

Rational Forecasts or Social Opinion Dynamics? Identification of Interaction Effects in a Business Climate Survey

Thomas Lux
Department of Economics
University of Kiel *
and
Kiel Institute for World Economy

May 29, 2008

Abstract

This paper develops a methodology for estimating the parameters of dynamic opinion or expectation formation processes with social interactions. We study a simple stochastic framework of a collective process of opinion formation by a group of agents who face a binary decision problem. The aggregate dynamics of the individuals' decisions can be analyzed via the stochastic process governing the ensemble average of choices. Numerical approximations to the transient density for this ensemble average allow the evaluation of the likelihood function on the base of discrete observations of the social dynamics. This approach can be used to estimate the parameters of the opinion formation process from aggregate data on its average realization. Our application to a well-known business climate index provides strong indication of social interaction as an important element in respondents' assessment of the business climate.

JEL classification numbers: C42, D84, E37

Keywords: business climate, business cycle forecasts, opinion formation, social interactions

Contact adress: Olshausenstr. 40, 24118 Kiel, Germany,
E-Mail: lux@bwl.uni-kiel.de

*Financial support from the EU STREP ComplexMarkets (contract number 516446) and from the Volkswagen foundation is gratefully acknowledged. Helpful remarks and comments by Simone Alfarano, Reiner Franke, Cars Hommes, Alan Kirman, Axel Leijonhufvud, Mark Salmon, Detlef Seese, Jamsheed Shorish and Enrico Zaninotti are gratefully acknowledged.

1 Introduction

Recent literature has started to consider the role of social interdependencies between individual decisions. The potential importance of social interactions has been highlighted in analyses of such diverse phenomena like human capital acquisition (Bénabou, 1996), drug addiction (Jones, 1994), social pathologies due to peer group effects (Glaeser, Sacerdote and Scheinkman, 1996) or herding in financial markets (Kirman, 1993; Lux, 1995). Early work on social spillovers that are not mediated by markets include Föllmer's analysis of economies with interdependent preferences and Schelling's (1971) study of the emergence of strict racial segregation as the overall outcome of the locational choice of individuals who only have a weak preference for neighbors of the same race. Existence and uniqueness of equilibria in large economies with both local or global interactions have been studied recently by Horst and Scheinkman (2006).

Empirical work on social interactions has mostly been based on an adaptation of the discrete choice framework allowing for social spillovers in agents' utility functions. Brock and Durlauf (2001a,b) provide an introduction into the econometric implementation of this approach. Extensions and applications of this approach can be found in Ioannides (2006) and Krauth (2006), among others. While the discrete choice approach typically studies social interactions in cross-sectional data and assumes that the configuration of choices represents a self-consistent equilibrium (i.e. expectations conditional on agents' beliefs concerning the behavior of others are rational), we are interested in a *dynamic* process of ongoing opinion formation within a group of agents. While our incorporation of social influences is very close (both in its spirit and its formal implementation) to Brock and Durlauf's more static approach to social interaction, we do not necessarily impose that agents have settled at an equilibrium. Another difference is that we do not model social interaction effects as due to spillovers in utility or payoff functions. Due to the nature of the time series we wish to model, we are profoundly ignorant about the relevant underlying incentives of agents. In fact, there might be no incentive component of any importance in our particular setting.

One area in which a dynamic process of opinion formation could arguably be of some relevance, is survey data on business expectations or so-called sentiment indices that are published by academic and private institutes in most developed countries. While these indices attract quite some public attention upon their regular compilation, they have only found scarce consideration in the macroeconomics literature. Due to the underlying motivation for collecting such data, much of the limited body of available literature focuses on the predictive power for macroeconomic activity of these surveys (cf. Hüfner and Schröder, 2002; Gelper et al., 2007, Taylor and McNabb, 2007). However, as far as we know, attempts at formulating behavioral models for the underlying data-generating process of these surveys are practically non-existent. It has been noted that there has been little effort to test positive models of expectation

or opinion formation on the base of the rich collection of survey data available in macroeconomics. While social interactions have been hypothesized to be of some importance in expectation formation (Carroll, 2003), such factors have to my knowledge not been incorporated explicitly in the small sample of papers testing positive models of expectation formation. The hypothesis underlying our present study is that these survey data might be viewed as the result of a social process of opinion formation among the respondents. If these data could be explained via social interactions, they would represent behavioral components of macroeconomic activity quite different from rational attempts at forecasting the future development of the business cycle. Rather than representing rational forecasts of future economic developments they could be interpreted as manifestations of *animal spirits*.

Fig. 1 gives an intuitive preview on our subsequent results. The figure contrasts the monthly observations of the ZEW Business Climate Index for the German economy compiled by the Centre for European Economic Research (German acronym: ZEW) at the University of Mannheim from about 350 respondents with a frequent measure of real economic activity (HP filtered industrial production). The business climate index is computed so that it is bounded by +1 and -1 from above and below (see sections 2 and 4 for details). Quite obvious, positive (negative) values are meant to indicate an optimistic (pessimistic) majority among respondents. The higher the absolute value, the more pronounced the positive (negative) outlook for the German economy. A glance at the lower panel shows a striking contrast to the real thing: while the output gap as measured by the residuals from the HP filter appears quite noisy, the climate index has much more obvious swings between low and high values. It appears that there is a much clearer image of the business cycle dynamics in the eyes of the observers compared to what can be extracted from real economic activity. The ZEW index is also characterized by very abrupt and drastic switches between more optimistic or more pessimistic majorities than any switch between positive or negative realizations of the output gap. The much lower noise ratio of the climate index is somewhat reminiscent of the higher volatility of stock prices compared to ex-post rational prices based on realized dividends, cf. the literature on excess volatility of financial markets (e.g. Shiller, 1981). Proponents of the excess volatility hypothesis argue that rational expectations of future dividends entering the present value model should be less volatile than their realizations. Similarly, one might argue that rational predictions of the business cycle should also be smoother than the subsequent realizations thereof as there will almost certainly be stochastic factors that are not known at the time when the forecast is formulated. However, since respondents are not requested to make a prediction of some measurable statistics of the business cycle but only issue qualitative opinions, it is not clear how to map the climate index into a prediction of GDP growth rates or industrial production.

FIG. 1 ABOUT HERE

The pronounced swings of our sentiment series is quite typical of such data. While one could, in principle, imagine that these swings are caused by the revelation of important news about the subsequent development, the hypothesis we are going to explore in this paper is that these swings are imprints of a process of social interaction among respondents. It is not difficult to imagine that respondents' changing assessments of the economic outlook are at least in part influenced by the evolution of the opinion of their peers. Interpersonal affects might come into play via private exchange of opinions but probably even more so via the influence of a 'social field' of the average mood of their peer group of which they learn through a variety of professional and private channels of communication. To test for the existence of such an interactive element in opinion formation, we will use a formalization close in spirit to that of Brock and Durlauf's discrete choice with social interactions, albeit without including any elements of utility or payoff maximization (there may be no such element in survey responses).

We adopt a framework of stochastic transitions between discrete alternatives along the lines of Weidlich and Haag (1983) and Lux (1995). While the basic goal is to identify potential interaction effects, this framework is general enough to allow us to also cover exogenous factors of influence on the opinion dynamics. Naturally enough, macroeconomic data would be our candidate explanatory variables. Including both these exogenous forces and an intrinsic feedback allows us to study their interplay in the formation of group expectations. There is another important issue we explore in our study: while we have a relatively constant number of respondents in our survey (about 350), it is not clear whether all these participants would, in fact, act as independent decision makers. This issue is quite subtle: apart from the overall hypothesized interaction effect, there might be coherence within subgroups of the entire pool of respondents that is so strong as to lead to entirely synchronized behavior. The behavior of such synchronized subgroups would simply collapse onto that of a single agent (and any member of the group would be a representative agent of it). The dynamics of the opinion formation process would look differently if certain subgroups would always move together. The framework to be formalized below allows us to cope with this phenomenon: first, we start by specifying the opinion dynamics for a given number of independent actors, equal to the average number of respondents in the survey. Since the number of agents explicitly appears as a variable in our model, we may, however, also adopt an agnostic view and let the model speak on the number of *effectively independent* agents. As it turns out, endogeneizing the number of active groups of agents allows a huge improvement in the goodness-of-fit of the model. Subsequent statistical analyses confirm that this specification covers the salient features of the data much better than alternative specifications. Further explanatory power is obtained by allowing for a 'momentum' effect in addition to the baseline social interaction. In contrast to these refinements of the social part of the dynamics, allowing for an additional feedback from macroeconomic data (e.g., industrial production) only improves slightly the goodness-of-fit with a more modest increase of the

likelihood. While the statistical properties (in terms of matching conditional and unconditional moments) of the model remain almost unchanged, the macro factor leads to synthetic replications from Monte Carlo simulations that are better able to match the particular patterns of ups and downs observed during the sample period. Our simple stochastic model also allows to compute confidence bounds for future observations from the transient density. We use these to assess whether the empirical series could be a likely realization of the process of social interaction given the initial condition and the macro influence. We also explore whether any single entry would be a probable realization conditional on the last month's entry and the contemporaneous macro feedback. As it turns out, in both cases the empirical data hardly ever move out of the pertinent 95 percent confidence intervals which nicely confirms the explanatory power of the model.

The rest of the paper is structured as follows: in section 2, the basic stochastic framework of social interactions will be introduced together with a review of its properties. Section 3 contemplates the problem of estimating the parameters of such a stochastic framework with an ensemble of interacting agents. Section 4 provides some results on Monte Carlo experiments with small samples to arrive at insights on the reliability and accuracy of our subsequent estimates. Section 5 then contains the application to the ZEW index of the business climate and section 6 provides a detailed analysis of the statistical properties of Monte Carlo replications of the estimated models to explore their explanatory power together with an assessment of their goodness-of-fit. Section 7 concludes.

2 A ‘canonical’ stochastic model of social interaction

As a simple formalization for the process of social opinion formation, we adapt an approach that goes back at least to Weidlich and Haag (1983) and that 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 aggregate outcome of this choice process at any point in time can be described via the difference between the number of individuals in the “+” and “-” groups:

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

or an equivalent *opinion index*:

$$x = \frac{n}{N} = \frac{n_+ - n_-}{2N} \quad \text{with } x \in [-1, 1]. \quad (2)$$

Agents' beliefs are either optimistic or pessimistic; they change their beliefs in continuous time, with a Poisson process describing the changes from the “+” to the “-” group or *vice versa* within the next instant. We denote the pertinent transition rates by w_{\uparrow} and w_{\downarrow} and assume that they are the same for all agents within each group.

For the sake of illustration, we follow the earlier literature quoted above by assuming an exponential functional form of the transition rates w_{\uparrow} and w_{\downarrow} :

$$w_{\uparrow} = v \exp(U), \quad w_{\downarrow} = v \exp(-U). \quad (3)$$

The function U might be labeled the ‘forcing function’ for transitions and is analogous to the utility function in a proper discrete choice setting. It is assumed to consist of a constant factor (bias) α_0 and a second component formalizing group pressure in favor or against homogeneous decisions, $\alpha_1 x$:

$$U = \alpha_0 + \alpha_1 x. \quad (4)$$

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 “+” (“-”) opinions 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). With this set-up the opinion dynamics is described as the aggregate outcome of $2N$ coupled Markov processes for agents' choices in continuous time. During small time increments Δt , the probability of an agent to switch from his previous group (decision) to the other alternative, is approximately equal to $w_{\uparrow} \Delta t$ and $w_{\downarrow} \Delta t$, respectively. Note that this formalisation implies that the strength of the social influence is the same for all agents. Interpersonal differences are covered in the stochasticity of the process, i.e. by the very fact that individual choices are not deterministic, but are only determined in expectation. In utility-based discrete choice models, this stochasticity is generated via the assumption of stochastic terms in individuals' utility functions that are random draws from an extreme-value distribution.

Models with the above basic ingredients have been thoroughly investigated in the literature. The basic features of the model can be summarized by the following findings¹:

- i) For $\alpha_1 \leq 1$, the group dynamics defined by (3) and (4) 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).

¹cf. Weidlich and Haag, 1983, chap. 2; Lux, 1995, essentially the same jump-Markov process is used by Blume and Durlauf (2003) as a dynamic process of strategy revisions in a discrete choice framework with social interactions.

- ii) For $\alpha_1 > 1$ and α_0 not too large, the stationary distribution has two maxima $x_+ > 0$ and $x_- < 0$. 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 $\bar{\alpha}_0$ given by:

$$\cosh^2(\bar{\alpha}_0 - \sqrt{\alpha_1(\alpha_1 - 1)}) = \alpha_1 \quad (5)$$

with $\cosh(\cdot)$ denoting the hyperbolic cosine, $\cosh(x) = \frac{1}{2}(e^x + e^{-x})$. One might note that these findings are perfectly analogous to those in models of discrete choice with social interactions, cf. Brock and Durlauf (2001b, propositions 1 through 3): Moderate influence of social interaction ($\alpha_1 \leq 0$) leads to a balanced distribution of the population on both alternative choices while strong interaction leads to the emergence of a majority in one alternative. A positive (negative) bias α_0 generates asymmetry as it introduces a preference for one of both alternatives.

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 :

$$\begin{aligned} \frac{d}{dt}\bar{x} &= v(1 - \bar{x})e^{\alpha_0 + \alpha_1\bar{x}} - v(1 + \bar{x})e^{-\alpha_0 - \alpha_1\bar{x}} \\ &= 2v[\tanh(\alpha_0 + \alpha_1\bar{x}) - \bar{x}] \cosh(\alpha_0 + \alpha_1\bar{x}). \end{aligned} \quad (6)$$

(6) 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 (6).

However, it is worthwhile to emphasize that in the case of a finite population eq. (6) is only an approximation to the mean value dynamics, i.e. it gives the most likely path that the average opinion takes from some initial condition. It is neither exact nor is it necessarily close to any particular realization of the process. Although the mean value equation becomes exact in the limit of an infinite population, with a finite pool of agents the stochastic elements of the dynamics would lead to non-negligible fluctuations in the composition of the opinion index. For example, in the case of multiple equilibria, switches between both modes of the distribution might occur due to the inherent stochastic fluctuations of the opinion dynamics while the mean value equation could only predict convergence towards the nearest equilibrium from any set of initial conditions.

3 Estimation: The Basic Framework

While the stochastic properties of population processes like the one depicted in sec. 2 have been studied in great detail (Weidlich and Haag, 1983; Aoki, 1996; Weidlich, 2000), this literature has not developed a systematic approach towards estimation of such models. In the following I will outline, how such models can be estimated via a fairly conventional maximum likelihood procedure.

The basic ingredient in our estimation procedure is the so-called Fokker-Planck equation for the time development of the transitional density of macroscopic observables of the process. The Fokker-Planck equation associated to a stochastic process is a parabolic partial differential equation that occupies a very prominent place in statistical physics (Risken, 1989; Frank, 2005). However, it seems that due to the different research perspectives in this discipline, 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 seems straight forward: if one 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 compute the likelihood of each observation conditional on the realization of the process in the previous period. 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 still resort to numerical approximations of the Fokker-Planck equation. Numerical integration of partial differential equations via finite difference or finite element methods is also a well developed field (Thomas, 1995) and has found important applications both in statistical physics 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 partial differential equation:

$$\frac{\partial f(x)}{\partial t} = \frac{\partial}{\partial x}(\mu(x, \theta)f(x)) + \frac{\partial^2}{\partial x^2}(g(x, \theta)f(x)). \quad (7)$$

If (7) refers to a Fokker-Planck equation, the unknown function $f(x, t)$ is the transitory density of x , and $\mu(x, \theta) = -A(x, \theta)$, $g(x, \theta) = \frac{1}{2}D(x, \theta)$ with $A(x, \theta)$ and $D(x, \theta)$ are the drift and diffusion functions of the process, and θ is 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. (7). Various methods for discretisation of the stochastic equation (7) can be used. Applying a finite difference approach, the first and second derivatives on both sides of eq. (7) could be approximated either via forward differences or 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.

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

In a forward discretization, (7) would have to be replaced by

$$\frac{f_j^{i+1} - f_j^i}{k} = \frac{\mu_{j+1}f_{j+1}^i - \mu_j f_j^i}{h} + \frac{g_{j+1}f_{j+1}^i - 2g_j f_j^i + g_{j-1}f_{j-1}^i}{h^2} \quad (8)$$

with $f_j^i := f(x_0 + j \cdot h, ik)$ and $\mu_j := \mu(x_0 + j \cdot h, \theta)$, $g_j := g(x_0 + j \cdot h, \theta)$. This forward approximation is also known as the *explicit* finite difference approximation as it provides a closed-form solution for the mesh points at time $i+1$. 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 $(i + \frac{1}{2})k$ 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 present application boundary conditions should prevent a leakage of probability mass to points outside the support of the transient density. The very natural condition to conserve mass within the support is, therefore:

$$f_{-\frac{1}{2}}^j = f(x_0 - \frac{1}{2}h, jk) = 0 \quad \text{and} \quad f_{N_x + \frac{1}{2}}^j = f(x_0 + (N_x + \frac{1}{2})h, jk) = 0. \quad (9)$$

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. In our setting, the no-flux boundary conditions guarantee conservation of probability mass within the underlying x -interval if (7) governs the dynamics of a transient density (i.e. if (7) is a Fokker-Planck equation).

The drift and diffusion term of the Fokker-Planck equation for our process are given by²:

$$A(x) = \frac{n_-}{2N}w_{\uparrow}(x) - \frac{n_+}{2N}w_{\downarrow}(x) = v(1-x)e^{\alpha_0+\alpha_1x} - v(1+x)e^{-\alpha_0-\alpha_1x} \quad (10)$$

which, of course, coincides with the right-hand side of (6), while the diffusion term is:

$$D(x) = \frac{1}{N}\left(\frac{n_-}{2N}w_{\uparrow}(x) + \frac{n_+}{2N}w_{\downarrow}(x)\right) = \frac{1}{N}(v(1-x)e^{\alpha_0+\alpha_1x} + v(1+x)e^{-\alpha_0-\alpha_1x}). \quad (11)$$

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. 2 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, $T = 3$ 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 with density $\Phi^N(x_0 + A(x)k, D(x)k)$ evaluated at grid points $-1 + jh; j = 0, 1, \dots, N_x$, 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.

FIG. 2 ABOUT HERE

On the base of the Crank-Nicolson (or any other finite difference approximation), we can estimate the parameters of a diffusion process with discretely

²The Fokker-Planck equation is obtained as a second-order approximation to the complete characterization of the probability flux over all states (the so-called Master equation, cf. Weidlich and Haag, 1983; Lux, 1995)

spaced observations via approximate maximum likelihood: The negative log-likelihood of a sample of observations X_0, \dots, X_T is

$$-\log f_0(X_0 | \theta) - \sum_{s=0}^{T-1} \log f(X_{s+1} | X_s, \theta) \quad (12)$$

where $f_0(X_0 | \theta)$ 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 $f(X_{s+1} | X_s, \theta)$ 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. In his Theorem 3, he shows that the grid size has to behave like $k(T) = T^{-\delta}$ 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. Appendix A provides an illustration of the second-order accuracy of our discretisation via some worked-out numerical examples.

4 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. 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, 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 order of magnitude of this sample size is also in line with the number of available monthly observations of the ZEW index in our sample (which is 176). The time scaling parameter v has been fixed in order to have a certain number of switches 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 $k = \frac{1}{8}$ ($k = \frac{1}{16}$) 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 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 1 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 equivalence 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. In Appendix A we also provide evidence for the alleged second-order accuracy of the Crank-Nicolson approximations which underscores its suitability for ML estimation.

TABLE 1 ABOUT HERE

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 through IV, with $N = 25$, $N = 50$ or $N = 175$ in both cases. The results are exhibited in Table B1 in the Appendix. Given the small sample size, the behavior of the estimates seems also quite satisfactory. We comment on a few particular observations in the Appendix.

5 Empirical Application: Interaction Effects in a Business Climate Index

Since we have focused on 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 of *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). A broader range of confidence indices is compiled by the European Commission for the member states of the European union (European Commission, 2007). Many of these indices are close to the simple structure of our ‘canonical’ model in that they very literally ask for whether respondents are optimistic (“+”) or pessimistic (“-”) concerning the prospects of their economy. The only difference to our above model is that these indices mostly also allow for a neutral assessment. To accommodate this additional possibility 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 these data³. Here we focus on the ZEW index

³As detailed in Weidlich and Haag (1983, c. 6) the above framework could easily be extended by allowing for a neutral valuation and various degrees of positive or negative sen-

as one particularly interesting example. What makes it particularly suitable for our purpose is that in contrast to many other sentiment indices it represents the average of binary resp. tertiary responses in a very direct way, i.e. without any further aggregation involved, and that it has a rather constant number of participants (about 350 respondents) while other indices exhibit more fluctuations in their number of respondents over time. The group of respondents is furthermore more homogeneous than in most other surveys as it consists mainly of leading professionals from the finance industry. This selection of respondents implies on the one hand, that there should be more communication within this group (directly and indirectly via targeted media) than in a more anonymous sample selected via randomized nation-wide telephone interviews. On the other hand, one could hypothesize that financial experts should be less prone to interaction effects which lends further interest to our results.

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 in Sec. 2. In contrast, the indices for the U.S. economy are computed as weighted averages over categorical answers to different questions while the second important index of the German business climate, Ifo, starts with sector-specific surveys and aggregates them to an overall business climate indicator. The available monthly record of the ZEW sentiment index (starting in December 1991 and running through July, 2006) had already been displayed in Fig. 1 above. 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. In the introduction, we had already compared 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. 1 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 appearance 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 2. 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

timents by slight changes of individuals' transition rates. Adopting the formalization of transition rates proposed by Weidlich and Haag, the macroscopic dynamics of the index would indeed remain unchanged.

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. Figs. 3 and 5 below). This is due to the fact that, in our framework, 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, not act like independent agents (independent in performing their movements, not independent in the sense that their movements between “+” and “-” were not influenced by other 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. Table 2 shows 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. The ‘effective’ number of agents in our estimation is only about 40 ($2N$) compared to the much higher official sample size of about 350. As concerns the other 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.

TABLE 2 ABOUT HERE

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:

$$U_t = \alpha_0 + \alpha_1 x_t + \alpha_2 y_t. \quad (13)$$

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. Various such macro feedbacks have been investigated. As typical macroeconomic data at monthly frequency we tried interest rates, industrial production and changes of unemployment rates. In our model 3 we report

the influence of industrial production (deseasonalized and HP filtered, as displayed in Fig. 1). Note that the direction of the feedback is not predetermined in our model, i.e. α_2 could turn out positive or negative. The outcome of the exercise shows that industrial production adds some explanatory power: we obtain a significantly negative coefficient together with lower values of the AIC and BIC criteria. For interest rates, in contrast (results are available upon request), the estimated coefficients α_2 are not significant and overall improvements compared to model 2 are smaller. Quite the same holds for various measures of unemployment (with the change over the past 12 months entering as regressor because of the non-stationarity of the raw data): parameter estimates oscillate between significant and insignificant depending on which measure is used, the AIC and BIC values are between those of models 2 and 3 and the parameter estimates of the interaction components are hardly affected. Remarkably, the coefficient for the influence of changes of unemployment is positive in all cases. Combining two or three macroeconomic factors leads to very modest improvements ($\log L \approx 649$). Mostly, at most the coefficient for IP remains significant, while again the parameters for the interaction components are barely affected. However, even for the model 3, the improvement compared to model 2 is much smaller than the increase in $\log L$, AIC and BIC achieved by adding N as a free parameter (the step from model 1 to model 2). What is perhaps puzzling is the negative sign of the feedback effect from industrial production (similarly we obtained counterintuitive positive coefficients for unemployment and somewhat more plausible negative ones for interest rates) which is in contrast to a positive contemporaneous correlation of about 0.28 between both series. It appears to depict some type of ‘contrarian’ behavior: if the economic data is indicating a boom phase, our respondents already appear to forestall the overheating of the economy and the subsequent downturn and *vice versa*.⁴ In the estimation exercise reported in Table 2, the realization of industrial production is that of the previous period (which, in fact, in the case of IP is not known at this time to survey participants since the first statistical estimates are only released somewhat later). We have also experimented with various leads and lags without much change of the results.

Models 4 and 5 in Table 2 depict another extension of our baseline model: here we include a kind of ‘momentum’ effect in the opinion dynamics. Eq. (13) is now modified to include the change of the climate index from the month $t - 1$ to the last observation:

$$U_t = \alpha_0 + \alpha_1 x_t + \alpha_2 y_t + \alpha_3 (x_t - x_{t-1}) \quad (14)$$

One may interpret this as respondents reacting not only to the net influence of their environment but being particularly sensitive to changes of the business

⁴At the time of writing this version (June 2007) the ZEW index declined ‘unexpectedly’ despite the German economy experiencing a long-term maximum of its growth rate.

climate themselves.⁵ While one could argue that they might try to extract information on trends, it would certainly be hard to come up with a fully rational explanation for why the change of the index should have an influence on its subsequent development. Note also that a priori both a positive as well as negative feedback (if any) could be imagined. In fact, the negative coefficient on industrial production might suggest a similar contrarian element for the perceived momentum. As it turns out (cf. Table 2), the momentum effect is significantly positive. It again leads to a remarkable improvement of the model, but does not affect previously estimated parameters by too much. Adding industrial production as an explanatory variable (model 5) again leads to a further increase of the likelihood which is, however, again much more modest compared to the gain obtained from model 4. Smaller gains would result from alternative macroeconomic factors. In summary it, therefore, appears that macroeconomic variables add only a very slight fraction of the explanatory power, while the major improvements are obtained via refinements of our social opinion formation process.⁶

6 Specification Tests

How closely do time series from the estimated models mimic the empirical behavior of the ZEW index? Fig. 3 exhibits three simulations over the same time horizon ($T = 176$ integer periods) of model 5 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) as well as the ‘momentum’ of the index itself 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 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. Model 2 to 4 are not too different in their appearance. 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 $N \approx 20$. Fig. 4 shows the mean and 95 percent confidence bounds from the transient density

⁵We have to be a bit careful about the interpretation of the momentum term in our stochastic process: in order to guarantee that the process has Markov properties, we assume that $\Delta x_t = x_t - x_{t-1}$ only becomes public knowledge at the time when the new survey result is available (at time t). Respondents are, therefore, assumed to not update this variable between surveys (i.e. they only become aware of the current ‘momentum’ at the time when the survey is released). In this way, we can use it as an independent variable in the transition rates without having to modify the structure of the Fokker-Planck equation. If, in contrast, agents would update Δx_t between integer time steps, we would have to deal with a continuous time dynamics with delays for which even finite difference approximations would become quite cumbersome.

⁶This holds at all stages of our estimation exercise: if we add industrial production as an explanatory variable in model 1 (with fixed $N=175$), the likelihood only increases to -722.9 with virtually unchanging parameters for the social dynamics.

computed for model 3 over the whole observation period given the first observation of the index as the initial condition and incorporating the feedback from 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 that simulations of models 2, 4, and 5 would lead to very similar patterns. However, for models 2 and 4 the sample paths would not be synchronous to the empirical series simply because there is no exogenous factor.⁷ As can be seen from Fig. 5, the 95 percent confidence interval from model 1 excludes the better part of the empirical record, so that this baseline model could be clearly rejected as a potential data-generating process. For model 5, we could not perform the same exercise since the discrete momentum effect is hard to capture in the Fokker-Planck equation. We can, however, resort to numerical simulations in this case which gave a 95 percent confidence interval (from 1000 repetitions) that improves slightly on the analytical results for model 3 in Fig. 4 (not shown here because it is almost undistinguishable from Fig. 3). Overall, our models 3 and 5, in fact, show how the fuzzy exogenous information in the lower panel of Fig. 1 could be translated into a much clearer image of the business cycle dynamics in the view of the respondents' sentiments (upper panel of Fig. 1) via the self-referential and self-reinforcing dynamics of the opinion formation process.

FIGURE 3 ABOUT HERE

FIGURE 4 ABOUT HERE

Since the estimated interaction parameter, α_1 , in models 2 through 4 is marginally below the bifurcation value of unity, the ups and downs of the sentiment index during the observation period would likely reflect shifts of unique equilibria that alternate between optimistic and pessimistic majorities. Note, however, that a standard, say 95 confidence interval for α_1 would not exclude the possibility $\alpha_1 > 1$ so that we could as well have an underlying bimodal process with switches between both modes triggered by exogenous forces together with the inherent volatility of the opinion dynamics.

As another specification test⁸ 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 on the realization in the previous

⁷While this synchronous behavior appears quite striking in simulated time series, the statistical improvement by models 3 and 5 compared to models 2 and 4 in terms of the 'distance' criteria in Table 3 is relatively modest.

⁸For a more direct test we could use the test for uniform residuals of out-of-sample density forecasts (Pedersen, 1994, Diebold et al., 1998). Using as in-sample data for parameter estimation the observations until the end of 2000 this relatively weak test does not reject models 2 and higher on the base of the remaining out-of-sample density forecasts.

period. Fig. 6 shows the 95 % confidence bounds for the subsequent period's realization from model 5 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.

FIGURE 5 ABOUT HERE

FIGURE 6 ABOUT HERE

Table 3 provides a statistical analysis of 1000 Monte Carlo replications of models 1 through 5 on the base of the estimated parameters displayed in Table 2. 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. The table shows the means and simulated 95 percent boundaries for the first four unconditional moments together with the relative deviation (the squared value of the mean divided by the variance) 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 moments 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 relatively 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 3 and 5 seems to have contributed to a better fit compared to models 2 and 4. Again, this provides a confirmation of our visual impression reported above.

Table 4 reports autocorrelations of the index 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. Interestingly, models 2 and 3 are only able to match about the first four lags while the autocorrelations remain much higher than the empirical ones for the longer lags. Inclusion of the 'momentum' effect leads to a better fit of the entire range of autocorrelations between 1 and 10 lags and also achieves a close agreement in the estimate of the parameter of fractional differentiation as given in the last row of Table 4. This statistics is the parameter for hypothesized hyperbolic decay of the autocovariances, $E[x_t x_{t-\tau}] \sim \tau^{2d-1}$ and it is estimated via the method proposed by Geweke and Porter-Hudak (1983). The motivation for inclusion of this statistics comes from the finding that various

survey data in the political arena are characterized by long-term dependence in the sense of hyperbolic decay of their autocovariances and autocorrelation functions (Box-Steffensmeier and Smith, 1998).⁹

TABLE 3 ABOUT HERE

TABLE 4 ABOUT HERE

7 Conclusion

Given the immense public attention devoted to survey measures of business climate or economic sentiment, there has been surprisingly little work trying to model these data. Of course, under a rational expectations perspective, the most interesting aspect would be to test unbiasedness of such survey expectations and to find out whether they have predictive power beyond that of other macroeconomic data. However, not all economists do firmly believe in the ubiquitous validity of the rational expectation hypothesis. If we go to the other extreme, business climate surveys might rather reflect Keynes' notorious animal spirits at work. In this paper we have adopted the latter viewpoint. However, rather than taking the state of prevailing animal spirits as given, we have proposed a positive model to explain the fluctuations in respondents' confidence in the economic development. As it turned out, this model appears to have significant explanatory power for the ups and downs of the consumer climate: the model's parameters for the conjectured social interaction are strongly significant and apparently this social component of the opinion dynamics is much more important for the goodness-of-fit of various variants of our model than added macroeconomic variables. In the absence of alternative explanatory models, we conducted a series of specification tests that on the whole suggest that the empirical record could have been envisaged as a particular sample path from our model. Alternative 'rational' explanations of the development of the business climate would have to show that the pronounced swings could be explained by the release of important bits of information within the pertinent time intervals.¹⁰ This is a problem similar to the identification of important news at the time of large changes of financial prices (cf. Cutler et al., 1989) and a casual search for such explanations did not reveal any plausible candidates for such information shocks. While we cannot exclude such explanations, we leave the burden of the proof to proponents of rational expectations and reiterate that the social contagion of animal spirits apparently provides us with a framework that explains the data well without having to rely on homogeneous information shocks.

⁹Alfarano and Lux (2007) show that models with multi-modal distributions might lead to time series with apparent long memory.

¹⁰One could certainly argue that the supposed interaction effects can be explained away by correlated information. In the absence of individual data, we cannot discriminate between unobserved endogenous factors and interaction effects (as shown by Lee, 2007, such an identification would require a sample with multiple groups of different sizes). In any case, an alternative explanation based on unobserved common shocks would be empirically void.

There are many directions into which research would fruitfully proceed from here: first, one should obviously study similar data sets from other countries to see whether interaction patterns are similar or not. We have already started such a comparative project and found quite similar results to those reported above in quite a number of cases. Second, if business cycles are, in fact, generated (at least partially) by animal spirits, the business climate measures would interact with objective economic quantities like industrial production.¹¹ It would, therefore, be worthwhile to include the opinion dynamics into a multi-variate setting of both objective measures of economic activity and more subjective survey indices. While conceptually not too difficult to imagine, such a framework would be computationally extremely demanding and would require the development of more efficient numerical algorithms. Third, one would also like to identify animal spirits in cases where no survey data exists. This would present one with the challenge of developing indirect methods of inference to identify hidden psychological states.

¹¹Note that causality tests cannot distinguish between rational expectations and animal spirits: while positive causation between contemporaneous survey expectations and future realizations of economic activity are routinely interpreted as evidence for rational expectations (if unbiased), they could as well be the imprint of a true causality running from survey expectations to subsequent economic development.

References

- Alfarano, S. and T. Lux, 2007, A noise trader model as a generator of apparent financial power laws and long memory, *Macroeconomic Dynamics*, **11** (Supplement 1), 80–101.
- Aoki, M., 1996, *New Approaches to Macroeconomic Modeling: Evolutionary Stochastic Dynamics, Multiple Equilibria, and Externalities as Field Effects*, Cambridge, University Press.
- Benabou, R., 1993, Workings of a city: Location, education and production, *Quarterly Journal of Economics* **CVIII**, 619–652.
- Blume, L. and S. Durlauf, 2003, Equilibrium concepts for social interaction models, *International Game Theory Review*, **5**, 193–209.
- Box-Steffensmeier, J. and R. Smith, 1998, Investigating political dynamics using fractional integration methods *American Journal of Political Science*, **42**, 661–689.
- Brock, W. and S. Durlauf, 2001a, *Interaction-Based Models*, Chap. 54 in: J. Heckman and E. Leamer, eds., *Handbook of Econometrics* vol. 5. Amsterdam: Elsevier.
- Brock, W. and S. Durlauf, 2001b, Discrete choice with social interactions, *Review of Economic Studies*, **68**, 235–260.
- Carrol, C., 2003, Macroeconomic expectations of households and professional forecasters, *Quarterly Journal of Economics* **CXVIII**, 209–248.
- Chen, P., 2002, 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**, 307–326.
- Cutler, D., J. Poterba and L. Summers, 1989, What moves stock prices?, *Journal of Portfolio Management* **4**, 4–12.
- Diebold, F., T. Gunther and A. Tay, 1998, Evaluating density forecasts with applications to financial risk management, *International Economic Review*, **39**, 863–883.
- European Commission, 2007, The Joint Harmonised EU Programme of Business and Consumer Surveys User Guide, http://ec.europa.eu/economy_finance/indicators/business_consumer_surveys/userguide_en.pdf.
- Föllmer, H., 1974, Random economies with many interacting agents, *Journal of Mathematical Economics*, **1**, 51–62.

- Frank, T., 2005, *Nonlinear Fokker-Planck Equations*, Berlin, Springer.
- Gelper, S., A. Lemmens and C. Croux, 2007, Consumer sentiment and consumer spending: decomposing the Granger Causal Relationship in the time domain, *Applied Economics*, **39**, 1–11.
- Geweke, J. and S. Porter-Hudak, 1983, The estimation and application of long-memory time series models, *Journal of Time Series Analysis*, **4**, 221–238.
- Glaeser, E., B. Sacerdote and J. Scheinkman, 1996, Crime and social interactions, *Quarterly Journal of Economics* **CXI**, 507–548.
- Jensen, B. and R. Poulsen, 2002, Transition densities of diffusion processes: Numerical comparison of approximation techniques, *Journal of Derivatives*, 18–34.
- Jones, A., 1994, Health, addiction, social interaction, and the decision to quit smoking, *Journal of Health Economic*, **13**, 93–110.
- Horst, U. and J. Scheinkman, 2006, Equilibria in systems with social interactions, *Journal of Economic Theory*, **130**, 44–77.
- Hüfner F. and M. Schröder, 2002, Prognosegehalt von ifo-Geschäftserwartungen und ZEW-Konjunktuerwartungen: Ein ökonometrischer Vergleich, *Jahrbücher für Nationalökonomie und Statistik*, **222**, 316–336.
- Hurn, A., J. Jeisman and K. Lindsay, 2006, *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.
- Iounnides, Y., 2006, Topologies of social interactions, *Economic Theory* **28(3)**, 559–584.
- Kirman, A., 1993, Ants, rationality, and recruitment, *Quarterly Journal of Economics* **CVIII**, 137–156.
- Kraft, M.,T. Landes and P. Weise, 1986, Dynamic aspects of a stochastic business cycle model, *Methods of Operations Research*, **53**, 445–453.
- Krauth, B., 2006, Social interactions on small groups, *Canadian Journal of Economics*, **39**, 414–433.
- Lee, L.-F., 2007, Identification and estimation of econometric models with group interactions, contextual factors and fixed effects, *Journal of Econometrics*, **140**, 333–374.
- Ludvigson, S., 2004, Consumer confidence and consumer spending, *Journal of Economic Perspectives*, **18**, 29–50.

- Lux, T., 1995, Herd behaviour, bubbles and crashes, *Economic Journal*, **105**, 881–896.
- Lux, T., 1997, Time variation of second moments from a noise trader/infection model, *Journal of Economic Dynamics and Control*, **22**, 1–38.
- Pedersen, A., 1994, *Uniform Residuals for Discretely Observed Diffusion Processes*, Department of theoretical Statistics, University of Aarhus.
- Poulsen, R., 1999, *Approximate Maximum Likelihood Estimation of Discretely Observed Diffusion Processes*, Manuscript, University of Aarhus.
- Risken, H., 1989, *The Fokker-Planck Equation: Methods of Solutions and Applications*, Berlin, Springer.
- Schelling, T., 1971, Dynamic models of segregation, *Journal of Mathematical Sociology*, **1**, 143–186.
- Seydel, R., 2002, *Tools for Computational Finance*. Berlin, Springer.
- Shiller, R., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review*, **71**, 421–436.
- Souleles, N., 2004, Expectations, heterogeneous forecast errors, and consumption: micro evidence from the michigan consumer sentiment surveys, *Journal of Money, Credit and Banking*, **36**, 39–72.
- Taylor, K and R. McNabb, 2007, Business cycles and the role of confidence: Evidence for Europe, *Oxford Bulletin of Economics and Statistics*, **69**, 185–208.
- Thomas, J., 1995, *Numerical Partial Differential Equations*, Berlin, Springer.
- Weidlich, W., 2000, *Sociodynamics: A Systematic Approach to Mathematical Modelling in the Social Sciences*. Amsterdam, Harwood Acad. Publ.
- Weidlich, W. and G. Haag, 1983, *Concepts and Methods of a Quantitative Sociology*, Berlin, Springer.

A Appendix A - Accuracy of Finite Difference Scheme

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:

$$v_1 = f - hc - kd - h^2l - k^2m + \dots \quad (A1)$$

$$v_2 = f - 0.5hc - kd - 0.25h^2l - k^2m + \dots \quad (A2)$$

$$v_3 = f - 0.25hc - kd - 0.0625h^2l - k^2m + \dots \quad (A3)$$

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

$$\frac{v_2 - v_1}{v_3 - v_2} \simeq 2 \frac{c + 1.5hl}{c + 0.75hl}. \quad (A4)$$

Hence, if the method is first-order accurate, $c \neq 0$ should be the dominating component and evaluating (A4) 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 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. 2. Table A1 exhibits results for selected grid points with (A4) 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.

$x t$	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 A1: Order determination for the Crank-Nicolson method applied to the interacting agent model. All parameter values and settings like in Fig. 2.

$x t$	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 A2: Order determination with a different initial value, $x_0 = 0.9$. All other parameters and settings as in Fig. 2 and Table 1.

Results become slightly worse if one considers more extreme starting points: Table A2 exhibits error ratios at selected grid points for the same model parameters and approximation scheme like in Table A1 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.

B Appendix B: Monte Carlo Runs with Endogenous N

Table B1 provides the results on our Monte Carlo runs with estimated parameters v, α_0, α_1 , and N . The basic message is that even for the small sample sizes of our study, the extended sets of parameters can be efficiently estimated. The average biases across the 200 replications are small in most cases except for a few outliers. One particular outlier is the case $\alpha_0 = 0, \alpha_1 = 0.8, N = 175$ for which N is strongly biased upwards. However, as the medians show this bias

might be due to some extreme realizations. Another interesting observation is that our algorithm has problems in disentangling the effects of a large bias and strong social interaction ($\alpha_0 = 0.2, \alpha_1 = 1.2$). This is perhaps not too surprising since with x fluctuating around a unique positive mode the factor $\alpha_1 x$ would exhibit only small fluctuations. Interestingly, however, this effect diminishes with increasing N .

		Crank-Nicolson ($k = 1/8$)				Crank-Nicolson ($k = 1/16$)			
		v	α_0	α_1	N	v	α_0	α_1	N
$a_0 = 0.000$	means	3.390	0.001	0.777	32.400	4.024	0.001	0.774	32.618
$a_1 = 0.800$	medians	3.268	0.000	0.807	26.011	3.401	0.000	0.812	26.726
$N = 25$	FSSE	2.471	0.007	0.172	19.703	2.565	0.007	0.180	20.063
	RMSE	2.634	0.007	0.173	21.000	2.756	0.007	0.181	21.414
$a_0 = 0.000$	means	4.766	0.002	0.736	80.047	4.840	0.002	0.713	80.058
$a_1 = 0.800$	medians	3.595	0.000	0.847	61.461	3.876	0.000	0.856	64.865
$N = 50$	FSSE	3.763	0.007	0.277	62.459	3.950	0.007	0.303	64.302
	RMSE	4.149	0.007	0.284	69.170	4.349	0.008	0.314	70.834
$a_0 = 0.000$	means	6.589	0.001	0.673	392.688	8.252	0.002	0.673	479.833
$a_1 = 0.800$	medians	3.385	0.000	0.824	191.185	4.362	0.000	0.863	243.992
$N = 175$	FSSE	7.136	0.005	0.359	433.211	10.354	0.009	0.475	588.831
	RMSE	7.972	0.005	0.380	483.861	11.587	0.009	0.490	661.749
$a_0 = 0.200$	means	1.696	0.220	0.759	28.370	27.620	0.221	0.758	28.118
$a_1 = 0.800$	medians	3.718	0.191	0.813	26.637	4.215	0.190	0.813	26.711
$N = 25$	FSSE	20.696	0.123	0.229	9.451	38.345	0.120	0.222	8.707
	RMSE	23.250	0.125	0.232	10.012	45.488	0.122	0.225	9.228
$a_0 = 0.200$	means	6.965	0.249	0.711	52.319	7.048	0.259	0.694	50.922
$a_1 = 0.800$	medians	3.216	0.208	0.790	50.401	3.410	0.208	0.789	48.795
$N = 50$	FSSE	16.054	0.158	0.281	19.450	14.627	0.168	0.300	18.396
	RMSE	16.497	0.165	0.294	19.539	15.141	0.178	0.318	18.379
$a_0 = 0.200$	means	4.035	0.279	0.662	185.145	4.985	0.294	0.637	184.206
$a_1 = 0.800$	medians	3.196	0.209	0.784	175.450	3.334	0.219	0.770	164.994
$N = 175$	FSSE	2.992	0.229	0.399	86.822	5.742	0.240	0.418	92.877
	RMSE	3.159	0.242	0.422	87.197	6.062	0.258	0.448	93.101
$a_0 = 0.000$	means	3.161	0.000	1.197	26.372	3.110	0.000	1.196	26.050
$a_1 = 1.200$	medians	3.001	0.000	1.197	26.168	3.065	0.000	1.197	26.021
$N = 25$	FSSE	0.746	0.009	0.015	3.975	0.598	0.009	0.015	3.563
	RMSE	0.762	0.009	0.015	4.196	0.606	0.009	0.015	3.705
$a_0 = 0.000$	means	3.097	0.004	1.187	51.587	3.099	0.004	1.187	51.468
$a_1 = 1.200$	medians	3.053	0.000	1.197	51.526	3.072	0.000	1.196	51.356
$N = 50$	FSSE	0.761	0.031	0.050	6.593	0.686	0.031	0.050	6.516
	RMSE	0.765	0.032	0.052	6.766	0.691	0.031	0.052	6.664
$a_0 = 0.000$	means	3.825	0.026	1.127	176.176	3.781	0.055	1.138	175.935
$a_1 = 1.200$	medians	2.967	0.011	1.172	159.542	2.999	0.011	1.174	158.175
$N = 175$	FSSE	7.201	0.356	0.513	242.764	3.902	0.723	0.666	242.663
	RMSE	7.230	0.357	0.517	242.159	3.970	0.724	0.667	242.057
$a_0 = 0.200$	means	3.410	0.669	0.619	19.130	4.556	0.656	0.635	19.296
$a_1 = 1.200$	medians	1.501	0.596	0.715	18.465	1.705	0.588	0.736	18.976
$N = 25$	FSSE	9.547	0.390	0.488	5.205	14.034	0.391	0.489	5.183
	RMSE	9.531	0.609	0.758	7.837	14.086	0.600	0.747	7.699
$a_0 = 0.200$	means	22.016	0.256	1.134	54.254	39.919	0.342	1.042	52.693
$a_1 = 1.200$	medians	10.118	0.173	1.244	54.290	20.864	0.210	1.186	53.523
$N = 50$	FSSE	30.421	0.338	0.414	12.290	40.372	0.367	0.449	12.193
	RMSE	35.811	0.341	0.419	12.976	54.633	0.393	0.481	12.457
$a_0 = 0.200$	means	5.686	0.318	1.059	171.394	5.994	0.341	1.031	170.926
$a_1 = 1.200$	medians	2.834	0.195	1.209	174.389	3.027	0.280	1.102	167.431
$N = 175$	FSSE	15.959	0.520	0.628	48.890	15.518	0.544	0.658	53.062
	RMSE	16.144	0.532	0.642	48.901	15.766	0.561	0.677	53.086

Table B1: Note: The table reports the statistics of 200 Monte Carlo runs of each parameter set with a sample size of $T = 200$.

	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 1: 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 (RMSE).

	ν	α_0	α_1	α_2	α_3	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 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 4 (momentum effect)	0.14 (0.05)	0.10 (0.06)	0.91 (0.14)		2.11 (0.76)	27.24 (9.63)	-627.5	1265.1	1265.4
Model 5 (momentum + IP)	0.12 (0.05)	0.11 (0.06)	0.86 (0.16)	-2.82 (1.65)	2.23 (0.81)	25.12 (8.95)	624.9	1261.9	1260.1

Table 2: Parameter Estimates for Stochastic Models of Interacting Agents.
Note: Details on the underlying models appear in the main text. The numbers in brackets are standard errors of parameter estimates.

	data	models				
		1	2	3	4	5
mean: (95%)	0.352 (-0.588, -0.631, -0.405)	0.349 (0.045, 0.552)	0.324 (0.035, 0.528)	0.391 (0.214, 0.521)	0.343 (0.173, 0.479)	
std. dev: (95%)	0.370 (0.068, 0.185)	0.355 (0.251, 0.484)	0.384 (0.293, 0.506)	0.314 (0.245, 0.391)	0.356 (0.284, 0.424)	
skewness: (95%)	-0.620 (0.029, 1.685)	-0.908 (-1.880, 0.097)	-0.991 (-1.878, 0.053)	-0.947 (-1.620, 0.247)	-0.857 (-1.571, -0.181)	
kurtosis: (95%)	-0.428 (-0.652, 3.717)	0.591 (-1.322, 4.347)	0.540 (-1.283, 3.296)	0.804 (-0.911, 3.585)	0.207 (-1.000, 2.315)	
rel. deviation: (95%)	0.905 (8.706, 82.360)	1.368 (0.024, 4.022)	1.030 (0.022, 2.792)	1.766 (0.342, 3.956)	1.049 (0.209, 2.439)	
distance: (95%)		0.952 (0.899, 0.986)	0.297 (0.220, 0.455)	0.233 (0.171, 0.318)	0.223 (0.165, 0.311)	

Table 3: Unconditional moments from 1.000 Monte Carlo simulations of models 1 through 5 (95 percent confidence intervals from the simulations are given in brackets)

ACF	data	models				
		1	2	3	4	5
1 (95 %)	0.935	0.630 (0.456 0.963)	0.923 (0.845 0.967)	0.939 (0.908 0.968)	0.904 (0.844 0.944)	0.930 (0.890 0.955)
2	0.830	0.404 (0.162 0.929)	0.853 (0.715 0.936)	0.880 (0.820 0.934)	0.796 (0.674 0.879)	0.848 (0.769 0.901)
3	0.709	0.266 (0.013 0.890)	0.789 (0.595 0.907)	0.819 (0.732 0.900)	0.691 (0.523 0.811)	0.762 (0.653 0.845)
4	0.584	0.175 (-0.080 0.857)	0.729 (0.496 0.883)	0.758 (0.652 0.866)	0.592 (0.393 0.751)	0.675 (0.541 0.784)
5	0.465	0.116 (-0.133 0.820)	0.673 (0.398 0.860)	0.696 (0.566 0.833)	0.499 (0.266 0.699)	0.589 (0.432 0.723)
6	0.363	0.075 (-0.171 0.784)	0.620 (0.319 0.840)	0.633 (0.478 0.797)	0.419 (0.169 0.638)	0.508 (0.335 0.662)
7	0.272	0.048 (-0.188 0.747)	0.571 (0.241 0.813)	0.571 (0.392 0.759)	0.355 (0.092 0.594)	0.434 (0.250 0.616)
8	0.186	0.032 (-0.197 0.703)	0.525 (0.184 0.793)	0.512 (0.317 0.722)	0.302 (0.049 0.553)	0.366 (0.171 0.565)
9	0.094	0.022 (-0.213 0.668)	0.482 (0.121 0.774)	0.454 (0.239 0.678)	0.251 (-0.013 0.516)	0.298 (0.088 0.514)
10	0.017	0.014 (-0.220 0.631)	0.442 (0.078 0.752)	0.398 (0.167 0.640)	0.196 (-0.084 0.468)	0.228 (0.002 0.463)
d	0.553	0.194 (-0.343 0.978)	0.826 (0.338 1.261)	0.923 (0.455 1.346)	0.551 (0.027 0.992)	0.668 (0.202 1.051)

Table 4: Autocorrelations and estimated parameter of fractional differentiation d from 1.000 Monte Carlo simulations (95 percent confidence intervals from the simulations are given in brackets).




Figure 1: ZEW Sentiment Index and Industrial Production.

Finite Difference Approximation of Transitional Density




Figure 2: 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 3: Simulated trajectories from models 5 and 1 (lower right-hand panel). The broken lines show the empirical data




Figure 4: Mean and 95 percent confidence interval for Model 3 (from Fokker-Planck Equation conditional on initial observations and external macroeconomic information)




Figure 5: Mean and 95 percent confidence interval from Model 1 (from Fokker-Planck Equation conditional on initial observations and external macroeconomic information.)




Figure 6: 95 Percent interval for one-step iterations of transient density of Model 5.