Proceedings of the 36th
Conference on Decision & Control
San Diego, California USA • December 1997

Modeling Monetary Policy Using SPSA-based Neural Networks

Mehdi Mostaghimi

Portfolio Analysis Department
Pfizer Central Research
Groton, CT 06340
mehdi_mostaghimi@groton.pfizer.com

Department of Economics and Finance
Southern Connecticut State University
New Haven, CT 06515
mostaghimi@scsu.ctstateu.edu

Abstract

Simultaneous perturbation stochastic approximation-based neural networks (SPSA-NN) is introduced for modeling economic policy for the first time. A simulation method is used to compare the performances of two monetary policy models for the U.S. economy pursuing a targeted objective. In the first method a one-step linear feedback policy is used, and in the second method a policy based on SPSA-NN is used. It is shown that SPSA-NN policy is much faster to learn the system, and once it is learned, the policy is quick to adjust to the changes.

1. Introduction

Monetary policy decision making is the determination of the amount of the money supply and of the interest rate in the economy by the Federal Reserve System (FR) in order to accomplish certain policy objectives. The Federal Reserve System is a part of the U.S. federal government. It, however, acts independently from the elected administration and Congress in its decision making. Thus, it sets its own goals and objectives, which may or may not be consistent with the ones of the administration, and formulates its own monetary policy to accomplish these objectives.

There are several views expressed in the literature on the objectives of the FR and how they are formed and what factors are considered in forming them (for example, see [5]). One view is that the FR pursues policies to serve its own political objectives and, therefore, responds only to the political pressures exerted by the administration and Congress (members of the FR board are nominated by the administration and approved by Congress).

A second view is based on much recent evidence indicating that there is a game going on between the FR and the private sector. The private sector guesses what the best policy of the FR will be and acts to benefit from it. Thus, monetary policy is set by the FR with this behavior of the private sector in mind.

A final view, one that most observers have, is that the FR makes policies to stabilize the economy. Thus, the objective is to smooth out the business cycles and to keep the inflation low and predictable. It is believed that the FR uses mathematical/econometric models for policy making, or more precisely, uses one-step feedback models describing the policy variables in terms of their own past values and the values of the other economic variables.

The present research is consistent with this final view. In econometric policy formulation, a strong assumption is usually made that the policy maker has perfect information on how the value of a policy variable is related to the rest of the economic variables, and this relationship is usually defined as linear. Historical evidence about the U.S. economy shows that both of these assumptions are too strong and may be incorrect. Cordell [1] has indicated that there have been three tentative monetary policy regimes in the United States since the end of the Second World War: 1947:2-1969:12, 1970:1-1979:9, and since 1979:10. If we assume a change in monetary policy is a direct result of a change in the structure/dynamic of the economic system, then at least these changes must have happened in the economy to warrant a policy change. In fact, by using a well-accepted modeling of the U.S. economy for monetary policy analysis (for example, [5]) and by estimating this model for each of Cordell's regime periods and for the entire period, the results (Table 1) clearly indicate that the economic system structure is changed, going from one period to another and for the entire period.

Since it is impossible to predict accurately the changing behavior of an economic system in advance, a one-step linear feedback policy cannot adjust quickly to a change in the dynamic structure of a system, thus resulting in inefficient policies and instability in the economy. A policy formulation which is flexible enough to respond quickly to system changes would keep the economy stable.

In this research, we propose to formulate monetary policy using simultaneous perturbation stochastic

0-7803-3970-8/97 $10.00 © 1997 IEEE

approximation-based neural networks (SPSA-based NN) developed by Spall and Cristion [8]. This methodology uses the nonlinear structure of the NN for policy decision making, thus allowing for continuous change in the dynamic of the system. The methodology is model free and does not need the knowledge of the system. Rather, it uses SPSA, which is based on the noisy measurements of the data, to approximate the loss function and the gradient required for system optimization. This is in contrast to the other model-free NN models (for example, [3]), which require an explicit model of the system, to act as a true system, for calculating the gradient required in the back-propagation algorithm used there.

A simulation method is used in this research to study the effectiveness of two monetary policies, one based on a linear one-step feedback and one derived by SPSA-NN feedback, to stabilize the U.S. economy near a target objective. For this study, we assume that the economic system is known as the one estimated by a linear system for the entire period of 1959:09-1995:09. Then, we change the system dynamic in 1970:01, once by a shift in the level of the system and once by adding nonlinearity. The effectiveness of a policy, the stability of the system, is measured by the mean squared errors of the system measurement from the target objective. Policies are compared with respect to the speed of their responses to a change in the system behavior.

We are not aware of any study done using NN for economic policy analysis and we are certain that the SPSA-NN was never used for such studies. SPSA-NN, however, has been used in modeling a wastewater treatment system [8] and quality and process control [4], among other applications.

A brief review of SPSA-NN methodology is given in the next section. It is followed by a section on monetary policy simulation for the U.S. economy and summary and conclusion.

2. Simultaneous Perturbation Stochastic Approximation-Based Neural Networks

Artificial neural networks (NN) are mathematical models designed to achieve a certain goal by estimating a set of parameters in a nonlinear input-output system [10]. An NN system usually consists of several sets of nodes connected by layers: input layer, output layer, and one or more hidden layers. The parameters (θ) of the model are associated with the layers and are estimated from the observations through a learning period, usually by a back-propagation algorithm.

Stochastic approximation (SA) is a set of algorithms for searching for the optimal value of a stochastic function. For a loss function $L(\theta)$, the minimal θ^* is reached when

$$g(\theta) \equiv \frac{\partial L}{\partial \theta} = 0 .$$

There are two SA methods for solving this equation; Robbins-Monro, which uses the noisy observations of the gradient $g(\theta)$, and Kiefer-Wolfowitz, which uses the noisy data on the loss function $L(\theta)$. In both methods, the recursive algorithm for finding θ^* is

$$\hat{\theta}_k = \hat{\theta}_{k-1} - \alpha_k \hat{g}(\hat{\theta}_{k-1}) ,$$

where α_k, the gain coefficient, is a positive scalar, and $\hat{g}$ is an approximation of the gradient function g. Under some conditions, $\hat{\theta}_k$ converges to θ^*.

When $g(\theta)$ is not available, the only SA method for solving this equation is Kiefer-Wolfowitz. A standard approach to this method is, at each iteration, to form an approximation based on the values of the loss function L(.) by positively and negatively perturbing each component of $\hat{\theta}_{k-1}$. Thus, for p-dimensional θ, a total of 2p values of L(.) is required for estimating each gradient.

An alternative approach to the above approximation of the gradient is the Spall simultaneous perturbation (SP) method [6]. In this method, two observations of L(.), independent of the dimension of p, are formed by producing $\hat{L}^{(\pm)}_k$ as a result of perturbing $\hat{\theta}_{k-1} \pm c_k\Delta_k$, where c_k is a positive scalar, and $\Delta_k=(\Delta_{k1},\Delta_{k2},...,\Delta_{kp})^T$ is a vector of independent mean zero random variables satisfying certain regularity conditions. Thus, the simultaneous perturbation approximation to $g_k(\hat{\theta}_{k-1})$ is

$$\hat{g}_k(\hat{\theta}_{k-1}) = \begin{bmatrix} \dfrac{\hat{L}_k^{(+)} - \hat{L}_k^{(-)}}{2c_k\Delta_{k1}} \\ \vdots \\ \dfrac{\hat{L}_k^{(+)} - \hat{L}_k^{(-)}}{2c_k\Delta_{kp}} \end{bmatrix} .$$

Spall [7] has shown that $\hat{g}_k(\hat{\theta}_{k-1})$ is an approximately unbiased estimator of $g(\hat{\theta}_{k-1})$, and that SPSA can achieve the same level of asymptotic accuracy in estimating θ^* as the standard method with only 1/p the number of measurements of L(.). This is especially important in NN problems since p can be a very large number.

Applications of SPSA-based NN to policy/control decision making is described in [9]. Consider a system output x_{k+1} at time k+1 given by

Table 1: Estimated Models for the Dependent Variable DIP

Period	Independent Variables (time lag) c	DIP (-1)	DPPI (-1)	DM (-2)	DM (-3)	DM (-5)	DM (-6)	DM (-10)	DM (-12)	DSP (-1)	DR	DR (-1)
1951.08-1995.09	.001	.335	.109		.181		.203			.023	.003	
1951.08-1969.12		.329				.321	.404			.064	.010	
1970.01-1979.09	.002	.426		.302	.375		.308		-.351		.001	
1979.10-1995.09			.196	.198			.131	.116	.132		.003	.001

$$x_{k+1} = \Phi_k(x_k, y_{1,k}, y_{2,k}, \ldots, y_{l,k}, u_k, w_k),$$

where $\Phi_k(.)$ is a generally unknown, nonlinear function governing the dynamic of the system, y_i; for i=1,2,...,l, are other variables influencing the system, u_k is the policy input applied to affect the system at time k+1, and w_k is a serially independent random variable. The goal of a policy maker is to choose the sequence of the policy variable $\{u_k\}$ in such a manner that the output of the system is close to a targeted policy objective $\{t_k\}$.

An NN is used to produce the policy variable $\{u_k\}$. Assuming the structure of NN is known (number of layers and nodes), finding the optimal policy u_k is equivalent to estimating the parameters of NN (θ_k) minimizes a loss function $L_k(\theta_k)$ measuring the performance of the system x_{k+1} relative to its target value t_{k+1}. Assuming no information is available about the dynamic of the system, an SPSA approximation is used to approximate the loss function. In such a method, the output of the NN is the policy variable u_k, and the inputs are the system output x_k and the system target t_k.

3. Monetary Policy Simulation

In this study, we will use a simulation method to compare the performances of two different monetary policies for accomplishing a certain objective in the U.S. economy for the period of 1951:09 - 1995:09.

For this study we need a good model of the U.S. economy relating a policy variable, in addition to other economic variables, to a variable measuring the performance of the economy. For this purpose we select a good model, according to the literature (for example, [5]), for monetary policy analysis. In this model, the performance of the U.S. economy is measured by the index of industrial production IP, which is shown as a function of its own past values, the interest rate R (three-month T-bill rate), the total money supply M1, producer price index PPI, and S&P composite stock price index SP. The selected policy variable is the interest rate R. For a more precise description of these variables and where to find them, see [5].

As Salemi [5] and the others have shown, the best linear model showing that the index of industrial production is statistically related to the other economic variables, especially to the policy variable R, is in the following form

$$DIP = b_0 + b_1 DIP(-1) + \sum_{i=2}^{k-1} b_i X_i + b_k DR,$$

where DIP = log(IP) - log(IP(-1)),
DR = R - R(-1),
DPPI = log(PPI) - log(PPI(-1)),
DM = log(M1) - log(M1(-1)),
DSP = log(SP) - log(SP(-1)), and

for example DIP(-1) stands for the one-period lag variable of DIP. The variables X_i for i=2,3,...,(k-1) are generic used to represent the rest of the economic variables, which are not shown explicitly in the model.

Using the data from CITIBASE, the best fit model, AIC and BIC for overall validation and t-test at 95% for individual significance, using E-View (TSP) econometric software for the estimation, for the period of 1951:09 - 1995:09 is shown in the first row of Table 1. This model is used as the baseline system in our study.

In this study, it is assumed that the overall goal or objective of a policy is to stabilize the economic system around a targeted growth rate equal to the average value of DIP=0.002911. Thus, the target value ODIP is set to equal 0.002911.

The policy variable used for accomplishing this objective is the change in the interest rate, DR, set mostly by the Federal Reserve System, FR. It is assumed that the final view described in the introduction is the one the FR pursues in making monetary policy decisions and, in that, FR uses a one-step feedback policy to decide the value of DR. Let's call the policy estimated by such a method EDR. Assume the FR uses two different methods to arrive at such a policy. In the first method, a linear feedback policy is used, where EDR is generated as a function of the system variables and the latest estimates of the system parameters (for example, see [6]). For the second method, EDR is estimated by the SPSA-based NN method of [9], described briefly in this paper. The SPSA-NN method is model free and only uses the latest values of DIP, the targeted value of ODIP, the previous value of DR, and the random perturbations generated by a distribution with a mean of zero and a variance equal to the variance of DIP for the entire period equal to 0.010367. It is assumed that the NN has two hidden layers, the first one with 20 nonlinear nodes and the second one with 10 nonlinear nodes.

It is assumed that the economic system is behaving linearly and this behavior is according to the model identified for the entire period. This assumption gives a great advantage to the linear feedback policy method in that the policy-generating method has perfect knowledge of the system behavior and, thus, its performance is expected to be the best. In reality, however, the economic system is not perfectly known to the policy maker and he/she/it must guess the structure of the system as well.

In order to see how different policy regimes adjust to sudden changes and nonlinearities in the system, it is assumed that the system changes, beginning in 1970:01, without the knowledge of the policy makers. In the first case, this change is equivalent to a sudden shift in the level of the system. It is assumed that this shift is equivalent to one standard deviation of DIP, equal to 0.012581, estimated from the data for the entire period. In the second case, a nonlinearity is introduced in the system in the form of the multiplicative DIP(-1)*DR(-1):

$$DIP = b_0 + b_1 DIP(-1)*DR(-1) + \sum_{i=2}^{k-1} b_i X_i + b_k DR.$$

In comparing this model with the original one, we see that the parameter b_1 is the coefficient of the product of DIP(-1)*DR(-1), rather than only of DIP(-1), for the period beginning 1970:01. Since the policy makers are not aware of these changes in the system behavior, it is of interest to see how each policy method is adjusting to this change.

The performance of a policy is measured by how close the system performance, RDIP (the estimated value of DIP given a policy variable EDR), is to the policy objective ODIP. A mean square errors (MSE) method is used to measure the performance of a policy. The simulation process is as follows: beginning from period 1951:09, for every period, the system parameters are estimated using econometric software (E-View 2.). For the linear feedback policy model, the policy variable EDR is estimated by (see [5]):

$$EDR(+1) = (ODIP - (\hat{b}_0 + \sum_{i=1}^{k-1} \hat{b}_i X_i)) / \hat{b}_k.$$

For the SPSA-NN policy, this policy variable is estimated by a software developed using Spall and Cristion's [9] methodology. RDIP is then estimated by substituting EDR for DR in the estimated mode.

Figures 1 and 2 show, respectively, the mean square errors of the performances of the linear feedback policy and the SPSA-NN policy. In each figure, the systems are identified by the last one or two letters of their names: B for the baseline system, NL for the nonlinear system from 1970:01, and S for a shift in the level from 1970:01.

Overall, by comparing Figures 1 and 2, we see that the linear feedback policy performs better than the SPSA-NN policy in that the former has lower mean square errors. This comparison, however, is misleading. We all have expected the linear feedback policy performs better than the one from SPSA-NN, given that it has perfect knowledge about the system structure and the SPSA-NN policy has no knowledge about the system. The performance of a policy must be evaluated relative to its prior knowledge of system. A policy model must also be evaluated based on how quickly and how accurately it responds to a change in the system behavior. A superior policy model is faster in learning the system behavior and is more accurate in responding to such changes.

In comparing Figure 1 with Figure 2, it is very clear that the SPSA-NN policy is much faster to learn the system, and once the system is learned, the policy is quick to adjust to changes. This can be clearly seen on the left-hand side of Figure 2 (MSERIPNNB), where at the beginning SPSA-NN has no information about the system so the mean squared errors are very high; then, almost immediately, within a few

Figure 1: Linear Feedback Policy: Mean Squared Errors
MSERIPB: Baseline Model
MSERIPNL: Nonlinear System from 1970:01
MSERIPS: A One Standard Deviation Shift to DIP from 1970:01

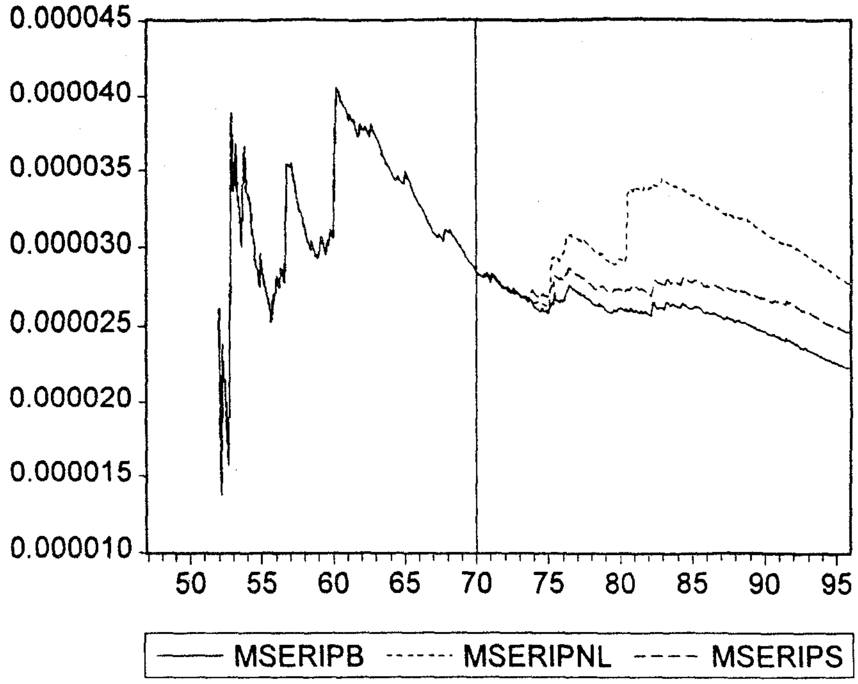

Figure 2: SPSA-NN Policy: Mean Squared Errors
MSERIPNNB: Baseline Model
MSERIPNNNL: Nonlinear System from 1970:01
MSERIPNNS: A One Standard Deviation Shift to DIP from 1970:01

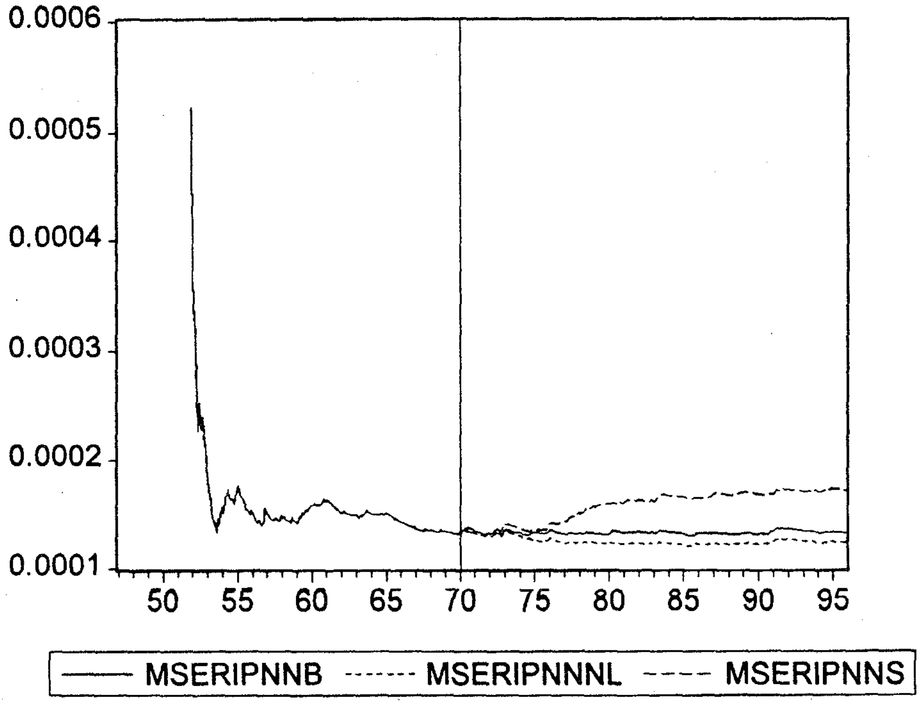

periods, it learns almost all that is possible about the system and stays on top for the rest of the period; mean squares errors stay low with minimum volatility. For the linear feedback, Figure 1 (MSERIPB), the policy maker knows a lot about the system right at the beginning, but as time passes and the system changes, the policy cannot adjust as fast; an increase in the level of mean square errors and sudden volatilities with high magnitudes result.

It seems that SPSA-NN is also faster to adjust when a structural change happens in the system. This can be seen, beginning in 1970:01, especially for the case of a nonlinear change in the system (MSERIPNL and MSERIPNNNL). In fact, the mean squared errors for the SPSA-NN policy have dropped below the one for the baseline. This indicates that when nonlinearities become present in the system, a characteristic of most economic systems, SPSA-NN actually performs better. The performance of the linear feedback policy for this case has greatly deteriorated, with higher mean square errors, for the obvious reason of the policy's inability to cope with a nonlinear change in the system. In the later years, beginning around 1983, the linear feedback policy begins to recover rapidly, decreasing mean square errors. The major reason for this is that the value of the coefficient of the nonlinear effect (b_1) is actually becoming smaller in the model (coefficient of DIP(-1) in Table 1 for period 1979:10-1995:09), indicating that the effect of the nonlinearity is very weak in the model. The SPSA-NN policy response for the case of a shift in the level of the system (MSERIPNNS) is not that satisfactory and it seems that the policy performance deteriorates with time; the gap between its mean square errors and the one for the baseline is widening. The same can also be said for the linear feedback policy, but the situation is not as severe.

A more accurate evaluation of the performance of a policy model is to compare it relative to its prior knowledge of the system [2]. It is shown that SPSA-NN performs very well when there are nonlinearities in the system.

4. Summary and Conclusion

In this research, we introduced the application of the simultaneous perturbation stochastic approximation-based neural networks (SPSA-NN) to modeling economic policy, in general, and to modeling monetary policy, in particular, for the first time. The monetary policy of a one-step linear feedback is compared with the monetary policy of SPSA-NN for the U.S. economy in accomplish a certain objective. It is learned that SPSA-NN is faster to learn the system and quicker to adjust to the changes in the system. As a result, the economic system is more stable around its targeted objective.

Acknowledgment: This research is partially supported by a grant from the Connecticut State University System. Special thanks to Professor Winnie Yu from Southern Connecticut State University for helping on computer implementation of the simulation.

References

[1] Cordell, L.R. (1989). "The Invisible Handshake: An Investigation of the Monetary Policy Process at the Federal Reserve System, 1966 to 1982," unpublished Ph.D. Dissertation, University of North Carolina, Chapel Hill, Dept. of Economics.

[2] Mostaghimi, M. (1997). "Monetary Policy Simulation Using SPSA-based Neural Networks," Working Paper, Department of Economics and Finance, Southern Connecticut State University, New Haven, CT 06515.

[3] Narendra, K.S. and K. Parthasarathy (1990). "Identification and Control of Dynamical Systems Using Neural Networks," *IEEE Trans. Neural Nets*, 1, 4-26.

[4] Rezayat, F. (1995). "On the Use of an SPSA-based Model-free Controller in Quality Improvement," *Automatica*, Vol. 31, No. 6, pp. 913-915.

[5] Salemi, M.K. (1995). "Revealed Preference of the Federal Reserve: Using Inverse-Control Theory to Interpret the Policy Equation of a Vector Autoregression," *Journal of Business & Economic Statistics*, Vol. 13, No. 4.

[6] Sargent, T.J. (1979). Macroeconomic Theory, Academic Press, New York.

[7] Spall, J.C. (1992). "Multivariate Stochastic Approximation using a Simultaneous Perturbation Gradient Approximation," *IEEE Transactions on Automatic Control*, AC-37, 332-341.

[8] Spall, J.C. and J.A. Cristion (1995). "Model-free Control of Nonlinear Stochastic Systems in Discrete Time," In *Proc. of the IEEE Conf. on Decision and Control*, pp. 2199-2204.

[9] Spall, J.C. and J.A. Cristion (1994). "Nonlinear Adaptive Control Using Neural Networks: Estimation with a Smoothed Simultaneous Perturbation Gradient Approximation," *Statistical Sinica*, 4, 1-27.

[10] Wasserman, P.D. (1989). Neural Computing: Theory and Practice, van Nostrand Reinhold, New York.