Categories of Criticism of the Rational Expectations Theory

Preston Miller, Clarence Nelson, Tom Supel

October 1975

Working Paper #: 49

Rsch. File #: 120.1

Categories of Criticism of the Rational Expectations Theory

By P. Miller, C. Nelson, and T. Supel

The spirit of this inquiry is as follows. The work of the past few years on the rational expectations-natural rate hypotheses by Wallace, Sargent, and others has resulted in the finding that monetary policy is powerless to exert any systematic counter cyclical influence on the economy. We take the view that judgments about whether and in what ways economic stabilization policy works in the real world must flow from the best available body of theory. That's very basic. Results of policy cannot be measured or tested directly; but various testable consequences of a general theory or model can. The success of a model's derived implications and predictions in standing up to observable facts helps us decide what is the "best available body of theory" on which to base our judgments about optimal policy making.

Our interpretation is that prior to the "rational expectations" onslaught, prevailing "best theory" (of course, not without its critics) could be represented by any of a variety of neo-Keynesian macro models including the FRB-MIT model we've used in our briefing work. Those models gave us an exploitable Phillips curve, and though we were often unconvinced that the projected quantitative tradeoff was realistic, we could carry on in the expectation that a succession of improvements in our model's equations and advances in techniques of optimal control would lead us to better and better policies.

The rational expectations class of models is a formidable contender to replace the body of prior "best theory." We accept that rational expectations has, in fact, dethroned the neo-Keynesian macro models we have known and loved (a love-hate kind of love to be sure).

But how strong a replacement is the new theory and how robust a "best theory" is it likely to be? Because of the severe practical consequences of accepting the rational expectations view, we consider it especially important to explore as many critical questions about the new theory as possible in order to assure ourselves that our grounds for so accepting the new view is the more solid.

II.

Why buy the "rational expectations" class of models in preference to the familiar class of neo-Keynesian macro models? Here we briefly note two directions for inquiry.

(1) The first essentially involves questions about validity, appropriateness, or realism of assumptions. (We reject, incidently, the Friedman 'positive economics' position that "testing" of assumptions is valueless in evaluation of theories.) Much of the form of the rational expectations innovation is essentially replacement of one set of (often implicit) behavioral premises about the manner in which economic agents make decisions by a set of somewhat more explicit premises deemed to be better because they are more representative of agents' actual know-how and cleverness. The behavioral functions in such a model are asserted to be more nearly consonant with individual optimizing decisions which must, in turn--and here an appeal to realism--take into account anticipated policy-maker responses to any projected scenario for the economy. Current macro models, in contrast, posit aggregate behavioral relations in an ad hoc way that are not necessarily consistent with "intelligent" individual optimizing behavior and specifically fail to incorporate any representation of savy about policy-maker reactions to any "unconditionally" projected set of outcomes.

Admittedly, rational expectationists have, in terms of formal theory, delivered only a halfway house. That is, the models till now tested and from which the contending consequences have been drawn are by-and-large modifications of traditional macro models, and still very much 'ad hoc'. The suggested program of rational expectations purists to build models of the economy from the ground up with consistent individual or micro optimizing assumptions is still very much a vision. Our view is that even at the ad hoc macro model level, rational expectations ideas have usefully revealed an important blind spot in representations of behavior (specifically about the role and process of expectations formation) in the class of models we have been depending upon. The question whether patching up these blind spots in some continuing ad hoc enterprise would be a possibly fruitful line of research cannot be finally answered until we know how well the rational expectations class of models stand up to critical review. One commonly occurring criticism that we will elaborate a bit later is, in effect, that rational expectations has gone too far in replacing a form of stupidity on the part of macro model agents with a form of superintelligence. If you put credence in that position, then from the point of view of "realism" of premises some middle ground is probably best.

(2) Perhaps the more important direction of inquiry in choosing between two competing bodies of theory for a "best" theory is that of deriving testable logical implications from the theories and subjecting them to empirical tests. All the better from an epistemological viewpoint if the implications are drawn so they conceivably can be shown to be false by the observed data and thus the theory conceivably can be rejected. Some of Sargent's work, not treated within the province of

this paper, raises questions about the possibility, in principle, of discriminating between the two major hypotheses on the basis of observed data. However, it appears to us that a complete and careful delineation of consequences of the two broadly different models does yield differences that aid us in discriminating between the two. We schematically represent our view of what is involved in Figure 1.

Figure 1

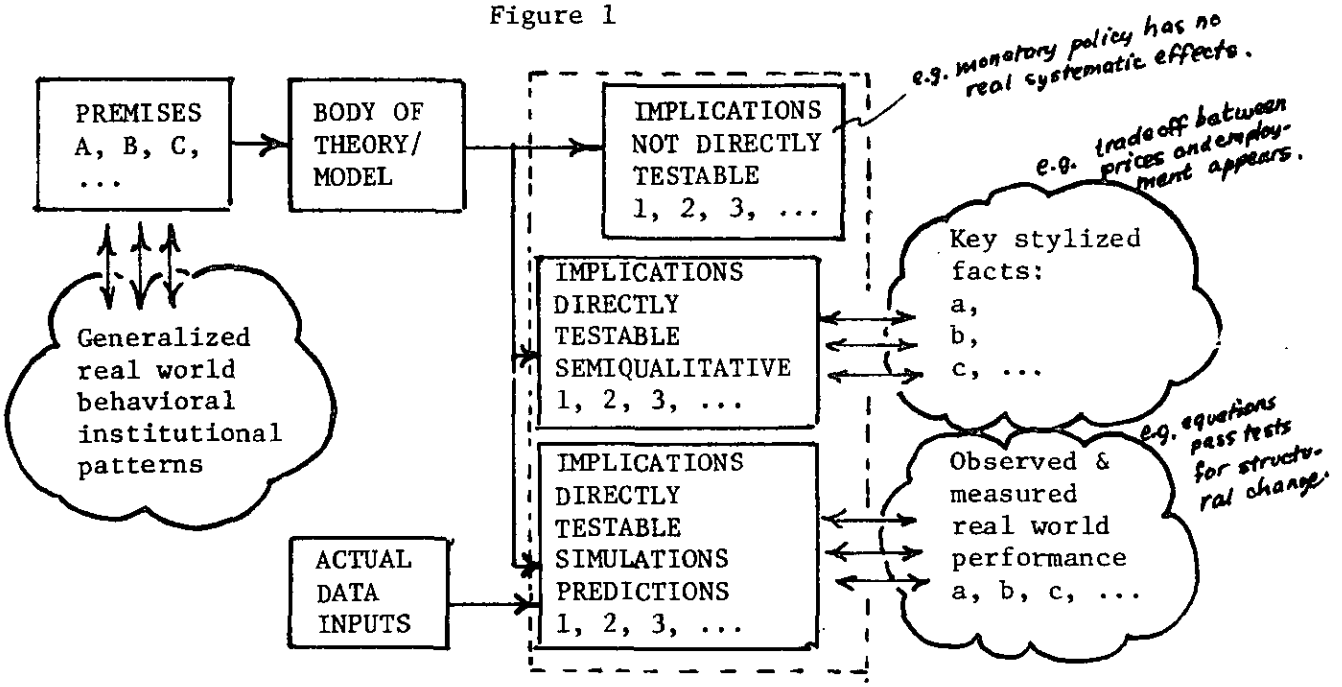

We entertain the notion of three (not totally unambiguous) categories of derivable results from theory. One set is not directly testable (not in principle, but as a practical matter, because we have not the ability to perform sufficiently rich and prolonged experiments with real world economies). The one specific result we cite in example is the derived result of overriding importance that monetary policy can have no systematic real effects. Another set consists of those statements our subject theory makes about its properties relative to a possible catalog of stylized facts of a qualitative or semiqualitative form that we gener-

alize from real world observational data. Wallace's earlier presentation in this series offered some elements to start such a catalog, including the following: (a) the observed "Phillips curve" correlation, (b) the existence of "involuntary" unemployment, and (c) the existence of appreciable periodic fluctuations in total output. To this list we may add (drawing on other's suggestions and our own): (d) procyclical investment behavior, (e) procyclical nominal interest rates, (f) "liquidity effect" (increase in real cash balances in early stages of inflationary episodes), (g) long persistence of over-/underemployment, and (h) steeper "Phillips curves" for countries with bigger real fluctuations. Just what elements ought to go into such a catalog deserves more critical thought, because the stylized facts so included provide important criteria for us to apply in screening out contending theories. The above list of elements is not intended as definitive in any sense, but it does portray the type of properties that a "best" model may be required to generate internally. At this point we merely observe that neo-Keynesian macro models pass some and fail some; that rational expectations models pass some that neo-Keynesian models fail; and that, to the best of our knowledge, only some tentative, augmented* rational expectations models on the current frontier of developing theory can even hope to "pass" all.

A third set of implications deductively generated by the model's dynamic process--in conjunction with some required input data such as that for exogenous variables--contains predicted or simulated time paths for variables endogenous to the model. Predictions over past historical periods, but out of the sample (estimation) period, can be

---

*The augmentation involves, for example, "search" theory and spatially separated labor markets.

directly tested against actual outcomes. Equivalently, somewhat more formally packaged tests for "structural change" can be made. It is now well-known that our existing stable of neo-Keynesian macro models have failed all strict tests that have been applied to them. But we are already to an extent persuaded that we should not rely on traditional macro models to provide scientific support for the notion that monetary policy can have predictable counter-cyclical effects. What we are missing is a demonstration that rational expectations models can perform better as predictors and simulators of actual outcomes.

III.

The foregoing comments serve to sketch the general grounds over which critical review of rational expectations may range. We turn now to a bit more specific an objective: to briefly and nontechnically summarize ideas from the recent literature that claim to, or may be interpreted as claiming to, reprieve activist monetary policy. We will attempt to represent these notions by at least heuristic reference to a specific simple macro model with built-in rational expectations used illustratively in some of the Wallace-Sargent papers.

The "structure" of the model is as follows:

$$(1) \quad y_t = a_0 + a_1 y_{t-1} + a_2(m_t - m_t^*) + u_t$$

$$(2) \quad m_t = g_0 + g_1 y_{t-1} + e_t$$

$$(3) \quad m_t^* = E(m_t | I_{t-1}).$$

In equation (1), y_t is some real variable of policy concern (say unemployment rate though it could be deviation of actual output from trend) and m_t is a policy instrument (say money supply though it could be inflation

rate). The a_i are fixed parameters and u_t is a random variable. The equation represents an economy whose real sector is driven by three active factors: (1) its own momentum (y_{t-1}); (2) the achieved value of the policy instrument (m_t) but only to the extent that the achieved value differs from what was expected by the agents for that period (m_t^*) viewed from the immediately prior period; (3) nonsystematic elements reflected through u_t.

Equation (2) represents the policy setting process involving a feedback rule with two parameters, g_0 and g_1, an observation on last period's unemployment rate and a random "miss," e_t. Where the policy objective is to minimize the variance of y, for example, the settings of g_0 and g_1 that will guarantee this achievement can be determined from knowledge of the parameters of equation (1). That follows, of course, from solving a straight-forward minimization problem within this model.

Equation (3) represents agents' expectations about the policy variable, m_t. Rational expectations are here portrayed as the mathematical expectation of m_t given knowledge of equation (2) and information about the true value of y_{t-1}. Since the policy maker is assumed here to have settled on values for g_0 and g_1, agents know everything in the policy maker's response function,

$$m_t = g_0 + g_1 y_{t-1} + e_t,$$

except the random element, e_t. Since e_t is known to be representable as randomly drawn from a distribution with mean zero and fixed variance, the forecast m_t^* under rational expectations is via equation (3):

$$m_t^* = g_0 + g_1 y_{t-1}.$$

That, of course, means that in equation (1) which tells us how the real economy is driven, the magnitude of the "surprise" element reflecting policy, $a_2(m_t - m_t^*)$, is equal to $a_2 e_t$, an unsystematic and unpredictable part of the policy maker's decision process.

This simple representation of a rational expectations' economy can be used to illustrate the nature of various types of criticism we can draw from the literature. The model could be elaborated and made empirically richer without altering the key lessons. We've sorted criticisms into three categories, and each category will be discussed in subsequent sections under separate authorship by the three participants in this paper. In the first category which will be discussed in Section IV are questions essentially about the validity of using equation (3)--mathematical expectations--to represent agents' forecasts of the policy variable. A contrasting point of view is that people aren't as clever or knowledgeable or well informed as such representation requires. If they aren't then the impotency of policy as modelled may incorrectly reflect characteristics of the real-world economy.

A second category of criticism, discussed in Section V, has to do with an assumption of price and wage flexibility implicit in the model. If for the moment we think of m_t as a price inflation variable and y_t as the (real) supply decision, then the representation of the model is that decisions of agents as per equation (3) utilize all of the information available up through time t-1 that is available to the policy maker. However, if agents are constrained--whether because of contractual obligations or for other reasons--from immediately making

the price or output setting that some of the I_{t-1} would otherwise dictate, then the policy innovation $a_2(m_t-m_t^*)$ in equation (1) need not wash out and some role for predictable policy effects on real variables would appear. This is not to say that agents are incapable of making unbiased forecasts of m_t on the basis of I_{t-1}. The point is that it may be an irrelevant exercise since they entered into contracts on the basis of forecasts made conditionally on I_{t-2}.

A third category of criticism seeks to raise fundamental questions about the validity of models in which only the "surprise" part of monetary variable settings, $(m_t-m_t^*)$, can have an effect on real variables. Money serves no role in the rational expectations model. Would the level of money--in addition to the surprise part--have an effect on real variables in a model where money is useful as a medium of exchange?

IV.

This section looks at some ideas from the literature that relate to the question: What if expectations are not formed in the neatly mathematical way specified in rational expectations models?

One competing possibility is that there are information differences among agents or between agents and the policy authority. We should take note there are two more or less distinct aspects to the cognitive processes in human decision making. One being knowledge (e.g., knowing about the structure of the economy or about sound decision algorithms or about optimal strategies or the equations of a system or even possessing the concept that a system exists) and the other being information (e.g., numerical or qualitative data about variables, basic

individual facts or observations, existential statements about states or conditions). Note, too, that knowledge or information as possessed may be partially or totally in error. Existence of differences in knowledge among participants in an economy could have quite different implications from existence of differences of information sets. Obviously, costs of acquiring knowledge and the amount of time involved can be very different from that associated with acquiring information.

Sargent [ ] treats information differences, discussing what he terms "partly rational" expectations in an article demonstrating the usual rational expectations results. The usual model, as presented, supposes that rational agents' expectations of m_t^*, say, are formed by preparing linear regressions on all past "conditioning variables"--call this the full information set I_{t-1}, for which we write $m_t^* = E(m_t | I_{t-1})$ as in equation (3). He then considers a more limited information set, say I'_{t-1}. If agents form expectations using $E(m_t | I'_{t-1})$, then, in general, the "surprise," $m_t - m_t^*$, will contain some systematic elements that the monetary authority knowing I_{t-1} can use in a limited way. The kind of play given to the monetary authority, Sargent emphasizes, is not such as to make it possible for the policy maker to achieve a chosen target for the real variable with any persistence, but is (on average) only the power to determine a greater or lesser variance for the real variables of the model over some run of time.

But there are other ways to conceive and model "partly rational" behavior. To say that all agents perform least squares regressions using complete or large information sets (or act equivalently) calls for a form of sophistication not thought to be representative of the real world by many economists.

Tobin [ ], commenting on Lucas's paper at the Board's 1970 Price Determination Conference, exemplifies these views:

> "Lucas' paper provides a rigorous defense of the natural rate hypothesis, and the study's rigor and sophistication have the virtue of making clear exactly what the hypothesis requires. The structure of the economy, including the rules guiding fiscal and monetary policy, must be stable and must be understood by all participants. The participants not only must receive the correct information about the structure but also must use all of the data correctly in estimating prices and in making quantitative decisions. These participants must be better econometricians than any of us at the Conference. If they are, they will always be--except for the unavoidable mistakes due to purely random elements in the time sequence of aggregate money demand--at their utility- and profit-maximizing real positions. These positions are invariant to any systematic changes in the sequence of aggregate money demand, either in the level of such demand or in any of its time derivatives.
>
> Once again, a pragmatist might conclude that he agrees with the natural rate hypothesis in principle but also believes that, in as long a run as can be of concern to policy-makers in an uncertain and changing world, a trade-off does exist for policy-makers as well as for statisticians."

Suppose only some agents in an economy are representable as rational in the foregoing sophisticated sense. Others are more limited, myopic, or even dumb. Expectations for the economy as a whole, then, are in some sense "mixed" rational and nonrational. Following the formulation in Gordon [ ] we may depict aggregate expectations formation relevant to equation (3) as a simple weighted-average of naive adaptive (as the forecast value, use last period's value) and rational expectations:

$$m_t^* = \lambda m_{t-1} + (1-\lambda) E(m_t | I_{t-1}).$$

The result of such a formulation is to restore more or less to monetary policy its short-run potency, depending on how close λ is to 1 or to 0. Rational expectationists may argue that decision makers who use $m_t^* = m_{t-1}$ are not making optimal forecasts and will, if businessmen, lose out to those who forecast as by $m_t^* = E(m_t | I_{t-1})$. That may certainly indicate a mode through which actions among some subsets of economic players will tend to move λ in the direction of zero. But that provides no basis for judging that λ is now (or indeed will be evermore) at zero. Empirically, by appeal to casual observation it is clear that some major sectors of decision makers (e.g., consumers) face no terminal consequences to persistent use of poor or wrong forecasts. Again, it seems clear that some, if not many, businesses survive over considerable periods guided by seat-of-the-pants judgments in which endemic systematic biases are neither ruled out by formal logic nor prohibited by sudden-death extinction. In fact, over some periods of time biased or suboptimal forecasts could out perform unbiased, optimal forecasts merely as a result of chance. It seems clear that a plausible Darwinian model of the business sector could be constructed that generates systematic bias in the aggregate even though any firm is eventually eliminated if its decision-process results (however arrived at) deviate greatly from those of optimal forecasting firms.

Cyert and DeGroot [ ] discuss a model in which learning takes place through feedback of information from the market. If firms do not know the form of the model of the process that determines price and, in fact, use a 'wrong' model (an "inconsistent" model) then learning from market information can lead to an equilibrium, though the process may converge slowly or even not at all. "Consistent" models, as they term

it, are those for which parameters may be unknown but the correct form is known. With Bayesian learning, consistent models lead progressively toward a rational expectations world. While that type of stylized learning experience can be fairly readily formalized, the step that carries firms from wrong models to the correct model is less tractable. They say only: "If firms have models that diverge drastically from reality, it seems reasonable to assume that management would recognize this condition and change the model. The firms would continue searching for a model that produced predictions that coincide more closely with actual observations."

Other authors have looked at possibilities for determinate monetary policy effects during a learning period in which the public adjusts its prior beliefs (assumed perhaps to be suddenly wrong, possibly because some structural change has occurred). Taylor [ ] talks of transitional expectations in such a setting. As is by now obvious, if the agents start with a false model yielding biased predictions and only gradually adjust this through learning as new data flows in, then the Phillips curve only gradually becomes unexploitable during the transition and monetary policy can systematically influence real variables during transition.

A premise underlying all of these learning-transition papers is that the natural resting point toward which the adaptive expectations process will move (regardless of whether the rate of transition is fast or slow) is rational expectations. That contention, we should note, is an empirical matter and it is not at all obvious that such a proposition is true or even approximately true. If it were, Taylor among others points out, the policy-during-transition issue is more or less hollow.

If the agents' forecasts are consistently biased merely because they do not have some of the information (or knowledge of new structure, say) that the policy authority possesses, then the policy authority has two policy options: (a) keep the information secret and work hard to exploit the gap, or (b) make all of the information public and relax. Whether the latter approach is in some higher sense effective, Taylor points out, depends on the cost of distributing and using information. The latter notion serves to remind us that even in a micro optimizing rational model world, costs in acquisition and use of information (and knowledge) may drive a wedge between m_t^* and full information $E(m_t|I_{t-1})$.

One final point, in a very quick and sketchy fashion: it seems possible (a point made by Hurwicz at last year's Rational Expectations Conference and never well grasped by this author) that risk aversive behavior can cause the pure rational expectations forecast to be discounted or compromised a bit. Illustratively, this would lead us to write equation (3) as:

$$m_t^* = \alpha E(m_t|I_{t-1})$$

where α may equal, say, 0.9. The effect could be to give the policy authority some leverage through which it could effect a stabilization policy. Why should "risk" introduce such a factor? We have come up with two hints, neither of which is thought out, and both of which may be wrong. The first is the casual observation that businessmen may purposely bias their forecasts if the costs associated with an overprediction are different from those associated with an underprediction. Using biased predictions then causes them to take actions more consistent

with their loss functions. The second is the result of a kind of game theoretic situation in which we pit agents (who aren't really sure whether the policy authority will stick to its g's or will instead try to "fool" them into doing the "right thing" as secretly determined by the authority) against the policy authority. The resolution of this game through some kind of mixed strategy in which $E(m_t|I_{t-1})$ only gets part weight by the agents is what supposedly accounts for the α-factor.

V

One of the fundamental premises underlying the theory of rational expectations is that prices and nominal wages are "flexible" in the current period. In the rationality models, flexible means that the current period values of these variables are determined endogenously rather than being predetermined by actions or events of earlier periods. The accordance of this premise with reality has been questioned by a number of persons. Phelps argues that many current period prices are set well in advance of the current period. For example, goods marketed in period t will be priced according to advertisements, catalogs, etc., that were printed in period t-1. And the firm had to issue printing orders on the basis of information known at the end of period t-2. Phelps also points out that this lag in price setting is of no particular importance to the rational expectations theory as long as the policy authorities are making their decisions on the basis of the same information as the firms. However, it seems reasonable to suppose that policy decisions are conditional on information available through period t-1. This difference in information, Phelps argues, will generally produce an optimal feedback-type policy rule.

Okun argues against price flexibility by asserting that search costs and tradition are very important elements of economic life. In a world of "customer" (as opposed to auction) markets, it is quantity which adjusts to demand shifts rather than price. And in a world of "career labor markets," wage policies are largely determined by tradition rather than current conditions.

Hall also argues that the issue of wage flexibility is of critical importance in understanding the behavior of unemployment. He focuses on wages in the "nonentrepreneurial" sector and claims that the sluggishness of wages in this sector causes the distribution of wages to widen when demand falls. Employing a search model of unemployment, Hall argues "that wage differentials do widen by enough during a contraction so that the extra unemployment associated with the contraction is the result of optimal behavior on the part of the unemployed." Hall is led to this search explanation of unemployment because other theories appear unable to adequately explain the persistence of unemployment. Once the unemployment rate moves away from its equilibrium value, it tends to stay away (on the same side) for a number of years.

In particular, Hall argues that the rational expectations theory is unable to explain the persistence of unemployment. If, in the above model, we let $a_1=0$ and interpret y and m as the unemployment rate and price level, respectively, then equation (1) can be viewed as a Phillips curve which may be written as

$$y_t = a_0 + a_2(m_t - m_t^*) + u_t$$

where a_0 is the mean of the unemployment rate. Let $v_t = m_t - m_t^*$ be the innovation in prices. Then the unemployment rate path may be described by

$$y_t = a_0 + a_2 v_t + u_t$$

Written in this form, it is clear that the rational expectations model can easily "explain" the unemployment rate at any point in time via

innovations or random shifts in the Phillips curve. However, there is nothing in the model to explain why y should remain on the same side of y^* for extended periods of time.

One possible explanation is that u is serially correlated. But this is not a satisfactory explanation, because u is a variable which represents that which is not explained by the theory, and relabeling u as persistence still leaves persistence unexplained. Thus, persistence of unemployment is another one of the stylized facts that is still without explanation.

In his remarks to the Conference on Rational Expectations, Phelps raised a number of other issues. Among them were:

(a) One premise of the theory of rationality is that the dispersion of the conditionally predicted inflation rate does not affect the short-term equilibrium values of real variables. Phelps argues that higher order moments may indeed matter, especially in a multiperiod setting. For example, a fixed rate money supply rule may cause a greater variance in the inflation rate in the current period than some feedback rule, but the variance 10 years out may be smaller. Thus, to the extent that variance affects agents' decisions, the particular policy rule may affect today's decisions.

(b) Even if one accepts the basic rational expectations model, there is still a question about the ability of policy to affect (in a systematic way) any real variables during the transition from one policy rule to another. The issue here is what happens when the policy makers adopt a new rule because of some unforeseen event such as an oil embargo or a large inventory decumulation. Even if

the new rule is announced beginning in period t, agents had made decisions for period t conditional on information available at the end of period t-1 when there was no oil embargo.

(c) In the rational expectations model, the proper value for the expected inflation rate is the reduced form value produced by the model being used. That is, rational expectations are rational relative to a model. The question being raised is, what happens when all agents are using different models--are they all modelling expectations, in some sense, optimally?

VI

The final criticism we raise applies to all macroeconomic models and is really a sad (funny?) commentary on the state of the art. The criticism is that we do not have a theory of money. Monetary models of the economy are really barter models with money grafted on; that is, money is not essential in describing the equilibria of the modeled "monetary" economies.

Before looking more closely at what it means for money to be essential, let us look at what this criticism implies about the Wallace-Sargent green paper model. Their modeled economy consists of n goods and m agents who maximize objective functions dependent only on real quantities of the n goods. Into this environment is added an $(m+1)^{th}$ agent--the monetary authority--which controls the quantity of an $(n+1)^{th}$ good-money. Money serves no purpose. It is not in agents' objective functions (which it shouldn't be), and it does not allow distributions of goods or production possibilities which are not also feasible under barter. Seen in this light it is not surprising Wallace and Sargent find that one monetary policy rule is as good as any other. What is surprising is Keynesians find that money matters.

Economists a long time ago recognized that money serves some purposes. At least three advantages are claimed for monetary economies over barter systems of exchange:

1. Use of money allows economies in the activity of exchange.
2. It eliminates the double coincidence of wants (i.e., individual who wants to trade lox for schmaltz herring does not have to find someone willing to trade schmaltz herring for lox).

3. It allows specialization of labor in production.

The problem economists face is in trying to formalize these intuitive notions. A model of the economy which attempts to formalize these intuitive notions should include costs to transacting, uncertainty about future prices, economies of scale to labor specialization, and absence of certain markets. The nonconvexities (increasing returns to scale, etc.) inherent in such an economy are in stark contrast to the nice properties exhibited by standard competitive models without money. The feature that some markets are absent implies that a complete theory of money should include a theory on the existence of markets. In other words the existence of markets should be endogenous in a model with money.

Since we are so far away from having a theory of money, does it make sense in macroeconomics to talk about the optimal quantity of money or about the optimal monetary policy rule? If the answer is "yes," then we should at least make an attempt to gauge the impact of different monetary policies on the usefulness of money as a medium of exchange.