1

The concept of rational expectations

Two divergent schools of thought originated at Carnegie-Mellon University in the late fifties and early sixties. During this period Herbert A. Simon had been refining his ideas on "bounded" rationality, a doctrine stressing man's* limited computational abilities in making decisions. At the same time, his colleague John F. Muth was working in another direction and developing the doctrine of rational expectations. Adding intrigue to this episode was the fact that Simon and Muth (along with Franco Modigliani and Charles Holt) were also collaborating on a book on inventory management and production control. Although the development of the diverse doctrines of bounded rationality and rational expectations by collaborators could be viewed as just a historical coincidence, it is more likely that an intense preoccupation with a common set of problems led the two researchers on different paths in search of a solution.

In their joint effort, Simon and Muth worked on problems of production scheduling and inventory management for the firm. These problems had been studied by many talented mathematicians, economists, and operations researchers who had developed general results for complex cases. The work of Holt, Modigliani, Muth, and Simon (1960) was aimed at deriving tractable, operational rules that could be easily applied in practice. As simplifying devices, they assumed that the various costs facing a company could be described by linear or quadratic cost functions. Why was this a useful simplification?

Simon (1956) had earlier shown that with quadratic costs and certain assumptions about the nature of uncertainty in the model, rules describing the optimal behavior for production and inventories for the firm would be linear functions of observable variables. In addition, in this framework, firms need to

* To avoid awkward wording, "man" and the pronoun "he" will sometimes be used generically, referring to both sexes.

consider only the *expected value* of future sales and could ignore other moments of the probability distribution of sales. This result, known as *certainty equivalence,* makes the calculation of inventory policies much easier by reducing the computational load for a typical decision maker.

Simon (1979) viewed the simplifications needed to obtain the certainty equivalence result in a positive light. He wrote, "the assumption of quadratic costs reduces the original problem to one that is readily solved. Of course, the solution, though it provides optimal decisions for the simplified world of our assumptions, provides, at best, satisfactory solutions for the real-world decision problem that the quadratic function approximates. In principle, unattainable optimization is sacrificed for in-practice, attainable satisfaction" (1979:499).

This approach, obtaining optimal solutions for a simplified world, is one strategy for easing the computational burden on decision makers. Another route is to give up explicit optimization and search for satisfactory solutions for a more realistic world. In Simon's terminology, both these strategies are examples of *satisficing* behavior, behavior that eschews complete optimization in favor of a more limited search for the best policy or decision.

Simon, therefore, viewed the certainty-equivalence results as a useful approximation for certain situations but did not elevate them to the plateau of a foundation for a general theory of economic behavior. In other environments different types of satisficing behavior might be more appropriate. In some situations the decision maker has to search out further alternatives rather than choose from among given alternatives. Because decision makers must decide when to choose, an extra level of complexity to problem solving is added. Certainty equivalence, therefore, was a useful, but by no means all-encompassing, paradigm for behavior.

John Muth viewed the problem somewhat differently. The certainty-equivalence results, which enabled the investigator to focus only on the *expected values* of variables that were uncertain, provided the key to an attack on another problem, which could loosely be termed the "interaction between expectations and reality." There were many examples of this problem in the literature at the time – perhaps the most well known concerned the "cobweb theorem" in agriculture.

The work of Nerlove (1958) and others stressed that farmers' planting decisions depended on the prices they expected to receive when the crop was marketed. In turn, the actual price for the crop depended on the amount finally harvested and the current level of demand. It was soon evident that different assumptions concerning the formation of price expectations could radically alter the actual price dynamics in the market. For example, if farmers based their price expectations on last year's price, there was a potential for dramatic instability in prices and production. Suppose that a spell of bad weather one year destroys part of the crop so that prices rise above normal. If farmers expect this high price to prevail, they will plant more than usual and when the resulting crop is harvested, prices will fall below normal. If this low price, in turn, is expected to persist, plantings will be lower than usual, leading to lower output and higher prices. Depending on the parameters of demand and supply curves, these price oscillations could either grow over time or dampen. Other price-expectation schemes would lead to different dynamic behavior for prices and production.

Another example of this phenomenon that was prominent in the literature of the fifties concerned the dynamics of hyperinflation. Cagan (1956) developed a simple model in which the velocity of money depended inversely on expected inflation, whereas expected inflation, in turn, was a function of past inflation. His model had the property that an autonomous increase in expected inflation would lead to an increase in velocity, which, in turn, would cause prices to rise. The increase in prices would then increase expectations of inflation, leading to a further increase in velocity. Depending on key parameters in the model, the burst of inflation caused by an autonomous change in expectations could either dampen or accelerate into a hyperinflation.

In both cases the explicit dynamics for prices depended on the precise nature of price expectations. Subsequent research demonstrated that these models were not robust to changes in the way that expectations were formed, and theoretical and empirical work rested critically on the exact specification of the price-expectation mechanisms.

Progress in economics, therefore, seemed to require a working, quantitative knowledge of how expectations of key variables were formed. Unfortunately, this verified theory of expectation formation neither existed at the time nor does it today. If anything, the psychological literature on expectations tends to suggest that people's expectations are intimately connected to their particular situation, and no general theory seems to work. Tversky and Kahneman (1974) presented evidence that decision makers are often subject to certain biases that naturally arise out of their circumstances. They argued, for example, that individuals, as a rule, judge distances by the clarity of objects in their field of vision. If objects are sharply perceived, then they are judged to be close at hand. This works well in most circumstances but, at times, can give very misleading results. Distances will be overestimated when visibility is poor, because objects will then be seen with less clarity. On the other hand, when visibility is exceptionally good, distances will be underestimated. Although these insights might be useful to predict behavior in certain situations, they are unlikely to become the foundation of a theory of aggregate behavior.

With the behavior of models so sensitive to the formulation of expectations, the lack of a general theory of expectations was an unsatisfactory state of affairs. It would be ludicrous for policy advice to be made contingent on what economists think the public believes.

In his path-breaking article "Rational Expectations and the Theory of Price Movements," John Muth (1961:315) recognized this problem. "To make dynamic economic models complete, various expectational formulas have been used. There is, however, little evidence to suggest that the presumed relations bear a resemblance to the way the economy works." Muth suggested that economists are often interested in how expectations might change in certain circumstances and thus should not be satisfied with fixed expectational formulas that do not allow for change when, for example, the structure of the system changes. If the underlying economic system changes, we would expect economic actors, at least after a certain amount of time, to change the way they form their expectations. Traditional models of expectation formation do not permit this adaptation.

Muth (1961:316) modestly put forward his hypothesis: "I should like to suggest that expectations, since they are informed predictions of future events, are essentially the same as the predictions of the relevant economic theory. At the risk of confusing this purely descriptive hypothesis with a pronouncement as to what firms ought to do, we call such expectations 'rational.'"

Muth noted that many economists, including Simon, thought that theories based on rational behavior were inadequate to explain observed phenomena. Muth argued the exact opposite point: Existing economic models did not assume enough rational behavior. One way to ensure this rationality was to insist that expectations of economic actors be consistent with the models used to explain their behavior. Muth's insight was that it was possible to require economic agents to form expectations of economic variables by using the very model that actually determined these variables. This ensured that the behavior of the model was consistent with individual actors' beliefs about the behavior of the economic system. Although this was Muth's basic point, a more formal and precise definition of rational expectations requires a review of several concepts from the theory of probability.

Conditional expectations and Muthian rationality

To understand the logic of rational expectations, one must be familiar with the concepts of probability densities, conditional probability densities, and the expectation operator. Although most of the economic examples use continuous random variables, for expositional purposes, it is useful to begin with discrete random variables. Let X be a random variable (such as a grade on a final exam) that can take on any of the values $X_1, \ldots, X_n$.

Let P_i be the probability that the random variable actually takes on the value X_i. The vector of probabilities $P_1, \ldots, P_n$ completely describes the information about the stochastic behavior of the random variable.

The expected value of a discrete random variable is the traditional measure of the central tendency of a probability distribution. It is defined as

$$\text{expected value} = E(X) = \sum_{i=1}^{n} P_i \cdot X_i \tag{1.1}$$

For continuous random variables, there are similar definitions. A continuous random variable can take on any value within a certain specified interval; for example, the amount of rainfall on a given day can be described as a continuous random variable ranging between zero and a positive number that reflects the heavens' rainfall capacity. A density function $f(X)$ describes the probability of various levels of rainfall. More precisely, $f(X)\,dX$ is the probability that the rainfall will be in the small interval dX around the level X. For continuous random variables, the expected value is defined as

$$\text{expected value} = E(X) = \int_a^b X f(X)\,dX, \tag{1.2}$$

where a and b are the lower and upper limits, respectively, of the random variable.

Conditional probability or conditional density functions are used extensively in the rational expectations literature. To understand this concept, imagine that a man approaches you on a train and wants to play a game of chance with a die that he removes from his pocket. You inspect the die and notice that it has six numbers and no obvious deformities. In your own mind you assign a probability of one-sixth to an occurrence of each number on the die.

Suppose, however, that an hour before the stranger approached you, a conductor on the train had wandered up and down the aisles warning the passengers about a stranger with a loaded die. In particular, the conductor told you the die came up on the number three half the time, the other numbers splitting the remaining probability. Warned with this information, you would now assign 50 percent probability to the occurrence of a three and 10 percent probability to the occurrence of each of the other numbers.

There are now two different probability distributions – one corresponds to the case in which you were warned and the other to the situation in which you were at the mercy of the stranger with the loaded die. In the former case your probability distribution can be described as *conditional* on being warned by the conductor. This is a simple example of the concept of a conditional probability distribution.

More generally, economic actors will make their probability assessments based on the information available to them at the time. Let I_{t-1} signify the information set that is available to economic actors at time $t - 1$. Then the notation $f(X_t|I_{t-1})$ will stand for the *conditional probability* density for the random variable X_t, given the information available at time $t-1$.

Corresponding to the conditional density is the *conditional expectation* defined as

$$\text{conditional expectation} = E[X_t|I_{t-1}] = \int_a^b X_t f(X_t|I_{t-1})\, dX_t. \quad (1.3)$$

The conditional expectation of a random variable is just the expected value of the variable formed by using the conditional density.

To understand some properties of conditional expectations, one should think of conditional expectations as forecasts of random variables. Associated with any forecast is a forecast error, ε_t, defined as

$$\text{forecast error} = \varepsilon_t = X_t - E[X_t|I_{t-1}]. \quad (1.4)$$

The forecast error has two important properties. First, the conditional expectation of the forecast error is zero. This follows directly from noting that, at time $t-1$, the conditional expectation (or the forecast) is known, so its conditional expectation is just the forecast itself. Using this fact, one finds that the conditional expectation of the forecast error is

$$E[\varepsilon_t|I_{t-1}] = E[X_t|I_{t-1}] - E[X_t|I_{t-1}] = 0. \quad (1.5)$$

The second property of the forecast errors is known as the *orthogonality property*. Forecast errors should not only have a zero expected value but also be uncorrelated with any information that is available to economic actors. If this were not the case, it would be possible to improve the forecast by incorporating this correlation into the forecast. To put it simply, an indication of a good forecast (and a property of conditional expectations) is that any subsequent forecast errors should be inherently unpredictable and hence unrelated to any information available at the time the forecast is formulated. Symbolically, the orthogonality principle can be expressed as

$$E[\varepsilon_t \cdot I_{t-1}|I_{t-1}] = 0. \quad (1.6)$$

Forecast errors derived from conditional expectations are uncorrelated with any information that is contained in the available information set. Shiller (1978) provides a further discussion of this important property.

Muth's rational expectations hypothesis essentially equates two distinct concepts; economic actors' subjective, psychological expectations of economic variables are postulated to be the mathematical conditional expectation of those variables. In other words, people's subjective expectations are, on average, equal to the true values of the variable.

This idea can be clarified by some notation. Let ${}_{t-1}X_t^{\varepsilon}$ be the subjective, psychological expectation for a variable X_t. Drawing on the previous notation, Muth's hypothesis asserts that

$$\text{subjective expectation} = {}_{t-1}X_t^{\varepsilon} = E[X_t | I_{t-1}]$$
$$= \text{conditional expectation.} \qquad (1.7)$$

Thus there is a connection between the beliefs of individual economic actors and the actual stochastic behavior of the system. This is the essence of the rational expectations approach.

To understand the implications of Muth's concept, one should distinguish between the problem of forecasting variables that are exogenous to the system and the problem of forecasting those that are endogenous to the system. For variables that are exogenous, forecasts or expectations by economic agents are important but do not affect the actual values of those variables. Exogenous variables are, by definition, those that are determined *outside* the existing system. On the other hand, expectations or forecasts of endogenous variables will affect the dynamics of the endogenous variables. The example of the farmers illustrates this point – their expectations about prices (an endogenous variable) affect the behavior of the endogenous variables (prices and quantities).

The hypothesis of rational expectations applies to both exogenous and endogenous variables but is most interesting for the latter. Expectations are rational if, given the economic model, they will produce actual values of variables that will, on average, equal the expectations. Expectations will diverge from actual values only because of some unpredictable uncertainty in the system. If there were no unpredictable uncertainty, expectations of variables would coincide with the actual values – there would be *perfect foresight*. The rational expectations hypothesis differs from perfect foresight because it allows for uncertainty in economic systems.

The example with the farmers should help to clarify the concept. A farmer with rational expectations performs the following thought experiment: What price can I expect so that if everyone else anticipates the same price, on average, we will all be correct? The farmer takes into account the anticipated supplies from his decisions and the decisions of similarly situated farmers, and calculates the price that will prevail if they all expect a particular price. In other words, they behave as if they possess a competitive, stochastic model of the market.

Farmers need not, of course, actually perform this thought experiment; as Milton Friedman has stressed, economic actors need only act *as if* they are maximizing utility or profits for our theories to work. Muth (1961:317) was quite explicit on this point: "[the hypothesis] … *does not* assert that the scratch work of entrepreneurs resembles the system of equations in any way."

It is not necessary for individuals to have identical expectations for economists to use the rational expectations hypothesis. Muth argued that individuals' expectations should, however, be distributed around the true expected value of the variable to be forecasted. In this way the average of individual forecasts would be the expected value of the true variable, although individuals could certainly differ in their beliefs.

Once the fiction of identical individuals is abandoned, it is possible to take another perspective on the rational expectations hypothesis and consider it from the point of view of arbitrage. In ordinary markets we do not require that all individuals respond to price signals in order to maintain a vibrant price system. Instead, we recognize that only a handful of individuals is required to arbitrage markets and ensure that, for example, coffee sells at the same transportation-cost-adjusted price in all locations. The same principle applies to the rational expectations hypothesis. If there is any economic profit to be gained from gathering and analysing information in order to predict the future, we would expect some individuals to pursue this strategy. If enough arbitrage activity takes place, the market might behave as if it is rational, even though many individuals in the market are simply passive.

The arbitrage perspective leads quickly to two points: (1) It suggests that the rational expectations hypothesis could be especially applicable to markets (such as financial markets) in which arbitrage is relatively costless. (2) It suggests that survey evidence that measures *average* rather than *marginal* beliefs or behavior might provide a misleading perspective on the applicability of the rational expectations hypothesis. The rational expectations hypothesis is not synonymous with arbitrage – it is certainly possible for a market to behave in a manner consistent with the rational expectations hypothesis even when arbitrage activities are costly. When arbitrage activity is relatively costless, however, this perspective provides an additional rationale for using the hypothesis.

The certainty-equivalence proposition was used extensively in Muth's article. Muth made the necessary assumptions to ensure that rational economic agents, for the most part, need concern themselves only with the mean or expected value of future variables and not worry about higher moments of the probability distribution, such as the variance. Muth, of course, recognized this as a simplification and thought that it was simply the price for obtaining some interesting and meaningful results. Subsequent work in the seventies (e.g., Lucas [1978a]) illustrated that it was possible to build models in which people's entire subjective-probability distributions would coincide with the true objective-probability distributions governing the system. This is the most general statement of the rational expectations hypothesis: The subjective-probability distributions of economic actors equal the objective-probability distributions in the system.

It is clear from the structure of Muth's (1961) article – theoretical analysis followed by some empirical applications – that he viewed the rational expectations hypothesis as a hypothesis of *positive* economics. The hypothesis could be wrong – it might be the case that expectation schemes do not have any consistent properties across economic models. If they do, however, powerful analytical results might be possible. In the sense that Muth was just proposing a hypothesis, his work could strike some as less dogmatic than Simon's. Simon's contention that expectations cannot embody too much rationality is as testable as Muth's conjecture that they might be more rational than commonly believed. Before putting the shackles on individuals' expectations and condemning them to a world of limitations and bounds, it might be useful, Muth argued, for economic science to explore precisely the opposite alternative.

A priori critiques of Muthian rationality

Much of the criticism and hostility toward Muth's concept of rationality that surfaces in some quarters can be traced to two factors. Sometimes the hypothesis has been employed in models that are thought to be simplistic. This, however, is not a criticism of the rational expectations hypothesis per se as much as it is a dissatisfaction with certain types of model building. On the other hand, some writers have rebelled at the idea that Muth's definition is the best, or even an adequate, definition of true rational behavior. The rational expectations hypothesis has been criticized for being, among other things, inconsistent with the subjectivist view of probability; an inadequate description of procedural rationality; and not a sufficiently general hypothesis to include learning and adaptive behavior.

The subjectivist critique of the rational expectations hypothesis, as outlined in Swamy, Barth, and Tinsley (1982), focuses on the central role that objective-probability distributions play in the theory. Subjectivists recoil at the notion of a true or objective probability distribution apart from the beliefs of particular agents. From the subjectivist point of view, probability beliefs are essentially the "bets" that an individual would be willing to make about the occurrence of a set of events. As Savage (1954) has shown, an internally consistent subjective probability measure can be derived for the individual if his choices among hypothetical bets satisfy certain axioms. Although different individuals might be willing to make different bets, as long as the bets of each individual satisfy the axioms, each individual can be said to be behaving in a consistent, coherent, and even "rational" manner. There is no need for individual probability beliefs to coincide with either each other or some outside "objective" standard. Thus the subjectivists would argue that there already exists a standard for rational behavior that bears little resemblance to the rational expectations hypothesis.

Robert E. Lucas, who engineered the rational expectations "revolution" in macroeconomics, has defended the rational expectations hypothesis from the subjectivist attack. Lucas did not dispute the subjectivist or Bayesian theory of the foundation of probability. However, he argued that

> the general hypothesis that economic agents are Bayesian decision makers has, in many applications, little empirical content: without some way of inferring what an agent's subjective view of the future is, this hypothesis is of no help in understanding his behavior. Even psychotic behavior can be (and today, is) understood as "rational" given a sufficiently abnormal view of relevant probabilities. (1977:15)

To employ economic theory in practice, Lucas argues, one must know what probability distributions agents actually use.

In situations in which it is possible to observe the frequencies of different events, it might make sense to assume that agents' subjective probability distributions mirror the observed frequencies. Lucas concedes that there could be situations in which observed frequencies provide little guide to decision makers. In these circumstances Lucas argues that *no* economic reasoning will be of any value. In short, economic science requires some theory about the content of subjective-probability beliefs; when it is possible, observed frequencies can be a useful guide to these beliefs. Predictive theories cannot be easily built on the principle that agents have subjective-probability distributions that cannot be related to objective events.

One of the implications of Muth's definition of rational expectations is that agents' expectations about variables should change when the conditional probability distribution governing the variables changes. This aspect of the theory has been challenged by Peter Rappoport (in an unpublished paper, "Rational Expectations and Rationality") on the grounds that it is inconsistent with a broader definition of rationality.

Rappoport takes as his paradigm for rationality the textbook story about the logic of scientific inference. According to this view, existing hypotheses are maintained until some evidence dictates that the hypothesis be rejected. In this spirit Rappoport argues for a reliance on classical hypothesis testing to determine if a particular expectation mechanism is no longer consistent with the data. Rappoport argues that because this theory is consistent with the logic of scientific progress, it deserves consideration as the primary theory of expectation formation.

This version of procedural rationality might, indeed, be an interesting theory of expectation formation, but it gains little independent support from its connection to the well-known falsification methodology of modern science. There are two grounds for skepticism. First, today it is generally believed that philosophy oversteps its bounds when it recommends substantive theories for a science; at best, philosophy is useful for clearing conceptual puzzles that can