

A Note on the Accuracy of Markov-Chain Approximations to Highly Persistent AR(1)-Processes*

Martin Flodén**

Stockholm School of Economics and CEPR

SSE/EFI Working Paper Series in Economics and Finance

No. 656

March 12, 2007

Abstract

This note examines the accuracy of methods that are commonly used to approximate AR(1)-processes with discrete Markov chains. The quadrature-based method suggested by Tauchen and Hussey (1991) generates excellent approximations with a small number of nodes when the autocorrelation is low or modest. This method however has problems when the autocorrelation is high, as it typically is found to be in recent empirical studies of income processes. I suggest an alternative weighting function for the Tauchen-Hussey method, and I also note that the older method suggested by Tauchen (1986) is relatively robust to high autocorrelation.

JEL classification: C60

Keywords: numerical methods; income processes; autoregressive process

1 The Problem

This note considers the accuracy of different methods that are commonly used to approximate autoregressive processes by Markov chains. We thus want to approximate the AR(1) process

$$z_{t+1} = (1 - \rho)\mu + \rho z_t + \varepsilon_{t+1}$$

by an n -node Markov chain $\{Z, \Pi\}$ where $Z = \{z^1, z^2, \dots, z^n\}$ and $\Pi = \{\pi_{i,j}\}$ where $\pi_{i,j}$ is the transition probability from z^i to z^j . Here $\varepsilon \sim N(0, \sigma_\varepsilon^2)$, and consequently the unconditional standard deviation of z is $\sigma_z = \sigma_\varepsilon (1 - \rho^2)^{-1/2}$.

Five alternative methods are used to approximate this AR(1) process. The first method follows Tauchen (1986). The nodes Z are equally spaced between $\pm 1.2\sigma_z \ln n$, and the transition probabilities Π are the probabilities $\pi_{i,j} = \Pr(z' \in [z^j - s, z^j + s] | z = z_i)$ implied by the AR(1) process. The step size s is half the distance between nodes, i.e.

*I thank Paul Klein for helpful discussions and the Wallander and Hedelius foundation for financial support.

**Department of Economics, Stockholm School of Economics, Box 6501, SE-113 83 Stockholm, Sweden. martin.floden@hhs.se, +46-8-7369263.

$s = (z^2 - z^1) / 2$, except at the extreme nodes where the intervals are $(-\infty, z^1 + s]$ and $[z^n - s, \infty)$.

The following three methods all build on Tauchen and Hussey's (1991). The nodes $\{z^i\}$ are then the Gaussian nodes for some normal distribution $N(\mu, \hat{\sigma}^2)$, and the difference between the approximation methods I use is the choice of $\hat{\sigma}^2$.¹

To summarize Tauchen and Hussey's method, let $\{z^i\}$ and $\{w^i\}$ denote the Gaussian quadrature nodes and weights for the normal distribution $N(\mu, \hat{\sigma}^2)$. Suppose now that $z_t = \mu$ and that $\sigma_\varepsilon^2 = \hat{\sigma}^2$. Then these Gaussian nodes and weights typically provide an excellent approximation of how z will develop in the next period. The problem is that if $z_t \neq \mu$, Gaussian quadrature would imply other nodes (and weights), but the Markov chain requires that the nodes are fixed. So, can we find nodes Z and probabilities Π that provide an approximation to the process for any $z_t \in Z$? Gaussian quadrature provides nodes $\{z^i\}$ and weights $\{w^i\}$ so that

$$\int g(\xi) f(\xi) d\xi \approx \sum g(z^i) w^i$$

where g is some function, $\xi \sim N(\mu, \hat{\sigma}^2)$, and f is the density function for ξ . Tauchen and Hussey (1991) note that if

$$z_{t+1} = (1 - \rho)\mu + \rho z_t + \xi_{t+1}$$

then

$$\int g(z_{t+1}) f(z_{t+1}|z_t) dz_{t+1} = \int g(z_{t+1}) \frac{f(z_{t+1}|z_t)}{f(z_{t+1}|z_t = \mu)} f(z_{t+1}|z_t = \mu) dz_{t+1},$$

i.e.

$$\int g(z_{t+1}) f(z_{t+1}|z_t) dz_{t+1} \approx \sum g(z^i) \frac{f(z_{t+1}|z_t)}{f(z_{t+1}|z_t = \mu)} w^i.$$

They therefore propose that the AR(1) process should be approximated by the nodes $\{z^i\}$ and the transition probabilities

$$\pi_{i,j} = \frac{f(z^j|z^i)}{f(z^j|z_t = \mu)} w^i.$$

As already mentioned, I consider three variants of the Tauchen-Hussey method. First, as suggested by Tauchen and Hussey (1991) I use $\hat{\sigma}^2 = \sigma_\varepsilon^2$. Most subsequent implementations of the Tauchen and Hussey method (e.g. Burnside, 1999) use this alternative, and Tauchen and Hussey (their section 3.3) argue that numerical evaluations support this choice. I refer to this specification as σ_ε in the tables below.

The second variant I consider is $\hat{\sigma} = \sigma_z$. Tauchen and Hussey (1991) also mention this as a possible choice, and Klein (2007) suggests this specification, which I refer to as σ_z in the tables. The third variant I consider is $\hat{\sigma} = w\sigma_\varepsilon + (1 - w)\sigma_z$ where $w = 1/2 + \rho/4$. The variance $\hat{\sigma}^2$ is then set to a weighted average of the conditional and unconditional

¹Different choices for the mean μ could also have been considered.

variances, and more weight is given to the conditional variance when the process is highly persistent. This variant is referred to as w in the tables.

Finally, the fifth method considered is outlined in Adda and Cooper (2003). This method first chooses n intervals such that z has equal unconditional probability to fall in each of the intervals. Second, one node is chosen for each interval, and this node is set to the expected value of z conditional on z being in that interval. Third, the transition probabilities are calculated as with Tauchen's (1986) method.

Using these five approximation methods, I consider the accuracy of approximations to different specifications of the AR(1) process

$$z_{t+1} = \rho z_t + \varepsilon_{t+1}.$$

The Markov chain approximations are often used in economics to model income processes, and I first consider the accuracy of approximations to three processes that have often been adopted in the literature. The first specification follows Aiyagari (1994) and sets $\rho = 0.60$ and $\sigma_\varepsilon^2 = 0.013$, while the second follows Hubbard et al. (1995, HSZ) and sets $\rho = 0.95$ and $\sigma_\varepsilon^2 = 0.030$. The final specification, suggested by Storesletten et al. (2000), is even more persistent and sets $\rho = 0.98$ and $\sigma_\varepsilon^2 = 0.020$. In addition to examining these processes, I also examine the methods' accuracy for a broader set of autocorrelations.

2 Results

Table 1 reports the autocorrelation, conditional standard deviation, and unconditional standard deviation implied by the Markov chains that are obtained with the different approximation methods with $n = 5, 9,$ and 15 nodes, and Table 2 reports the accuracy of approximations for a broader set of autocorrelation parameters for $n = 5$ and 9 nodes.² The tables first show that processes with relatively low persistence (for example Aiyagari's process with $\rho = 0.60$) are well approximated with all methods even with a small number of nodes. For these process, we however see that Tauchen and Hussey's quadrature based methods consistently deliver better approximations than Tauchen's and Adda and Cooper's methods. The approximations are less precise when persistence is high. Tauchen's method appears to be relatively robust and never completely fails to capture any of the three moments.

The standard Tauchen-Hussey implementation based on the conditional variance provides a decent approximation of the autocorrelation and the conditional variance. But with high persistence, the unconditional variance is much larger than the conditional, and this method chooses nodes that are close to the mean of the process. Consequently, as demonstrated in Table 1, the method generates an unconditional variance that is much lower than the true unconditional variance. This result is not surprising when one considers that the method was designed to generate an accurate approximation of the conditional development of z (starting from the unconditional mean of the process). When persistence is high, the unconditional distribution is very different from the conditional, and the resulting approximations are imprecise.

²Computer code is available at www.hhs.se/personal/floden.

The Tauchen-Hussey method based on the unconditional variance has the opposite problem; it chooses nodes far from the mean of the process, and with a small number of nodes the implied transition probabilities from one node to another are miniscule. As expected, the resulting approximation is relatively accurate for the unconditional variance, but typically at the price of a much too high autocorrelation and a too low conditional variance.

The third implementation of the Tauchen-Hussey method is a compromise between the other two and aims at delivering more robust approximations when the autocorrelation is high. I have experimented with different weighting schemes before choosing $w = 1/2 + \rho/4$. These weights deliver approximations that are relatively accurate and robust, except when the autocorrelation is very high and only a small number of nodes are used.

Adda and Cooper's method is also relatively robust at high persistence levels, and in particular it then generates good approximations of the autocorrelation and the unconditional variance. The approximation to the conditional variance is less precise. An advantage with this method is that the ergodic distribution is uniform over the grid in contrast to the other methods that tend to put more of the ergodic distribution at the center of the process. This uniformity facilitates Monte Carlo simulations, in particular when behavior in the extreme nodes is important (for example if wealth is concentrated among a small fraction of the population, and equilibrium prices are determined by the aggregate wealth holdings).

The different implications of these methods are exemplified by Table 3 and Figures 1-2. Table 3 shows the Markov chain approximations for the HSZ process ($\rho = 0.95$, $\sigma_\varepsilon^2 = 0.030$) when using $n = 5$ nodes. In particular, the table demonstrates that the implementation with the unconditional variance generates a wide grid with extreme persistence in each node. Figures 1-2 illustrate some aspects of the approximations of the HSZ process when using $n = 15$ nodes.³ Figure 1 shows the implied autocorrelation ρ conditional on the node. As expected, the figure clearly shows that all approximations have problems approximating the true autocorrelation at the extreme nodes where the process only can move in one direction. Figure 2 shows the implied standard deviations σ_ε conditional on the node and illustrates a main difference between the Tauchen (1986) and the Tauchen-Hussey (1991) methods; the former method generates approximations that are good on average while the latter method generates good approximations when the process is close to its mean.

3 Conclusions

This note has compared the accuracy of methods that are often used to approximate AR(1) processes by discrete Markov chains. The nodes generated by the Tauchen and Hussey's (1991) method are the Gaussian quadrature nodes that provide a good approximation of the dynamics of the process conditional on some current state of the process. When the process has high persistence, the conditional transition dynamics will differ substantially depending on the current state, and the most common implementation of the Tauchen-Hussey method has problems generating accurate approximations. In particular it generates a too narrow grid and cannot approximate the unconditional variance

³The implications of Adda and Cooper's method are not reported in the graphs.

in the autoregressive process. I suggest using another weighting function that generates a broader grid with the Tauchen-Hussey method, but I also note that the Tauchen (1986) method is relatively robust to high persistence.

References

- Adda, Jerome, and Russell Cooper (2003), *Dynamic Economics*, MIT Press
- Aiyagari, Rao (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving”, *Quarterly Journal of Economics*, 109, 659-684
- Burnside, Craig (1999), “Discrete State-Space Methods for the Study of Dynamic Economies”, in R. Marimon and A. Scott (eds.): *Computational Methods for the Study of Dynamic Economies*, Oxford University Press
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes (1995), “Precautionary Saving and Social Insurance”, *Journal of Political Economy*, 103, 360-399
- Klein, Paul (2007), “Discrete Approximations of an AR(1) Process with Gauss-Hermite Quadrature”, mimeo, University of Western Ontario
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2000), “Consumption and Risk Sharing over the Life Cycle”, NBER Working Paper No. 7995
- Tauchen, George (1986), “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions”, *Economics Letters* 20, 177-181
- Tauchen, George, and Robert Hussey (1991), “Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models”, *Econometrica* 59, 371-396

Table 1: Approximated AR(1) Processes

True	$n = 5$			$n = 9$			$n = 15$		
	Tauchen		Tauchen-Hussey	Tauchen		Tauchen-Hussey	Tauchen		Tauchen-Hussey
	$1.2 \ln n$	σ_ε	σ_z	$1.2 \ln n$	σ_ε	σ_z	$1.2 \ln n$	σ_ε	σ_z
Aiyagari's Process, $\rho = 0.60, \sigma_\varepsilon^2 = 0.013$									
ρ	0.6000	0.5844	0.5992	0.6024	0.6000	0.5982	0.6000	0.6000	0.6000
σ_ε	0.1140	0.1167	0.1137	0.1138	0.1139	0.1165	0.1140	0.1140	0.1140
σ_z	0.1425	0.1430	0.1418	0.1425	0.1424	0.1451	0.1425	0.1425	0.1425
z^n/σ_z		1.9313	2.2856	2.8570	2.4856	2.6367	1.4091	4.5127	3.9261
Hubbard, Skinner & Zeldes' Process, $\rho = 0.95, \sigma_\varepsilon^2 = 0.030$									
ρ	0.9500	0.9577	0.9073	0.9998	0.9524	0.9503	0.9394	0.9945	0.9496
σ_ε	0.1732	0.1843	0.1576	0.0101	0.1410	0.1982	0.1670	0.0561	0.1692
σ_z	0.5547	0.6037	0.3275	0.5622	0.4792	0.6205	0.4303	0.5556	0.5407
z^n/σ_z		1.9313	0.8921	2.8570	1.4079	2.6367	1.4091	4.5127	2.2238
Storesletten, Telmer & Yaron's Process, $\rho = 0.98, \sigma_\varepsilon^2 = 0.020$									
ρ	0.9800	0.9952	0.9261	1.0000	0.9895	0.9861	0.9619	1.0000	0.9815
σ_ε	0.1414	0.0838	0.1258	0.0000	0.0679	0.1466	0.1332	0.0014	0.1177
σ_z	0.7107	0.7938	0.2782	0.7261	0.5468	0.8448	0.3868	0.7173	0.6657
z^n/σ_z		1.9313	0.5685	2.8570	1.1521	2.6367	0.8980	4.5127	1.8198

Note: The first Tauchen-Hussey column uses the conditional variance and hence sets $\hat{\sigma} = \sigma_\varepsilon$, the second column uses the unconditional variance and sets $\hat{\sigma} = \sigma_z$, and the final column uses a weighted average of σ_ε and σ_z as described in the text.

Table 2: Accuracy of Approximations

ρ	Tauchen		Tauchen-Hussey, σ_ε		Tauchen-Hussey, σ_z		Tauchen-Hussey, w		Adda-Cooper				
	$\hat{\rho}/\rho$	$\hat{\sigma}_\varepsilon/\sigma_\varepsilon$	$\hat{\rho}/\rho$	$\hat{\sigma}_\varepsilon/\sigma_\varepsilon$	$\hat{\rho}/\rho$	$\hat{\sigma}_z/\sigma_z$	$\hat{\rho}/\rho$	$\hat{\sigma}_\varepsilon/\sigma_\varepsilon$	$\hat{\rho}/\rho$	$\hat{\sigma}_\varepsilon/\sigma_\varepsilon$	$\hat{\sigma}_z/\sigma_z$		
$n = 5$													
0.50	0.9686	1.0139	1.0012	0.9994	0.9990	1.0007	0.9998	0.9999	1.0000	0.9999	0.9310	0.9737	0.9471
0.60	0.9739	1.0234	1.0035	0.9972	0.9950	1.0039	0.9977	0.9996	0.9999	0.9993	0.9471	0.9888	0.9471
0.70	0.9798	1.0392	1.0085	0.9905	0.9793	1.0174	0.9819	0.9982	0.9997	0.9969	0.9665	1.0112	0.9471
0.80	0.9857	1.0692	1.0206	0.9869	0.9272	1.0573	0.8828	0.9948	0.9990	0.9869	0.9881	1.0494	0.9471
0.90	0.9944	1.1260	1.0543	0.9689	0.9379	1.0911	0.4217	1.0008	0.9986	0.9379	1.0060	1.1403	0.9471
0.92	0.9984	1.1287	1.0662	0.9638	0.9276	1.0807	0.2644	1.0054	0.9993	0.9086	1.0076	1.1793	0.9471
0.94	1.0044	1.1029	1.0804	0.9581	0.9162	1.0629	0.1159	1.0107	1.0010	0.8568	1.0076	1.2386	0.9471
0.96	1.0119	0.9902	1.0970	0.9518	0.9036	1.0416	0.0206	1.0162	1.0046	0.7506	1.0050	1.3425	0.9471
0.98	1.0155	0.5925	1.1170	0.9450	0.8898	1.0204	0.0001	1.0217	1.0097	0.4804	0.9989	1.5788	0.9471
$n = 9$													
0.50	0.9961	1.0174	1.0151	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.9790	0.9892	0.9762
0.60	0.9971	1.0216	1.0179	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.9896	0.9966	0.9762
0.70	0.9979	1.0287	1.0228	1.0000	0.9996	1.0003	0.9998	1.0000	1.0000	1.0000	1.0004	1.0078	0.9762
0.80	0.9985	1.0431	1.0333	0.9995	0.9981	1.0043	0.9922	0.9994	1.0000	0.9999	1.0086	1.0270	0.9762
0.90	0.9990	1.0849	1.0652	0.9956	0.9852	1.0374	0.8150	0.9965	0.9998	0.9969	1.0092	1.0733	0.9762
0.92	0.9992	1.1041	1.0804	0.9936	0.9786	1.0471	0.6745	0.9970	0.9997	0.9937	1.0082	1.0935	0.9762
0.94	0.9997	1.1303	1.1032	0.9907	0.9696	1.0501	0.4578	0.9995	0.9996	0.9857	1.0069	1.1248	0.9762
0.96	1.0015	1.1517	1.1378	0.9867	0.9575	1.0401	0.1867	1.0041	0.9997	0.9600	1.0055	1.1817	0.9762
0.98	1.0062	1.0368	1.1888	0.9816	0.9417	1.0204	0.0096	1.0093	1.0015	0.8322	1.0037	1.3244	0.9762

Note: The table shows moments implied by the respective Markov chains ($\hat{\rho}$, $\hat{\sigma}_\varepsilon$, and $\hat{\sigma}_z$) relative to the true moments implied by the AR(1) process, (ρ , σ_ε , and σ_z).

Table 3: Approximations of HSZ's Processes

row	Z	Π				
Tauchen						
1	-1.0713	0.8920	0.1080	0.0000	0.0000	0.0000
2	-0.5357	0.0445	0.8735	0.0820	0.0000	0.0000
3	0.0000	0.0000	0.0610	0.8780	0.0610	0.0000
Tauchen-Hussey: σ_ε						
1	-0.4948	0.7376	0.2473	0.0150	0.0002	0.0000
2	-0.2348	0.1947	0.5555	0.2328	0.0169	0.0001
3	0.0000	0.0113	0.2221	0.5333	0.2221	0.0113
Tauchen-Hussey: σ_z						
1	-1.5848	0.9999	0.0001	0.0000	0.0000	0.0000
2	-0.7520	0.0000	0.9998	0.0002	0.0000	0.0000
3	0.0000	0.0000	0.0001	0.9998	0.0001	0.0000
Tauchen-Hussey: w						
1	-0.7809	0.9207	0.0792	0.0001	0.0000	0.0000
2	-0.3706	0.0476	0.8486	0.1037	0.0001	0.0000
3	0.0000	0.0000	0.0873	0.8252	0.0873	0.0000
Adda-Cooper						
1	-0.7765	0.8232	0.1701	0.0067	0.0000	0.0000
2	-0.2950	0.1701	0.5919	0.2233	0.0146	0.0000
3	0.0000	0.0067	0.2233	0.5401	0.2233	0.0067

Note: The table shows the Markov chains generated by the different methods to approximate the process $\rho = 0.95$ and $\sigma_\varepsilon^2 = 0.030$ when $n = 5$ nodes. Since Z and Π are 'symmetric', only the first three rows are reported.




Figure 1: Approximation of the HSZ process with 15 nodes, true (red) ρ and ρ conditional on node (blue, marked)




Figure 2: Approximation of the HSZ process with 15 nodes, true (red) σ_ϵ and σ_ϵ conditional on node (blue, marked)