Lecture Notes 10

Assessing the accuracy of approximation

Evaluating simulation algorithms is an important step in determining approximate solutions to a model as it aims at evaluating the accuracy of the approximation. In other word, it should answer the question

> *What error an individual would do if he/she were to use the approximate decision rule ("rule of thumb") instead of the true decision rule?*

This question is far from being simple to answer, as we first need to qualify the error we are talking about. Further, accuracy is not the only dimension along which a solution method should be tested. An approximation method must be done along at least three dimensions:

1. accuracy of the solution,
2. speed of computation,
3. ease of utilization by any potential user.

If both 1 and 2 can be more or less easily tested, 3 is without any doubt the most difficult to assess as it does not call for a quantitative answer. Therefore, we will leave this issue.

Speed is of course primarily function of the hardware, but also of the chosen programming language. Further, you have to keep in mind that beside the execution CPU time, there is often the time spent on developing the code! Anyway, accuracy is indeed the most important issue and the remaining of these lecture notes will only deal with that issue. Testing the accuracy of a particular approximation method usually boils down to evaluate the approximate *decision rule* (or *policy function*) and the trajectories that this function delivers. At this point, we should distinguish two types of problems.

The first one contains the few models for which an analytical expression for the policy function exists (Burnside's [1998] , or the optimal growth model with log utility and complete depreciation). In this case, it is possible to make a systematic comparison of the approximated policy functions with the exact ones on the entire set of value of interest for the state variables. As we are comparing two functions, it is an exercise in numerical analysis, without any statistical dimension.

In the second type of problem — which is the one we usually face — no *a priori* knowledge of the exact form of the policy function is available. Although it may be instructive to compare the policy functions provided by different algorithms, such a comparison cannot *per se* determine the best method. However, it is still possible to test whether the simulated trajectories satisfy the restrictions implied by the rational expectation hypothesis and its statistical consequence: the error of forecast must be a martingale difference sequence. Another approach is more economic in the sense that it attempts to give an economic interpretation of the residuals of the Euler equation, and therefore endows the economist with an objective economic criterion to gauge the accuracy of the solution.

However, it is necessary to keep in mind what the simulated model will be used for. Some uses seem rather robust to moderate inaccuracy in the policy function. It seems to be the case for impulse response functions, for example. In other cases, high accuracy is necessary: Price dividend ratios

in asset pricing models seem extremely sensitive. When the model is used for policy analysis purposes, it is not enough to evaluate the quality of an algorithm via statistical means for the baseline parameters only. By its very nature, a dynamic stochastic simulation samples more often points near the center of the ergodic space for its state variables. However, a policy shift modifies the ergodic space for the economy and may bring it in a region which was little explored in the simulation with the baseline parameters.

Throughout these lecture notes, we will apply each method to the standard optimal growth model we have been dealing with so far.

$$\max_{\{c_{t+\tau}\}_{\tau=0}^{\infty}} E_t \sum_{\tau=0}^{\infty} \beta^{\tau} \frac{c_{t+\tau}^{1-\sigma} - 1}{1-\sigma} \text{ with } \beta \in (0,1) \text{ and } \sigma \in \mathbb{R}^{+} \backslash \{1\} \tag{10.1}$$

subject to

$$k_{t+1} = \exp(a_t) k_t^{\alpha} - c_t + (1-\delta) k_t \tag{10.2}$$

and to the side conditions $k_t > 0$ and $c_t > 0$ for all t. The stochastic process for a_t is given by

$$a_t = \rho a_{t-1} + \varepsilon_t \tag{10.3}$$

where $|\rho| < 1$ for stationarity purposes and ε is a gaussian white noise process with mean 0 and constant standard deviation σ_{ε}.

Following Taylor and Uhlig [1990], the previous stochastic growth problem was solved for different cases in order to provide with a systematic comparison of different solving methods. For comparison purposes, we consider the same parameterization as theirs imposing $\alpha = 0.33$ and $\rho = 0.95$ for all cases under study. Only β, σ and σ_{ε} will be changed along the different exercises. Their values are reported in table 10.1.

These experiments allow to evaluate the sensitivity of the accuracy of the approximation method to *(i)* the degree of impatience of the individuals which determines the overall persistence in the economy *(ii)* the coefficient of risk aversion of the individuals (σ), and *(iii)* the overall degree of uncertainty in

Table 10.1: Parameter choices for the experiments

β	σ	σ_ε
0.95	0.5	0.1
0.95	1.5	0.1
0.98	0.5	0.1
0.98	1.5	0.1
0.95	0.5	0.02
0.95	1.5	0.02
0.95	3.0	0.02
0.98	0.5	0.02
0.98	1.5	0.02
0.98	3.0	0.02

the economy (σ_a) which turns out to be very large in the four first experiments (the traditional RBC literature usually considers a value of σ_a lower than 0.01).

10.1 Comparing the methods

This section proposes a set of tools, widely used in the literature to assess the relative performances of different approximation methods. In this section, we will not deal at all with tools that aim at selecting, either from a statistical or an economic point of view, a method over the others, but will rather propose tools that permit to evaluate the differences arising from the use of one method rather than another. Therefore, this essentially will amount to state the major features of the outcome of each approximation in terms of qualitative (form of the decision rule, IRF...) and quantitative (moments, distributions ...) implications of the model. Two approximate solutions to this model are compared a log–linear and a linear approximation.

10.1.1 Qualitative implications

In this section, we discuss some tools that may be used to gauge the relative qualitative properties of different approximation methods. We essentially dis-

cuss two issues *(i)* the form of the decision rules and *(ii)* the impulse response functions that the model delivers when a given approximation method is used.

Decision rules

A first way to assess differences in different approximation methods is to get an idea of the form of the decision rules that the approximation furnishes. This indeed constitutes a first approach to detect potential failures in the form of the decision rules.

A first experiment that may be of interest is to plot the "deterministic" decision rule. What we mean by "deterministic" version of the decision rule is the function that relates — in our benchmark experiment — k_{t+1} or c_t to k_t ($\exp(a_t)$) for a given value of a_t (k_t). For instance, figure 10.1 reports the decision rules for capital stock and consumption in the case of a linear and log–linear approximation. The upper panels report the link between k_t and k_{t+1}, c_t

Figure 10.1: Deterministic decision rule

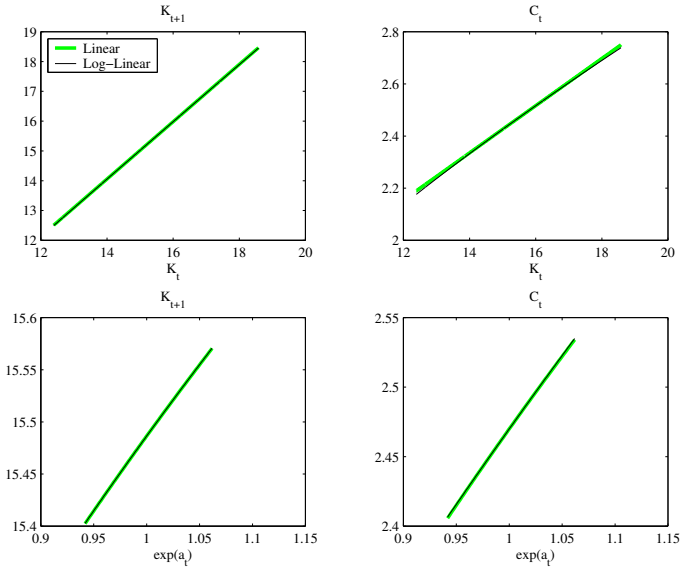

Note: This graph was obtained for $\beta = 0.95$, $\sigma = 1.5$ and $\sigma_\varepsilon = 0.1$.

given that a_t is set to its mean, while the lower panels draw the decision rule of

k_{t+1} and c_t as a function of the technology shock $\exp(a_t)$ given that the capital stock is held at its steady state level. As can be seen from figure 10.1, our two experiments (linear and log–linear approximation) do not differ that much in terms of "deterministic" decision rule, implying that their major qualitative properties are very similar. This is mainly due to the certainty equivalence property they impose and to the fact that for low volatilities, as we have already seen in these lecture notes. The linear and log–linear approximations lead to very similar results. Beside these certainty equivalence comparison of the decision rules, one may be interested in comparing the stochastic behavior of these decision rules. Therefore, we report in figure 10.2 the scatter plots for consumption and capital, and next period capital and capital, taking the uncertainty into account. The length of each series was 5000. It is important

Figure 10.2: Stochastic decision rule

Linear case

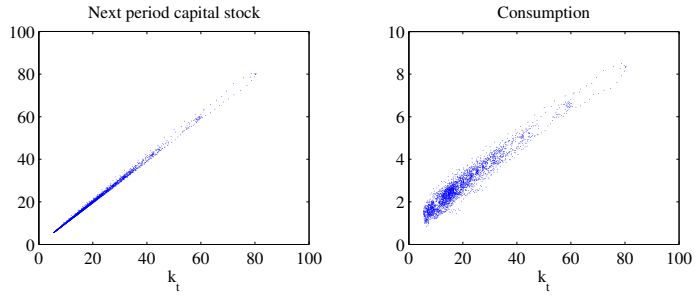

Log–linear case

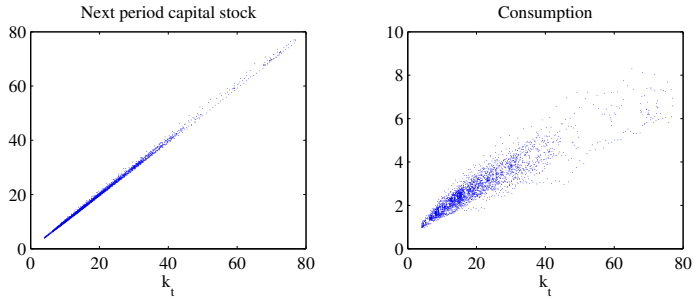

Note: This graph was obtained for $\beta = 0.95$, $\sigma = 1.5$ and $\sigma_\varepsilon = 0.1$.

to note that consumption was simulated relying on its approximation, then

using the approximate solution, we generated the capital stock using the law of motion of capital equation. In others, only consumption was approximated, output and the capital stock where generated using their nonlinear expression such that they can be taken to be "exact". Then, differences arise between the two approximation methods. First, the log–linear approximation yields a higher dispersion in the consumption decision rule relative to the linear. Conversely, it leads to a more concentrated decision rule for the capital stock — and therefore for investment. These differences essentially come from the fact that in this procedure, the overall volatility of the economy matters. Indeed, as we keep the nonlinear features of the model when simulating, the certainty equivalence is broken. Therefore, figure 10.2 indicates that the (log–)linearity form of the decision rule may have greater implications than what was suggested by the former analysis, especially in terms of moments and the form of the distribution generated by the model. Further, it appears that the greater volatility in consumption can be found when the capital stock is high, indicating that — as should be expected — one of the two methods performs worse in matching the tails of the true distribution.

Impulse response functions:

The computation of the IRF is conducted in the lines of Koop, Pesaran and Potter [1996]. In case of an arbitrary shock of magnitude σ, given state variables at time 0, an IRF for horizon h is defined as:

$$IRF(h, \sigma, 0) = \frac{E[y_{t+h}|\sigma; \mathscr{I}_0] - E[y_{t+h}|\mathscr{I}_0]}{E[y_{t+h}|\mathscr{I}_0]}$$

where $\mathscr{I}_0$ is the information set available at the time 0. Conditional expectations, involved in IRF computations, are computed using Monte-Carlo integration. Then IRF are obtained as follows:

1. Given an initial value for x, x_0, we draw R replications of a vector of innovations for x, $\mathcal{E}^r = \{\varepsilon_s\}_{s=0}^{H}$.

2. We then compute R realizations of the price–dividend ratio, y_t, using the approximated solution to the model. They are denoted $y_h^r(x_0, \mathcal{E}^r t)$ for $h = 0, \ldots, H$ and $r = 1, \ldots, R$.

3. We compute R realizations of the price–dividend ratio, y_t using the same draws plus one additional arbitrary shock σ, which affects ε_t. They are denoted $y_h^r(x_0, \sigma, \mathcal{E}^r)$, for $h = 1, \ldots, H$ and $r = 0, \cdots, R$.

4. We compute the IRF as the relative difference:

$$IRF^r(h, \sigma) = \frac{y_h^r(x_0, \sigma, \mathcal{E}^r) - y_h^r(x_0, \mathcal{E}^r)}{y_h^r(x_0, \mathcal{E}^r)} \quad h = 0, \cdots, H$$

5. We finally form the averages for each simulated data:

$$IRF^R(h, \sigma) \quad = \quad \frac{1}{R} \sum_{r=1}^{R} IRF^r(h, \sigma) \text{ for each } h = 0, \cdots, H$$

   Thus for R large, we have:

$$\lim_{R \to \infty} IRF(h, \sigma, 0)^R = IRF(h, \sigma, 0)$$

It is worth noting that for the approximation methods we are using in our simple application, this procedure may be simplified as they both make the certainty equivalence assumption. Indeed, this procedure relies on simulations in order to account for differences in the steady state levels that might arise when the decision rules are nonlinear. But furthermore, this procedure accounts for differences that might occur in the case of small versus big shocks, or positive versus negative shocks. Such differences may be huge when the model we consider is highly nonlinear, as it would be the case if the utility function were to display high curvature, or if we were to consider a model with irreversible investment as in Christiano and Fisher [2000].

Figure 10.3 reports the impulse response function of consumption and physical capital in face a one standard deviation positive technological shock. The most striking result that stands out of this graph is that the IRF do

Figure 10.3: Stochastic decision rule

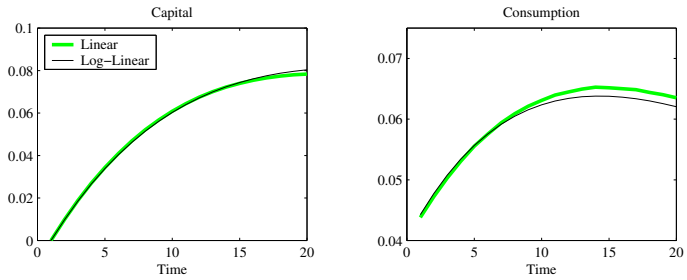

Note: This graph was obtained for $\beta = 0.95$, $\sigma = 1.5$ and $\sigma_\varepsilon = 0.1$.

not differ much, even if the volatility of the shock is quite high compared to standard calibrations. Thus, even if we decision rules differ, as the previous section as shown, the IRF are quite similar. This result partly come from the fact that the linearity that each of these methods imposes — and its corollary certainty equivalence assumption — implies the great similarity in the IRF. But, it seems to be a more general result the IRF would still be similar for more elaborated methods. This seems to indicate that for a model, which does not generate that much nonlinearities, accuracy of the approximate decision rules may not be of a great deal as far as IRF are concerned. Nevertheless, this would not be true any more with a model in which nonlinear features are important (for instance a model of irreversible investment, or a model with borrowing constraint). Therefore, IRF analysis may be a good way to detect failures in the approximation if they are not in accordance with what is predicted by the theory.

10.1.2 Quantitative features

We now try to characterize tools that can be used to assess the differences in the quantitative features associated to different approximation methods. We essentially discuss two issues *(i)* the form of the decision rules and *(ii)* the impulse response functions that the model delivers when a given approximation method is used.

Distributions

A first way to discuss the differences across different technics is to have a look at the distribution of the variables of interest they generate. Therefore, in figure 10.4, we report the non parametric estimator of the distribution of consumption and the capital stock, obtained simulating a sample of 5000 data points. At a first glance, it appears that the model generates a distribution

Figure 10.4: Distribution of variables

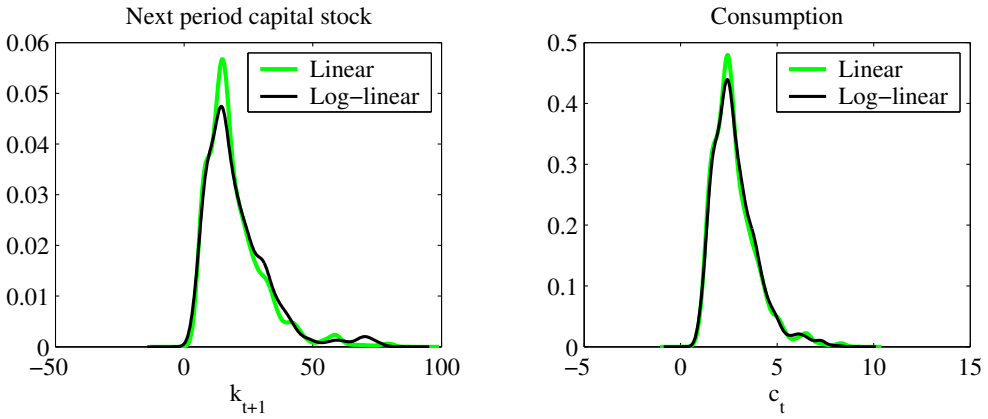

Note: This graph was obtained for $\beta = 0.95$, $\sigma = 1.5$ and $\sigma_\varepsilon = 0.1$.

that resembles a log–normal. Further, although we ran the same experiment for each model (in particular we used the same seed to draw the innovations of the technological shock), it turns out that the distributions slightly differ. In particular, a simple graphical inspection of the distribution of the capital stock indicates that the linear approximation generates a more concentrated distribution than does the log–linear.

In order to provide with a better understanding of the differences arising in the distributions, it is interesting to compare the moments generated by each method.

Moments comparison

The comparison between the moments generated by different approximation methods is a convenient way to understand the sources of the differences arising across the distributions generating by each method. Therefore, we report in table 10.2 the mean, standard deviation, minimum and maximum, the skewness and kurtosis of the distribution of capital and consumption for each method we consider. The moments were obtained relying on 1000 *Monte–carlo simulations* of 2500 data points. The same set of shocks was used for both approximations. It turns out that the linear model generated a slightly higher mean of the capital stock than the linear approximation. This might reflect the fact that the use of the log–linear rather than linear approximation as a rule of thumb to form expectations generates a precautionary saving motive that leads the agents to accumulate at a greater pace. This implies that, on average, output is greater when the log–linear approximation is used, such that the implied wealth effect enables the agents to consume more. Thus, the average level of consumption is also greater when a log–linear approximation is used. Nevertheless, the difference between the two methods is not statistically significant at the 5% level.

Table 10.2: Moments of the distribution

	Capital		Consumption	
	Linear	log–linear	Linear	log–linear
Mean	19.4117	19.3671	2.8258	2.8164
Std. Dev.	9.4605	9.6937*	1.0965	1.0807*
Maximum	63.0216	59.2960	7.4915	7.4246*
Minimum	7.5715	4.6400*	0.8243	0.9117*
Skewness	2.6066	1.8265*	1.5467	1.1721*
Kurtosis	6.2228	5.1871*	5.0187	4.4666*

A $\star$ indicates that the moment is significatively different from the linear case at the 5% level

The differences are statistically significant when we focus on the standard

deviation of the capital stock and consumption level, although tiny in the latter case. This confirms the differences we reported inspecting the stochastic decision rules for the capital stock and the consumption level. The differences in the amplitude of fluctuations seem to be more pronounced in the case of capital, as shown by the minimum and the maximum values of the capital stock. The log–linear approximation generates a higher minimal and a lower maximal level of capital than the linear case, thus narrowing the support of the distribution. This result is a by–product of the concavity of the logarithm transformation that tends to limit the amplification of the technological shocks, contrary to the linear approximation that transmits the technological shocks one for one. Further, the linear approximation generates more skewness and kurtosis in the distribution of the capital stock — recall that when simulating only consumption was approximated. The explanation of this phenomenon may lie in the fact that the model is almost log–linear and that physical capital and consumption are almost log–normally distributed. Thus, a log–linear approximation should fare better. On the contrary, a linear approximation may have a hard time accounting for the log–normality thus overestimating both skewness and kurtosis. Therefore, this simple comparison provides some information on the potential failures of the linear approximation method.

Besides these results, it may be important to gauge the differences in the approximation method relying on moments that really depend on the level accuracy of the solving method. For instance, in the case of optimal growth model — and more generally of any general equilibrium model in which agents have access to financial markets — studying asset pricing implications provide with a great deal of information on the properties of competing approximation methods. For instance, Christiano and Fisher [2000] have shown that asset pricing moments are largely affected by small deviations from the true policy rule. For instance, they report experiments[1] were the bias in consumption

[1]In the case of conventional PEA, Table 3, p.1207.

with respect to the true solution — obtained using dynamic programming — is about 0.2%, but were the bias in the equity premium reaches 78%. This can be easily understood as asset pricing are — in this type of model — strongly related to the intertemporal marginal rate of substitution. If we consider a CRRA utility function, as the marginal utility of consumption is essentially related to the inverse of consumption to a given power, a small approximation error on consumption translates into a huge error in the marginal utility. The intertemporal marginal rate of substitution is then strongly biased, which affects asset pricing moments a great deal. Then, even if accuracy on consumption is acceptable, it may be the case that this will not be acceptable in terms of asset pricing. This actually reveals one important feature of the quest of accuracy: all depends what the economist is interested in. Macroeconomic aggregates are usually found to be not very demanding in getting a high level of accuracy, on the contrary asset pricing must be handled carefully.

An implication of the former discussion is that small differences in the approximate solutions for aggregate variables may translate in huge differences in asset pricing implications. Therefore, asset pricing may carry an important piece of information on the quantitative properties of different methods.

CPU Time

Another important quantitative feature that may incline the choice of the economist in favor of a particular method is the CPU time. Although this feature nowadays becomes less and less important due to the large increase in power of computers, this may constitute a crucial point whenever one is willing to perform a large number of different experiments — namely test different calibrations. One criterion that may go beyond the obvious need for accuracy, is therefore the need to minimize CPU time. But besides the simple CPU time, one should take care of the amount of time needed to code and optimize the algorithm, which may — under certain circumstances — be more costly than the CPU time.

10.2 Accuracy tests

This section is no more concerned with the comparison of the outcome of different rules, but now deals explicitly with the problem of evaluating the accuracy of a particular rule and thus tackle the question of the quality of the approximation.

10.2.1 A difference criterion

This approach can only be used when the true solution to the problem is available. In this case, the accuracy of the approximation method can be directly checked against the "true" decision rule. This is undertaken relying on the two following criteria

$$E_1 = 100 \times E \left| \frac{y_t - \widetilde{y}_t}{y_t} \right|$$

and

$$E_\infty = 100 \times \max \left\{ \left| \frac{y_t - \widetilde{y}_t}{y_t} \right| \right\}$$

where y_t denotes the true solution to the problem and $\widetilde{y}_t$ is the approximation of the true solution by the method under study. E_1 represents the average relative error an agent makes using the approximation rather than the true solution, while E_∞ is the maximal relative error made using the approximation rather than the true solution. In our case, we take the dynamic programming solution to be the true one. It should however be stressed that this solution is not always available or may be too costly to implement.

10.2.2 An economic approach

In this section, we follow Judd [1992] and assess the accuracy of the approximate solution by evaluating the degree of rationality the agents would reach relying on the approximate solution as a rule of thumb to form expectations. In others, a solution will be said to be accurate when the intertemporal error the agent would make using the approximate solution would be small enough,

which remains to be defined. This error may be measured, in the optimal growth model, by the residual of the Euler equation, u_t

$$u_t = c_t - \left[\beta E_t c_{t+1}^{-\sigma}\left(\alpha \exp(a_{t+1}) k_{t+1}^{\alpha-1} + 1 - \delta\right)\right]^{-\frac{1}{\sigma}}$$

For interpretability purposes, the Euler residual is normalized by the level of consumption, such that u_t/C_t expresses the degree of irrationality of the agents in terms of consumption goods. In others, it furnishes the loss in terms of consumption a agents would suffer from by using the approximate solution rather than the true solution. Then the following criterion may be defined:

$$\mathcal{E}_1 = \log_{10}\left(E\left|\frac{u_t}{c_t}\right|\right) \tag{10.4}$$

$$\mathcal{E}_2 = \log_{10}\left(E\left\|\frac{u_t}{c_t}\right\|^2\right) \tag{10.5}$$

$$\mathcal{E}_\infty = \log_{10}\left(\max\left|\frac{u_t}{c_t}\right|\right) \tag{10.6}$$

The first criterion provides with a scale free measure of the average *intertemporal error* an agent would make using the approximate solution, the second criterion furnishes a similar information, although slightly different as it can be interpreted in terms of volatility, and the last criterion gives the maximal error the agent would make. In order to understand the criterion let's assume that $\mathcal{E}_1$=-6. This tells us that using the approximate rule, the agent would make a mistake of 1\$ whenever she consume 1,000 000 \$, which should be considered as small enough — from an economic point of view — to consider the approximation as accurate. Therefore, this criterion is totally ad hoc, in the sense that the definition of *small enough* is left to the willingness of the economist.

Table 10.3 report theses criteria for the log–linear approximation method. The numbers were computed for values of capital ranging within the interval $[(1-\Delta_k)k^\star; (1+\Delta_k)k^\star]$ were $k^\star$ is the deterministic steady state and $\Delta_k = 0.2$, and value of the technology shock that insures that 99.99% of the distribution

of a_t is covered. The integral involved by the expectation was evaluated using a 20 nodes Gauss–Hermite quadrature.

Table 10.3: Judd's criteria

β	σ	σ_ε	$\mathcal{E}_1$	$\mathcal{E}_2$	$\mathcal{E}_\infty$
0.9500	0.5000	0.1000	-1.9756	-3.8514	-1.2016
0.9500	1.5000	0.1000	-2.0121	-3.8908	-1.2366
0.9800	0.5000	0.1000	-2.4244	-4.7454	-1.6484
0.9800	1.5000	0.1000	-2.3496	-4.5811	-1.5817
0.9500	0.5000	0.0200	-3.2966	-6.4510	-2.4559
0.9500	1.5000	0.0200	-3.3500	-6.5215	-2.5076
0.9500	3.0000	0.0200	-3.4302	-6.6201	-2.5829
0.9800	0.5000	0.0200	-3.7580	-7.3709	-2.9216
0.9800	1.5000	0.0200	-3.6917	-7.1988	-2.8560
0.9800	3.0000	0.0200	-3.7325	-7.2316	-2.8936

As expected, the method performs badly in face of a high volatility, as, for instance, when $\sigma_\varepsilon = 0.1$ an agent that would use the log–linear rule to compute her expectations would make an average error of 1\$ for each 200\$ of consumption purchases, which might not be acceptable from an economic point of view. In others, the solution is not acceptable from an economic point of view. The intertemporal error rises dramatically whenever we look at the maximal error, as a 1\$ intertemporal error for each 150\$ consumption purchases, which cannot be accepted from an economic point of view. The error decreases with the volatility of the technological shock.

10.2.3 A statistical approach

The *statistical approach* proposed by Den Haan and Marcet [1994] relies on *Monte–Carlo simulations* of the model. The idea is to test for the *martingale difference property* of the residuals of the Euler equation. Intuitively, if the candidate solution delivered by an algorithm is accurate, it should satisfy the rational expectation hypothesis. In particular, the expectation errors should

be innovations, meaning that they should not be correlated with variables belonging to the information set of the individuals that formulate expectations. This is precisely what this evaluation method tests.

Let us assume that we deal with a model of the form

$$E_t F(y_{t+1}, x_{t+1}, y_t, x_t, \varepsilon_{t+1}) = 0 \tag{10.7}$$

where $F : \mathbb{R}^{n_y} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_e} \longrightarrow \mathbb{R}^{n_x+n_y}$. E_t is the standard conditional expectations operator. ε_t is the set of innovations of the structural shocks that affect the economy. Let us now assume that the model has been solved by an approximation method that yields

$$y_t = g(x_t; \theta) \tag{10.8}$$

$$x_{t+1} h(x_t, y_t, \varepsilon_{t+1}) = h(x_t, g(x_t; \theta), \varepsilon_{t+1}) \tag{10.9}$$

Given a sequence of $\{\varepsilon_t\}_{t=1}^T$ and an initial condition x_0, it is easy to perform dynamic simulations for x_t and y_t from equations 10.8 and 10.9. Let $\widetilde{y}_t$ and $\widetilde{x}_t$ denote simulated realizations from the approximate solution. Let us now define the residuals by

$$u_{t+1} = F(\widetilde{y}_{t+1}, \widetilde{x}_{t+1}, \widetilde{y}_t, \widetilde{x}_t, \widetilde{\varepsilon}_{t+1}) \tag{10.10}$$

This testing procedure has primarily been developed by Den Haan and Marcet [1994] and borrows from the econometric literature on GMM estimation. Let us first define a set of q simulated instrumental variables, denoted $\widetilde{z}_t$, where $\widetilde{z}_t \in \widetilde{I}_t \subset I_t$. These *simulated instruments* must correspond to the predetermined and exogenous variables that enter into the information set at period t. They are thus essentially obtained from the simulated variables $\widetilde{x}_t$. From (10.7), the residuals u_{t+1} must satisfy $q \times n$ conditions:

$$E\left[u_{t+1} \otimes \widetilde{z}_t\right] = 0 \tag{10.11}$$

The simulated counterparts of (10.11) can be defined as

$$g_T = \frac{1}{T} \sum_{t=1}^{T} u_{t+1} \otimes \widetilde{z}_t \tag{10.12}$$

As pointed out by Den Haan and Marcet [1994], g_T converges to zero almost surely as T goes to infinity if the solution S is exact. But, because of sampling errors, g_T will not be exactly equal to zero under the null hypothesis that the approximation method furnishes the "true" solution. Therefore, Den Haan and Marcet [1994] introduce a test statistic similar to the Hansen's [1982] *J–statistics* to decide whether g_T is significantly different from zero:

$$J_T = T g_T' \Omega_T^{-1} g_T \tag{10.13}$$

where Ω_T is a consistent estimate of the variance–covariance matrix:

$$V = \lim_{j \to \infty} \sum_{-j}^{j} E\left[(u_{t+1} \otimes \widetilde{z}_t)(u_{t+1-j} \otimes \widetilde{z}_{t-j})'\right]$$

Under stationary and ergodic assumptions for the process $\{u_{t+1} \otimes \widetilde{z}_t\}_{t=1}^{t=T}$ and assuming that g_T is asymptotically distributed as a gaussian stochastic variable, we have

$$\sqrt{T} g_T \to \mathcal{N}(0, \Omega_T) \quad \text{as} \quad T \to \infty \tag{10.14}$$

Therefore, the statistic J_T is asymptotically distributed as a chi-square with $q \times n$ degrees of freedom. It is worth noting that no loss of degrees of freedom occurs as all parameters are known with certainty.

We now apply it to the optimal growth model. In each case, 500 time series of 3000 observations were used. For each experiment, the set of instruments was $z_t = \{1, k_t, a_t, k_{t-1}, a_{t-1}\}$, so that the global specification test statistic admits a chi–square distribution with 5 degrees of freedom. Following Den Haan and Marcet [1994], we report the percentage of draws that are in the lower and upper 5% tails for the test. For large value of σ_ε, the log-linear approximation does not provide accurate solutions (first line of table 10.4). In each case, the empirical size of the statistics differs significantly from the theoretical size. For instance, when using the theoretical critical value at the 5% confidence level, the test leads to non–rejection of H_0 (the residuals are a martingale difference sequence) in only 0.27% of the cases for the first parameterization.

Another way of looking at the results is to have a look at the empirical critical value at the 5% significance level. The critical value for a chi–square distribution with 5 degrees of freedom at the 5% significance level is 1.1455, in all the cases, the empirical value lies above it implying that the test always leads to rejection of a martingale difference sequence for the residuals. Therefore, the approximation does not fully satisfy the rational expectation hypothesis. This may also be seen looking at the last column of the table, which reports the average value of the statistics across the simulations. For instance, the first row of the table indicates that the log–linear approximation is ways out of the true decision rule as Q_T=10.12, while the critical value at the 5% significance level is 2.94, which therefore leads to strong rejection of the martingale difference property of the Euler residuals: the approximation does not fulfill the restrictions imposed by the rational expectations hypothesis under this parameterization.

This is particularly true when the degree of impatience is large and the relative risk aversion is low. In each case under study, the approximation cannot fulfill the set of orthogonality conditions, with the possible exception of a low variance of shocks, a low degree of impatience and a low intertemporal substitution parameter. Except for the cases of large shocks and low value of the risk aversion, the log–linear approximation performs well. In most cases, it appears that the empirical size of the J-statistics is close to the theoretical size of the chi-square distribution with five degrees of freedom. This result is illustrated by figure 10.5 that reports the theoretical cdf and the empirical associated to the cases ($\beta = 0.95$, $\sigma = 1.5$ and $\sigma_\varepsilon = 0.1$) and ($\beta = 0.95$, $\sigma = 1.5$ and $\sigma_\varepsilon = 0.02$). In the first case, the empirical cdf departs significantly from the theoretical cdf (a $\chi^2(5)$) for every size. Conversely, the two cdf are very close when the standard error of the shock decreases.

Other tests may be developed in these lines. For instance, one may just check that the Euler residuals are iid, or one may test all orthogonality conditions in isolation in order to detect the source of the rejection of the test.

Table 10.4: Den–Haan and Marcet's test

			Empirical size		Empirical values	
β	σ	σ_ε	lower 5%	upper 5%	5% S.L.	Stat.
0.95	0.5	0.1	0.27	61.27	2.9444	10.1154
0.95	1.5	0.1	2.73	91.32	1.5429	5.5875
0.98	0.5	0.1	0.37	63.97	2.8089	9.9574
0.98	1.5	0.1	3.00	90.33	1.5024	5.9418
0.95	0.5	0.02	3.95	92.70	1.3043	5.4108
0.95	1.5	0.02	3.79	94.30	1.2884	5.2368
0.95	3.0	0.02	3.28	92.40	1.3428	5.5097
0.98	0.5	0.02	4.56	92.84	1.2227	5.4512
0.98	1.5	0.02	3.67	94.17	1.3434	5.3233
0.98	3.0	0.02	3.89	93.64	1.3975	5.4558

Note: The critical value for a chi–square distribution with 5 degrees of freedom at the 5% significance level is 1.1455.

Figure 10.5: Global specification test

$\sigma_\varepsilon = 0.1$ $\sigma_\varepsilon = 0.02$

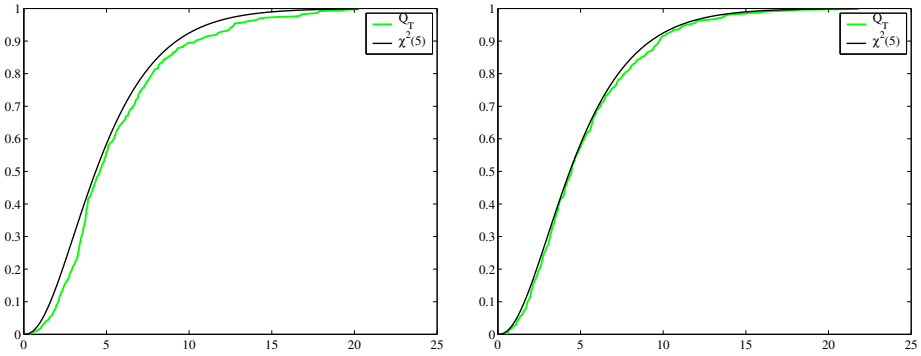

Bibliography

Burnside, C., Solving asset pricing models with Gaussian shocks, *Journal of Economic Dynamics and Control*, 1998, *22*, 329–340.

Christiano, L.J. and J.D.M. Fisher, Algorithms for solving dynamic models with occasionally binding constraints, *Journal of Economic Dynamics and Control*, 2000, *24* (8), 1179–1232.

Haan, W. J. Den and M. Marcet, Accuracy in Simulations, *Review of Economic Studies*, 1994, *61*, 3–17.

Hansen, L.P., Large Sample Properties of Generalized Method of Moment Models, *Econometrica*, 1982, *50*, 1269–1286.

Judd, K., Projection Methods for Solving Aggregate Growth Models, *Journal of Economic Theory*, 1992, *58*, 410–452.

Koop, G., M.H. Pesaran, and S.M. Potter, Impulse Response Analysis in Nonlinear Multivariate Models, *Journal of Econometrics*, 119-147 1996, *74*.

Taylor, J.B. and H. Uhlig, Solving Nonlinear Stochastic Growth Models : A Comparison of Alternative Solution Methods, *Journal of Business and Economic statistics*, 1990, *8*, 1–17.

Index

Contents

List of Figures

List of Tables