Annals of Economic and Social Measurement, 3/1, 1974

QUANTITATIVE ANALYSIS FOR DECISIONS AT THE FEDERAL RESERVE

BY JAMES L. PIERCE*

The purpose of this paper is to describe how policy analysis and advice is made at the Federal Reserve Board as well as to discuss the implications of the process for optimal control applications. An attempt is made to highlight those areas where control applications might make the greatest contribution to improving the policy-making process. Some of the problems that have been encountered in using control techniques at the Fed--problems with structural models and specification of objective functions -- are pointed out along with some of the insights that these applications have provided.

The way in which our work at the Fed bears on control theory applications can probably best be illustrated by discussing how policy analysis and advice is actually made at the Federal Reserve. To begin, it might be instructive to provide a description of the policy-making process itself.

Every month in Washington, there is a meeting of the Federal Open Market Committee (FOMC), which is the basic policymaking body of the Federal Reserve. This Committee is composed of the seven members of the Board of Governors plus five Federal Reserve Bank presidents. The presidents of the remaining seven Federal Reserve Banks also attend these meetings and freely enter into the discussion, but do not vote. Thus, there are twelve voting members on the Committee. Immediately, the severe problems involved in specifying an objective function for monetary policy can be seen; there are twelve individuals with twelve separate sets of preferences trying to reach a single decision. Somehow, however, decisions do get made. The FOMC decides upon open market operations—the purchases and sales of Government securities, which are made daily in New York--that affect directly commercial bank reserves in the economy. This is the primary vehicle through which monetary policy operates.

In addition to the FOMC, the seven members of the Federal Reserve Board formally meet together several times a week and--among their other duties—decide upon reserve requirements and approve discount rate changes, two additional monetary policy instruments. But because these actions are carefully coordinated with the activities of the Committee, for purposes of this discussion, the FOMC will be considered to be the ongoing policymaking body.

The staff periodically makes several kinds of presentation to the FOMC. They will all be discussed in some detail here because each of them may provide some insights for control applications.

Three or four times a year, depending on the behavior of the economy, a major effort is made to prepare quarterly forecasts of the economy for the FOMC. These usually run 4 to 6 quarters into the future. The forecasts are not extended any further because generally not much credence is put in longer-run forecasting. I

* Associate Director, Division of Research and Statistics, Board of Governors of the Federal Reserve System. The opinions expressed here are my own and do not necessarily reflect those of the Federal Reserve Board or its staff. I wish to thank Mary Ann Graves for her valuable help and advice in converting the original taped version of this talk into an intelligible form for publication.

must admit, as someone who has tried to do forecasting, I share these sentiments. This basic forecast is conditional on explicit assumptions concerning monetary policy.

To begin its analysis, the staff makes a conditional forecast that contains some kind of simple policy assumption. For instance, it may be assumed that a particular growth path for the money stock will obtain during the period, or maybe a simple pattern in interest rates is taken as the policy. Given this basic assumption, we run our econometric model—which is our own version of the FMP (FRB-MIT-Penn) model—and obtain a set of results after judiciously adjusting the constant terms in equations that have not been tracking well. These adjustments can be justified as a means of building in prior information to make a conditional forecast for a specific future time period.

At the same time, another forecast is being made by judgmental means—i.e. without aid of a formal model. Judgmental forecasters at the Fed usually have a very good feel for what is going on in the economy, and they often make better short-term forecasts than the models do. The judgmental forecast is compared to the model forecast. The differences between the two forecasts often lead to insights and revisions in each forecast. A consensus forecast is then arrived at that is a blend of the forecasts obtained from the two methods.

The consensus forecast provides figures for such target variables as GNP and its major components, the unemployment rate, and the rate of inflation. At this point, the quarterly model is adjusted in terms of intercept shifts in individual equations so as to force its sectors and totals to agree quarter by quarter with the consensus forecast. Once these adjustments are made, we run policy alternatives off of the adjusted form of the model.

Because we normally do not have to adjust the model very far in order to make it agree with the consensus forecast, these adjustments have virtually no effect on the multipliers in the model, even though the model is nonlinear. Thus, the policy alternatives applied to the adjusted form of the model give us, for all intents and purposes, the same marginal effects as would occur if we had not adjusted the model to begin with. Its multipliers are left intact, but the adjustments allow us to talk about a common level of the target variables.

The policy alternatives may be for different growth paths of a monetary aggregate, such as the money stock, or for different patterns of an interest rate. In the interests of clarity, this discussion will focus on the money stock. Alternative growth paths of the money stock can be handled very conveniently with our neo-Keynesian model. We usually run several alternatives—an easier alternative, a tighter alternative, and then different time paths in between. This gives the FOMC some feel for what the tradeoffs are among important target variables such as employment, output, and the inflation rate.

We also do alternative simulations for those sectors about which we feel particularly uncomfortable or uncertain. One of our biggest problems is predicting fiscal policy over the policy horizon. As you know, in the United States, monetary and fiscal policy are determined separately. The Fed has to predict fiscal policy just as anybody else does (including the President and Congress), and we sometimes make large errors in the predictions. Currently, for example, we have to guess whether the President's impounding scheme will or will not work, how Congress

will react, and whether more or less money ultimately will be spent. When uncertainty about fiscal policy is particularly pressing, we work out various assumptions about the fiscal sector to examine the implications for some given monetary policy. In addition to the bothersome fiscal sector, there are times when we feel uneasy about a particular co-determined variable, like inventory investment. At that time, we may try different patterns of inventory investment, again, to see what the implications of variability in that sector are for the selection of a particular policy.

By the end of this process, the staff has compiled a large number of alternatives. However, there is a real limit to the number of alternatives that can be presented to any audience. A welter of detail seems to cause more confusion than assistance. While the entire exercise is quite useful to the staff, the presentation to the FOMC must be more limited. Thus the staff determines what are the most crucial issues at the time, and presents the results relevant to those issues.

On the basis of the staff's presentation and its own evaluation of events, the FOMC sets a policy for the next six months or more in terms of, say, the growth path of the money stock. In principle, the FOMC's longer-run strategy can be updated each month; in practice, however, it does not change that often—for no other reason than new and useful information does not arrive on a monthly basis. While we do not go through a full-scale forecasting and policy alternative analysis each month, we do update our forecast for each meeting based on the policy trajectory previously chosen by the Committee and on any new data. As events warrant, the FOMC does change its basic policy trajectory.

I would like to point out that while we have never formally presented an optimal control solution to the FOMC, we have produced such solutions for a certainty version of our model. The outcome has been very useful in designing policy alternatives and as background for policy advice. In the cases studied so far, it has been possible to come very close to the optimal solution with some very simple policy moves. One of our big fears had been that in the optimal solution, the money stock or interest rates might explode right off the charts. Of course the policymakers would not believe this sort of result, and probably rightfully so. In fact, however, we have been able to come very close to an optimal trajectory with quite a gradual and smooth movement in the path of both interest rates and the money stock. Thus, even in a rudimentary form, we have found these exercises to be very helpful.

While the FOMC meets once a month, its open market operations go on daily. Thus, every month, the Committee must decide upon its "operating strategy," that is, it has to give instructions to the trading desk in New York, where securities are actually bought and sold, as to what to do during the month that will lapse before the next Committee meeting. These instructions are composed of two parts. First, the Committee states its decision regarding its longer-run strategy, say, some growth path in the money stock. Then it must decide about how to return the targeted variable—whether the money stock or an interest rate—to the desired path when it strays off course. For instance, the money stock is very rarely exactly on course because it cannot be determined with complete accuracy. The FOMC's decision on how quickly to bring the money stock back towards the desired longer-run path is not independent of its decision concerning the variability in short-term

interest rates that it will tolerate — or vice versa. Thus, there is a tradeoff between control of the money stock, or any other monetary aggregate, and variability of short-term interest rates.

Each month the staff presents short-run operating strategy alternatives to the FOMC in very much the same manner that it presents longer-run GNP projections. Namely, we have a model, a form of which was described by Pindyck and Roberts at this conference, that is used to generate monthly forecasts of the money market. The forecasts are conditional on income estimates derived from the quarterly forecast and on a specific pattern of bank reserve growth. The forecasts from the monthly model are blended with independent judgments (along with the results of other models) to arrive at a set of alternatives to present to the FOMC. For instance, if the money stock is off track, then the alternatives are in the form of various growth paths of bank reserves required to move the money stock from its current value back to the desired long-run path over different time horizons. Predictions of the implication for short-term interest rates of the various alternatives are also presented. The more quickly that the money stock is brought back to the long-run path, the greater will be the movement in interest rates, other things equal. Since there is some concern about interest movements in their own right, the estimated tradeoff between hitting the money stock target and movements in interest rates are also presented. On the basis of this presentation and its own evaluation of the money market, the FOMC decides on its short-run operating strategy.

In general, the longer-run movements in the money stock, or some other monetary aggregate, are viewed as stabilizing the economy (income, employment, and inflation), and the shorter-run movements, in part, are viewed as stabilizing the money market (e.g., short-term interest rates). One can think of the situation in terms of spectral analysis: the high frequency movements in the money stock are really those used to stabilize the money market and the low frequency ones are use to stabilize the economy. There is a conflict between these movements in the money stock, but the conflict need not be very great so long as the money stock tends to fluctuate evenly around the long-run trajectory. Sometimes this is not so, and the money stock gets too far off the trajectory to allow it to be returned in any short period of time without unacceptable movements in interest rates. As a result, longer-run stabilization policy itself may be affected. Unfortunately for our purposes, the money stock is not solely set by the Federal Reserve but rather is co-determined in the economy: it is the path of bank reserves that is under policy control. While the stock of money is highly influenced by Federal Reserve policy, it is still difficult both to predict and control. Furthermore, the money stock is not the only intermediate target for monetary policy: interest rates, credit conditions, etc. also compete for the FOMC's consideration. When these other factors are given heavier weight, the money stock can stray even farther from its predetermined growth path—perhaps, with no cost to the objectives of policy.

I would like to suggest some implications of these procedures for optimal control applications. The first concerns the use of an intermediate target variable as a means for obtaining the ultimate ends of monetary policy. The FOMC really makes two decisions: it makes the decision as to what is the preferred time path of the economy, and then makes the decision as to what intermediate target, say, the

growth path in the money stock, will be consistent with this goal. In other words, it aims at the money stock—or at other times interest rates—as a vehicle for accomplishing its ends with respect to the real sectors of the economy.

It is not obvious that this is the appropriate thing to do. One is certainly entitled to ask: why not go directly from the true instruments of policy (open market operations, the discount rate, and reserve requirements) to the real economic targets? Why go through this intermediate vehicle at all? Clearly, this sort of procedure would not make any sense in a world of certainty, where we knew the exact relationships between the instruments and the ultimate targets. To the contrary, however, we operate in a world with a high degree of uncertainty. The rationale for using an intermediate target lies in the fact that its data are more frequently available than are data on the real sectors: movements in the intermediate target can provide early information on how the real sectors are responding to policy. In addition, it must be under some degree of policy control and it must be causally related to the ultimate objectives of policy. Thus, the difficulty or ease with which a target for the money stock can be hit in a particular situation, presumably, indicates what is happening in the real sectors. While this idea of using an intermediate target has appeal, no one has proved that it is appropriate.

One very useful application of optimal control procedures would be to analyze the conditions under which it is desirable to use an intermediate target for monetary policy. Furthermore, if these conditions are likely to exist, what is the best intermediate target to use? If the conclusion were to be that it is never or hardly ever appropriate to use an intermediate target variable, then it is important to know the costs incurred by pursuing one.

The next issue I want to discuss is uncertainty in general. It is difficult to overemphasize the degree of uncertainty with which policy decisions must be made. A high degree of uncertainty concerning future values of exogenous variables is one of the reasons why it is difficult to make reliable forecasts very far into the future. The forecasting errors in the exogenous variables become so large, or at least the variance around some expected value becomes so large, that the worth of our GNP projections diminishes greatly as the forecast horizon is extended.

Another area of uncertainty has to do with our models. I want to stress this because users of control theory often tend to take models as given and work out solutions without seriously questioning the reasonableness of the models. This tendency is not very harmful when one is working on technique. However, there is a real danger of giving more credence to model results than they deserve, especially if a particular policy trajectory is highly influenced by the choice of a model.

The problem lies not only with uncertainty concerning the true value of model parameters, but also with the structure of models themselves. I cannot state with much certainty that we have a good approximation to the economy with our models at the Federal Reserve Board. I have even more doubt about other models that are used for policy analysis. My particular concern involves whether or not we have correctly approximated the impact of monetary policy in the models. For example, we have found that with some relatively minor changes in the specification of our quarterly model—changing just three or four equations—we can importantly alter its policy multipliers; I believe this is true of other models as well. A couple of examples should make the point.

We have a wide variety of money demand functions from which we can choose, all estimated with about the same R^2, with about the same standard error of estimate, and all about as intuitively sensible or unsensible as the other. However, the results from the full model are quite different just because of the different interest elasticities of money demand implied by these functions. The way that this and other structural models work is to impose clearing in the money market: short-term interest rates must move sufficiently to equate money demand with money supply. Other things equal, the smaller the short-run interest elasticity of money demand, the greater the movement in interest rates required to induce the public to change its holdings of money balances by a given amount. For a given change in the growth in the money stock, different money demand functions with their different elasticities will imply different changes in short-term interest rates. The changes in short rates then feed to longer-term rates, to wealth, and to spending. Thus, the short-run impact of a change in the money stock predicted by the model will depend rather crucially on the money demand function selected. At this time, we simply have no reliable guidelines to help us choose among the functions available.

Another example is provided by the model's consumption function. Theoretically, consumption should depend upon wealth and it does in our model. Unfortunately, as soon as we have consumption dependent upon wealth, we have to predict the impact of monetary policy on the stock market, because the major component of the variance in wealth is the variance in stock prices. Given a choice, I think any of us would far prefer to just go ahead and forecast consumption than to forecast the stock market. However, because we would be losing a major channel through which monetary policy works by leaving wealth out, we leave it in. We then have all kinds of *ad hoc* procedures for explaining the stock market. The way that the stock market is specified to adjust to monetary policy variables is crucial to estimating the impact of these variables on consumption and, hence, GNP. Again we have no reliable means of selecting the "correct" specification.

Another question I want to raise is: what is an acceptable way of evaluating a model? How do we know when it is right? There are many models that explain their sample period well and also do pretty well outside the sample period, but which differ drastically in their specifications and also in their implied policy multipliers. Unfortunately, there is no very reliable method of choosing among alternative models. In a related vein, we have learned that individual structural equations might look sensible, but when they are put together in a full model, they can give some very strange results. Thus, it becomes very difficult to know when an equation is good or not. Should it be judged as an individual equation or in terms of how it contributed to the full model?

An excellent case in point has to do with stability of models. Should one impose stability on a model or not? By stability, I mean if the model is shocked with, say, a permanent change in the growth rate of the money stock, does the model go through explosive cycles? The real world might be like that, but then again it may not. It is not possible to know because the real world never gets shocked in the way that models are shocked. Instability *per se* does not bother control theorists because they point out that the system can be stabilized even though it is structurally unstable. Perhaps that may be what happens in the economy. Perhaps the economy is inherently unstable but policy at least has been good enough on average to keep

it from exploding. It is difficult to determine where the truth lies. The problem is clearly an important one because the kind of stability conditions imposed will have implications for a model's dynamic policy multipliers. It seems fair to say that there is no firm basis for knowing what kind of stability conditions to impose on our models, if any.

There is an additional problem that plagues any user of models for forecasting and policy analysis. There is no solid basis for establishing a practical method of incorporating prior information into models. This information comes from such sources as recent performance of the model, judgmental assessments of the economy, and special survey information. Adjustments of a model's constant terms is a rough and ready way to deal with this problem, but the method is clearly deficient. Slope coefficients as well as intercepts should be adjusted when appropriate. The development of better procedures for incorporating prior information would provide a major contribution to policy analysis.

At this point I would like to offer a few comments about the objective function used in the formulation of monetary policy. As mentioned earlier, it is particularly difficult to talk about "the" objective function because there are really twelve objective functions on the FOMC, and there has to be a majority of people on the Committee to agree on policy. The different ways that the Committee members grapple with uncertainty often condition disagreements more than their basic underlying objectives. For instance, if someone places a high weight on avoiding high inflation rates, he will be very worried that somehow future inflation rates have been underestimated. He will be willing to pay a relatively high penalty, in terms of higher expected unemployment rates, in order to avoid a bad draw in the sense of getting inflation rates greater than anticipated. The same sort of argument applies to a member who is worried about unemployment. Being central bankers and being in a position where decisions have to be made, the Committee members are risk averse—they are willing to trade off expected value for decrease in variance.

Their aversion to risk often takes the form of restricting movements in policy instruments. This occurs not because instrument stability is necessarily valued *per se*, but rather stems from a fear of going outside the range of experience. These particular restrictions, then, do not belong in the objective function. By restricting movements in its instruments, the FOMC has in a sense solved its own control problem. I think there is too much tendency on the part of researchers doing control applications simply to assume that there should be a penalty cost on movements in the instruments, and then justify this assumption by observing that the instruments, in fact, have not moved very much in the real world. This procedure precludes us from ever being able to demonstrate whether or not the movements have been too restrictive.

For monetary policy, there really is no cost (in an economic resource sense) of large movements in the reserve instrument. It is no more costly to buy a billion dollars of Government securities per unit than to buy a thousand dollars worth, and there are probably great scale economies. Thus, movements in bank reserves do not belong in the objective function. It is true that policymakers worry about short-term variability of interest rates and, at times, about the level of rates. These interest rate considerations should appear in the objective function, or at least as side conditions in a control problem.

Three additional observations concerning objective functions may be in order at this point. First, it is very difficult to convince a policymaker to move an instrument in what he views to be the wrong direction. That is to say, if income is expanding very rapidly and the models are predicting that it is going to fall in the future unless he eases up, it is very difficult to get him to ease up because that sort of policy recommendation is contrary to what is going on currently. I must say that until our models do a lot better, his wariness may be justified. Again, the problem is one of how to handle risk: what if the model were wrong? What if the economy were expanding very rapidly, the policymaker eases up, but economic expansion becomes more rapid? The cost of the error to the policymaker would be very large.

Second, I have observed over time that risk aversion on the part of the policymakers leads to risk aversion on the part of those giving advice. The reason is very simple: if advice is followed and it turns out to be wrong, the policymakers probably will not listen next time. Thus, the people giving advice also have a loss function that further compounds the risk aversion. As a result, policy advice often goes only part way, trying to point policy in what appears to be the right direction. Contrary to the opinion of some observers, policy advisers are not in the habit of recommending "fine tuning" of the economy.

Third, for optimal control studies, we need to know how crucial the weights are in the loss function. We have done a few optimal control experiments in this area using our quarterly model; they indicate that at least for some initial conditions, the choice of weights is not very important. In these experiments, wide variations in the relative weights assigned to the unemployment rate and the rate of inflation resulted in surprisingly similar optimal policy trajectories. The reason that this result occurs is that the inflation rate responds much less rapidly to changes in policy variables than does the unemployment rate. In the longer run, however, the effect of monetary policy is much more powerful on the rate of inflation. Thus, even if the unemployment rate receives a relatively large weight in the objective function, an attempt to bring it quickly back to target will set in train forces leading ultimately to a relatively large rise in inflation above its target. Thus, so long as the inflation rate receives a weight in the objective function, it will reduce the incentive to move the unemployment rate quickly to its desired value, although some movement is desirable. In the longer run, small changes in the unemployment rate are associated with relatively large changes in the rate of inflation, so again the inflation rate must enter importantly in computing the loss.

This would be a very powerful result if it held for a large number of initial conditions. It would indicate that researchers would not have to worry so much about getting the correct weights in the objective function. The result would also be a demonstration of the robustness of the technique of optimal control. If it turned out that the technique depended crucially on these weights, however, then it becomes much weaker because we really have no way of knowing what the weights are.

I would like to conclude my remarks with some observations on the complexity of models used for control applications. There seems to be a great desire among economists to work with the newest and biggest models. It is particularly difficult to work with new models, especially if they are large, because their properties are not well known and because their sheer size leads to severe technical

problems. The Federal Reserve only has about two or three instruments, so we are not terribly interested in looking at results where sixty different instruments can be varied. It would be extremely productive if we could talk in a more meaningful way about movements in real output, employment, and inflation in response to variations in one instrument. There probably is a very substantial payoff to working with smaller models that describe the behavior of those two or three target variables and their relation to a policy instrument. The use of small models would clear away a great deal of the pure technical problems, particularly in stochastic control problems. It is not at all clear to me that larger models are needed in order to carry out this kind of experiment. In fact, we are currently engaged in efforts to come up with a scaled down version of our own quarterly model. Hopefully, we will then be able to do control problems more efficiently than we have been able to do in the past.

I would like to conclude by saying that I think the work on optimal control is very promising, and our applied work at the Fed indicates that control techniques can and will provide important contributions toward solving stabilization problems.

Board of Governors of
the Federal Reserve System