary equilibrium in which markets for future goods are replaced by some form of expectations, themselves functions of current prices and quantities, has indeed been developed, though its empirical content is necessarily meager if the formation of expectations is left unanalyzed. But the

true neoclassical spirit is being denied in such a model. Although we are not usually explicit about it, we really postulate that when a market could be created, it will be. I sometimes think that welfare economics ought to be considered an empirical discipline. Implicitly, if an opportunity for a Pareto improvement exists, then there will be an effort to achieve it though some social device or another. In our theories and to a considerable extent in practice, the cheapest way in many cases is the creation of a market; and markets do emerge. If a market is impractical for one or another of the reasons we usually call "market failure," then very likely some other social device will at least be tried: government intervention; codes of professional ethics; or economic organizations with some power intermediate between the competitive firm and the government.

Thus, the failure of markets for future goods must be regarded as an analytic problem as well as a presupposition. It seems to me there are two basic causal factors. One is that contracts are not enforceable without cost and forward contracts are more costly to enforce than contemporaneous contracts; the other is that because of the many uncertainties about the future, neither buyers nor sellers are willing to make commitments which completely define their future actions. Let me take these two points up in turn.

The ability to make enforceable contracts is a necessary but not sufficient condition for a market. However, there is no way to insure complete enforceability. An individual may make a contract which he cannot in fact fulfill. Penalties may indeed be imposed on failure to live up to one's agreement, but they are not a substitute for compliance from the viewpoint of the other party, and there is always a degree of cost in enforcing the penalties. The laws of bankruptcy are a social recognition that complete enforceability is impossible and that it is even socially desirable to set limits on the penalties for failure. However, when the exchange of values for values is simultaneous or nearly so, the contracts may almost be self-enforcing. If a good has been sold and not paid for, it can be recovered; if there is a continuing relation of buyer and seller, a failure to settle bills can be met by refusal to make further deliveries, in which case the loss is minimized. With contracts extending into the distant future, on the contrary, the possibility of failure to comply becomes greater, partly because the self-enforcement aspects become weaker, partly because unexpected changes may intervene to make even a sincerely intended compliance difficult or impossible.

The outstanding examples of forward contracts are credit instruments. The buyer, who is taking the risks of default, is motivated to protect himself by seeking more information about the seller. The lender wants to know the borrower's assets, the prospects for changes in them, possibly even what he is going to do with the money. This very individualized information-seeking relation is quite far from the arm's length impersonal model of a market. The so-called capital markets are in many structural aspects very different from our model markets. It is of course an empirical question how far their behavior departs from the model. But the recurrent theme of credit rationing and availability doctrines, the essential imperfections of the credit market which underlie monetarist theories of cyclical fluctuations suggest that the incomplete enforceability of credit contracts and the protective steps taken by lenders are significant factors in explaining the working of the market.

While the enforceability question explains why those forward contacts that are made do not constitute a perfect market, we need more to understand why even

these are so limited in their coverage of future goods transactions. There are forward contracts in money, some commodities, real estate, but very little else. The explanation lies in uncertainties of both buyers and sellers about prices and quantities and about technology and tastes. Using uncertainties about prices and quantities as an explanation for market failure is a circular argument, though not necessarily a fallacy. That is, if all markets for future goods existed and cleared all transactions, then there would be no price-quantity uncertainties. But this much is true; if some markets for future goods do not exist, then the agents have uncertainties which are relevant to their behavior on markets for complementary or substitute goods. As Hicks showed a long time ago, complementarity and substitution can occur over time as well as simultaneously. If, as I will argue in a minute, uncertainty can tend to destroy markets, then we can conclude that the absence of some markets for future goods may cause others to fail.

To illustrate, the demand for capital goods at any point of time is dependent on the prices and sales of the product at future points of time. Therefore the demand for future capital goods will depend on expectations about the product at some still more removed time. If we assume only that we will not have markets for products at some distant point of time, then the resulting uncertainty will reflect itself in a failure of the market for capital goods in the nearer future, which will in turn create still further uncertainties.

Thus, if some markets for future goods are nonexistent, there will be uncertainties on the other markets; in addition, demand and supply conditions for the future are uncertain because of technological and taste shifts. Assume that both buyers and sellers are risk averters. Then without going into details it is reasonable to conclude that both demand and supply will have a downward bias as compared with the situation in which uncertainty is absent. A buyer will be unwilling to contract for purchase of a good if a superior or cheaper substitute may be available; and the seller will be unwilling to accept a price sufficiently low to be suitable to the buyer, particularly if he thereby precludes himself from a possible opportunity to shift his resources to other closely related goods. It would seem possible, at least, that there will be no price at which transactions in future goods will take place. From a theoretical viewpoint, one might say that the market is in a strange sort of equilibrium; there is some shadowy sort of price at which supply and demand are equated at zero. But this price is not performing much of a signalling function.

There is one ultra-neoclassical approach to the market treatment of uncertainties, in which I take some pride. That is the notion of a contingent market. Instead of letting uncertainty ruin existing markets, we can take it explicitly into account by buying and selling commitments to be carried out only if some uncertain event occurs. We could in principle imagine agreements to transact which will hold if and only a given conceivable technological innovation does not take place, with a second market for transactions valid if the innovation does take place. Then we can restore the possibility of markets.

Such contingent markets are not entirely unknown; insurance contracts are the purest example, and equity markets and cost-plus contracts provide more muddied illustrations. But they are relatively rare. Why this should be so follows again from the general problem of information costs and dispersal. If contracts are contingent on the occurrence of some event, then it must be verified whether or not the event occurred. But this is information, and as the example of a technological

innovation suggests, it is information likely to be much more easily available to one party than to the other. Hence, the range of possible contingent contracts becomes limited to those for whom the events are easily verifiable for both parties. The implications of these limits are known in the insurance literature as adverse selection and moral hazard, and they are of immediate practical significance in such matters as health insurance. But more broadly, they so limit the scope of contingent markets in practice that, as argued before in connection with markets for future goods, they prevent the emergence of even technically possible markets because of the large unresolved uncertainties.

I hope enough has been said to indicate the widespread implications of costly, dispersed information for the process by which future-oriented economic decisions are made. Let me remark, briefly in view of the length of time I have already taken, that informational costs and values play a key role in modifying the structure even of contemporaneous transactions. The individual optimizing agent is supposed to know at least his technology or tastes and the prices he faces. We have already argued a good deal of uncertainty with respect to the future economic implications of present economic choices. But in addition there is the possibility that technological information, which would be useful to him, exists somewhere in the world but outside his firm. There are grounds for engaging in the active pursuit of information. We begin to enter the realm of diffusion of innovations, to which some sociologists as well as economists have contributed. The interesting points here are the biasses in the information channels, some of which, at least, can be explained in terms of differential costs of acquiring information. For example, the well-documented role of personal influence in accepting innovations can be interpreted as due to a perceived high reliability of such information; in economic terms, this means more information per unit of expenditure of time or money.

The terms of trade with the outside world should not be regarded as freely given to the firm. In a world with a large number of commodities, even knowing the prices of relevant commodities involves the costly acquisition of certain kinds of information. This remark has given rise to a large literature on search in recent years. One implication which has been only slightly explored is that the concept of the market begins to weaken, and Jevons's Law of Indifference becomes more of an equilibrium condition than a statement valid about a market even in disequilibrium. At a moment of time, prices of what would usually be thought of as the same commodity bought or sold by different firms can differ because buyers or sellers may not, in their ignorance and in the presence of costs of search, find it worthwhile to shop further. Obviously, the important application of this principle may be to the labor market. Clearly, there are important informational differences between the employees currently working for a firm and potential substitutes elsewhere, although these are interchangeable in pure neoclassical theory. Indeed, there are differences both in the information the firm possesses about its employees as compared with alternatives and the information which employees have about the economic opportunities and the specific production conditions of the firm as compared with outsiders. It appears that considerations of this type must play some role in understanding the continued possibility of unemployment and particularly the sluggish response of wages to market disequilibria.

I am far from exhausting the implications of an information-economical viewpoint for the economic world. I look forward to exciting developments in the next decade.