Solving Nonlinear Stochastic Growth Models: more comparisons of alternative solution methods.[1]

by

Harald Uhlig

Dep. of Economics,

University of Minnesota

and

Institute for empirical macroeconomics

Minneapolis

[1]This project was performed at the Institute for empirical macroeconomics, which is jointly funded by the Federal Reserve Bank of Minneapolis and the University of Minnesota.

The models studied by macroeconomists are becoming increasingly complex and richer in structure. Electronic computing power becomes ever cheaper. Consequently, the demand arises for numerical techniques capable of solving these macroeconomic models, taking account of the often nonlinear and stochastic nature of the problem. These solution methods can be used to guide or support the theoretical intuition on these models or even replace them in cases where a purely theoretical analysis seems hopeless or infeasible.

This paper intends to extend the comparisons of different alternative solution techniques for nonlinear rational expectation models as performed by Taylor for the Nonlinear Rational Expectations Modelling Group. Descriptions of the different techniques can be found in Taylor resp. in papers produced by the authors themselves.

I. The task

The researchers were asked to apply their techniques to find solutions and simulations for three different problems. This paper discusses only the first problem and the different simulation results.

The model is the same as described in Taylor: there is a one good growth economy in which agents maximize

$$E_t[\sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\tau}}{1-\tau}]$$

subject to

$$c_t + k_t - k_{t-1} = \theta_t \, k_{t-1}^{\alpha}$$

and the positivity constraints $c_t > 0$ and $k_t > 0$. θ_t is assumed to follow the stochastic process

$$\log(\theta_t) = .95 \log(\theta_{t-1}) + \epsilon_t,$$

where ϵ_t are independently and identically drawn from a $N(0,\sigma^2)$–distribution. It seems clear that this model has a unique stationary distribution on $\mathbb{R}$ as solution, describable as a continuous decision rule

$$k' = f(k,\theta)$$

and

$$c = g(k,\theta)$$

on $\mathbb{R}_+ \times \mathbb{R}_+$. Therefore, it is in principle possible to measure the quality of the different solutions by some distance between the decision rule computed by that technique and the true decision rule. The following 10 cases were considered:

	β	τ	σ
1	.95	.5	.1
2	.95	1.5	.1
3	.98	.5	.1
4	.98	1.5	.1
5	.95	.5	.02
6	.95	1.5	.02
7	.95	3.0	.02

8	.98	.5	.02
9	.98	1.5	.02
10	.98	3.0	.02.

Cases 1 through 4 were already considered in Taylor.

There are three quite separate issues when comparing these solutions:

1. What properties can we find out about the economic model at hand for a particular setting of parameters.
2. How do these properties compare across cases of different parameters.
3. Given a model, what do the solutions tell us about the different solution techniques.

The research is aimed at the third issue: ultimately, we are interested in developing techniques to deal with more complex and/or more interesting models. The growth model (although considered in macroeconomic modelling) serves only as a vehicle to develop intuition on how the different techniques might behave under different circumstances. The interesting question is not so much wether the solutions demonstrate appealing results apart from the specific model at hand, but rather on how close (in some sense) the solutions comes to the theoretically correct solution and wether this behaviour is robust against the change of the model. The issue is to characterize a class of models for which a particular solution methods can be expected to yield reasonable results with respect to some metric. Economic intuition about

particular test models (like the growth model in this paper) can give guidance and intuition with regard to these properties, i.e. questions 1 and 2 above can serve as a vehicle to answer 3.

II. The simulations

At the time this paper was printed, simulations by the follwing authors were available:

B. Ingram,

A. Marcet,

G. Tauchen,

J. Gagnon,

W. Coleman,

and several solutions, provided by L. Christiano, labeled

LogLQ with normal disturbances,

LogLQ with discrete Markov–disturbances,

LinLQ – normal,

LinLQ – Markov,

dynamic programing (only for cases 5 – 10).

We hope to include simulations by E. McGrattan, P. Labadie and C. Sims later on.

The solution methods are described in Taylor resp. in papers by the various authors themselves.

III. Numerical results

A complete list of numerical results for all 10 cases can be found in the numerical appendix. We calculated[2]

- contemporaneous covariance matrices of consumption and capital,
- auto– and cross– correlations of consumption and capital for up to 10 lags,
- univariate autoregressions for consumption
- univariate autoregressions for capital
- regressions of consumption resp. capital on both consumption and capital

These statistics were also computed in Taylor for the cases 1 through 4. Furthermore, we computed

- contemporaneous covariance matrices of first differences in consumption and capital
- regressions of first differences in consumption and capital on both consumption and capital, including R^2.

The later calculation can be used for testing the random walk hypothesis for consumption with the simulations[3], since R^2 and the F–statistic are equivalent (given degrees of freedom and length of time–series, which should not matter much here, since our simulations are usually 1000 entries long or longer). An R^2 close to zero supports the random walk hypothesis.

[2]The calculations were performed with RATS. To a substantial part, the program is due to A. Hossain.

[3]An issue about unit roots does not arise if the solution method properly calculates a stationary distribution for the variables.

The differences among the models can be quite substantial. An extraction of the R^2–statistics from the regressions with 4 lags on past data[4] yields the following table:

case	1	2	3	4	5	6	7	8	9	10
Ingram	.71	.77	.84	.85	.56	.84	.91	.67	.83	.95
Marcet	.42	.05	.03	.03	.05	.05	.05	.03	.03	.03
Tauchen	.16	.37	.34	.27	.50	.38	.33	.35	.28	.27
Gagnon	.23	.06	.08	.09	.17	.04	.04	.06	.06	.06
LogLQ –No	.36	.04	.12	.04	.42	.05	.02	.23	.04	.01
LogLQ –Ma	.23	.10	.14	.04	.32	.05	.02	.27	.03	.01
Coleman	.41	.05	.34	.03	.44	.06	.03	.36	.04	.02
LinLQ –No	.16	.06	.08	.04	.33	.05	.02	.28	.03	.01
LinLq –Ma	.35	.05	.23	.02	.35	.05	.02	.23	.02	.01
Dyn.Progr					.37	.04	.02	.29	.02	.02

Since the random–walk hypothesis might be considered an important issue in this model, the different solution techniques seem to be rather far apart in that they deliver different answers to the same question.

[4] I.e. we estimated the regression equation

$$c_t - c_{t-1} = \alpha + \sum_{j=1}^{4} \beta_j c_{t-j} + \sum_{j=1}^{4} \gamma_j k_{t-j} + \epsilon_t$$

In the next table, we computed $\sigma^2_{\Delta k} / \sigma^2_{\Delta c}$, where Δ denotes first differencing of the time series and σ denotes the variance. This ratio measures the relative volatility of investment vs consumption–increases.

case	1	2	3	4	5	6	7	8	9	10
Ingram	48.6	130.	350.	49.9	1.3	111.	357.	1.3	79.8	1078
Marcet	29.8	14.3	76.7	66.3	10.9	10.9	10.9	63.5	63.6	63.5
Tauchen	7.8	2.2	1.7	1.6	2.6	2.1	2.1	1.6	1.7	1.7
Gagnon	19.8	3.8	2.2	2.8	17.1	4.6	3.4	7.8	3.7	3.2
LogLQ –No	23.5	9.1	48.5	36.2	29.2	11.0	7.9	131.	60.2	46.1
LogLQ –Ma	19.3	21.4	48.6	30.5	26.5	10.2	9.1	139.	55.6	41.2
Coleman	24.3	11.7	101.	54.1	26.1	11.6	10.2	106.	53.0	55.2
LinLQ –No	7.8	3.6	22.3	12.4	24.5	9.9	7.9	136.	48.6	48.9
LinLq –Ma	24.4	9.1	115.	44.1	24.4	9.1	7.3	115.	44.1	37.3
Dyn.Progr					28.0	10.0	8.2	149.	52.7	46.7

Again, the results show large differences. Note in particular in both tables the dependencies of the results on the parameters τ and (for the second table) the discount factor β.

IV. Graphical results

Often, regression results and second order properties of a simulation can look reasonable although one might not want to accept a simulation if one could actually "see" how it is moving through some higher–dimensional state space. It can happen, that the time series calculated by the method lie on some rather restricted, possibly nonlinear subspace or that they move between and then stay around certain "islands" in the fashion of strange attractors.

Therefore, looking at various plots of the simulated time series is another method besides calculating regressions and correlations to further our understanding and intuition on how the different methods behave. I now describe a few graphs given in the appendix of this paper. These description might not have an important impact in analyzing economic questions of the growth model at hand. The point however is, that one might expect these solution methods to exhibit similar or related behaviour in other models. In these models then it would be bold to expect that these behaviours never have an impact on answering particular, important economic questions. The purpose of this paper is to use the growth model to analyze the solution methods and not the other way around.

In set 1 of the graphs, simulations to case 4 are shown. I appeal to the naked eye for the obvious differences. All these plots were done within the same bounds. To get a better idea, on how in particular Ingrams solution and the LinLQ – Markov solution behave, a plot in individual bounds is added.

One might suspect from these plots, that some of these simulations move in a particularly remarkable fashion. e.g. Ingrams solution seems to exhibit some spiraling motion, whereas Gagnons solution shows periods of drops in consumption together with periods of extending the capital stock. This is in fact true and can be seen better in the plots of set 2, in which subsequent data points are connected by an arrow–pointed line. We can clearly see the "island–circling" behaviour in Ingrams as well in Marcets solution and the phase–change pattern in Gagnons simulation.

Taking first differences yields investment and consumption–increases (possibly negative). The plots of set 3 exhibit the behaviour of the different methods, again in common bounds. The pattern in Tauchens solution is rather striking, but also the saddle–shaped form in the LogLQ – Normal simulation appears somewhat strange.

For that reason, we reproduced Tauchens solution and the LogLQ – Normal solution for all 10 cases in set 4. Except for case 1, the striking pattern appears in all of Tauchens simulations. The saddle–shape of the LogLQ – Normal solutions seems to disappear more or less in cases 5 through 10, which use a variance of the error–term in the autoregression equation for $\log(\theta_t)$, which is one fifth of the variance in the cases 1 through 4. We appended some plots for case 3 of the actual movement of the time–series, indicated by arrows connecting subsequent data points.

The output–vs–capital plots in set 5 show quite different behaviours across different

solution methods: the "island–to–island–hopping" in Ingrams case, the meandering of Gagnons solution or the strange pattern in the two LinLQ – solutions, which we again reproduced for cases 5 and 10 for comparison.

Plots of theta–vs–other series (set 6) usually look rather similar (see simulation of Marcet and LinLQ – Normal) except for the simulation by Tauchen and through dynamic programming.

Finally in set 7, we included plots of histograms of consumption for cases 4 and 8 (these are essentially steady state distributions). The differences between the histograms can be quite striking, e.g. with respect to properties like single–peakedness or multi–peakedness etc (we chose a rather coarse histogram–grid to generate smoothness of the distribution function without loosing too much information about its shape).

V. Conclusions

It appears that even in such a simple model, the different solution methods can yield quite different results with respect to their econometric behaviour as well as their pattern of movement as demonstrated by the graphs. It seems essential to get a better intuition for where these differences come from and how big they can be before relying on conclusions drawn from these methods e.g. for policy analysis.

On a technical note, it seems a good idea to use the same shock–series and initial conditions for all simulations to do a direct comparison. Also would be nice to actually generate 3D–plots of the decision rule for this problem, defined over a sufficiently fine grid of θ's and k's. Furthermore, one would like to compare costs of computation for the different methods by e.g. calculating the solutions on the same computer and measuring the time.

As for other models, some questions are still wide open:

- what happens in models that already exhibit a multiplicity of solutions (finitely many / a continuum). E.g. what equilibrium gets picked in a stochastic OLG–model with money ?
- what happens in models that don't show stationarity (like e.g. growth models with $\log(\theta_t) = \log(\theta_{t-1}) + \epsilon_t$. Is it always necessary to transform the model into a stationary one?
- Is it possible on a theoretical basis to prove results on e.g. stability or convergence of the method (and the convergence–speed) for some class of sufficiently regular models. How do the different methods compare?
- It might be interesting to select a particularly interesting economic issue and carefully construct a stochastic model for it. The different techniques can then be used to analyze this model. One can then compare the different answers and policy conclusions generated by these techniques.

Surely some of these questions have been answered already for some of the solution

methods. It might be nice to see a tabulation of those results that are available.

Taylor, J.B., "Solving nonlinear stochastic growth models: a comparison of alternative solution methods", preliminary draft Aug. 1988

case 04, capital(t) vs. consumption(t)

simulation Ingram

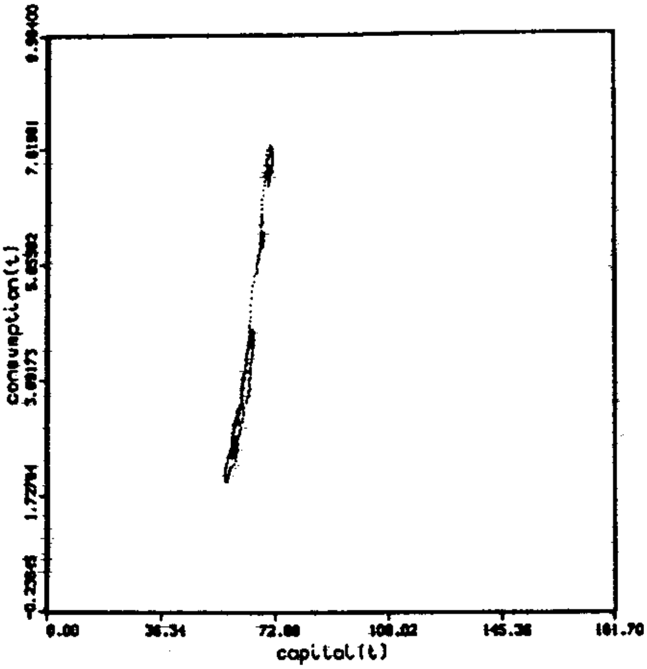

case 04, capital(t) vs. consumption(t)

simulation Touchen

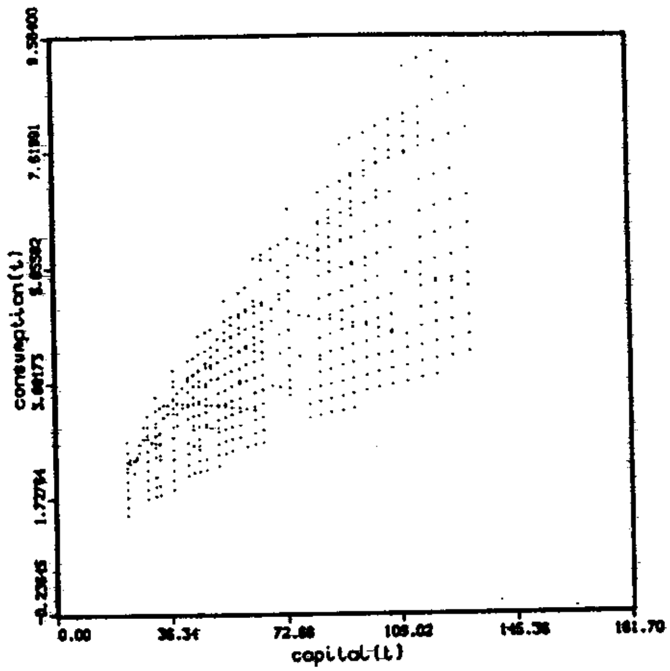

case 04, capital(t) vs. consumption(t)

simulation Marcel

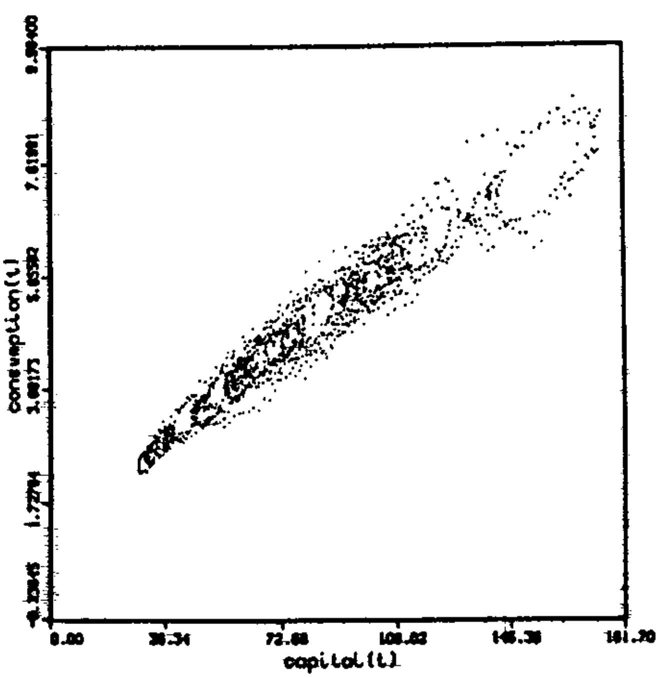

case 04, capital(t) vs. consumption(t)

simulation Gagnon

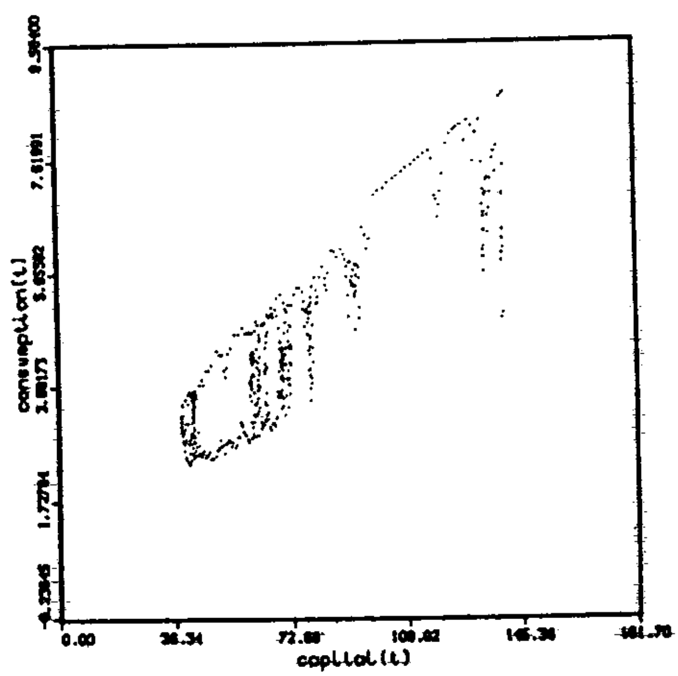

case 04, capital(t) vs. consumption(t)

simulation LogLQ - Normal

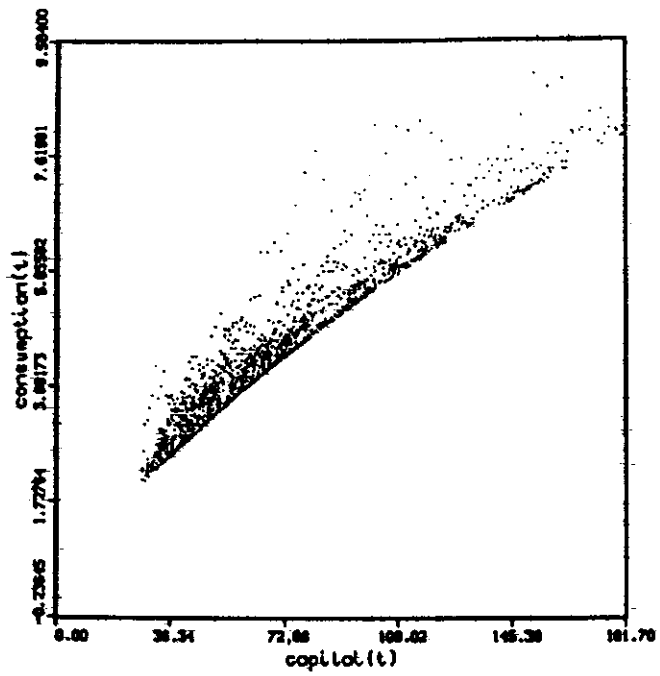

case 04, capital(t) vs. consumption(t)

simulation Coleman

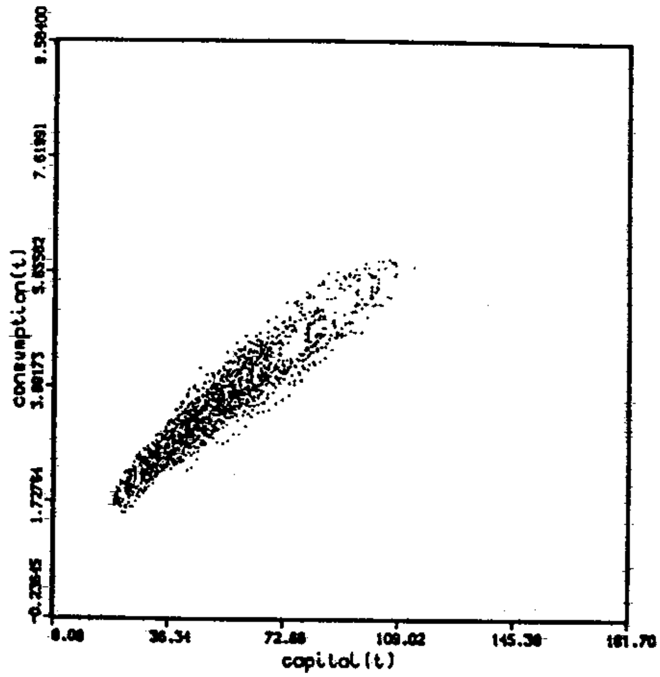

case 04, capital(t) vs. consumption(t)

simulation LogLQ - Markov

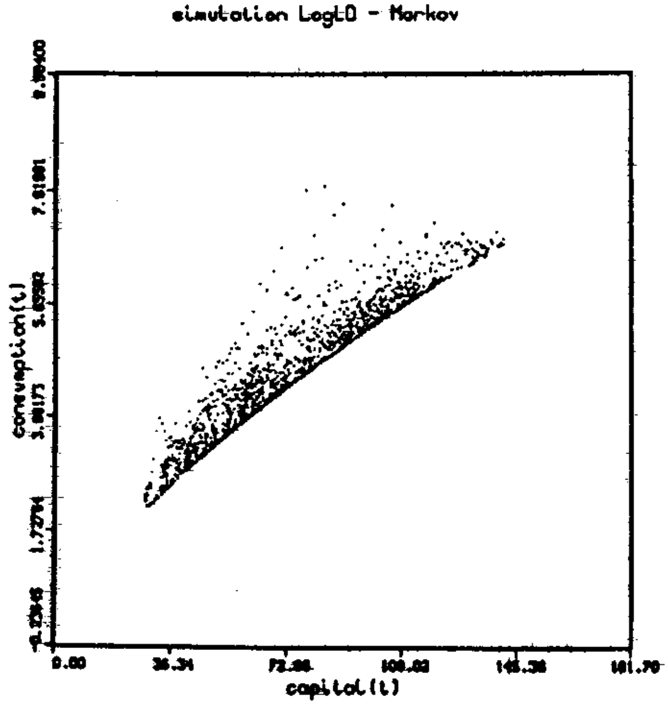

case 04, capital(t) vs. consumption(t)

simulation LinLQ - Normal

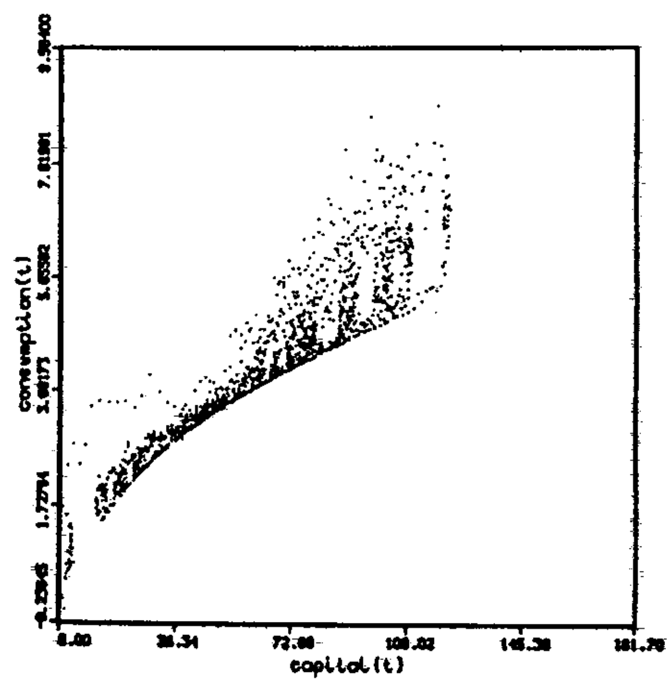

case 04, capital(t) vs. consumption(t)

simulation LinLQ - Normal

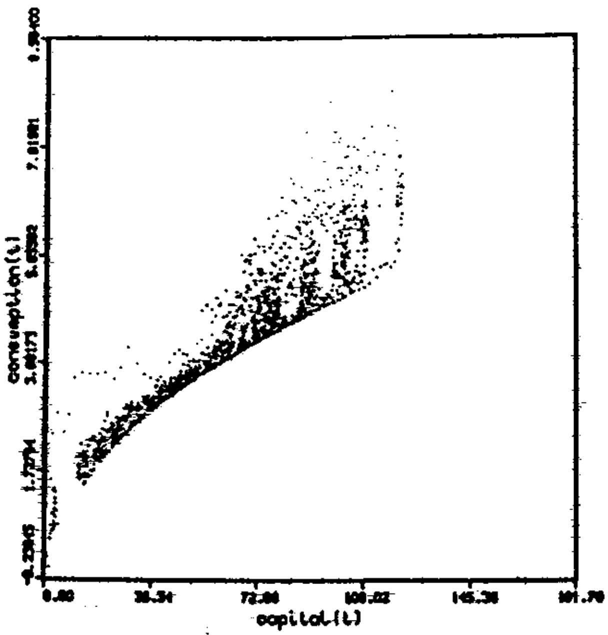

case 04, capital(t) vs. consumption(t)

simulation LinLQ - Markov

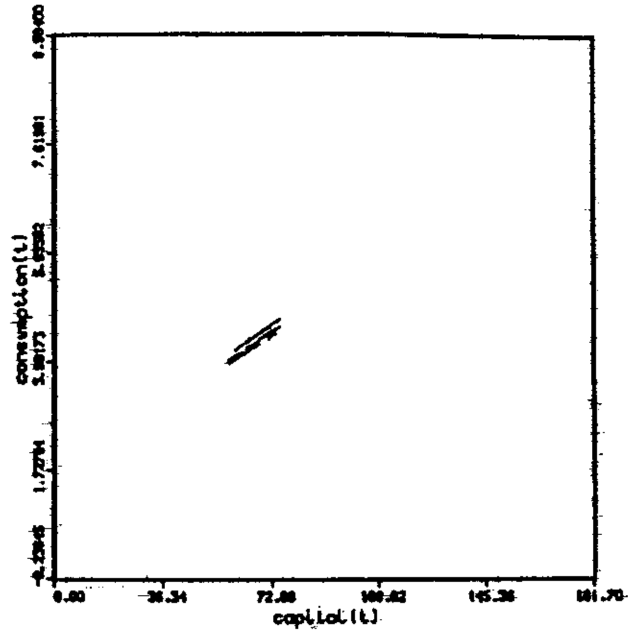

case 01, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

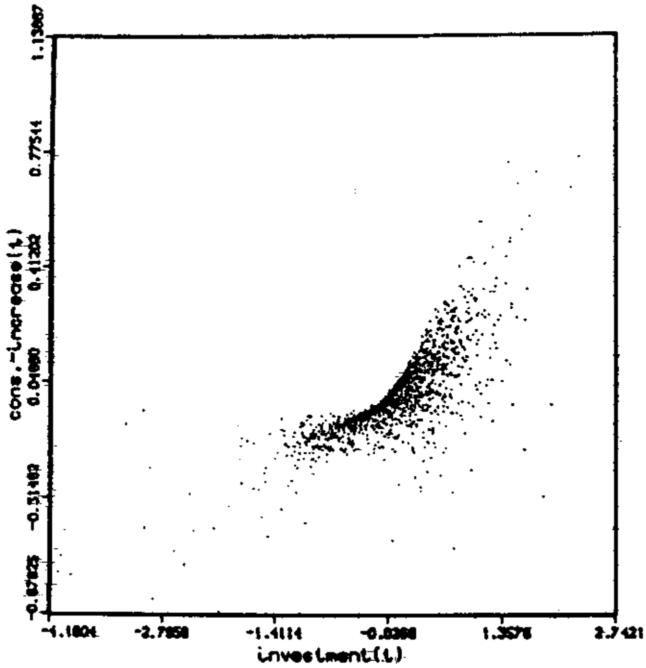

case 02, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

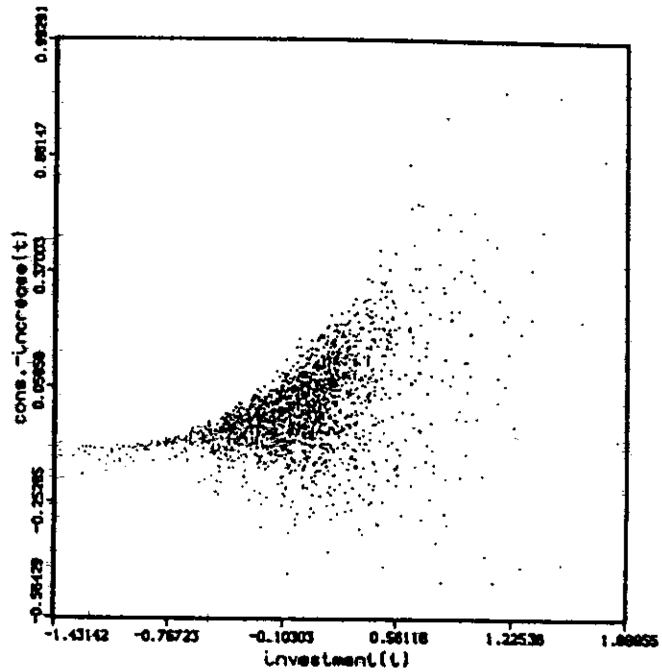

case 03, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

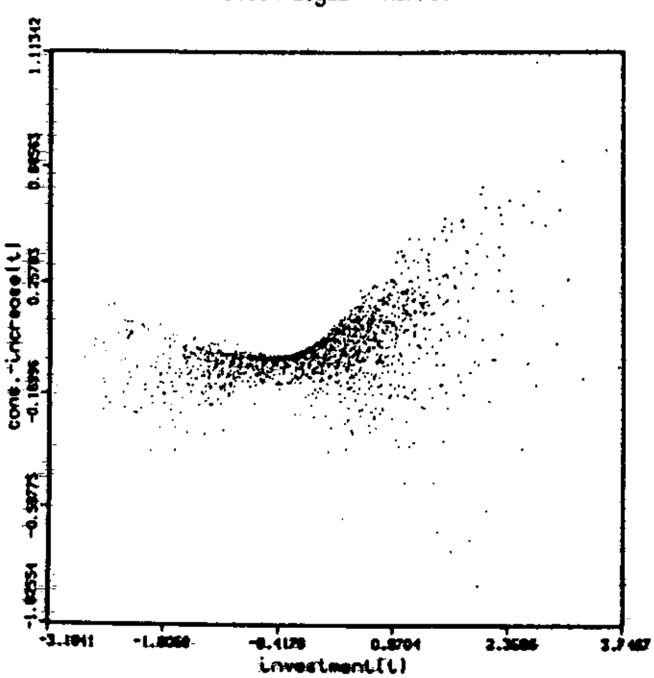

case 04, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

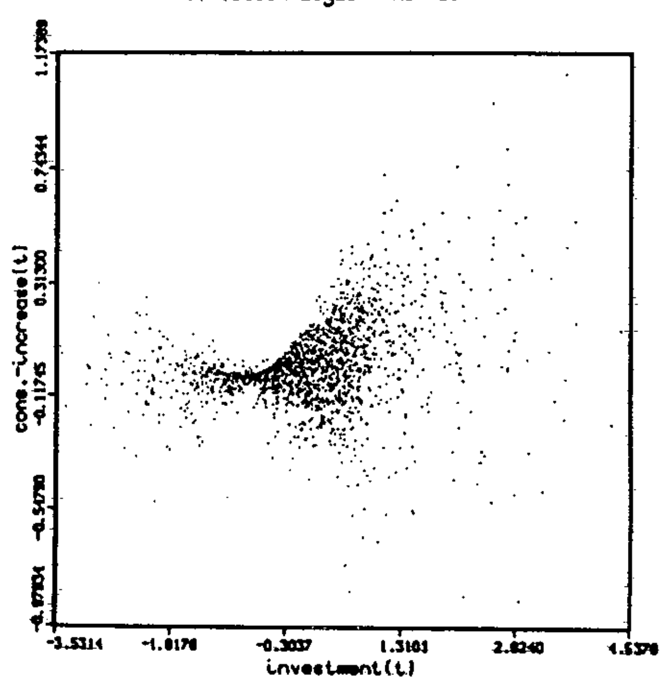

case 05, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

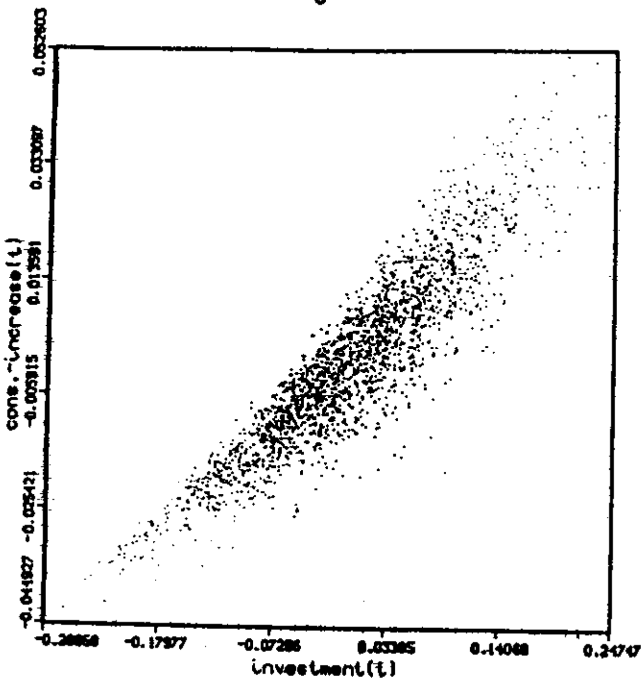

case 06, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

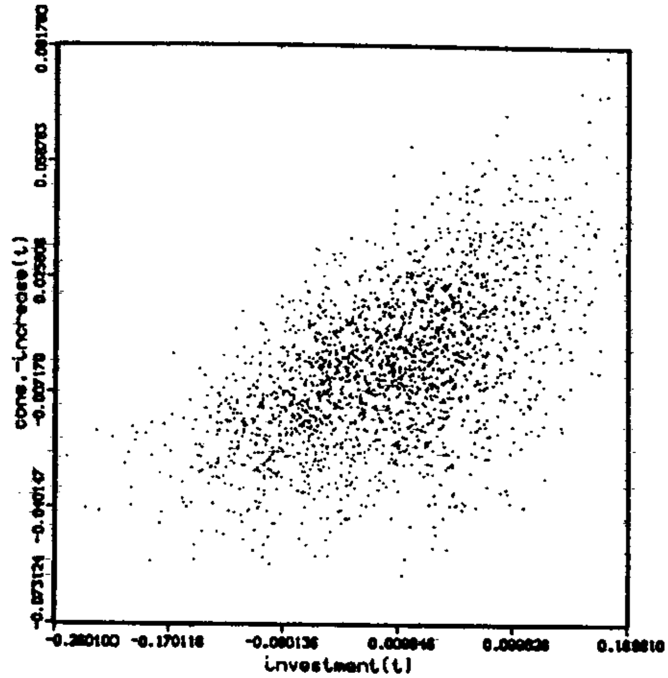

case 07, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

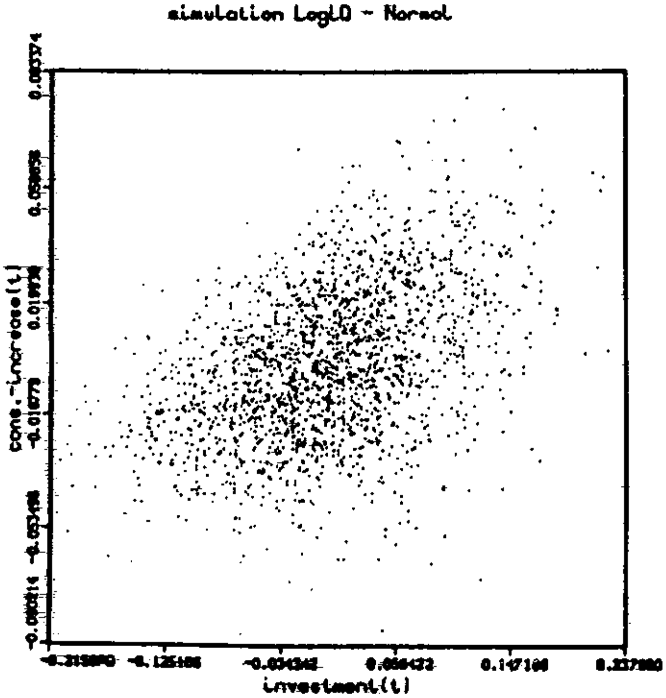

case 08, investment(t) vs. cons.-incr.(t)

simulation LogL0 - Normal

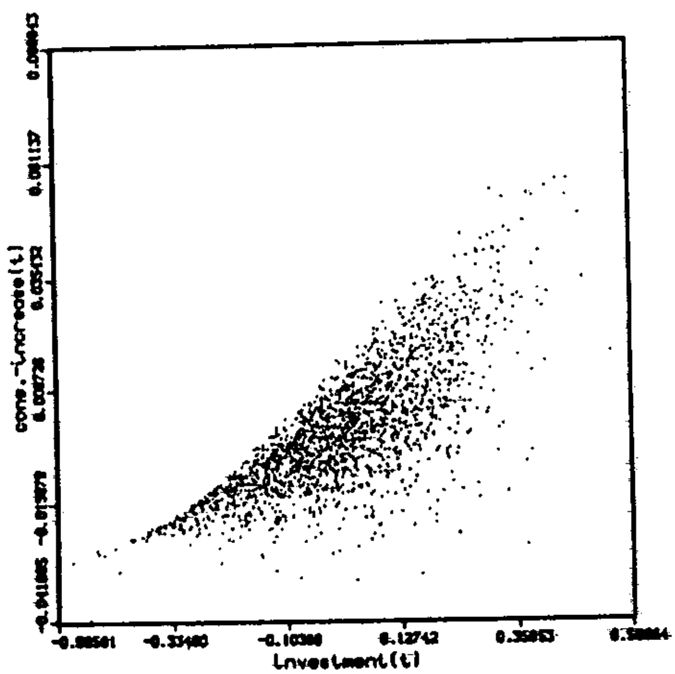

case 09, investment(t) vs. cons.-incr.(t)

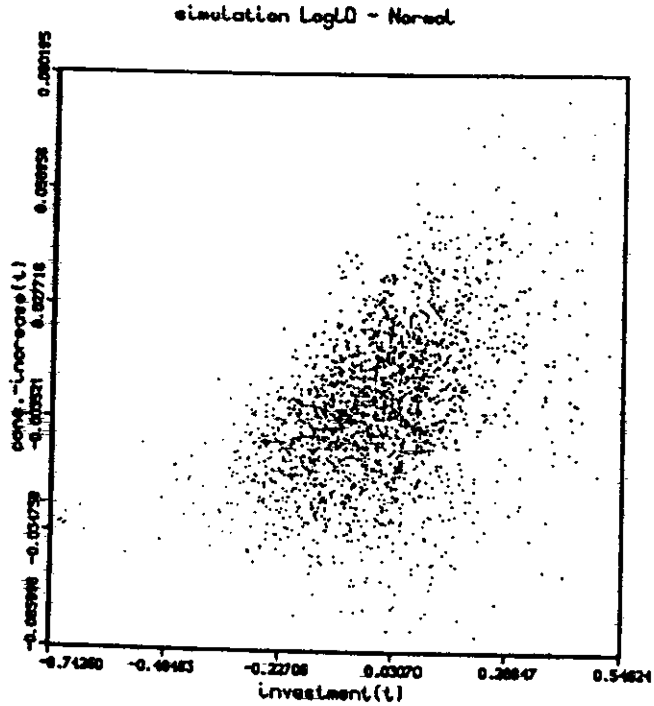

:ase 10, investment(t) vs. cons.-incr.(t)

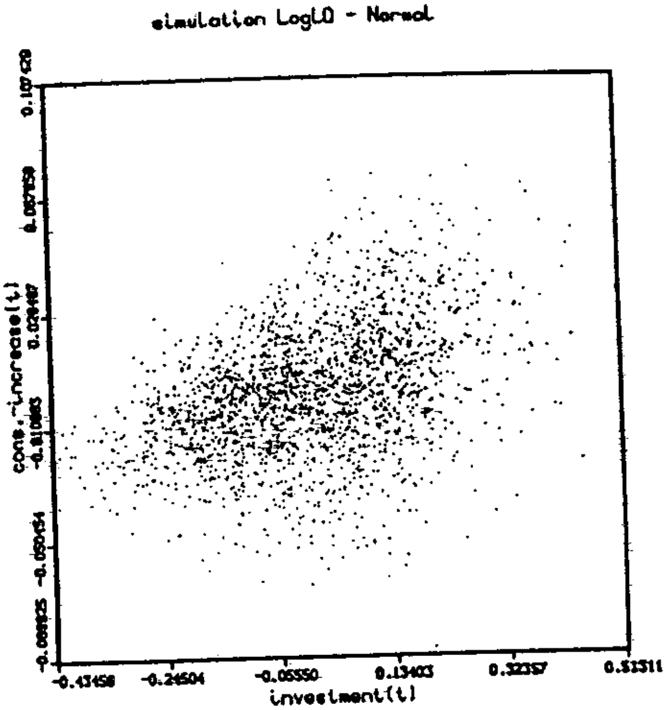

case 01, investment(t) vs. cons.-incr.(t)

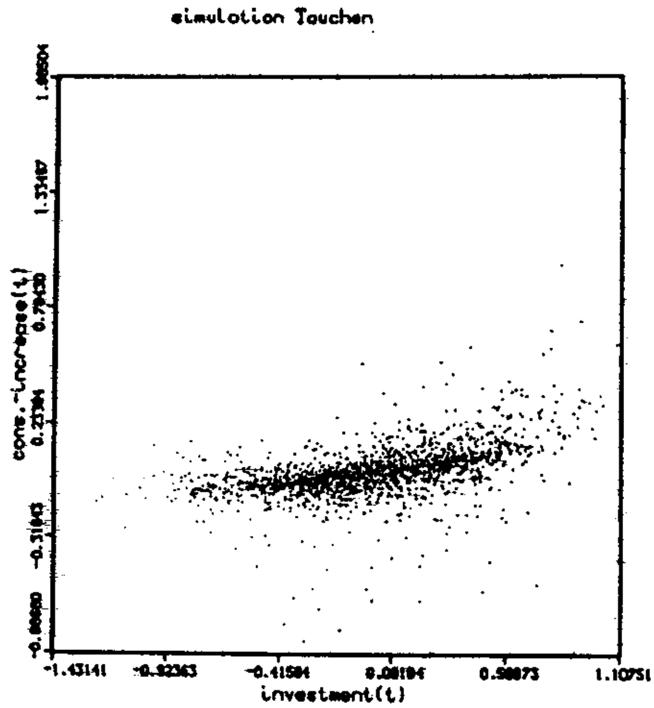

case 02, investment(t) vs. cons.-incr.(t)

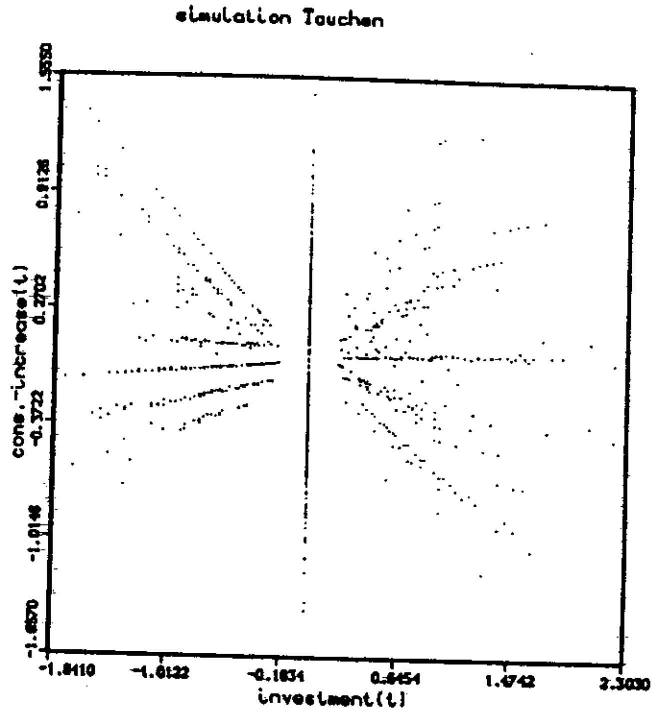

case 03, investment(t) vs. cons.-incr.(t)

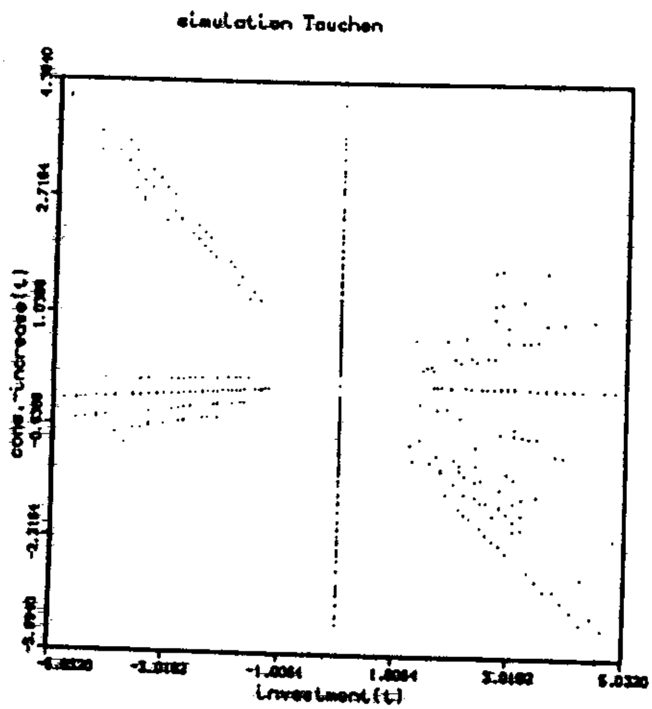

case 04, investment(t) vs. cons.-incr.(t)

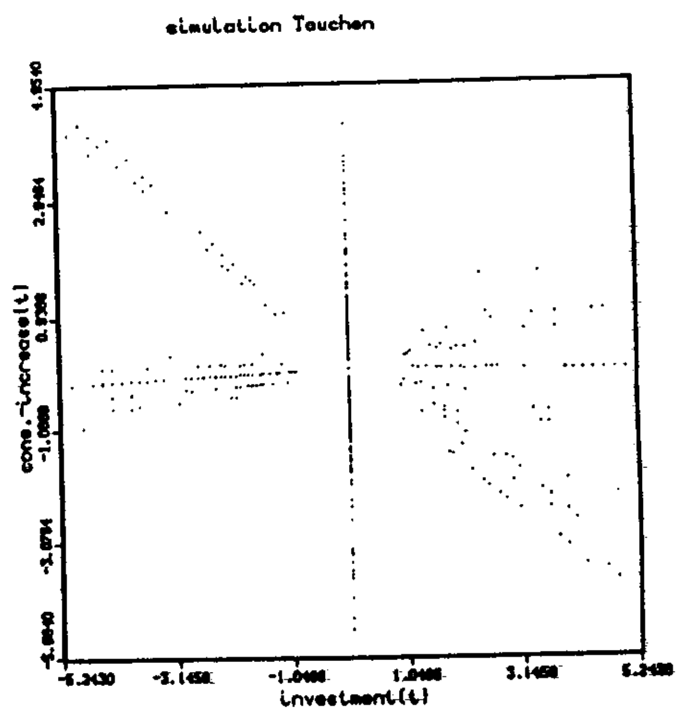

case 05, investment(t) vs. cons.-incr.(t)

simulation Touchen

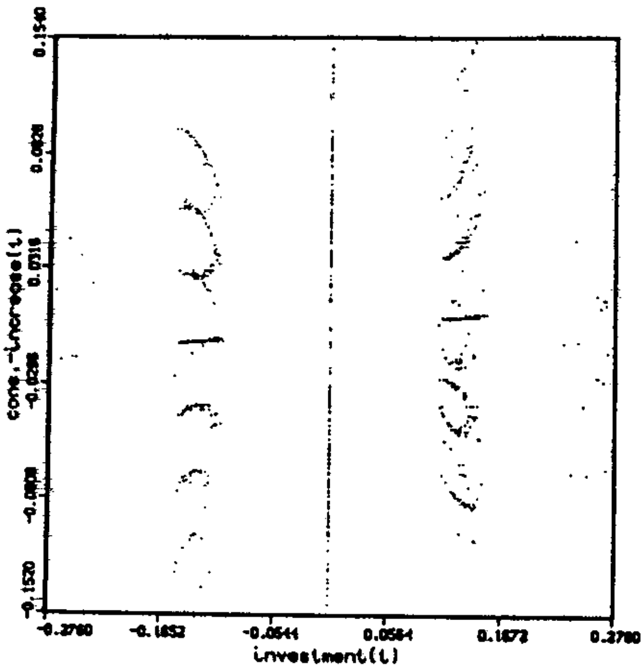

case 06, investment(t) vs. cons.-incr.(t)

simulation Touchen

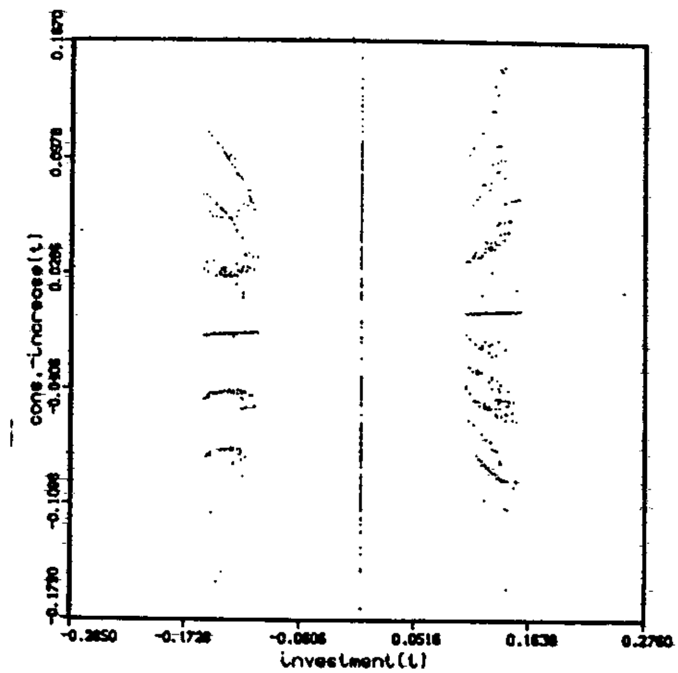

case 07, investment(t) vs. cons.-incr.(t)

simulation Touchen

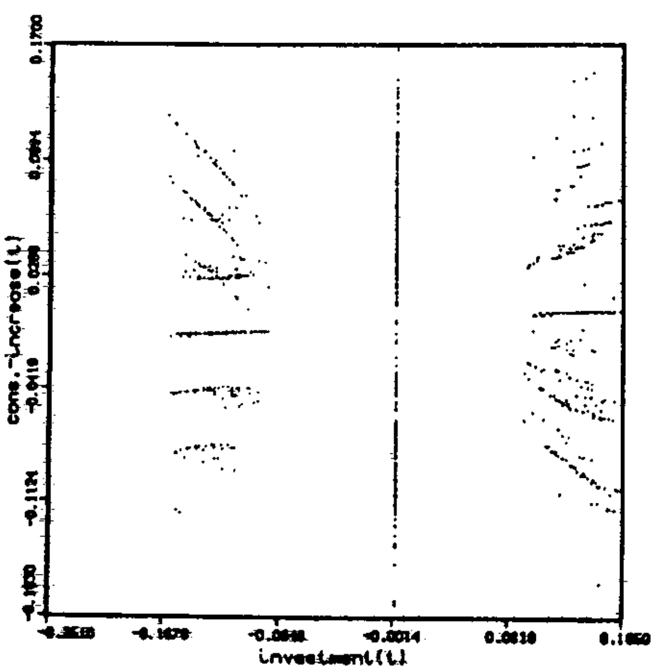

case 08, investment(t) vs. cons.-incr.(t)

simulation Touchen

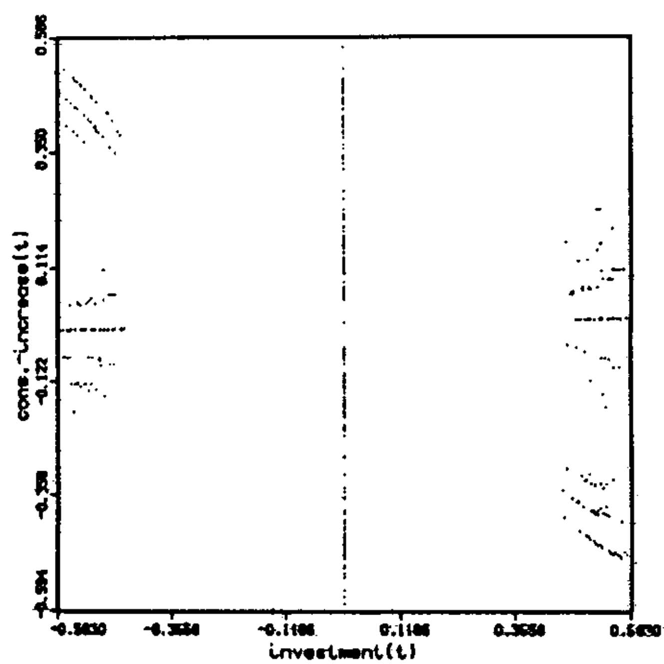

case 09, investment(t) vs. cons.-incr.(t)

simulation Touchen

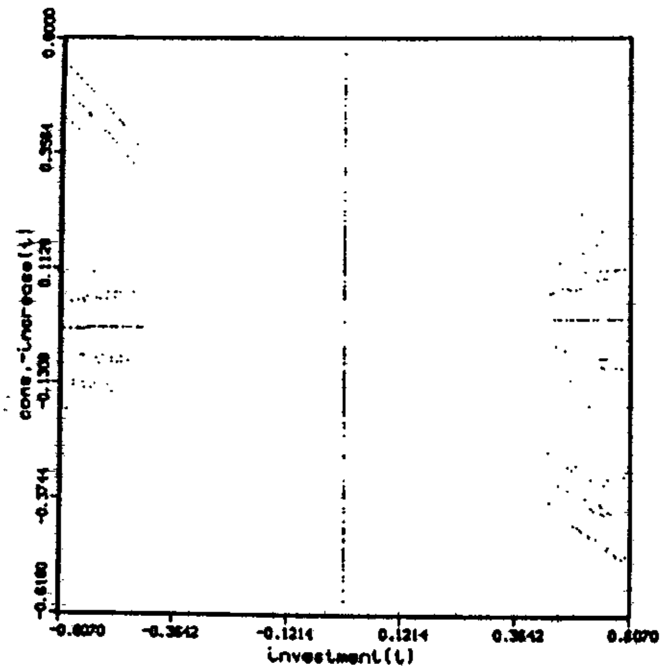

case 10, investment(t) vs. cons.-incr.(t)

simulation Touchen

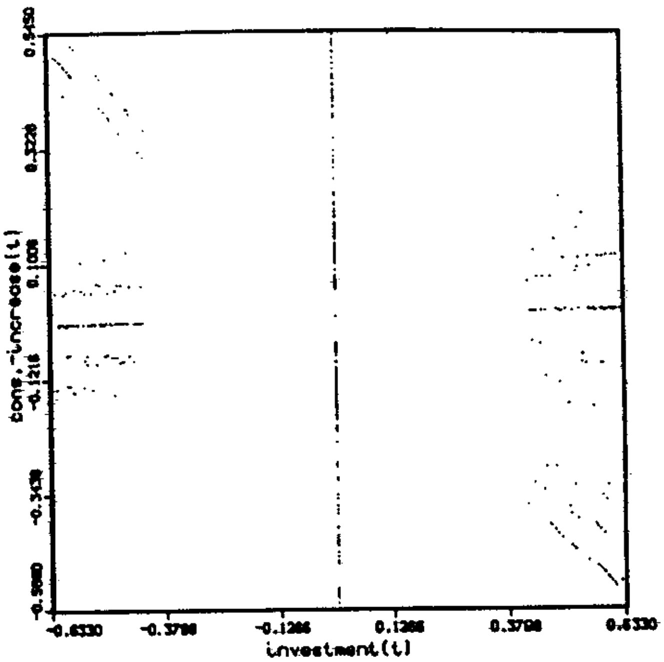

case 03, investment(t) vs. cons.-incr.(t)

simulation LogLQ - Normal

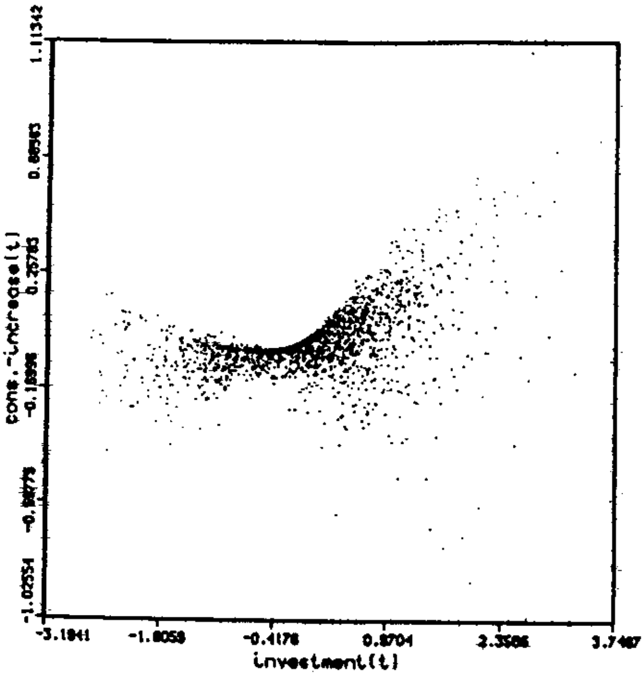

case 03, investment(t) vs. cons.-incr.(t)

simulation LogLQ - Normal

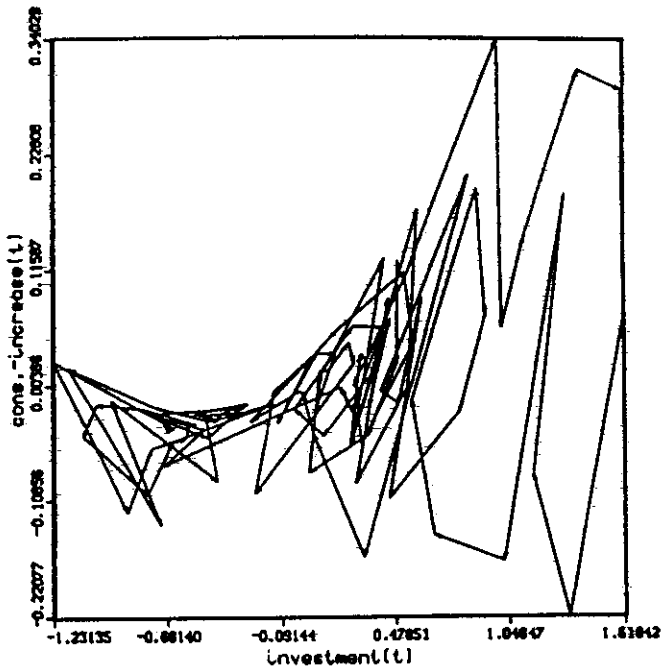

case 03, investment(t) vs. cons.-incr.(t)

simulation Tauchen

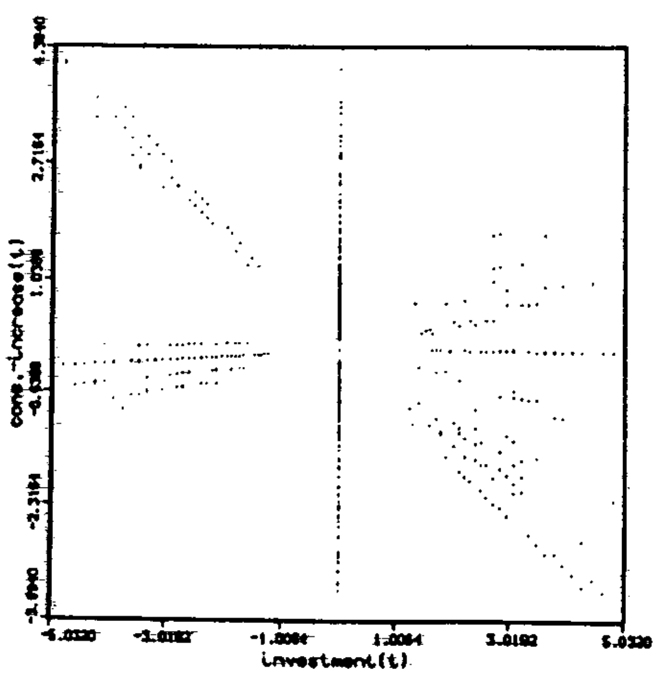

case 03, investment(t) vs. cons.-incr.(t)

simulation Tauchen

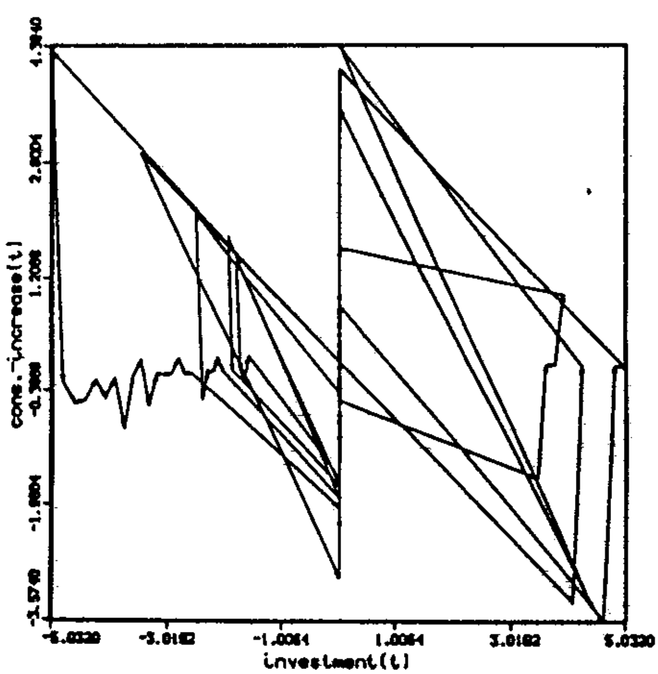

case 04, output(t) vs. capital(t)

simulation Ingram

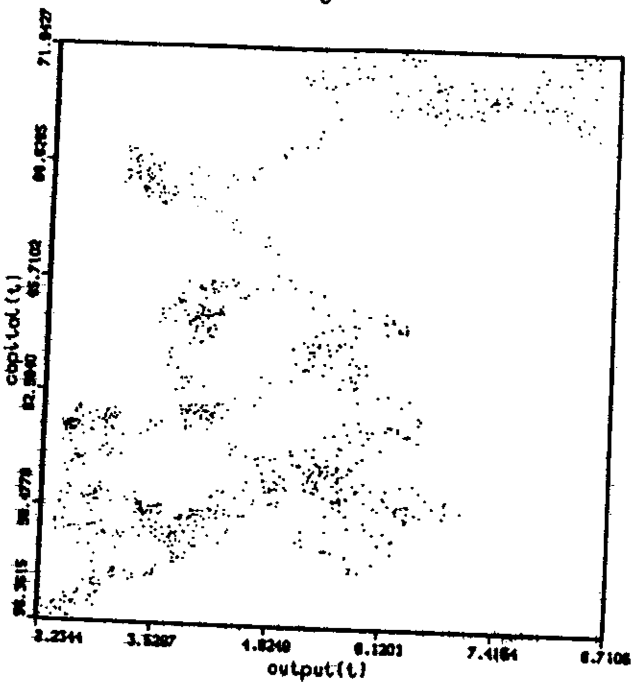

case 04, output(t) vs. capital(t)

simulation Tauchen

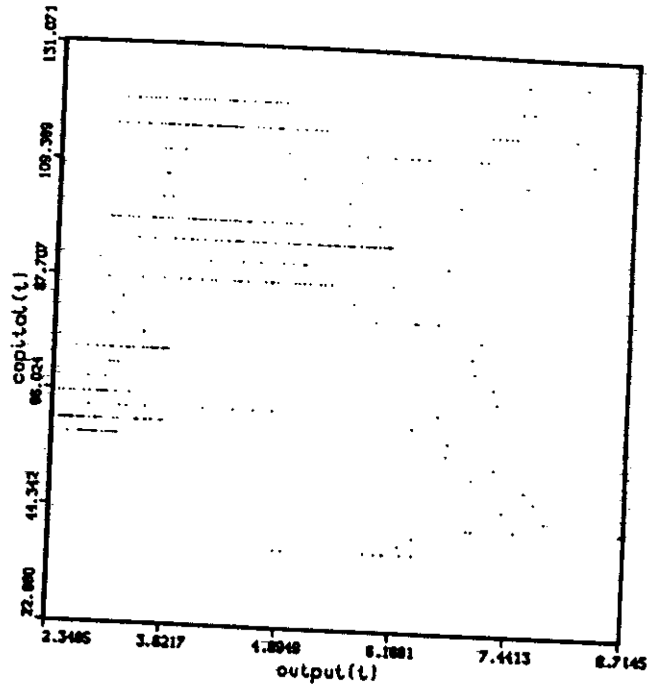

case 04, output(t) vs. capital(t)

simulation Marcel

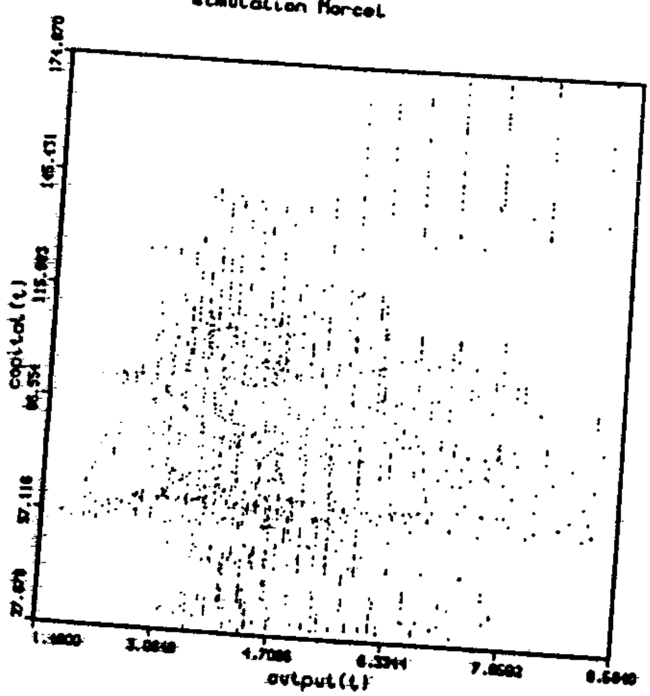

case 04, output(t) vs. capital(t)

simulation Gagnon

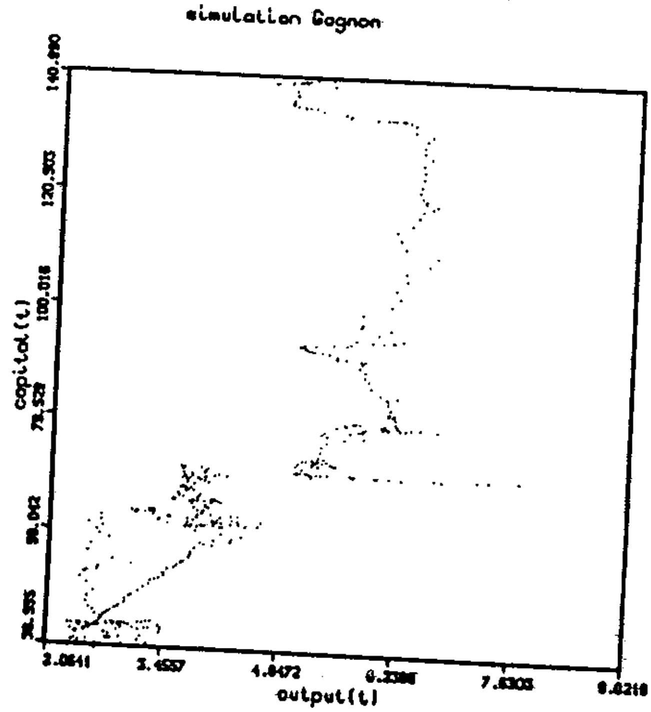

case 04, output(t) vs. capital(t)

simulation LogLQ - Markov

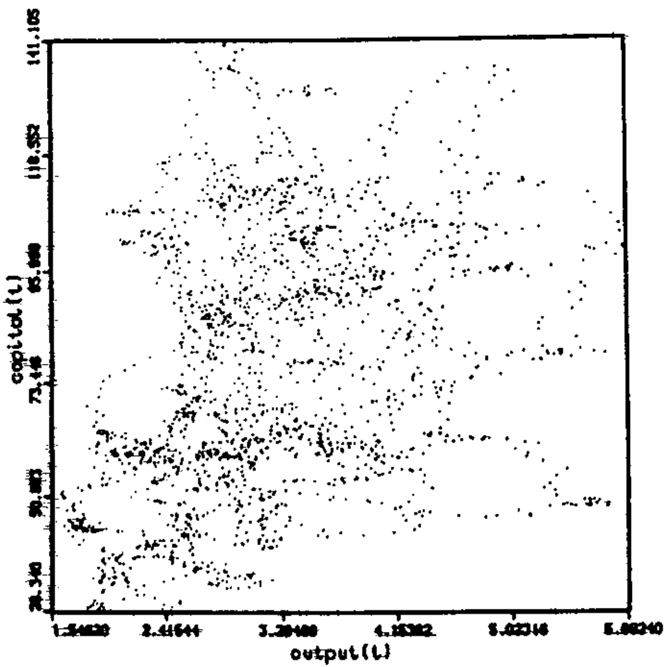

case 04, output(t) vs. capital(t)

simulation LinLQ - Normal

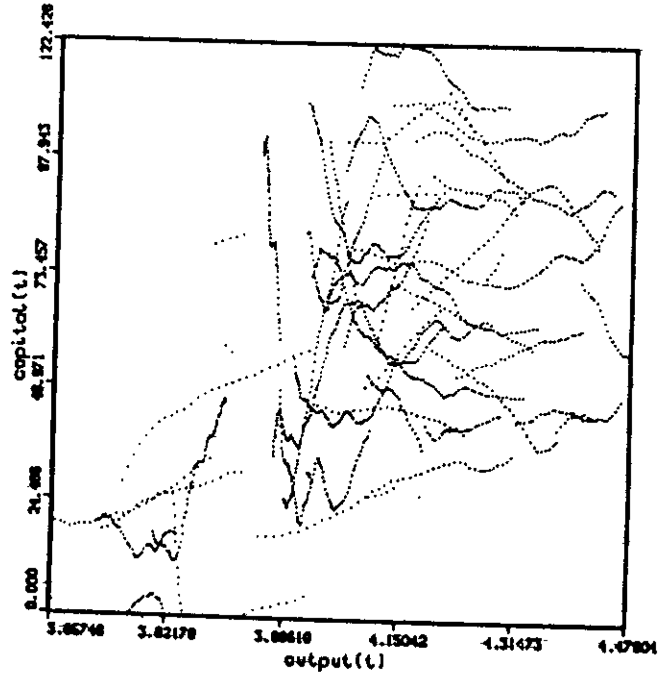

case 04, output(t) vs. capital(t)

simulation LogLO - Normal

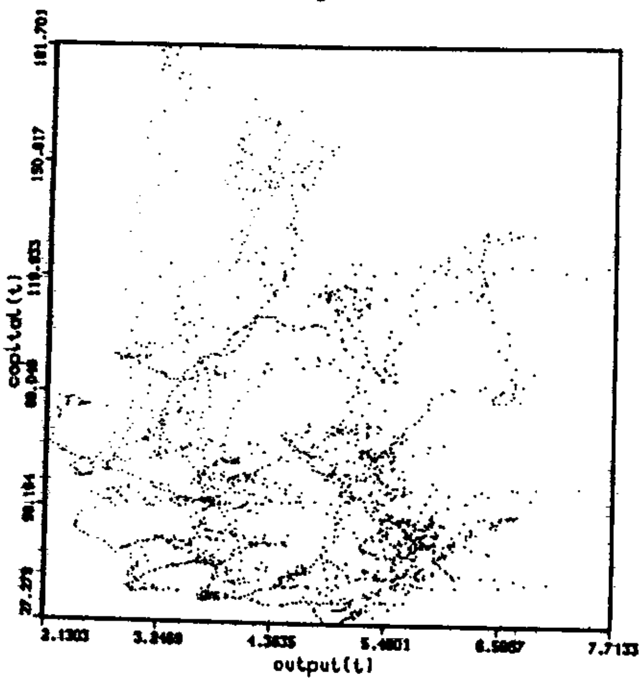

case 04, output(t) vs. capital(t)

simulation Coleman

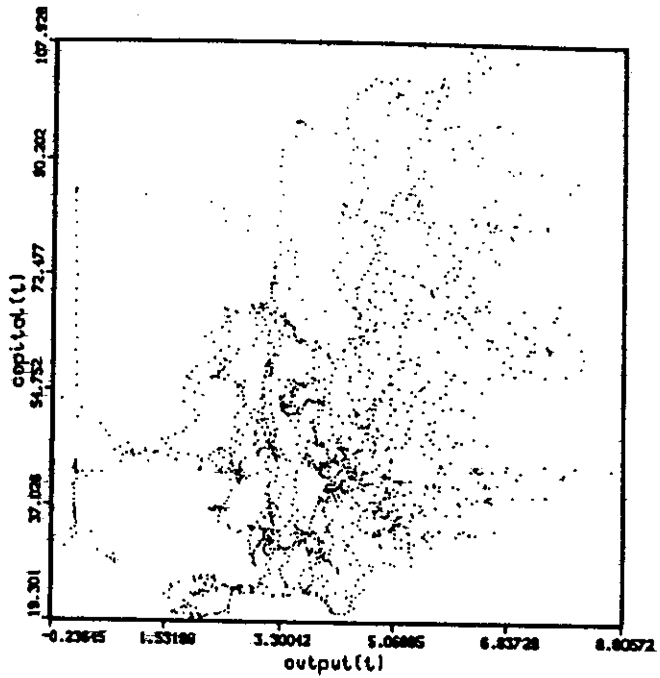

case 05, output(t) vs. capital(t)

simulation LinLQ - Normal

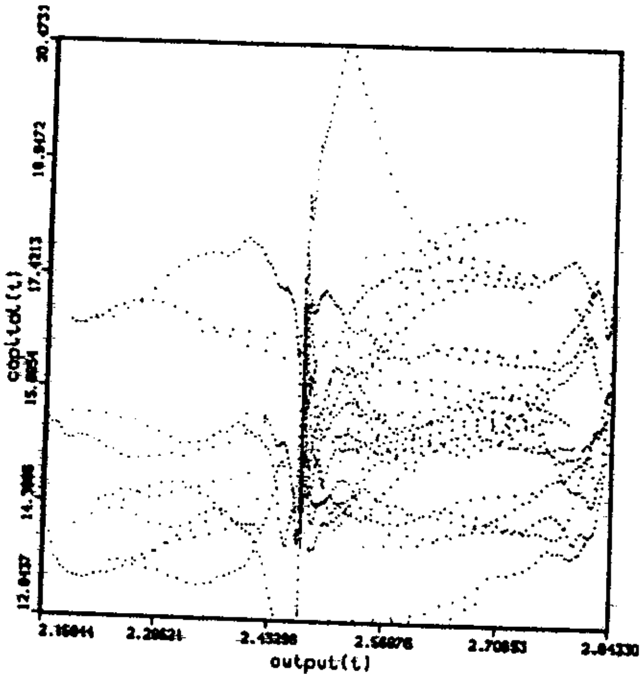

case 05, output(t) vs. capital(t)

simulation LinLQ - Markov

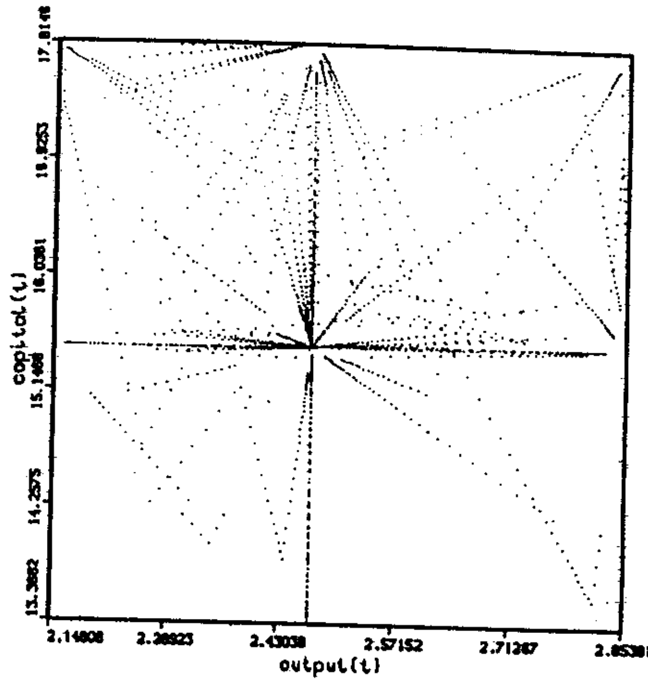

case 10, output(t) vs. capital(t)

simulation LinLQ - Normal

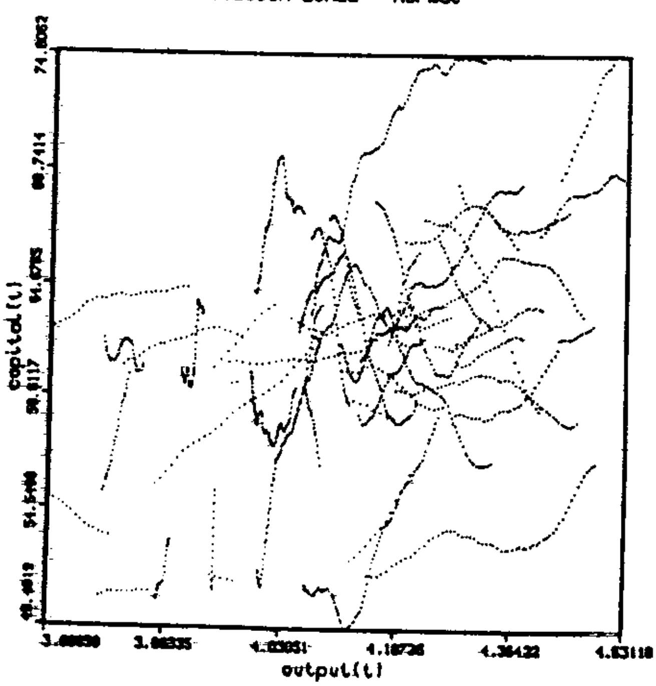

case 10, output(t) vs. capital(t)

simulation LinLQ - Markov

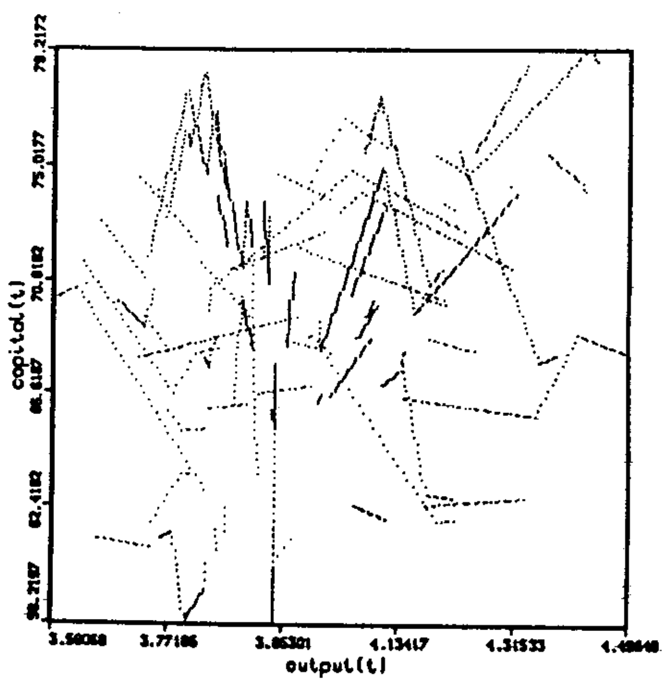

case 08, theta(t) vs. cons.-incr.(t)

simulation Marcet

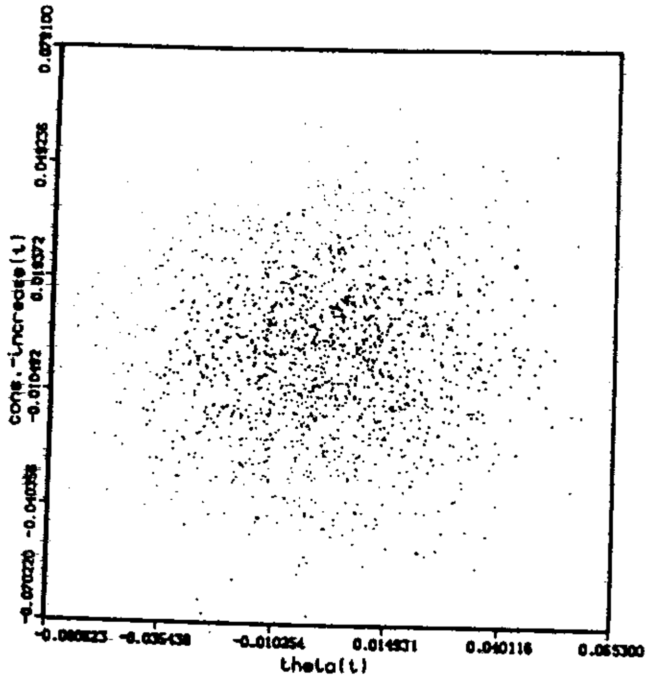

case 08, theta(t) vs. cons.-incr.(t)

simulation LinLQ - Normal

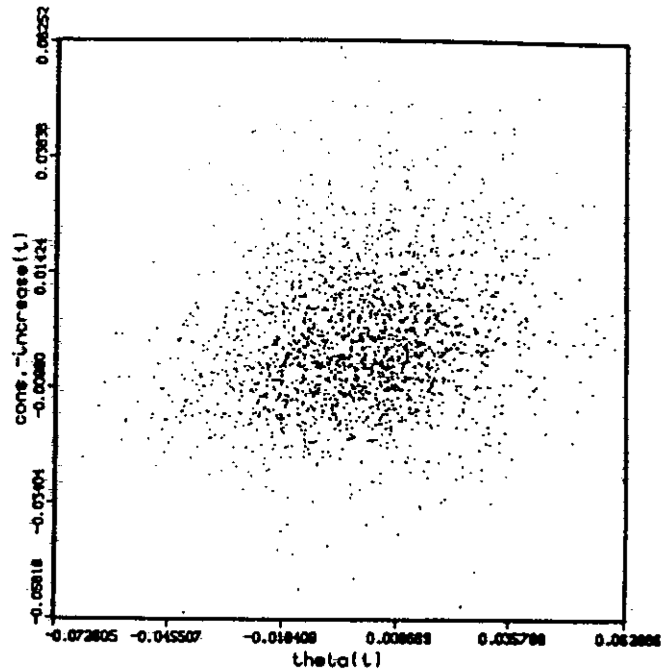

case 08, theta(t) vs. cons.-incr.(t)

simulation Touchen

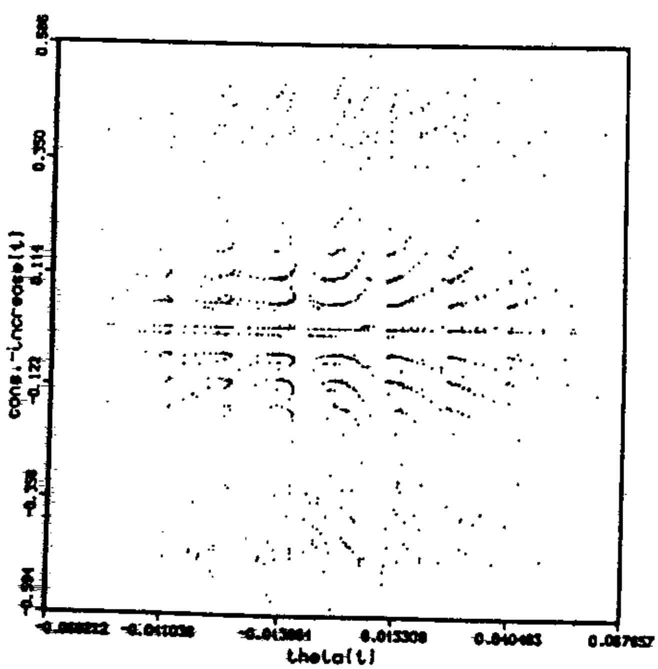

case 08, theta(t) vs. cons.-incr.(t)

simulation Dynamic Programing

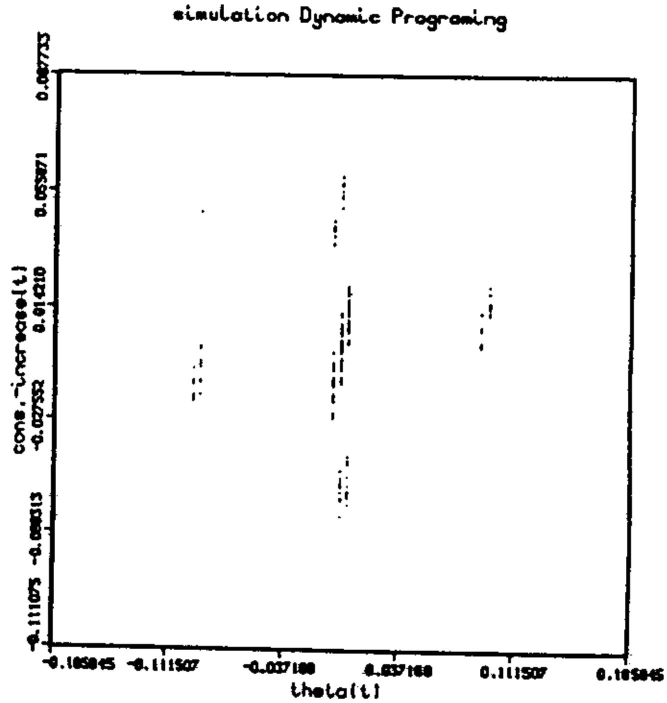

case 08, consumption(t), histogram

simulation Ingram

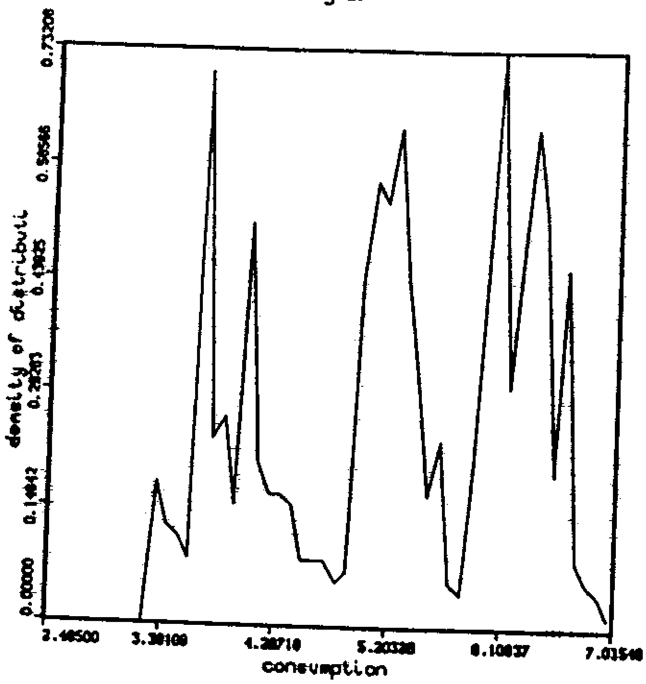

case 08, consumption(t), histogram

simulation Tauchen

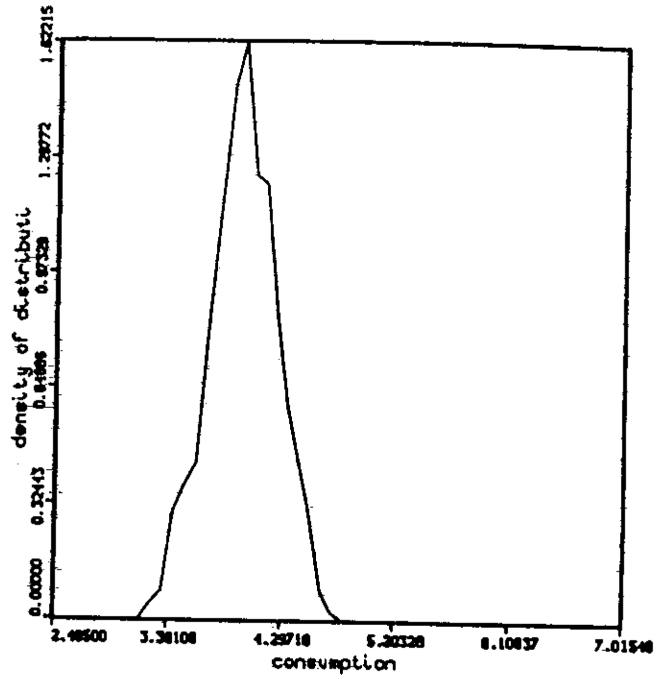

case 08, consumption(t), histogram

simulation Marcet

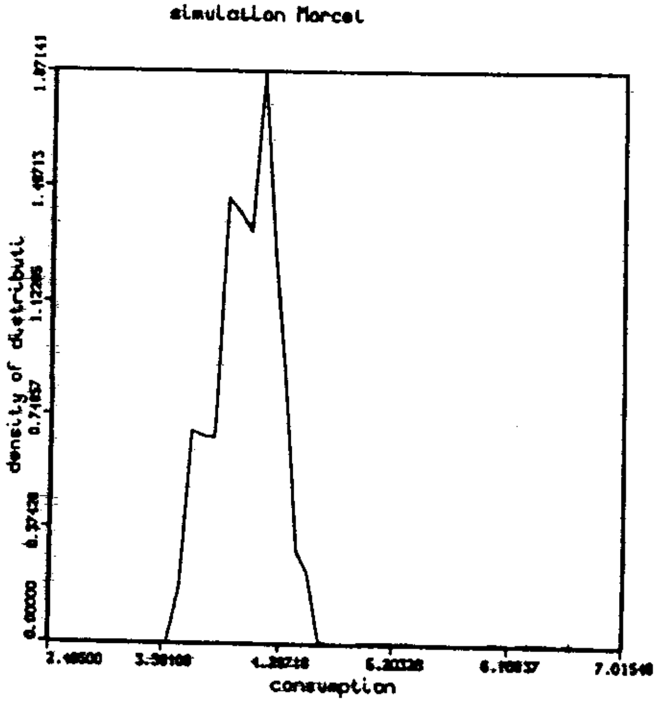

case 08, consumption(t), histogram

simulation Gagnon

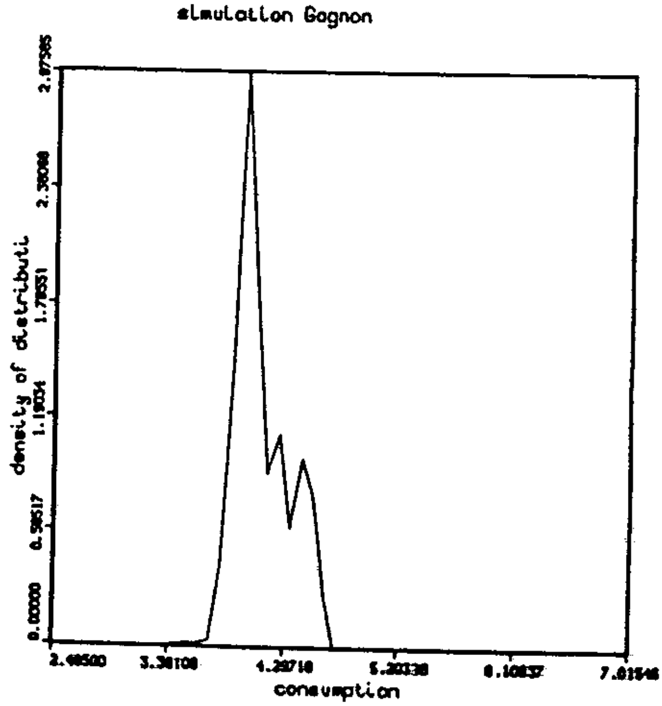

case 08, consumption(t), histogram

simulation LogLQ - Normal

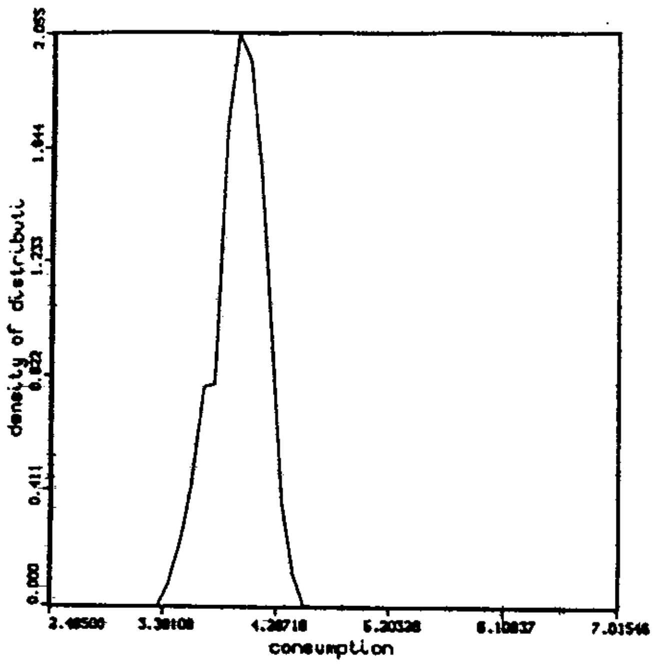

case 08, consumption(t), histogram

simulation Coleman

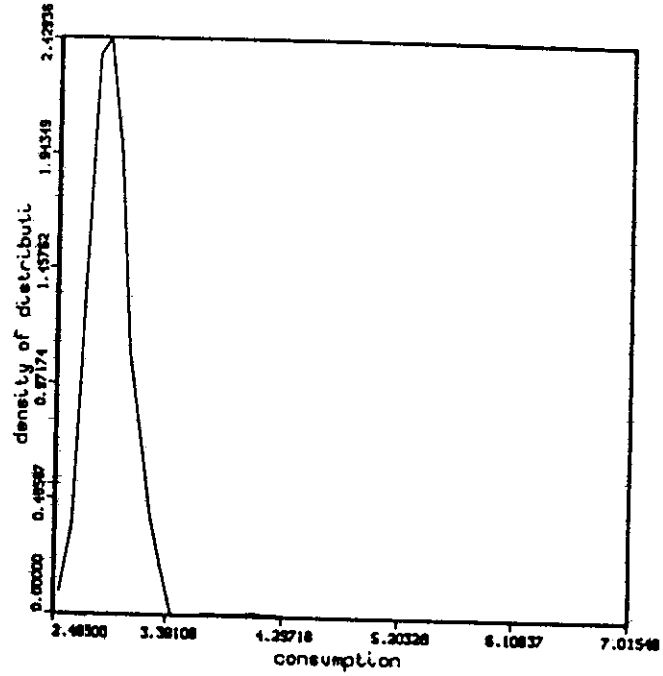

case 08, consumption(t), histogram

simulation LogLQ - Markov

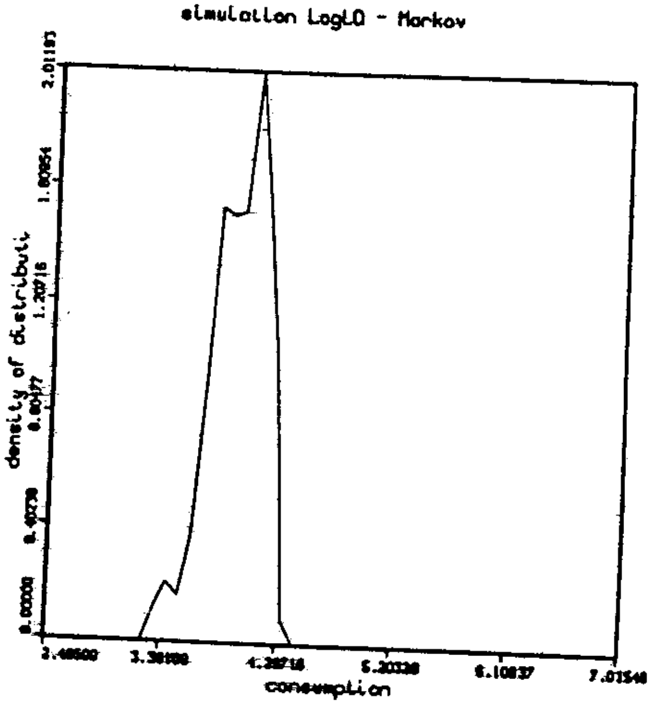

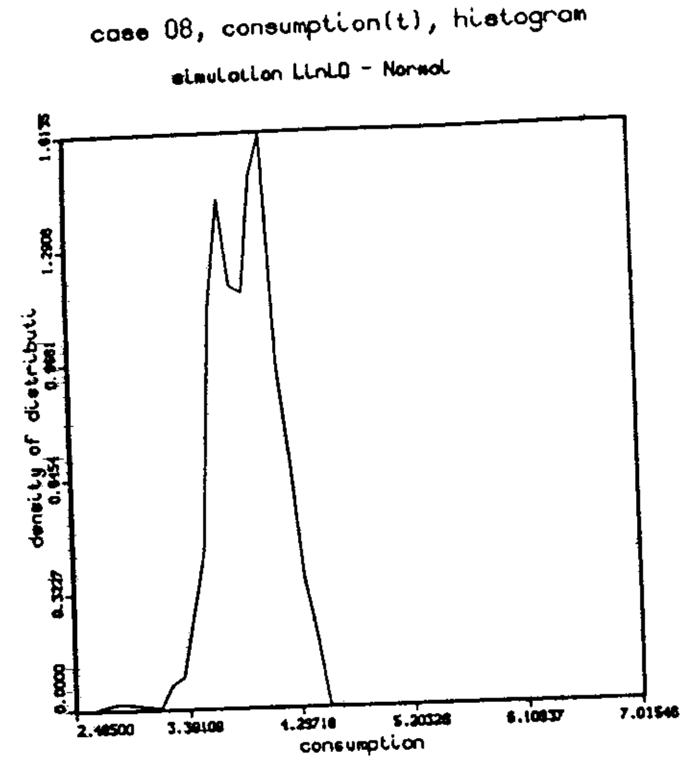
case 08, consumption(t), histogram
simulation LinLQ - Normal
density of distributi
0.0000
0.3227
0.6454
0.9681
1.2908
1.6135
2.48500
3.39108
4.29718
5.20328
6.10937
7.01546
consumption

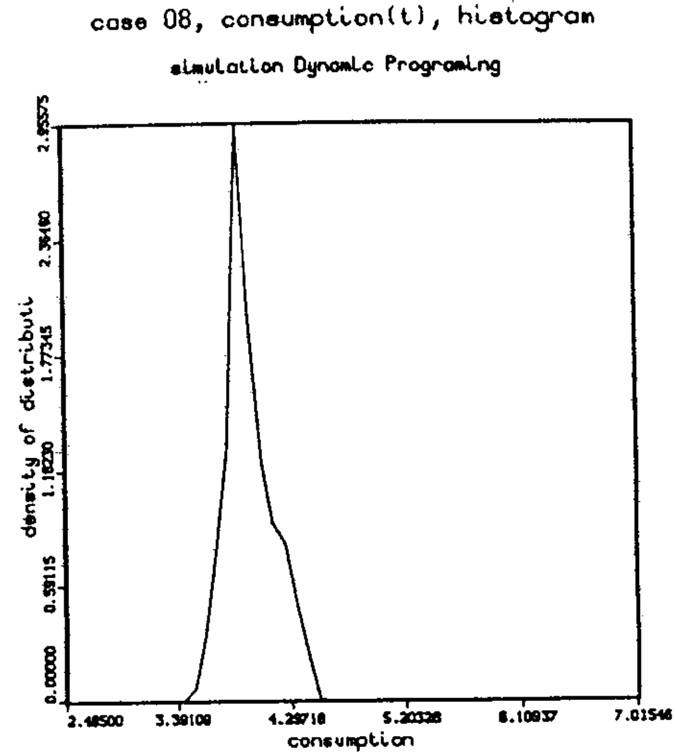
case 08, consumption(t), histogram
simulation Dynamic Programing
density of distributi
0.00000
0.59115
1.18230
1.77345
2.36460
2.95575
2.48500
3.39108
4.29718
5.20328
6.10937
7.01546
consumption

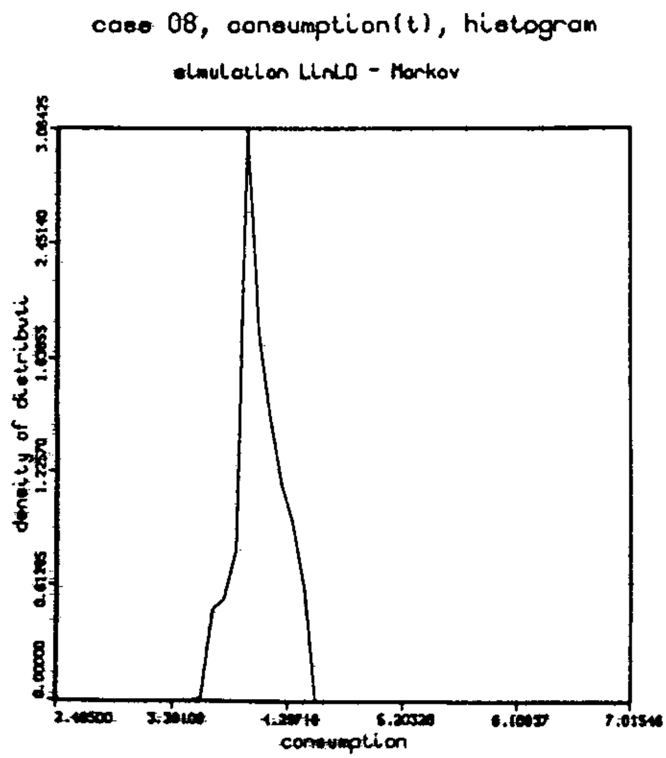
case 08, consumption(t), histogram
simulation LinLQ - Markov
density of distributi
0.00000
0.61285
1.22570
1.83855
2.45140
3.06425
2.48500
3.39108
4.29718
5.20328
6.10937
7.01546
consumption