Parameter Estimation for Stochastic Models of Interacting Agents: An Approximate ML Approach

Thomas Lux
Department of Economics
University of Kiel *

November 29, 2006

Abstract

Simple models of interacting agents can be formulated as jump Markov processes via suitably specified transition probabilities. Their aggregate dynamics might then be analyzed by the Master equation for the change of the probability distribution over time, or the Fokker-Planck equation that is obtained by a power series expansion and governs the probability distribution for fluctuations around an equilibrium. With such information on the transient density of the process, maximum likelihood estimation of its parameters becomes feasible. Even if the Fokker-Planck equation can not be solved explicitly, one can resort to numerical approximations like the Crank-Nicolson method for approximate ML estimation. We explain this algorithm with a simple model of interacting agents and show that the approximate ML procedure works well and has desirable accuracy even in the case of bimodal limiting distributions. We illustrate possible applications by estimating the parameters of this model for a popular business climate index for the German economy.

- preliminary version -

*Financial support from the EU STREP ComplexMarkets (contract number 516446) is gratefully acknowledged.

1 The basic framework

Many agent-based models with simple rules of agents' adaption to external influences allow for a representation of their aggregate quantities via Markovian stochastic processes. A comprehensive formalization of the time development of the probability distribution over their configuration space can be attempted via the Master equation formalism or approximations to it such as the Fokker-Planck equation for the dynamics of the transitional density (cf. the seminal methodological contributions by Weidlich and Haag, 1983; Aoki, 1996; and Weidlich, 2002). In the following we will particularly focus on the fertilization of the Fokker-Planck equation for estimation of the parameters of the underlying hypothetical model. The Fokker-Planck equation associated to a stochastic process is a parabolic differential equation for the change in time of the transitory density of the process. Although the Fokker-Planck equation occupies a very prominent place in statistical physics (Risken, 1989; Frank, 2005), it seems that due to the different research perspectives in this field, it has never been used as a tool for estimation of parameters of physical models.

Nevertheless, the use of the Fokker-Planck equation for parameter estimation is straight forward: if on has available discrete observations of a diffusion process and if the Fokker-Planck equation of the hypothesized process could be solved explicitly, the time-dependent solution to the transient density at the times of observations could be used to estimate the parameters via a standard maximum likelihood approach. Unfortunately, in models of interacting agents, a closed-form solution to the Fokker-Planck equation is usually not available. In this case, however, we could resort to numerical approximations of the Fokker-Planck equations. Numerical integration of partial differential equations via finite difference of finite element methods is also a well developed field (Thomas, 1995) and has found important applications both in statistical physics (Park and Petrosian) and financial mathematics (Seydel, 2002, part III). A well-known area of application is the pricing of American options and exotic options for which no closed-form solutions of the modified Black-Scholes equation exist. The only application within an estimation framework can be found in a different branch of computational finance, namely diffusion processes of the term structure of interest rates. The first to propose approximate ML estimation on the base of a numerical integration of transitory densities has been Poulsen (1999) whose approach has been compared to alternative methods by Jensen and Poulsen (2002). Hurn *et al.* (2006) propose refinements using finite elements rather than finite differences.

In order to set the stage for the presentation of this methodology, consider a parabolic stochastic differential equation:

$$\frac{\partial f(x)}{\partial t} = \frac{\partial}{\partial x}(\mu(x,\theta)f(x)) + \frac{\partial^2}{\partial x^2}(g(x,\theta)f(x)) \tag{1}$$

If (1) refers to a Fokker-Planck equation, the unknown function $f(x,t)$ is the transitory density of x, and we can write $\mu(x,\theta) = -A(x,\theta)$, $g(x,\theta) = \frac{1}{2}D(x,\theta)$ with $A(x,\theta)$ and $D(x,\theta)$ the drift and diffusion functions of the process, and θ a set of unknown parameters that one wants to estimate.

If no closed-form solution for $f(x,t)$ is available (which will mostly be the case), one can study the time development of the density via numerical integration of eq.(1). Various methods for discretisation of the stochastic equation (1) can be used. If one uses an evenly spaced discrete grid, one speaks of *finite difference methods*, if one uses a flexible grid, one speaks of refined finite element methods.

In principle, the first and second derivatives on both sides of eq. (1) could be approximated either via forward differences of backward differences (called explicit or implicit methods). Higher accuracy of the approximation can be achieved by combining both forward and backward differences by computing central differences around intermediate grid points. This approach is known as the Crank-Nicolson method and will be adopted in what follows.

To concretize the finite difference approximation, consider a 'space' grid with distance h between adjacent knots: $x_j = x_0 + j \cdot h$; $j = 0, 1, ..., N_x$ and similarly equally spaced points along the time axis between $t = 0$ and the final time T: $t_i = i \cdot k$ with $i = 0, ..., N_t$ and $k = \frac{T}{N_t}$.

In a forward discretization, (1) would have to be replaced by

$$\frac{f_j^{i+1} - f_j^i}{k} = \frac{\mu_{j+1} f_{j+1}^i - \mu_j f_j^i}{h} + \frac{g_{j+1} f_{j+1}^i - 2g_j f_j^i + g_{j-1} f_{j-1}^i}{h^2} \tag{2}$$

with $f_j^i := f(x_0 + j \cdot h, ik)$ and $\mu_j := \mu(x_0 + j \cdot h, \theta)$, $g_j := g(x_0 + j \cdot h, \theta)$. Replacing the forward difference on the left-hand side by the backward difference $f_j^i - f_j^{i-1}$, we obtain the implicit finite difference approximation. While the forward and backward approximations are of local accuracy (at the mesh points) $O(k) + O(h^2)$, higher accuracy can be obtained by taking the average of both the forward and backward difference approximation. This is known as the Crank-Nicolson method and can be shown to have local accuracy $O(k^2) + O(h^2)$. Note that the Crank-Nicolson approach effectively approximates the continuous-time diffusion at intermediate points $(i + \frac{1}{2})k$ and $(j + \frac{1}{2})h$ rather than those on the grid itself.

Because of the necessity of restricting the approximation to a finite interval, boundary conditions have to be imposed in order to prevent transitions to inaccessible states. In the Crank-Nicolson approach this requires the assumption

$$f^j_{-\frac{1}{2}} = f(x_0 - \frac{1}{2}h, jk) = 0 \qquad \text{and} \qquad f^j_{N+\frac{1}{2}} = f(x_0 + (N_x + \frac{1}{2})h, jk) = 0 \quad (3)$$

(cf. Giuliani). While such simple Dirichlet boundary conditions preserve the local second order accuracy, more complex derivative boundary conditions in certain applications would require a careful analysis of the errors brought about by their discretization. It is important to note that the no-flux boundary conditions guarantee conservation of probability mass within the underlying x-interval if (1) governs the dynamics of a transient density (i.e. if (1) is a Fokker-Planck equation).

The second order accuracy of the Crank-Nicolson scheme can be checked in applications by trying different step sizes h or k. Denote by v_1, v_2 and v_3 the approximations of the continuous solution f, using step sizes k and h, $\frac{h}{2}$ and $\frac{h}{4}$, respectively. Then by expanding the error of the approximation in a Taylor series:

$$v_1 = f - hc - kd - h^2 l - k^2 m + ... \quad (4)$$

$$v_2 = f - 0.5hc - kd - 0.25h^2 l - k^2 m + ... \quad (5)$$

$$v_3 = f - 0.25hc - kd - 0.0625h^2 l - k^2 m + ... \quad (6)$$

It follows that the quotient of the differences of these approximations yields:

$$\frac{v_2 - v_1}{v_3 - v_2} \simeq 2\frac{c + 1.5hl}{c + 0.75hl} \quad (7)$$

Hence, if the method is first-order accurate, $c \neq 0$ should be the dominating component and evaluating (7) at the grid points, we would expect to see values close to 2. On the contrary, prevalence of values around 4 all over the place would be seen as a confirmation of the theoretical second-order accuracy of the Crank-Nicolson scheme. The same operation can be performed in the time direction as well using differences k, $\frac{k}{2}$, and $\frac{k}{4}$. We will see an illustration of this experimental determination of the order of accuracy in our application below.

On the base of the Crank-Nicolson (or any other finite difference approximation), we can estimate the parameters of a diffusion process with discretely spaced observations via approximate maximum likelihood.

The negative log-likelihood of a sample of observations $X_0, \ldots, X_T$ is

$$-log f_0(X_0 \mid \theta) - \sum_{s=0}^{T-1} log f(X_{s+1} \mid X_s, \theta) \tag{8}$$

where $f_0(X_0 \mid \theta)$ is the density of the initial state (which in practical applications will be skipped because of its negligible influence and the possible lack of a closed-form solution for the stationary density) and $f(X_{s+1} \mid X_s, \theta)$ is the value of the transitional density at $s+1$ conditioned on the previous observation at time s, X_s. This continuous density is approximated by our finite difference scheme. Poulsen (1999) shows that the pertinent estimator is consistent, asymptotically normal and can be asymptotically equivalent to full ML estimates, at least under the Crank-Nicolson approximation scheme.

2 Finite difference approximation of a canonical' interaction model

For illustration of the above framework, we use an approach that goes back at least to Weidlich and Haag (1983) and had been used in a macroeconomic setting by Kraft, Landes and Weise (1986) among others and in behavioral finance models by Lux (1995, 1997). The model deals with a binary choice problem and stochastic transitions of agents between both alternatives due to exogenous factors and group pressure. Let the two groups have occupation numbers n_+ and n_- ,respectively, with the overall population size being $2N$ (multiplication by 2 simply serves to avoid the case of an odd number of individuals).
The socio-economic configuration at any time can be described via the difference between group occupation numbers:

$$n = \frac{1}{2}(n_+ - n_-) \tag{9}$$

or an equivalent opinion index:

$$x = \frac{n}{N} = \frac{n_+ - n_-}{2N} \qquad \text{with } x \in [-1, 1]. \tag{10}$$

A simple stochastic process of individual moves between groups can be built upon Poisson probabilities to jump from the "+" to the "−" group or *vice versa* within the next instant. We denote this transition rates by $w_\uparrow$ and $w_\downarrow$ and assume that they are the same for all agents within our group.
For the sake of our illustration, we follow the earlier literature assuming an exponential functional form of $w_\uparrow$ and $w_\downarrow$:

$$w_\uparrow = v\, exp(U), \qquad w_\downarrow = v\, exp(-U) \tag{11}$$

The forcing function U is assumed to consist of a constant factor (bias) α_0 and a second component formalizing group pressure in favor or against homogeneous decisions, $\alpha_1 x$:

$$U = \alpha_0 + \alpha_1 x. \tag{12}$$

The parameters of the model are, thus: v which determines the frequency (time scale) of moves between groups, α_0 which generates a bias towards the choice of "+" ("−") if positive (negative) and α_1 which formalizes the degree of group pressure (if it is positive, if negative it would rather imply a tendency of non-conformity). Models with this basic ingredients have been thoroughly investigated in the literature. The basic outcome of the model can be summarized by the following findings:

i) For $\alpha_1 \leq 1$, the group dynamics defined by (11) and (12) is characterized by a stationary distribution with a unique maximum. If $\alpha_0 = 0$, this maximum is located at $x^* = 0$. It shifts to the right (left) for $\alpha_0 > 0$, (< 0).

ii) For $\alpha_1 > 1$ and α_0 not too large, the stationary distribution has two maxima. If $\alpha_0 = 0$, the distribution is symmetric around 0. It becomes asymmetric if $\alpha_0 \neq 0$ with right-hand (left-hand) skewness and more concentration of probability mass in the right (left) maximum if $\alpha_0 > 0$, (< 0) holds.

iii) If $|\alpha_0|$ becomes very large, the smaller mode vanishes and the stationary distribution becomes uni-modal again. This happens if $|\alpha_0|$ increases beyond the bifurcation value $\overline{\alpha_0}$ given by:

$$\cosh^2(\overline{\alpha_0} - \sqrt{\alpha_1(\alpha_1 - 1)}) = \alpha_1 \tag{13}$$

In most applications, the first step towards an analysis of the above group dynamics consists in the derivation of a quasi-deterministic law of motion for the first moment of x:

$$\frac{d}{dt}\overline{x} = v(1 - \overline{x})e^{\alpha_0 + \alpha_1 \overline{x}} - v(1 + \overline{x})e^{-\alpha_0 - \alpha_1 \overline{x}} \tag{14}$$

(14) is exact in the limit of an infinite population size and provides a first-order approximation of the dynamics of x for finite populations.

One easily recovers that the features of the unconditional distribution ((i) to (iii)) are reflected in the existence and stability of steady states of (14).

A more comprehensive description of the dynamics can be obtained via the Fokker-Planck equation. Its drift component can be shown to be:

$$A(x) = \frac{n_-}{2N} w_\uparrow(x) - \frac{n_+}{2N} w_\downarrow(x) = v(1-x)e^{\alpha_0+\alpha_1 x} - v(1+x)e^{-\alpha_0-\alpha_1 x} \quad (15)$$

which, of course, coincides with the right-hand side of (14), while the diffusion term is:

$$D(x) = \frac{1}{N}\left(\frac{n_-}{2N} w_\uparrow(x) + \frac{n_+}{2N} w_\downarrow(x)\right) = \frac{1}{N}\left(v(1-x)e^{\alpha_0+\alpha_1 x} + v(1+x)e^{-\alpha_0-\alpha_1 x}\right) \quad (16)$$

This is certainly a case in which the conditional density can not be solved for explicitly due to the high degree of non-linearity of both the drift and diffusion components. For numerical integration, we can, however, resort to the Crank-Nicolson scheme as introduced above. Fig. 1 shows an example with a strongly peaked initial distribution which evolves into a bi-modal distribution over time. Underlying parameters are: $v = 3, \alpha_0 = 0, \alpha_1 = 1.2, N = 50$ for the parameters of the agent-based model, $h = 0.0025$ and $k = 0.01$ for the discretization in space and time, $T = 2$ for the time horizon of the numerical integration and a space grid extending from -1 to 1 in accordance with the support of the variable x has been used. The initial condition, $x_0 = 0$, has been approximated by a Normal distribution density $\Phi^N(x_0 + A(x)k, D(x)k)$ evaluated at grid points $-1 + jh; j = 0, 1, \ldots, N_x$, in the x direction for the first time increment k. This avoids the problems of a Dirac δ-function as initial condition and can be interpreted as a first-order Euler approximation using the known drift and diffusion functions for the initialization of the approximation.

We proceed by checking the theoretical second-order accuracy of our approximation. We perform the order determination separately in each direction for the approximation exhibited in Fig. 1. Table 1 exhibits results for selected grid points with (7) applied to both the h and k distances. As can be seen, the expected dominance of values close to 4 is nicely confirmed and we can convince ourselves that the algorithm has no problem in tracking the transition from uni-modality to bi-modality with the required degree of accuracy.

$x \mid t$	0.25	0.5	0.75	1	1.25	1.5	1.75	2
-0.75	4.0	4.0	3.9	4.0	4.0	4.0	4.0	3.9
-0.5	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0
-0.25	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.25	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.5	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.75	4.0	4.0	3.9	4.0	4.0	4.0	4.0	3.9

h-ratio

$x\|t$	0.25	0.5	0.75	1	1.25	1.5	1.75	2
-0.75	4.0	3.9	4.0	4.0	4.0	4.0	4.0	4.0
-0.5	3.9	4.0	4.0	4.0	4.0	4.0	4.0	4.0
-0.25	3.9	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0	6.1	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.25	3.9	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.5	3.9	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.75	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0

k-ratio

Table 1: Order determination for the Crank-Nicolson method applied to the interacting agent model. All parameter values and settings like in Fig. 1

$x\|t$	0.25	0.5	0.75	1	1.25	1.5	1.75	2
-0.75	6.1	4.2	4.0	4.0	4.0	4.0	4.1	4.1
-0.5	5.0	4.1	4.0	4.0	4.1	4.1	4.2	4.2
-0.25	4.4	4.0	4.0	4.1	4.2	4.5	4.9	4.8
0	4.1	4.0	4.1	7.2	3.4	3.5	3.3	2.1
0.25	4.0	4.2	3.7	3.7	3.6	3.2	5.9	4.4
0.5	4.1	3.6	5.7	4.2	4.1	4.1	4.0	4.0
0.75	4.4	4.0	4.0	4.0	4.0	4.0	4.0	4.0

h-ratio

$x\|t$	0.25	0.5	0.75	1	1.25	1.5	1.75	2
-0.75	13.3	4.1	4.0	4.0	4.0	4.0	4.0	4.0
-0.5	5.8	4.0	4.0	4.0	4.0	4.0	4.0	4.0
-0.25	4.3	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.25	4.0	4.1	4.0	4.0	4.0	4.0	4.0	4.0
0.5	4.1	4.0	4.0	4.0	4.0	4.0	4.0	4.0
0.75	3.9	4.0	4.0	4.0	4.0	4.0	4.0	4.0

k-ratio

Table 2: Order determination with a different initial value, $x_0 = 0.9$. All other parameters and settings as in Fig. 1 and Table 1.

Results become slightly worse if one considers more extreme starting points: Table 2 exhibits error ratios at selected grid points for the same model parameters and approximation scheme like in Table 1 but with $x_0 = 0.9$ rather than

$x_0 = 0$. As can be seen, the approximation suffers somewhat at small t for values very far from the initial value. This deviation from second-order accuracy is likely due to the initialization via the Euler approximation (which is not second order accurate) but this effect gets nicely washed out with increasing time horizon.

3 Monte Carlo simulations of approximate ML estimation

We now turn to estimation of model parameters on the base of the numerical approximation to the Fokker-Planck equation. Poulson (1999) has demonstrated that this approximate likelihood approach is consistent and asymptotically normal and that it can be asymptotically equivalent to 'true' ML estimation under certain conditions. In his Theorem 3, he shows that the grid size has to behave like $h(t) = T^{-\delta}$ with $\delta > \frac{1}{4}$ which will be guaranteed in our applications. He also points out that - in contrast to simulated ML approaches - there is no stochastic approximation error and the accuracy of the approximation is directly controlled by the user. In order to study the performance of the method we conduct a small simulation experiment on the base of our canonical interaction model. Because of the time needed for approximate ML with numerical integration of the transient density we have to restrict this Monte Carlo study to a few selected parameter values. The following sets of parameters have been chosen:

- set I: $v = 3$, $\alpha_0 = 0$, $\alpha_1 = 0.8$
- set II: $v = 3$, $\alpha_0 = 0.2$, $\alpha_1 = 0.8$
- set III: $v = 3$, $\alpha_0 = 0$, $\alpha_1 = 1.2$
- set IV: $v = 3$, $\alpha_0 = 0.2$, $\alpha_1 = 1.2$

In all scenarios, $N = 50$, i.e. the population size is equal to 100 ($2N$). Our choice of parameters is governed by our interest to compare the performance in situations with uni-modal and bi-modal distributions, as well as situations with and without a bias term $\alpha_0 \neq 0$.

Because of the computational demands of this method, the sample size has been restricted to $T = 200$ observations at discrete integer time intervals which have been extracted from a true multi-agent simulation with small time increments $\Delta t = 0.01$. The time scaling parameter v has been fixed in order to have some chance of switching between both modes in the bi-modal case as otherwise we would not expect the estimation procedure to detect a bi-modal distribution (whether this conjecture really holds, might be checked in subsequent Monte Carlo experiments). The Crank-Nicolson finite difference discretization is applied with widths $k = \frac{1}{8}$ and $h = 0.02$ in the time and space direction, respectively (note that in the space direction $h = 0.02$ corresponds exactly to the discreteness of the index x for our setting with $N = 50$).

In order to have at least a certain benchmark for comparison of accuracy of the parameter estimates, we compare the resulting estimates with those obtained under $k = 1$. The later can be interpreted as an Euler approximation since it approximates the transient density by a Normal distribution (with mean and standard deviation taken from the drift and diffusion functions of the Fokker-Planck equation) which in the Crank-Nicolson approach is used only for the initialization of the iterations. This Euler approximation does, of course, not yield consistent estimates and so we would expect it to be inferior to the Crank-Nicolson-ML approach. In order to get some insight into the dependence of the parameter estimates on the step size used in the Crank-Nicolson approximation, we also compare results obtained with time increments $k = \frac{1}{8}$ and $k = \frac{1}{16}$.

Table 3 shows our results exhibiting the mean estimates, finite sample standard errors and root-mean squared errors for all underlying parameters. The main message is that we can estimate the parameters v, α_0 and α_1 quite accurately even for our relatively small sample of 200 observations. In all cases, the Crank-Nicolson estimates are by far better than those obtained on the base of the Euler approximation, in terms of bias and standard error. One also infers that estimated parameters become somewhat less reliable in the cases of parameter sets II and IV as compared to I and III, respectively. The reason is probably that a positive bias interferes with the effects of interaction so that the variability of estimated parameters across samples increases. Nevertheless, the overall bias and standard error still remain reasonable even in those cases with $\alpha_0 = 0.2$ (with the exception perhaps of the estimates of v for parameter set IV). In contrast, Euler estimates appear essentially useless in these cases. As concerns the influence of the density of the grid, we observe only minor differences between the Crank-Nicolson approximations with $k = \frac{1}{8}$ and $k = \frac{1}{16}$. In fact, results do not uniformly improve when reducing the time increments: while one obtains slight improvements for the parameters α_0 and α_1, the estimates of v seem to deteriorate. The near equivalance of both settings together with seemingly reasonable biases and standard errors suggests the conclusion that using finer grids would probably not improve significantly the quality of the parameter estimates.

Another set of Monte Carlo experiments is motivated by realizing that the number of agents (the system size) N appears as a variable in the diffusion part of the Fokker-Planck equation. Neglecting the issue of discreteness of N, we can, in principle, also use our approach to arrive at an estimate of the number of active agents instead of imposing a predetermined value of N. In our pertinent Monte Carlo experiments, we use again parameter sets I and IV (i.e. a symmetric setting with weak herding and an asymmetric one with strong herding), with $N = 50$ or $N = 500$ in both cases. The results are exhibited in Table 4... (to be added)

		Euler ($k = 1$)			Crank-Nicolson ($k = 1/8$)			Crank-Nicolson ($k = 1/16$)		
		v	α_0	α_1	v	α_0	α_1	v	α_0	α_1
set I	mean	0.999	-0.001	0.642	2.980	-0.000	0.793	3.023	-0.000	0.794
	FSSE	0.091	0.007	0.052	0.567	0.005	0.028	0.585	0.005	0.028
	RMSE	2.003	0.007	0.166	0.564	0.005	0.028	0.583	0.005	0.028
set II	mean	0.439	0.578	0.123	2.992	0.216	0.772	3.547	0.211	0.782
	FSSE	0.048	0.097	0.171	1.046	0.057	0.105	1.422	0.038	0.069
	RMSE	2.561	0.390	0.698	1.041	0.059	0.108	1.517	0.039	0.071
set III	mean	1.019	0.000	1.173	2.884	0.000	1.196	2.952	0.000	1.196
	FSSE	0.126	0.024	0.034	0.457	0.009	0.015	0.499	0.009	0.015
	RMSE	1.913	0.024	0.043	0.469	0.009	0.016	0.499	0.009	0.015
set IV	mean	0.232	1.741	-0.698	1.369	0.262	1.123	1.748	0.245	1.144
	FSSE	0.026	0.350	0.426	0.245	0.127	0.159	0.439	0.127	0.159
	RMSE	2.768	1.580	1.945	1.326	0.141	0.175	1.326	0.135	0.168

Table 3: Approximate ML Estimates: the table displays the mean parameter estimates over 200 Monte Carlo replications together with their finite sample standard errors (FSSE) and root mean squared errors (RSME).

4 Estimation of Interactive Opinion Formation: The Case of the ZEW Business Climate Index

	ν	α_0	α_1	α_2	N	logL	AIC	BIC
Model 1 (baseline)	0.78 (0.06)	0.01 (0.01)	1.19 (0.01)			-726.9	1459.8	1464.1
Model 2 (end. N)	0.15 (0.07)	0.09 (0.06)	0.99 (0.14)		21.21 (9.87)	-655.9	1319.7	1322.0
Model 3 (feedback from i)	0.14 (0.07)	0.15 (0.17)	0.94 (0.21)	-1.09 (2.55)	19.55 (10.13)	-655.7	1321.5	1321.8
Model 4 (feedback from IP)	0.13 (0.06)	0.09 (0.07)	0.93 (0.16)	-4.55 (2.53)	19.23 (8.78)	-650.4	1310.9	1311.1
Model 5 (no interaction)	0.04 (0.01)	0.45 (0.12)		-11.93 (3.98)	5.75 (1.59)	-654.2	1316.4	1318.7

Table 5: Parameter Estimates for Stochastic Models of Interacting Agents. Details on the underlying models appear in the main text. The numbers in brackets are standard errors of parameter estimates.

Since we have focused an a very simple interaction scheme, it is not obvious that its structural features should be easily applicable to economic data. Weidlich and Haag (1983, c.5) and Kraft, Landes and Weise (1986) had proposed simple business cycle models with, for example, investment decisions being driven by an opinion process like the one outlined in Sec. 2. Such models could be estimated using the above methodology. We leave this more demanding multi-variate application to future research and turn to a particular type of uni-variate time series in which interaction effects could arguably play some role. Various surveys for *business climate* or *sentiment* are regularly conducted in many countries that seem to receive much more attention by the public than by academic researchers. The leading examples are the Michigan Consumer Sentiment Index and the Conference Board Index for the U.S. economy, which have been reported monthly since the end of the 70ties (Ludvigson, 2004, Souleles, 2004). In Germany, similar surveys are conducted by the Ifo Institute (Ifo Business Climate Index) and the Center for European Research (ZEW) at the University of Mannheim (denoted the ZEW Index of Economic Sentiment). Among these indices, the ZEW Sentiment index comes closest to the simple structure of our 'canonical' model in that it very literally asks for wether respondents are optimistic ("+") or pessimistic ("-") concerning the prospects of the German economy over the next six months. The only difference to our model is that ZEW also allows for a neutral assessment. We might assume that neutral subjects can be assigned half and half to the optimistic and pessimistic camp which, then, would allow us to apply our model directly to their data. The index is, in fact, reported as the percentage of optimists minus pessimists so that it can be directly used as the opinion index x introduced in Sec. 2. In contrast, the indices for the U.S. economy are computed as weighted averages over categorial answers to different questions while Ifo starts with sector-specific surveys and aggregates them to an overall business climate indicator. The ZEW index is, therefore, probably the only one that delivers us with an aggregate of very pure binary (resp., ternary) assessments. Fig. 2 displays the entire available monthly ZEW series (starting in December 1991 and running through July, 2006). Despite quite a number of differences in the data collection process, its development is broadly parallel to that of the Ifo index. What is striking is the very pronounced cyclical behavior of the ZEW index with very sudden movements upward and downward and a certain stagnation at times at a high or low plateau. One could, in fact, argue that the dynamics of the ZEW index is reminiscent of a bi-modal stochastic dynamics switching between a high positive and a moderately negative equilibrium. It might be worthwhile to contrast this series with what it is designed to predict, the cyclical component in economic activity. This cyclical component appears in the lower panel of Fig. 2 in the form of residuals of monthly industrial production from the Hodrick-Prescott filter, which is widely seen as the state-of-the-art approach for disentangling trend components and cyclical components in economic activity. Somewhat surprising, the *perception* of the business cycle dynamics as reflected in the survey allows a much more clear cut categorization of its phases than the much more random nature of filtered IP.

The ZEW surveys are based on about 350 respondents so that we might take this information as a parametric restriction on N (assuming N=175). We, then, have to estimate the parameters v, α_0 and α_1 in a baseline application of our interacting-agents framework. Results are shown in Table 5. Interestingly, the crucial parameter α_1 is significantly larger than unity indicating bi-modality of the limiting distribution. Despite the impression of a dominance of positive assessment over the whole sample period (quite in contrast to stereotypes of German "angst") the bias term α_0 turns out to be not significantly different from 0. Unfortunately, simulations of the estimated model show, that it most likely would get stuck within one mode over a time horizon of the length of our sample (176 observations) and would on average at most switch only once from one mode to the other (cf. Fig. 3 and 5 below). Transitions between modes are governed by chance fluctuations and become more and more unlikely the higher the number of agents. *Vice versa*, frequent switches would only occur for a relatively small size of the underlying population. In order to reconcile our observation of a relatively large number of apparent switches of the mood of the respondents with the 'official' system size of 350 respondents, we could argue that the 'effective' system size is smaller than the official number. This would happen if some respondents would actually move broadly synchronously and would, therefore, nor act like independent agents. While we cannot check this assertion due to the anonymity of the data, we could let the index itself speak on the underlying effective system size by adding N to the list of parameters estimated via approximate ML. This approach is somewhat similar to that of a recent paper by Chen (2002) who argues that the relative extent of fluctuations of macroeconomic data is neither in agreement with a representative-agent framework nor with the assumption of a system size identical to the number of individuals or companies in an economy. On the base of simple stochastic models, he argues that the relative deviation (the square of the means divided by the variance of the data) yields an estimate of the implicit number of degrees of freedom. While this rough non-parametric approach somewhat surprisingly provides us with $\frac{\mu^2}{\sigma^2} \approx 0.9$, our parametric model arrives at an estimate of $N \approx 20$ (i.e. 40 independent agents or groups of agents). Note that this added flexibility leads to a relatively large increase in the log likelihood and is preferred over the baseline model by both the AIC and BIC criteria. As concern estimated parameters, α_0 still is insignificant, while the interaction coefficient falls marginally below 1 indicating uni-modality albeit with possibly large excursions into extreme configurations. Remarkably, the estimate of the parameter v decreases from 0.78 to 0.15 when proceeding from model 1 to model 2. The likely reason is that the first estimation would have to come up with a higher mobility of the population (higher propensity to change opinion) in order to compensate for the stagnatory tendency of the larger imposed population of model 1.

We have remarked in sec. 2 that our framework allows to incorporate exogenous effects on the opinion formation process. In order to do so we simply could expand the influence function U by introducing additional factors that could be of importance to the assessment of the business cycle by the respondents of the survey:

$$U = \alpha_0 + \alpha_1 x + \alpha_2 y \tag{17}$$

Most naturally, y could be macroeconomic data of the same frequency itself (i.e. monthly), although our framework could also accommodate data of higher or lower frequency. Two such macro feedbacks have been investigated in our models 3 and 4: the interest rate (3 month FIBOR rate) and industrial production (HP filtered, as displayed in Fig. 2). Note that the direction of the feedback is not predetermined in our model, i.e. α_2 could turn out to positive or negative. The outcome of the exercise shows that industrial production adds more explanatory power than the inclusion of the interest rate: for the former we obtain a significantly negative coefficient together with lower values of the AIC and BIC criteria. For the interest rate, in contrast, the estimated coefficient α_2 is not significant and overall improvements compared to model 2 are smaller. However, even for the preferred model 4, the improvement compared to model 2 is much smaller than the increase in $logL$, AIC and BIC achieved by adding N as a free parameter (the step from model 1 to model 2). Furthermore both the inclusion of the interest rate and of industrial production have only very negligible effects on the parameters of the interacting agent model. What is perhaps puzzling is the negative sign of the feedback effect from industrial production which is in contrast to a positive contemporaneous correlation of about 0.28 between both series. It appears to depict some type of 'contrarian' behaviour: if the economic data is indicating a boom phase, our respondents seem already forestall the subsequent overheating of the economy and the subsequent downturn and *vice versa*.

We tried also whether we could get rid of interaction, but the results of the pertinent exercise (model 5) are not too plausible: a very large positive predisposition ($\alpha_0 = 0.48$) together with a small population ($N = 6$) and a low degree of flexibility ($v = 0.03$) is needed to compensate for the missing interaction effect. AIC and BIC rank this variant inferior to model 4.

Table 6: Unconditional moments from 1.000 Monte Carlo simulations

	data	models				
		1	2	3	4	5
mean:	0.352	-0.588	0.349	0.335	0.324	0.368
(95%)		(-0.631, -0.405)	(0.045, 0.552)	(0.100, 0.517)	(0.035, 0.528)	(0.149, 0.540)
std. dev:	0.370	0.098	0.355	0.376	0.384	0.334
(95%)		(0.068, 0.185)	(0.251, 0.484)	(0.265, 0.508)	(0.293, 0.506)	(0.244 ,0.429)
skewness:	-0.620	0.615	-0.908	-0.932	-0.991	-0.900
(95%)		(0.029 ,1.685)	(-1.880, 0.097)	(-1.778 ,-0.068)	(-1.878, 0.053)	(-1.764, -0.054)
kurtosis:	-0.428	0.575	0.591	0.390	0.540	0.859
(95%)		(-0.652, 3.717)	(-1.322, 4.347)	(-1.308 ,3.775)	(-1.283, 3.296)	(-0.869, 3.398)
rel. deviation:	0.905	49.705	1.368	1.116	1.030	1.298
(95%)		(8.706, 82.360)	(0.024, 4.022)	(0.055, 3.418)	(0.022, 2.792)	(0.225, 3.238)
distance:		0.952	0.335	0.333	0.297	0.290
(95%)		(0.899, 0.986)	(0.255, 0.493)	(0.259, 0.462)	(0.220, 0.455)	(0.228, 0.378)

Table 7: Autocorrelations from 1.000 Monte Carlo simulations

		models				
ACF	data	1	2	3	4	5
1	0.935	0.630	0.923	0.936	0.939	0.925
(95 %)		(0.456 0.963)	(0.845 0.967)	(0.876 0.969)	(0.908 0.968)	(0.881 0.956)
2	0.830	0.404	0.853	0.875	0.880	0.851
		(0.162 0.929)	(0.715 0.936)	(0.764 0.940)	(0.820 0.934)	(0.762 0.908)
3	0.709	0.266	0.789	0.816	0.819	0.780
		(0.013 0.890)	(0.595 0.907)	(0.658 0.913)	(0.732 0.900)	(0.661 0.859)
4	0.584	0.175	0.729	0.761	0.758	0.710
		(-0.080 0.857)	(0.496 0.883)	(0.556 0.886)	(0.652 0.866)	(0.561 0.809)
5	0.465	0.116	0.673	0.707	0.696	0.643
		(-0.133 0.820)	(0.398 0.860)	(0.467 0.863)	(0.566 0.833)	(0.477 0.759)
6	0.363	0.075	0.620	0.655	0.633	0.578
		(-0.171 0.784)	(0.319 0.840)	(0.382 0.841)	(0.478 0.797)	(0.391 0.710)
7	0.272	0.048	0.571	0.605	0.571	0.515
		(-0.188 0.747)	(0.241 0.813)	(0.313 0.822)	(0.392 0.759)	(0.301 0.658)
8	0.186	0.032	0.525	0.558	0.512	0.455
		(-0.197 0.703)	(0.184 0.793)	(0.246 0.801)	(0.317 0.722)	(0.226 0.611)
9	0.094	0.022	0.482	0.513	0.454	0.399
		(-0.213 0.668)	(0.121 0.774)	(0.190 0.782)	(0.239 0.678)	(0.153 0.568)
10	0.017	0.014	0.442	0.470	0.398	0.345
		(-0.220 0.631)	(0.078 0.752)	(0.133 0.760)	(0.167 0.640)	(0.079 0.519)
d	0.553	0.194	0.826	0.875	0.923	0.792
		(-0.343 0.978)	(0.338 1.261)	(0.351 1.266)	(0.455 1.346)	(0.290 1.221)

5 Some specification tests

How close do time series from the estimated models get to empirical behavior of the ZEW index? Fig. 3 exhibits three simulations over the same time horizon of model 4 together with the empirical data. For these simulations, we have used time increments $\Delta t = 0.01$ for the ongoing opinion formation between integer time steps and have injected the knowledge of the current exogenous factor (HP-filtered industrial production) at integer time steps. As it can be seen, the visual appearance of the three Monte Carlo runs is pretty similar to that of the index itself and the significant feedback from industrial production seems to direct the simulations towards a pattern that is broadly synchronous with the ups and downs of the empirical record. Fig. 4 shows the mean and 95 percent confidence bounds from the transient density computed for model 4 over the whole observation period given the first observation of the index as the initial condition and incorporating the feedback from the news about industrial production. Since the empirical record stays within the 95 % bounds for practically the entire time horizon, we may conclude that we have no reason to reject the hypothesis that the empirical data could have emerged as one particular sample path from our stochastic model. We note however that simulations of models 2, 3, and 5 would lead to very similar patterns. However, for models 2 and 3 the sample paths would not be synchronous to the empirical series simply because there is no exogenous factor (model 2) or its influence is apparently weak (as it holds for the interest rate in the case of model 3). In contrast, model 1 yields a very different pattern as shown in the lower right panel of Fig. 3 since with the higher 'official' number of respondents shifts between equilibria become less frequent than with $N \approx 20$. As can be seen in Fig. 5, the 95 percent confidence interval from the transient density of model 1 reflects the low probability of regimes switches and remains narrowly concentrated around the initial state up to above period 120. Note that the bias term now includes the exogenous component, i.e. α_0 has to be replaced by $\alpha_0 + \alpha_2 y$ (with y here given by filtered IP). The ups and downs of the sentiment index during the observation period would, then, mostly reflect switches between unique equilibria that alternate between optimistic and pessimistic majorities. Our model 4, in fact, shows how the fuzzy exogenous information in the lower panel of Fig. 2 is translated into a much clearer image of the business cycle dynamics in the view of the respondents' sentiments (upper panel of Fig. 2) via the self-referential and self-reinforcing dynamics of the opinion formation process.

As another specification test we try to assess whether the abruptness of the up and down movements of the index is captured by our model. For this purpose we compute a series of one-period iterations of the transient density and extract the 95 percent confidence intervals conditional as the realization in the previous period. Fig. 6 shows the 95 % confidence bounds for the subsequent period's realization which apparently is never left by the empirical record. Upon close investigation one might, however, find some of the downturns are getting close to the lower boundary while the ups are pretty much in the center of the 95 percent bound.

Table 5, 6 and 7 provide a statistical analysis of 1000 Monte Carlo replications of models 1 through 5 on the base of the estimated parameters displayed in Table 5. In order to get an impression of how closely we match the statistical features of the data, we compare a selection of conditional and unconditional moments. Table 6 shows the means and simulated 95 percent boundaries for the first four unconditional moments together with the relative deviation as defined in Chen (2002) and the mean absolute distance between the entries of each simulation and the 176 empirical observations. As we can see, for the first to third moment as well as the relative deviation, models 2 to 5 are all pretty close to the empirical numbers while model 1 (using the 'official' number of 350 active agents) is far off the mark in all cases. This confirms the visual impression reported above that the patterns of all models with an endogenous number of effective agents are relatively similar while model 1 stands out by its tendency of getting frozen in the lower mode due to the negative initial condition and the high level of persistence caused by the large number of 350 agents. For the remaining statistics, we first see that kurtosis is poorly matched by all models, which might however be attributed to the volatility of this measure for small samples. The distance between the empirical observations and synthetic data again shows the greatest discrepancy for model 1 compared to all others while the feedback from industrial production in models 4 and 5 seems to have contribute to a better fit compared to models 2 and 3. Again, this provides a confirmation of our visual impression reported above.

Table 7 reports autocorrelations of the simulated series for lags 1 to 10. A glance at smaller lags again indicates that ACFs from models 2 to 5 are all very close to their empirical counterpart while model 1 has a much lower degree of dependence. Unfortunately, all models are only able to match about the first four lags while the autocorrelations remain much higher than the empirical ones for the longer lags.

References

Aoki, M., *New Approaches to Macroeconomic Modeling: Evolutionary Stochastic Dynamics, Multiple Equilibria, and Externalities as Field Effects*, Cambridge, University Press 1996

Chen, P., Microfoundations of Macroeconomic Fluctuations and the Laws of Probability Theory: The Principle of Large Numbers versus Rational Expectations Arbitrage, *Journal of Economic Behavior and Organization* **49**, 2002, 307-326

Jensen, B. and R. Poulsen, Transition Densities of Diffusion Processes: Numerical Comparison of Approximation Techniques, *Journal of Derivatives* 2002, 18-34

Hurn, A., J. Jeisman and K. Lindsay, *Teaching an Old Dog New Tricks Improved Estimation of the Parameters of Stochastic Differential Equations by Numerical Solution of the Fokker-Planck Equation*, Manuscript, Queensland University of Technology, Brisbane 2006

Kraft, M.,T. Landes and P. Weise, Dynamic Aspects of a Stochastic Business Cycle Model, *Methods of Operations Research* **53**, 1986, 445-453

Ludwigson, S., Consumer Confidence and Consumer Spending, *Journal of Economic Perspectives* **18**, 2004, 29-50

Lux, T., Herd Behaviour, Bubbles and Crashes, *Economic Journal* **105**, 1995, 881-896

Lux, T., Time variation of second moments from a noise trader/infection model, *Journal of Economic Dynamics and Control* **22**,1997, 1-38

Souleles, N., Expectations, Heterogeneous Forecast Errors, and Consumption: Micro Evidence from the Michigan Consumer Sentiment Surveys, *Journal of Money, Credit and Banking* **36**, 2004, 39-72

Poulsen, R., *Approximate Maximum Likelihood Estimation of Discretely Observed Diffusion Processes*, Manuscript, University of Aarhus, 1999.

Østerby, O.,*The Error of the Crank-Nicolson Method for Linear Parabolic Equations with a Derivative Boundary Condition*, Manuscript, University of Aarhus, 1998

Giuliani, P., *Fully Implicit and Crank-Nicolson Methods*, Manuscript, University of St. Andrews

Thomas, J., *Numerical Partial Differential Equations*, Berlin, Springer 1995

Frank, T., *Nonlinear Fokker-Planck Equations*, Berlin, Springer 2005

Risken, H., *The Fokker-Planck Equation: Methods of Solutions and Applications*, Berlin, Springer 1989

Weidlich, W. and G. Haag, *Concepts and Methods of a Quantitative Sociology*, Berlin, Springer 1983

Park, B. and V. Petrosian, Fokker-Planck Equations of Stochastic Acceleration: A Study of Numerical Methods, *Astrophysical Journal Supplement* **103**, 1996, 255

Weidlich, W., *Sociodynamics: A Systematic Approach to Modeling in the Social Sciences*. London, Taylor & Francis, 2002.

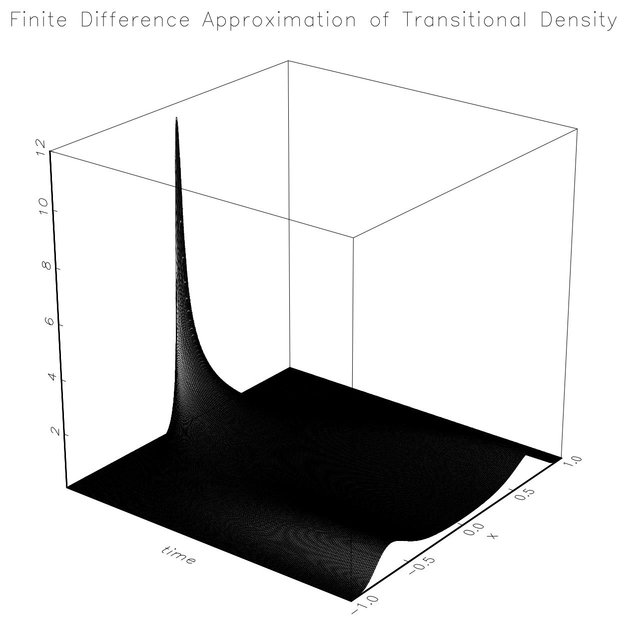

Figure 1: An illustration of the development of the transient density in the bimodal case. The initial state x_0 has been approximated by a Normal distribution with small standard deviation and mean x_0.

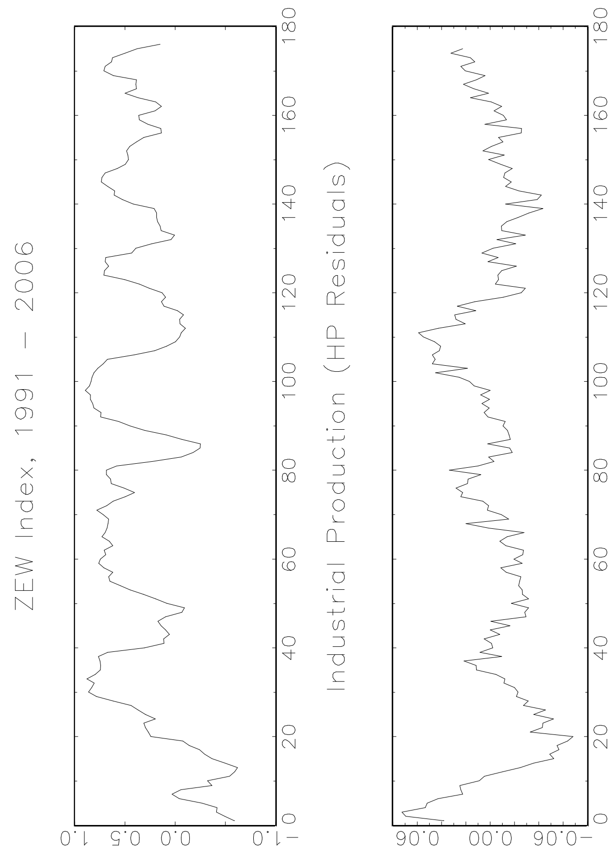

Figure 2: ZEW Sentiment Index and Industrial Production.

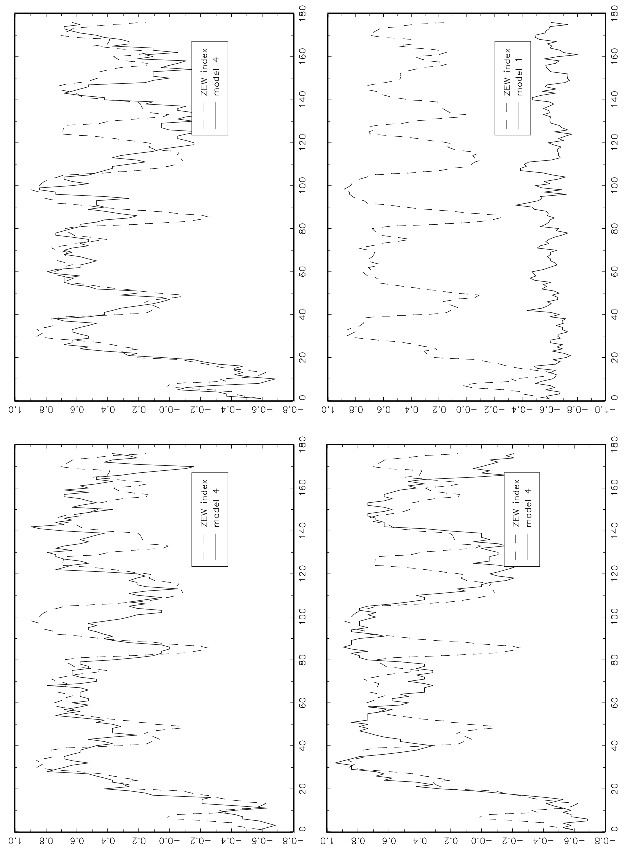

Figure 3: Simulated Trajectories from Models 4 and 1 (lower right-hand panel.)

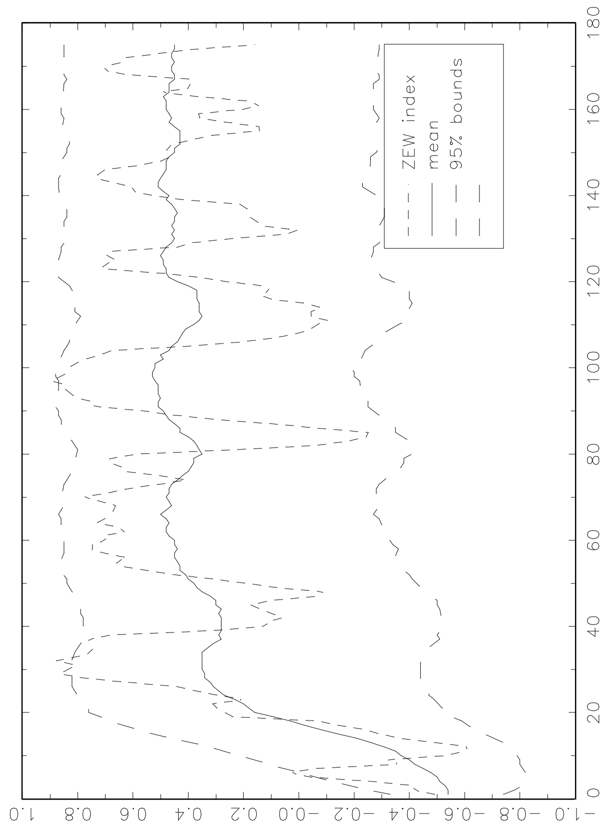

Figure 4: Mean and 95 Percent Confidence Interval for Model 4 (from Fokker-Planck Equation)

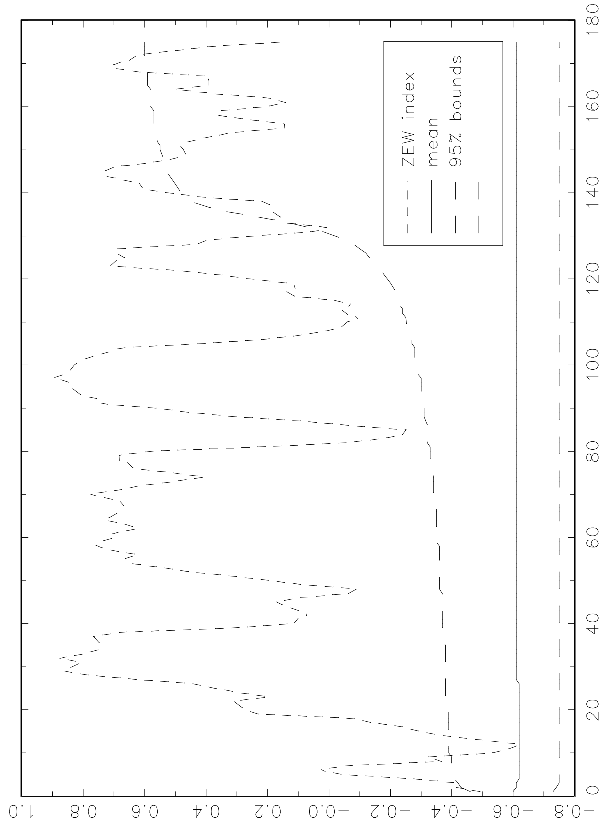

Figure 5: Mean and 95 Percent Confidence Interval from Model 1 (from Fokker-Planck Equation)

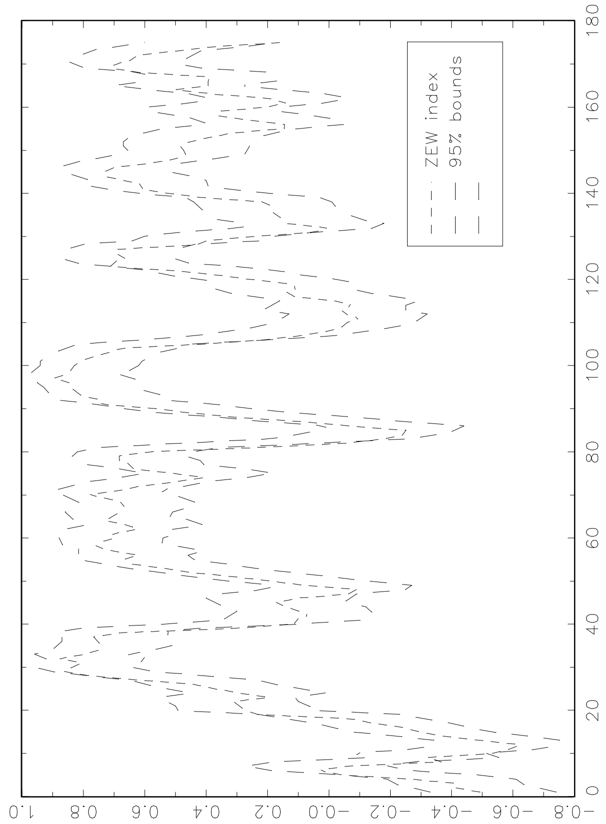

Figure 6: 95 Percent Confidence Interval for One-Step Iterations of Transient Density of Model 4.