

Parameterized Expectations Algorithm and the Moving Bounds

Lilia MALIAR and Serguei MALIAR

Departamento de Fundamentos del Análisis Económico, Universidad de Alicante,
Campus San Vicente del Raspeig, Ap. Correos 99, 03080, Alicante, Spain
(malialr@merlin.fae.ua.es and maliars@merlin.fae.ua.es)

The Parameterized Expectations Algorithm (PEA) is a powerful tool for solving nonlinear stochastic dynamic models. However, it has an important shortcoming: it is not a contraction mapping technique and thus does not guarantee a solution will be found. We suggest a simple modification that enhances the convergence property of the algorithm. The idea is to rule out the possibility of (ex)implosive behavior by artificially restricting the simulated series within certain bounds. As the solution is refined along the iterations, the bounds are gradually removed. The modified PEA can systematically converge to the stationary solution starting from the nonstochastic steady state.

KEY WORDS: Nonlinear models; Numerical solutions methods; Optimal growth; Parameterized expectations algorithm.

1. INTRODUCTION

The Parameterized Expectations Algorithm (PEA) is a non-finite state space method for computing equilibria in nonlinear stochastic dynamic models (e.g., Wright and Williams 1982; Miranda and Helmlinger 1988; Marcet 1988; den Haan and Marcet 1990; Christiano and Fisher 2000). The method is as follows: approximate the conditional expectation in Euler's equation by a parametric function of state variables and find the parameters, which minimize the distance between the expectation and the approximating function.

Several properties make the PEA an attractive tool for researchers in the area of economic dynamics. First, if a low-degree polynomial approximation delivers a sufficiently accurate solution, the cost of the algorithm does not practically depend on the dimensionality of the state space. Second, the PEA can be applied for analyzing not only the optimal economies but also the economies with externalities, distortions, liquidity constraints, and so on. Finally, the algorithm is fast and simple to program. For an extensive discussion of the method and its applications, see Marcet and Lorenzoni (1999) and Christiano and Fisher (2000).

The main drawback of the PEA is that it is not a contraction mapping technique and thus does not guarantee a solution will be found. In fact, if the assumed decision rule happens to be far from the true solution, the algorithm is likely to diverge. To achieve convergence, one has to wisely choose initial values for the parameters in the approximating function as well as a procedure for updating the parameters on each iteration.

To systematically find a good initial point for iterations one can use homotopy: “[S]tart with a version of the model which is easy to solve, then modify these parameters slowly to go to the desired solution.... It is often possible to find such ‘known’ solutions and to build a bridge that goes to the desired solution” (Marcet and Lorenzoni 1999, p. 156). One can also start from a solution that is previously computed by another numerical method, such as the log-linear approximation (see Christiano and Fisher 2000). It is evident, however, that the need to search for an initial point can seriously complicate implementing the PEA in practice.

This paper describes a simple modification that enhances the convergence property of the PEA. We consider the version of the algorithm developed by Marcet (1988) that evaluates the expectations by using Monte Carlo simulation. Our idea is to rule out the possibility of (ex)implosive behavior by artificially restricting the simulated series within certain bounds. As the solution is refined along the iterations, the bounds are gradually removed. We call this modification “moving bounds.”

Introducing the moving bounds resolves the problem of finding a good initial guess in the sense that the modified PEA is able to converge even if the initial guess is not very accurate. In our example, the modified PEA can systematically find the stochastic solution starting from the nonstochastic steady state. It is also important to mention that the practical implementation of the moving bounds is simple: one only has to automatically insert several lines in the original PEA code. In the remainder of the paper, we formally describe the modification of the moving bounds and provide an illustrative example.

2. THE MOVING BOUNDS

We modify the PEA described in Marcet and Lorenzoni (1999) to include the moving bounds. Consider an economy, which is described by a vector of n variables, z_t , and a vector of s exogenously given shocks, u_t . It is assumed that the process $\{z_t, u_t\}$ is represented by a system

$$g(E_t[\phi(z_{t+1}, z_t)], z_t, z_{t-1}, u_t) = 0, \quad \text{for all } t, \quad (1)$$

where $g: R^m \times R^n \times R^n \times R^s \rightarrow R^q$ and $\phi: R^{2n} \rightarrow R^m$; the vector z_t includes all endogenous and exogenous variables that are inside the expectation, and u_t follows a first-order Markov process. It is assumed that z_t is uniquely determined by (1) if the rest of the arguments is given.

We consider only a recursive solution such that the conditional expectation can be represented by a time-invariant function $\Phi(x_t) = E_t[\phi(z_{t+1}, z_t)]$, where x_t is a finite-dimensional subset of (z_{t-1}, u_t) . If the function $\Phi(\cdot)$ cannot be derived analytically, we approximate $\Phi(\cdot)$ by a parametric function $\psi(\beta, x)$, $\beta \in R^v$. The objective will be to find β^* such that $\psi(\beta^*, x)$ is the best approximation to $\Phi(x)$ given the functional form $\psi(\cdot)$,

$$\beta^* = \arg \min_{\beta \in R^v} \|\psi(\beta, x) - \Phi(x)\|.$$

This can be done by using the following iterative procedure.

- Step 1. Fix upper and lower bounds, \underline{z} and \bar{z} , for the process $\{z_t(\beta), u_t\}$. For an initial iteration $i = 0$, fix $\beta = \beta(0) \in R^v$. Fix initial conditions u_0 and z_0 ; draw and fix a random series $\{u_t\}_{t=1}^T$ from a given distribution. Replace the conditional expectation in (1) with a function $\psi(\beta, x)$ and compute the inverse of (1) with respect to the second argument to obtain

$$z_t = h(\psi(\beta, x_t(\beta)), z_{t-1}, u_t). \quad (2)$$

- Step 2. For a given $\beta \in R^v$ and given bounds \underline{z} and \bar{z} , recursively calculate $\{z_t(\beta), u_t\}_{t=1}^T$ according to

$$\begin{aligned} z_t(\beta) &= \underline{z} && \text{if } z_t(\beta) \leq \underline{z}, \\ z_t(\beta) &= \bar{z} && \text{if } z_t(\beta) \geq \bar{z}, \\ z_t(\beta) &= h(\psi(\beta, x_t(\beta)), z_{t-1}, u_t) && \text{if } \underline{z} < z_t(\beta) < \bar{z}. \end{aligned}$$

- Step 3. Find a $G(\beta)$ that satisfies

$$G(\beta) = \arg \min_{\xi \in R^v} \|\phi(z_{t+1}(\beta), z_t(\beta)) - \psi(\xi, x_t(\beta))\|.$$

- Step 4. Compute the vector $\beta(i+1)$ for the next iteration,

$$\beta(i+1) = (1 - \mu)\beta(i) + \mu G(\beta(i)), \quad \mu \in (0, 1).$$

- Step 5. Compute $\underline{z}(i+1)$ and $\bar{z}(i+1)$ for the next iteration,

$$\begin{aligned} \underline{z}(i+1) &= \underline{z}(i) - \underline{\Delta}(i), \\ \bar{z}(i+1) &= \bar{z}(i) + \bar{\Delta}(i), \end{aligned}$$

where $\underline{\Delta}(i)$ and $\bar{\Delta}(i)$ are the corresponding steps.

Iterate on Steps 2–5 until $\beta^* = G(\beta^*)$ and $\underline{z} < z_t(\beta^*) < \bar{z}$ for all t .

To perform Step 3, one can run a nonlinear least squares regression with the sample $\{z_t(\beta), u_t\}_{t=1}^T$, taking $\phi(z_{t+1}(\beta), z_t(\beta))$ as a dependent variable, $\psi(\cdot)$ as an explanatory function, and ξ as a parameter vector to be estimated. We will not discuss the choice of the functional form for the approximation, the parameter number, the simulation length, and so on, as all of these are extensively analyzed in the previous literature (e.g., Marcet and Lorenzoni 1999). Here, we focus only on the issue of convergence.

Thus, our modification is to artificially restrict the simulated series, $\underline{z} < z_t(\beta^*) < \bar{z}$. If \underline{z} and \bar{z} are set to $-\infty$ and $+\infty$,

respectively, the modified version is equivalent to the original one. The role of the bounds is discussed below.

Unlike the traditional value-iterative methods, the PEA does not have the property of global convergence. To be precise, if the approximation $\psi(\beta, x)$ happens to be far from the true decision rule, $\Phi(x)$, then the simulated series $\{z_t(\beta), u_t\}_{t=1}^T$ become highly nonstationary; as a result, the regression does not work appropriately and the algorithm diverges. Hence, one has to initially choose and subsequently update β such that $\psi(\beta, x)$ remains sufficiently close to the true decision rule, $\Phi(x)$. The need to fulfill this requirement can complicate the use of the PEA in practice; for example, one has to search for an initial guess by using homotopy or the log-linear approximation.

We approach the problem from a different perspective. Specifically, rather than trying to ensure that $\psi(\beta, x)$ always remains close to $\Phi(x)$, we attempt to enhance the convergence property of the PEA and, consequently, to prevent the algorithm from failing if the approximation happens to be far from the true decision rule. The moving-bounds method exploits the fact that under the true decision rule, $\Phi(x)$, the process $\{z_t, u_t\}_{t=1}^T$ is stationary. The bounds artificially induce the stationarity of possibly (ex)implosive simulated series $\{z_t(\beta), u_t\}_{t=1}^T$ by not allowing such series to go beyond a fixed range $[\underline{z}, \bar{z}]$. This range is small, initially. However, it increases at each subsequent iteration. The bounds, therefore, play a stabilizing role at the beginning, when the approximation $\psi(\beta, x)$ is probably not accurate. As the PEA converges to the stationary solution, the bounds gradually lose their importance and eventually become completely irrelevant.

In practice, there is no need to impose bounds on all of the simulated series. It is sufficient to restrict only the series for endogenous state variables, which are calculated by recursion and thus have a natural tendency to (ex)implode. In general, the remaining variables will be continuous functions of state variables and thus will be restricted automatically. Also, in some applications, there is no need to readjust (move) the bounds on each iteration. It is possible to fix the bounds \underline{z} and \bar{z} at the beginning so that the algorithm will eventually converge.

We discuss a possible choice of the moving bounds parameters in the next section. As an initial guess, we use the nonstochastic steady state. An advantage of this approach is that the initial point is computed in a simple and systematic manner. The drawback is that the nonstochastic steady state solution can be far from the true stochastic solution, and, hence, the convergence can be slow. It is important to mention that using the steady state as an initial guess is not feasible within the original PEA framework.

3. AN EXAMPLE

To illustrate the application of the moving-bounds method, we consider the simplest one-sector stochastic growth model,

$$\max_{\{c_t, k_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \delta^t \frac{c_t^{1-\gamma} - 1}{1-\gamma}, \quad \text{s.t. } c_t + k_t = (1-d)k_{t-1} + \theta_t k_{t-1}^\alpha,$$

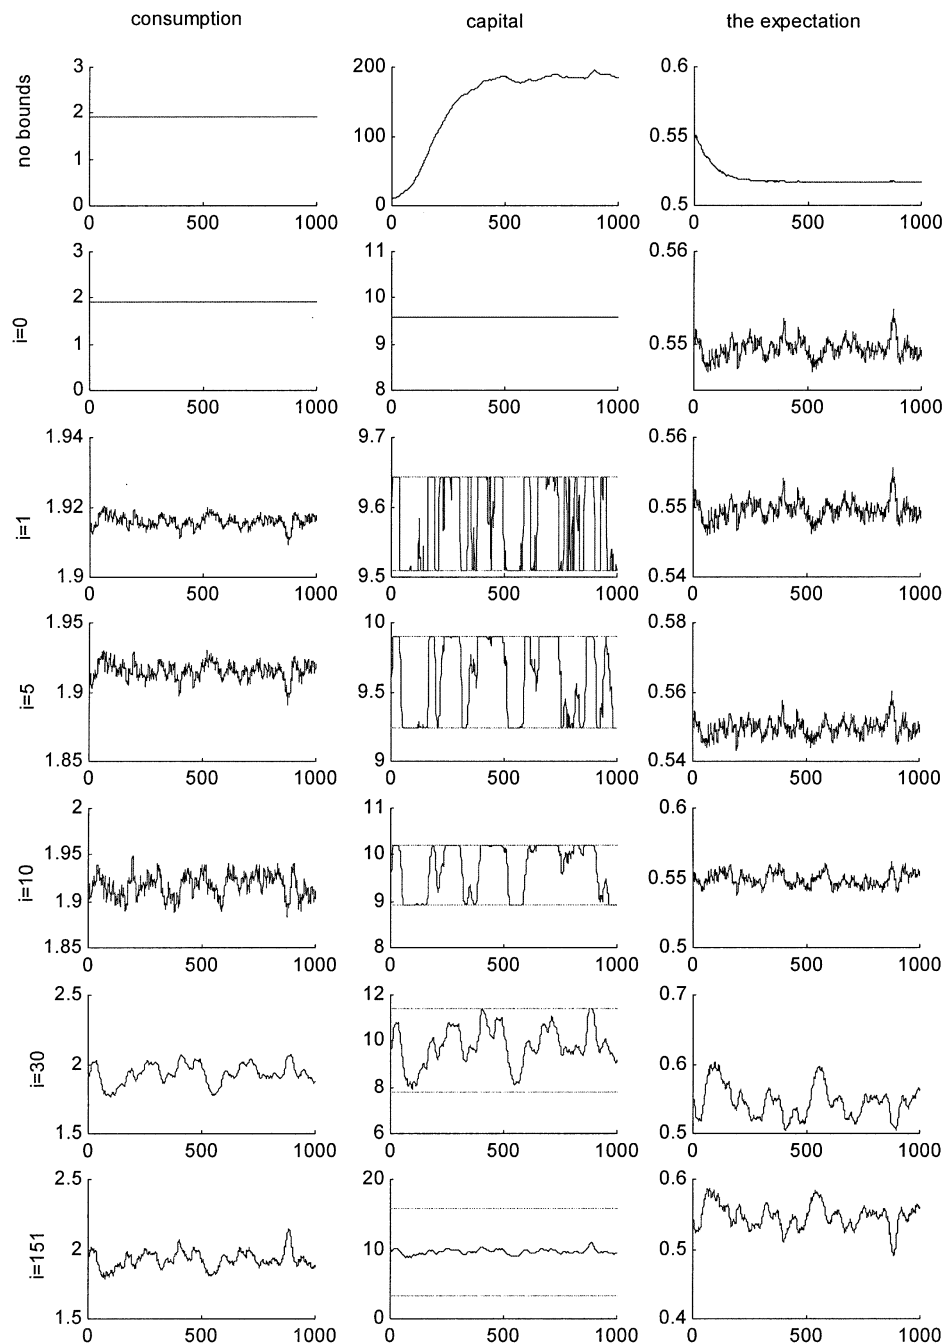


Figure 1. The Realization for Consumption, Capital, and the Expectation for a Single Stochastic Simulation (exploding capital, 1,000 periods).

where $\log \theta_t = \rho \log \theta_{t-1} + \epsilon_t$ with $\epsilon_t \sim N(0, \sigma^2)$, the initial condition (k_{-1}, θ_0) is given, and $\alpha \in (0, 1)$. If the utility is logarithmic, $\gamma = 1$, and there is full depreciation of capital, $d = 1$, the model allows for an analytic solution: $c_t = (1 - \alpha\delta)\theta_t k_{t-1}^\alpha$. In general, the closed-form solution to this model is not known.

The previous paper by den Haan and Marcet (1990) solves this model under $\gamma = 1$ by using the PEA and the homotopy approach. They start from the solution to the model under $d = 1$ and change d from 1 to 0 in 10 steps; they calculate the solution for each step and employ it as the initial guess for the next d .

We show how to solve the model by using the modified version of the PEA. The program is written in

MATLAB and is available on both the ASA FTP data archive, <ftp://www.amstat.org/>, and the authors' websites, <http://merlin.fae.ua.es/maliarl> and <http://merlin.fae.ua.es/marliars>. Following den Haan and Marcet (1990), we approximate the conditional expectation by

$$E_t[c_{t+1}^{-\gamma}(1-d + \alpha\theta_{t+1}k_t^{\alpha-1})] \cong \exp(\beta_0 + \beta_1 \log \theta_t + \beta_2 \log k_{t-1}),$$

where $\beta = (\beta_0, \beta_1, \beta_2)$ is a vector of coefficients to be found. We can calibrate β as

$$\beta_0 = \ln[c_{ss}^{-\gamma}(1-d + \alpha\theta_{ss}k_{ss}^{\alpha-1})], \quad \beta_1 = 0, \quad \beta_2 = 0.$$

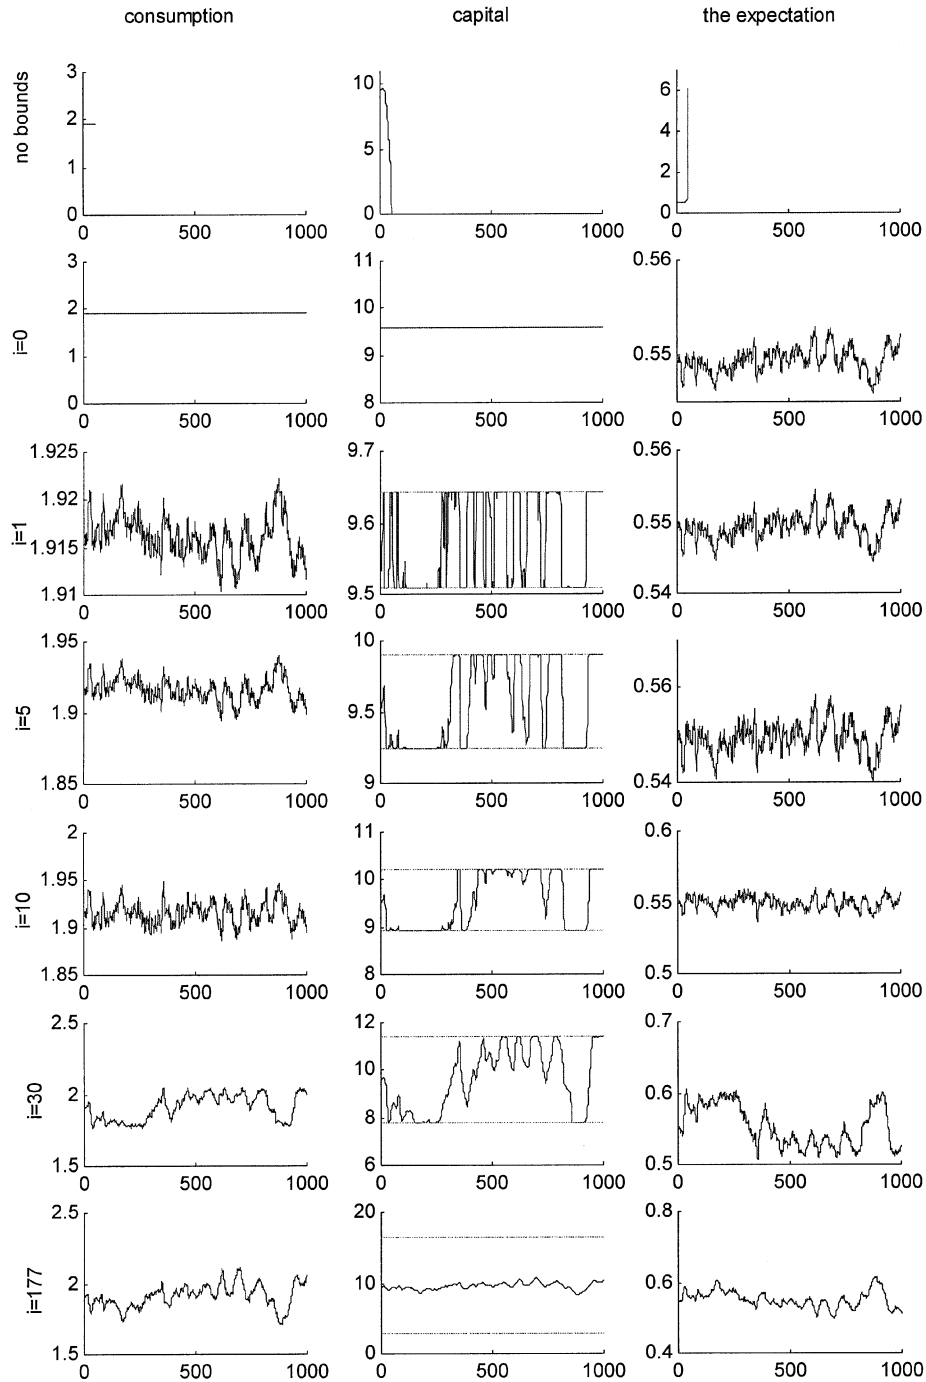


Figure 2. The Realization for Consumption, Capital, and the Expectation for a Single Stochastic Simulation (imploding capital, 1,000 periods).

However, to illustrate the convergence ability of the modified PEA, in the supplied MATLAB program, we draw β_1 and β_2 from the normal distribution, $N(0, 1)$. The algorithm has no difficulty in converging starting from such a random initial condition.

The moving-bounds parameters are

$$\begin{aligned} \underline{k}(i) &= k_{ss} \exp(-ai), \\ \bar{k}(i) &= k_{ss}(2 - \exp(-ai)), \end{aligned}$$

where $a > 0$, i is the number of iterations performed, and the variables with the subscript ss are the steady-state values.

Under this choice, on the first iteration ($i = 0$), the simulated series coincide with the steady-state solution, $k_t(\beta) = k_{ss}$, for all t . On the subsequent iterations, the lower and upper bounds gradually move, approaching 0 and $2k_{ss}$, respectively. The parameter a determines the pace at which the bounds are moved.

To simulate the model, we fix the model's parameters as follows:

α	δ	γ	d	ρ	σ	k_{-1}	θ_0
0.33	0.95	1	0.02	0.95	0.01	k_{ss}	1

We choose the updating PEA parameter $\mu = 0.5$. We set the moving-bounds parameter to $a = 0.007$, which corresponds approximately to having $\underline{k} = 0.5z_{ss}$ and $\bar{k} = 1.5k_{ss}$ after 100 iterations. We fix the length of simulation to $T = 1,000$ periods. The convergence criterion used is that the L^2 distance between vectors β obtained in two subsequent iterations is less than 10^{-5} .

By construction, the moving-bounds modification may help the PEA to converge, although it may not affect the final solution. Therefore, we will not provide any results regarding the properties of the solution; the discussion in den Haan and Marcet (1990, 1994) applies to the modified PEA without changes. We shall just illustrate how the moving-bounds method works in practice.

Figures 1 and 2 show two examples of the stochastic simulations. The series plotted in each row are consumption, capital, and the value of the expression inside the conditional expectation, respectively. The first row, "no bounds," corresponds to the first iteration of the original PEA, when no restrictions are imposed on the simulated series. The subsequent rows, " $i = 0$," " $i = 1$," and so on, are the simulated series obtained after i iterations are performed. The last row shows the final time series solution.

As we can see, when no restrictions are imposed, capital series become highly nonstationary. In the first case (Fig. 1), capital explodes quickly to almost 20 steady-state levels, whereas in the second case (Fig. 2), capital implodes to 0 in less than 100 periods. These graphs illustrate the problem of the initial point in the original PEA framework. To be specific, our initial guess of the steady-state solution here proved to be inaccurate and led to nonstationary series that may not be used in the regression. It is not surprising, therefore, that the original PEA might have difficulty in converging.

As follows from subsequent graphs, the poor initial guess does not create a problem for the modified PEA. Initially, at " $i = 0$," the bounds make the simulated series coincide with the steady state. On the next iteration, " $i = 1$," the possible range for capital increases; the capital series start fluctuating and hitting the bounds. On subsequent iterations, the solution refines and the range for capital continues to increase; the bounds are touched less and less frequently and, eventually, are never in operation. At this point, the task of the moving bounds is completed, but the iterations continue until the required accuracy in the fixed point is achieved.

One can easily check that our simple program is capable of finding the solution under any meaningful values of

the model's parameters. Furthermore, the property of convergence is not affected by a choice of the updating procedure; for example, one can assume full updating by setting $\mu = 1$. Finally, the algorithm has no difficulty in converging when the simulation length increases to 10,000 or even to 100,000 periods.

4. CONCLUSION

This paper suggests a simple modification that enhances the convergence capability of the PEA. Specifically, the modified PEA does not suffer from the problem of the poor initial guess and can systematically converge starting from the nonstochastic steady-state solution. In the example considered, the property of convergence proved to be robust to all meaningful changes in both the model's and the algorithm's parameters. We discuss only one example; however, we find the moving-bounds modification to be useful in several other applications.

ACKNOWLEDGMENTS

We thank Jeffrey Wooldridge, the associate editor, and an anonymous referee for valuable comments. We acknowledge the previous help of Albert Marcet in clarifying some issues related to the PEA. This research was partially supported by the Instituto Valenciano de Investigaciones Económicas and Ministerio de Ciencia y Tecnología, BEC 2001-0535. All errors are ours.

[Received March 2000. Revised May 2001.]

REFERENCES

- Christiano, L., and Fisher, J. (2000), "Algorithms for Solving Dynamic Models with Occasionally Binding Constraints," *Journal of Economic Dynamics and Control*, 24, 1179–1232.
- den Haan, W., and Marcet, A. (1990), "Solving the Stochastic Growth Model by Parameterizing Expectations," *Journal of Business and Economic Statistics*, 8, 31–34.
- den Haan, W., and Marcet, A. (1994), "Accuracy in Simulations," *Review of Economic Studies*, 6, 3–17.
- Marcet, A. (1988), "Solving Nonlinear Stochastic Models by Parameterizing Expectations," unpublished manuscript, Carnegie Mellon University.
- Marcet, A., and Lorenzoni, G. (1999), "The Parameterized Expectation Approach: Some Practical Issues," in *Computational Methods for Study of Dynamic Economies*, eds. R. Marimon and A. Scott, New York: Oxford University Press, pp. 143–171.
- Miranda, M., and Helmlinger, P. (1988), "The Effect of Commodity Price Stabilization Programs," *American Economic Review*, 78, 46–58.
- Wright, B., and Williams, J. (1982), "The Economic Role of Commodity Storage," *Economic Journal*, 92, 596–614.