Complex Systems **8** (1994) 75–89

Contrarians and Volatility Clustering

E. R. Grannan
Department of Physics, University of California,
Irvine, CA 92717, USA

G. H. Swindle
Department of Statistics and Applied Probability,
University of California, Santa Barbara, CA 93106, USA

Abstract. We introduce a new origin of volatility clustering in economic time series generated by systems of interacting adaptive agents. Each agent is assigned a random subset of a fixed collection of predictors. At every time step each agent generates an action based upon its assigned predictors. Some agents are contrarians, that is, they act at variance with the natural action suggested by a predictor. Agents that perform poorly are replaced. At each time step the signal value is generated solely by the cumulative actions of the agents on the current history of the time series. We observe numerically that under the dynamics induced by the removal of poor performers, even when contrarians are introduced at a very low density, the system evolves to a state in which contrarians comprise nearly half of the population. Furthermore, the time series generated by these systems exhibits volatility clustering. Elimination of either the contrarian behavior or the removal of poor players precludes volatility clustering.

1. Introduction

Many time series in economics and finance exhibit a phenomenon called clustered volatility, which refers to signals that display high volatility during some intervals of time and low volatility at other times. Volatility clustering has been observed in stock returns and related derivative securities [1, 2], interest rates [3], and foreign exchange rates [4]. A variety of time series models have been constructed and used in these and other works in an effort to characterize such signals statistically. However, progress in understanding the origins of clustered volatility has been limited. A thorough literature survey describing existing empirical evidence of clustered volatility and related statistical analyses is contained in [5].

In this paper we consider a class of adaptive systems which, through the actions of competing agents, self-organize to states in which time series with clustered volatility are generated. Like the systems in [6] discussed below, each new value of the signal is generated by the actions of players making decisions based on a set of prescribed predictors applied to the current history of the signal. However, the systems which we consider here have two additional features. First, contrarians—players which act "against" the natural action indicated by a predictor—are introduced. Also, players which perform poorly are replaced. This second feature provides for an evolutionary dynamics, which, in conjunction with contrarians, yields interesting behavior. In particular, we observe that even when contrarians are introduced at a low rate the systems equilibrate with densities of contrarians just under a half. Furthermore, the resulting signals exhibit regions of high and low volatility, that is, clustered volatility. To the best of our knowledge, this is a new mechanism for the origin of this phenomenon.

In addition to the purely academic pursuit of understanding the origin of clustered volatility, there are countless practical financial applications of models which generate realistic signals. Statisticians and econometricians frequently use time series models to capture relevant features of financial signals such as clustered volatility, and then apply the results to price derivative securities. However, such methods are intrinsically limited to the modeler's ability to select a tractable class of time series models that estimate particular attributes of the financial signal to a reasonable degree of accuracy. Adaptive systems such as those discussed below could potentially be calibrated to model financial signals directly.

The systems that we examine here are motivated by the work of W. B. Arthur. In [7], Arthur suggests the importance of the formation of beliefs of economic agents in an inductive procedure, and sketches a framework in which one could construct a population of "investors," each formulating its current hypothesis about the present state of a market by applying a set of predictors on the history of a financial signal. The signal itself is generated by the actions of the investors, which are boundedly rational. This train of thought is one of a variety of suggested departures from equilibrium theory and the associated assumption of rational expectations, in which investors or agents solve optimization problems given prescribed probability distributions on future events. For a survey of work related to bounded rationality in economic systems see [8].

An exceedingly simple system that captures much of the essence of [7], was introduced in [6]. Arthur describes the model in essentially the following way. A bar has 50 seats and a population of 100 people who consider patronizing the bar each night. Every individual is "happy" if he chooses to visit the bar on a night in which fewer than 50 people arrive (so that the bar is not too crowded), and is "unhappy" if 50 or more people attend. A pool of M (say $M = 20$) *ad hoc* predictors is constructed. For example, the predictors could be: (1) the same as last night; (2) the same as two nights ago; (3) a linear interpolation of the last three nights; and so on. Each individual is randomly assigned a subset of L ($1 \leq L \leq M$) predictors. The key feature of the system is that the attendance (signal) at time $n+1$ is generated solely by

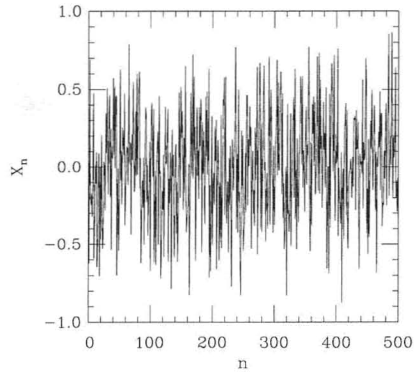

Figure 1: Sample signal X_n of 500 time steps for systems with $L = 5$ predictors held by each player out of a possible 40 after equilibration of the system without contrarians ($p_{\text{con}} = 0$) and without removal ($\theta = \infty$).

the responses of the players to information contained in the signal up to time n. Specifically, to generate a signal at a given time step, all predictors are evaluated over a window of, say, the last 100 time steps, and ranked according to the least square error in the predictions over this window. At each time step, every player uses the prediction generated by the best predictor (lowest error) *of the L that it has been assigned.* The decision of each player is to visit the bar if its best predictor predicts an attendance of fewer than 50 or conversely, to stay home if its best predictor anticipates attendance of 50 or more. Aggregate attendance is calculated, and the procedure is iterated.

This system clearly generates a stationary process. It is also clear that if $L = M$, then each player is in possession of all predictors, and the signal will bounce between full attendance and zero attendance since the actions of all players must coincide. Arthur's insight is that if $L < M$, then the inhomogeneity of beliefs can yield a more interesting time series. A realization of the signal generated by a process which is essentially identical to the system described above with $M = 40$ predictors and each player holding $L = 5$ is shown in Figure 1 (a signal value of zero in the figure corresponds to the attendance of half the players in the bar problem). Another observation by Arthur is that the system does not reach a trivial stationary state, insofar as the set of predictors being used at each time step does not reach a fixed point. This nontrivial aspect of the stationary state is a consequence of the fact that a good predictor is actively used by a large fraction of the population, which results in poor future performance.

The systems we consider here will add two important features to the above system: contrarians and an evolutionary dynamics. These two elements are essential to the appearance of volatility clustering.

The fact that the signal becomes trivial when every player possesses all predictors motivates the addition of contrarian strategies. Namely, if a given predictor has been the most successful at recent predictions (i.e., its error is smallest among all predictors), then in the original system the majority of the population is using this predictor to make decisions. The exact fraction depends upon L and M in a simple way. It would clearly be advantageous for a player to use this predictor as a contrarian, that is, to make a decision at variance with what the predictor suggests. This would result in visiting the bar during nights of low attendance, and staying home when it is crowded, at least while the system remains in the state where a large fraction of players are following the same predictor in the usual protrarian way.

The distinction between protrarians and contrarians is crucial to what follows. Allowing contrarian behavior is *not* equivalent to doubling the predictor set. Specifically, a contrarian using a predictor is not the same as a protrarian using the reflection of the predictor, because the reflected predictor will have a high error when the original predictor has a low error. When a player is currently using its best predictor as a protrarian, it will attend if this predictor suggests a low turnout; if the best predictor is being used as a contrarian, the player will attend if this predictor suggests a high turnout. The effect on the success of the predictor is manifestly different in the two cases. The actions of protrarians undermine the success of a predictor, while contrarians enhance the success. In these systems, if the best predictor happens to be used by nearly the same number of contrarians and protrarians, it can remain the best predictor for some time. Without contrarians, this is not possible.

The second feature added to these systems is an evolutionary dynamics: players that perform poorly (e.g., too often attend on busy nights and stay home on quiet nights) are replaced (see [9] for a discussion of the evolution of strategies in a model of stock traders). It is this attribute that allows for interesting dynamics in the number of protrarians and contrarians using given predictors. Specifically, the density of contrarian strategies converges to a value much higher than the density at which it was introduced. The signals generated by these systems are interesting in that clearly discernible regions of high and low volatility arise. Figure 2 contains a realization of such a time series for a system with $M = 40$ predictors, where each player holds $L = 25$ predictors. This value of L is much larger than that used later in this paper, and was selected here solely to make the volatility clustering clearly discernible.

In the next section, we describe the precise dynamics of the two systems. The third section describes simulation results. Here we show realizations of the processes as the number of predictors assigned to each player varies, and as the probability of generating contrarians is adjusted. We will also show how the density of contrarian strategies converges to limiting values which are relatively independent of the rate at which contrarians are generated. In addition, we discuss the distributions of the signal values and block variances. Most importantly, we show plots of the conditional variance of the next signal value given the previous value. These plots establish the appearance of clustered volatility when contrarians are present.

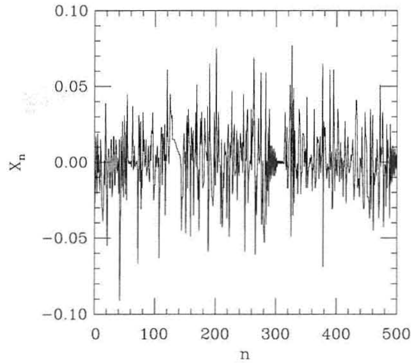

Figure 2: Sample signal X_n of 500 time steps for systems with $L = 25$ predictors held by each player out of a possible 40 after equilibration of the system with contrarians ($p_{\text{con}} = .5$) and with removal ($\theta = 100$).

2. The systems

The systems we consider generate signals with values in the interval $[-1, 1]$. There will be a total of M possible predictors Φ_j, $1 \leq j \leq M$, which are functions that map the history of the signal to the interval $[-1, 1]$:

$$\Phi_j : (\ldots, X_{n-1}, X_n) \rightarrow [-1, 1] \tag{1}$$

These functions generate the predictions at each time step. Typically the Φ_j's used in our simulations are generated randomly from some restricted class of functions (this results in a quenched disorder in the system). Each predictor Φ_j is also accompanied by an error denoted by e_j, the evolution of which will be described shortly.

The number of players in the game will be denoted by N. Each player is assigned L predictors. We will view each player as being defined by a vector of length L, $(i_1, \ldots, i_L)$, where the magnitude of the entries prescribes the predictors and the sign indicates protrarian/contrarian modes. So, for example, if $L = 3$ and a player is defined by the vector $(-3, 5, 8)$, then it has predictors 3, 5, and 8 at its disposal. It is a contrarian with respect to 3 and a protrarian with respect to 5 and 8. Initially, each player is generated by assigning L randomly selected predictors, and making each of the predictors contrarian with probability p_{con}. Here p_{con} is a parameter which denotes the probability that a newly assigned predictor will be used in contrarian mode. Note that p_{con} need not be, and in fact is typically very different from, the stationary density of predictors used by the players in contrarian mode, as will be seen later.

Based upon the actions of its predictors on the history $(\ldots, X_{n-1}, X_n)$ each player i will adopt a state $s_i(n + 1)$ that is either 1 or -1 (analogous to visiting the bar or staying home). This is done in the following way. Each player i selects the predictor with the lowest error from its available predictors $\Phi_1^i, \ldots, \Phi_L^i$. This predictor, denoted by $\hat{\Phi}_n^i$, generates a prediction for the next value of the signal: $\tilde{X}_{n+1}^i = \hat{\Phi}_n^i[(\ldots, X_{n-1}, X_n)]$. The player then acts in the following way:

(A) If player i is *protrarian* with respect to $\hat{\Phi}_n^i$, then $s_i(n+1)$ takes the *opposite* sign of the prediction $\tilde{X}_{n+1}^i$.

(B) If player i is *contrarian* with respect to $\hat{\Phi}_n^i$, then $s_i(n+1)$ takes the *same* sign of the prediction $\tilde{X}_{n+1}^i$.

Remark. Actions (A) and (B) are analogous to the concept of each player visiting the bar (protrarian mode) or staying home (contrarian mode) when its active predictor predicts a less crowded night.

The value of the signal at the next time $t = n + 1$ is the average

$$X_{n+1} = \frac{1}{N} \sum_{i=1}^{N} s_i(n+1). \tag{2}$$

Once this new signal value is generated, two things must be done:

(i) Players doing well during this time step are rewarded.

(ii) The updated predictor errors are calculated.

Regarding the first task (i), player i is considered to have "won" if the sign of $s_i(n+1)$ is different than X_{n+1}. This is analogous to visiting the bar during a time of low attendance or staying home during a busy time. An important feature of the system is that players that perform poorly are replaced. This can be done in a variety of ways without significant changes in the properties of the system. The method used here is to prescribe a threshold integer θ. As the system evolves, a counter for each individual is updated that denotes the difference between the number of "wins" and the number of "losses," bounded in magnitude by θ. When an individual's counter drops below $-\theta$ it is replaced with a randomly chosen set of L predictors, again with probability p_{con} that a predictor being used is in contrarian mode. The counter is then reset to zero.

To accomplish the second task (ii), the errors of the predictors are updated according to the rule

$$e_j(n+1) = \alpha e_j(n) + \beta_j(n+1) + z_j(n+1) \tag{3}$$

where $0 < \alpha < 1$ and

$$\beta_j(n+1) = \begin{cases} 0 & \text{if } \operatorname{sign}\{\Phi_j[(\ldots, X_{n-1}, X_n)]\} = \operatorname{sign}(X_{n+1}). \\ 1 & \text{otherwise} \end{cases} \tag{4}$$

Therefore, predictor Φ_j is hurt by predicting a value of the wrong sign. The parameter α specifies the memory of the predictors. The last term in (3), $z_j(n+1)$, is a sequence of independent random variables with a uniform distribution on $[0, \delta]$ which eliminates degeneracies arising from the simple form of β_j in (4). The width of the distributions δ can be arbitrarily small without affecting the results (we take $\delta = .01$ in our simulations). The Markovian form of the error update is particularly convenient both in simulation and analysis.

3. Results from simulations

For our simulations we restrict our attention to a class of *Linear Predictor Models* in which the predictors Φ_j are linear functions mapping a finite range R of the history of the signal to a real number corresponding to the projected value of the signal in the next time step:

$$\Phi_j[(\ldots, X_{n-1}, X_n)] = \sum_{k=1}^{R} \gamma_{jk} X_{n-k}. \tag{5}$$

The coefficients γ_{jk} are generated randomly. For example, in the simulations discussed below these coefficients are independently selected from a normal distribution with mean zero and standard deviation 0.4 as the system is initialized, and remain fixed for the duration of the simulation. This particular space of predictors contains many of the ad hoc predictors used by Arthur.

Unless otherwise stated, the following parameters are fixed:

(i) The number of players: $N = 1001$.

(ii) The number of predictors: $M = 40$.

(iii) The range of the predictors: $R = 4$.

(iv) The threshold for removal: $\theta = 100$.

Remark. The qualitative nature of the results does not seem to depend in any significant way on the precise values of these parameters. Furthermore, the phenomena described below are also observed in systems where the predictors are selected from different function spaces. For example, we have examined a class of systems with predictors selected from a class of nonlinear functions which only act on the last value of the signal ($R = 1$), and the qualitative behavior of these systems appears similar to that of the Linear Predictor Models discussed here.

Those parameters which have a significant effect on the behavior of the system are:

(i) p_{con}: the probability of a newly assigned predictor being used in contrarian mode.

(ii) L: the number of predictors assigned to each player.

The qualitative nature of the results is as follows. In the absence of contrarians (i.e., with $p_{\mathrm{con}} = 0$), the signal does not exhibit clustered volatility. In the presence of contrarians ($p_{\mathrm{con}} > 0$), the distribution of the number of predictors in contrarian mode increases from an initial density of p_{con} to a limiting value which is just under a half, and is rather independent of the value of p_{con}. The time series generated by systems with $p_{\mathrm{con}} > 0$ show clustered volatility (described in more detail below) to a degree which seems to increase with L.

Next we discuss the detailed results.

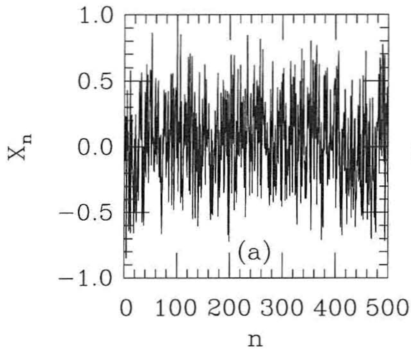

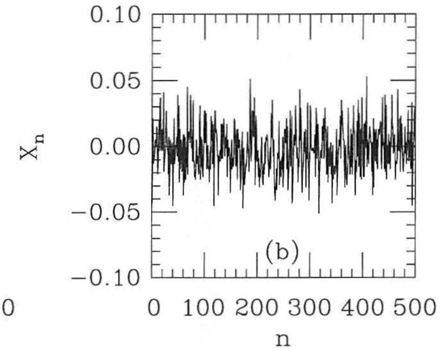

Figure 3: Sample signals X_n of 500 time steps for systems with $L = 5$ predictors held by each player out of a possible 40 after equilibration of the system. 3(a) is for the system with $p_{\text{con}} = 0$ (without contrarians); 3(b) is with $p_{\text{con}} = .5$. Note that the signal values X_n are reduced considerably with the addition of contrarians.

3.1 Sample signals

Figures 3 and 4 contain realizations of the time series generated by the system for several parameter values. In Figures 3(a) and 3(b), we show the system where each player has $L = 5$ predictors out of 40 possible with $p_{\text{con}} = 0$ and $p_{\text{con}} = .5$, respectively. Note that in each case the signal seems roughly centered around zero. However, the standard deviation for the system with $p_{\text{con}} = 0$ is approximately 0.3458 while the standard deviation takes a much smaller value of 0.0201 when $p_{\text{con}} = .5$. Figures 4(a) and 4(b) show the same plots for the system with $L = 15$. Once again the signals are roughly centered at zero, but have very different standard deviations: 0.5835 and 0.0238, respectively. In each case, visual comparison seems to point to more structure in the presence of contrarians.

3.2 Density of contrarians and equilibration

Before discussing detailed statistics of the stationary processes generated by these systems we first mention an interesting feature which was used to ascertain if the systems have reached their stationary states. In Figures 5(a) and 5(b), we show the fraction of predictors used in contrarian mode as a function of time for the systems with $L = 5$ and $L = 15$, respectively, for various values of p_{con}. The density of contrarians at time t, denoted by $c(t)$, is calculated by averaging the fraction of predictors used in contrarian mode over consecutive time intervals of sizes $2^k \times 10^4$ ending at t. An interesting feature is that the density converges to a limiting density which can be significantly different from p_{con}. When $p_{\text{con}} = .5$, the system *decreases* to an equilibrium density of contrarians which is approximately .489. The fact that the limiting density is less than .5 is a real feature and not merely a conse-

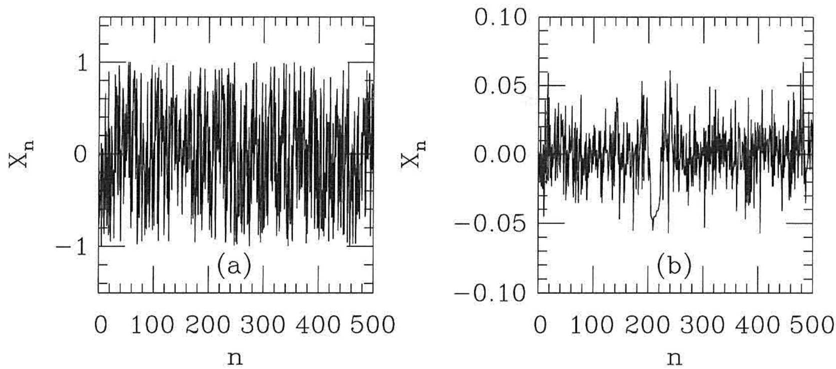

Figure 4: Sample signal X_n of 500 time steps for systems with $L = 15$ predictors held by each player out of a possible 40 after equilibration of the system. 4(a) is with $p_{\rm con} = 0$; 4(b) is with $p_{\rm con} = .5$. As in Figure 3, the amplitude of X_n is reduced considerably by contrarians. In addition, comparison with Figure 3 shows that larger values of L are associated with a higher variance in X_n.

quence of a finite number of players. When $p_{\rm con} = .1$ the density increases to approximately .481. This increase is due to the fact that if protrarians abound, it is generally lucrative to be a contrarian. Consequently, when $p_{\rm con}$ is small but positive, prior to equilibration, players removed due to poor performance are predominantly protrarian.

We use the convergence illustrated in Figure 5 as a criterion for the convergence of the system to its stationary state. When the density of contrarians stabilizes to within a given tolerance (.001) of the limiting value we begin generating our sample statistics.

3.3 Distributions

We now discuss the distributions of the signal values and the local variance. Figures 6(a) and 6(b) contain the empirical densities $f(x)$ of the signal values x for $L = 5$ with $p_{\rm con} = 0$ and $p_{\rm con} = .5$, respectively. Figures 7(a) and 7(b) show the same densities for $L = 15$. Note that the variance of the signal is reduced considerably by the presence of contrarians. The rapid change in monotonicity of the density shown in 6(b) and 7(b) is due to the discrete nature of the system and to the quenched disorder in the predictors. It is conceivable that averaging over the initial distribution of the random predictors yields a smooth density. Figure 8(a) shows the stationary density $f(v)$ of a local variance v when $p_{\rm con} = 0$ for $L = 5$ and $L = 15$. By this we mean that the distribution of the variance v of the time series over windows of 20 time steps is calculated. Figure 8(b) is similar, but with $p_{\rm con} = .5$ for the two systems. Note that the increase in L from 5 to 15 results in a noticeable broadening of the distribution of the local variance. Also note

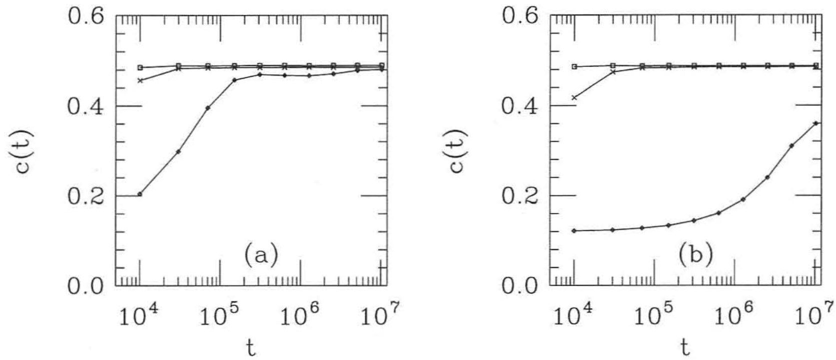

Figure 5: The ratio $c(t)$ of the number of predictors held in contrarian mode to the total number of predictors held (LM) averaged over consecutive time intervals of length $2^k \times 10^4$ for integer k ending at t. Figure 5(a) is for $L = 5$. Squares show the trace when $p_{\text{con}} = .5$, crosses when $p_{\text{con}} = .3$, and diamonds when $p_{\text{con}} = .1$. Note that for these different values of p_{con}, the equilibrium density of contrarians is nearly the same and is less than .5. Figure 5(b) is the same as 5(a) for a system with $L = 15$. Note the much slower rate of convergence when $p_{\text{con}} = .1$.

that the addition of contrarians reduces the variance significantly.

3.4 Conditional variances

One measure of how the variance of X_n depends upon the history up to time $n-1$ is

$$\text{Var}[X_n \mid X_{n-1}] = E[X_n^2 \mid X_{n-1}] - \left(E[X_n \mid X_{n-1}]\right)^2 \tag{6}$$

where $E[X_n \mid X_{n-1}]$ denotes the conditional expectation of X_n given X_{n-1}. The conditional variance (6) is a function of X_{n-1} and a simple quantitative measure of how the volatility of the upcoming value is conditionally dependent on the history of the signal.

Remark. The simplest time series models which utilize such conditional variances are the ARCH models [10] in which

$$X_n = \sigma_n Z_n \qquad \text{and} \qquad \sigma_n^2 = \beta_0 + \sum_{i=1}^{q} \beta_i X_{n-i}^2 \tag{7}$$

where Z_n are i.i.d. mean zero random variables. The main point is that the variance of the upcoming signal value depends on the history of the signal. If all of the β_i's in (7) are positive and if the recent history of the time series is volatile (that is, high values of X_{n-i}^2), then the variance of the upcoming value is high. Conversely, conditioning on low values of X_{n-i}^2 results in a

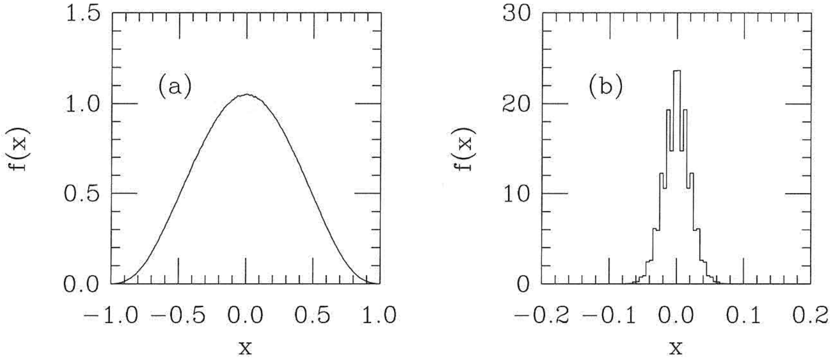

Figure 6: The stationary probability density $f(x)$ of the signal values for $L = 5$. 6(a) is with $p_{\text{con}} = 0$ and 6(b) is with $p_{\text{con}} = .5$. The lack of smoothness of the density in 6(b) is an attribute of the discrete nature of the system, both in the number of players and in the number of strategies.

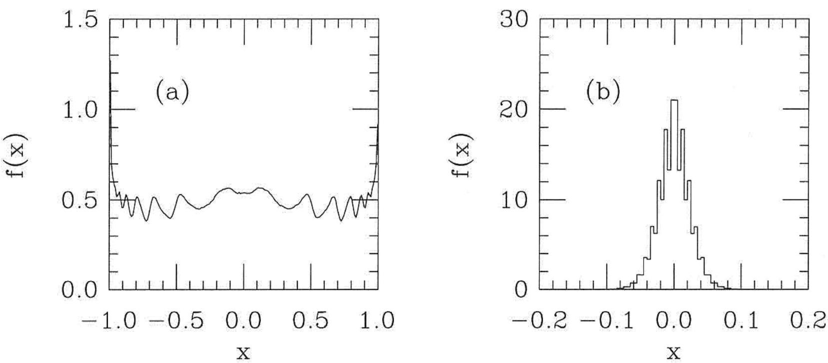

Figure 7: The stationary probability density $f(x)$ of the signal values for $L = 15$. 7(a) is with $p_{\text{con}} = 0$ and 7(b) is with $p_{\text{con}} = .5$. Note the striking change in appearance of $f(x)$ when $p_{\text{con}} = 0$ as L is increased from 5 to 15 (compare with 6(a) and 7(a)).

diminished variance for the upcoming value. Such an ARCH process would, therefore, exhibit clustered volatility. In general, one must consider the variance of X_n conditioned upon the entire history of the signal and, in fact, a major consideration in time series modeling is determining how far back the conditioning is pertinent.

The conditional variance with which we work is conditioning only on the last signal value—this will be more than adequate to illustrate interesting behavior. In Figures 9(a) and 9(b), we plot the conditional variance for a system with $L = 5$, and $p_{\text{con}} = 0$ and $p_{\text{con}} = 0.5$, respectively. Here $v_c(x) = \text{Var}[X_n \mid X_{n-1} = x]$ as defined in (6). Figures 10(a) and 10(b) are

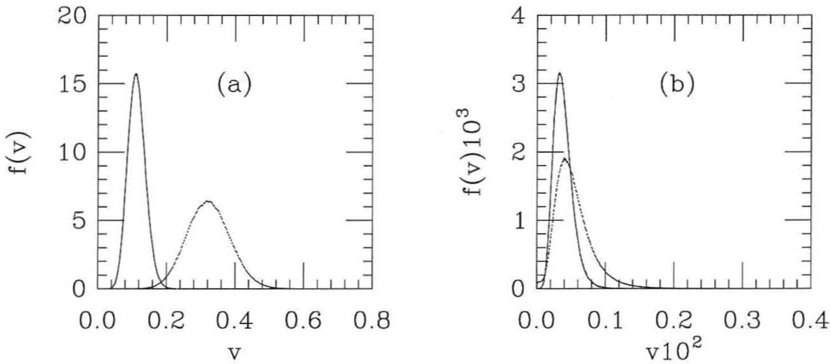

Figure 8: The stationary probability density $f(v)$ of the signal variance over blocks of 20 time steps. 8(a) is with $p_{\text{con}} = 0$: the solid plot is for $L = 5$ and the dashed plot is for $L = 15$. 8(b) is the same with $p_{\text{con}} = .5$.

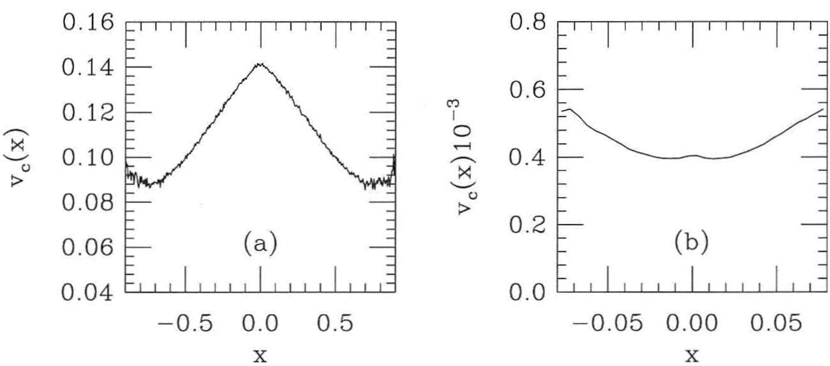

Figure 9: The conditional variance $v_c(x) = \text{Var}[X_n \mid X_{n-1} = x]$ for $L = 5$. 9(a) is with $p_{\text{con}} = 0$ and 9(b) is with $p_{\text{con}} = .5$. The addition of contrarians results in a clearly enhanced volatility at extreme values of x, that is, clustered volatility. Without contrarians, 9(a) shows a reduced volatility for extreme values of x.

similar, but for a system with $L = 15$. The contrast between the (a) and (b) figures is striking. In the presence of contrarians, it is clear that extreme values of X_{n-1} are associated with higher values of the variance of X_n. This is a manifestation of clustered volatility: if one conditions on the current signal value being far from the unconditional mean (which is approximately zero), then the signal is currently "volatile," and the upcoming value is of higher variance than that without the conditioning. In the absence of contrarians, the opposite is true, with a suppressed variance when the current value is far from the origin.

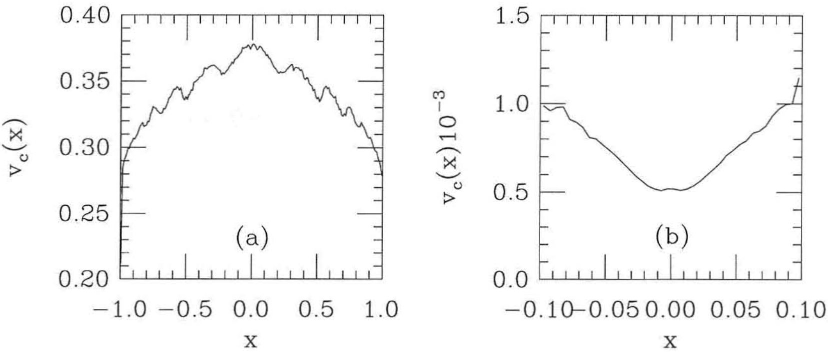

Figure 10: The conditional variance $v_c(x)$ for $L = 15$. 10(a) is with $p_{\text{con}} = 0$ and 10(b) is with $p_{\text{con}} = .5$. The addition of contrarians and the large value of L results in an even more striking display of clustered variance than in 9(b).

3.5 Scaling

Next we examine the partial sums of the signals: $S_n = \sum_{i=1}^{n} X_i$. In particular, we observe the scaling of $\text{Var}(S_n)$ for large n. Simulations indicate the usual diffusion scaling $\text{Var}(S_n) \sim \sigma n$. The value of σ varies as L varies, and with the addition of contrarians. For $L = 5$, $p_{\text{con}} = 0$ results in $\sigma \approx .122$ and $p_{\text{con}} = .5$ yields $\sigma \approx .0005$. For $L = 15$, we have $p_{\text{con}} = 0$ when $\sigma \approx .308$ and $p_{\text{con}} = .5$ when $\sigma \approx .0009$.

3.6 Conclusions

The results presented here show that a variety of nontrivial time series can be generated from relatively simple systems. Of particular note is the appearance of clustered volatility when contrarian behavior and evolutionary dynamics is introduced. While the presence of clustered volatility in many financial signals is well established empirically, plausible mechanisms for its origin have been elusive.

The systems and results discussed here are not to be taken as particularly realistic models of real economic agents. Rather, the fact that these simple adaptive systems can exhibit behavior which is observed in real economic signals and which is largely unexplained makes further work in this area compelling. Analysis of such adaptive interacting systems will presumably require techniques outside of the realm of equilibrium theory. In these systems the agents do not reach a fixed point in the strategies that they adopt—the agents' "perceptions" (primitive as they are) of their optimal strategies are evolving and highly coupled. The stationary behavior of these systems is characterized not by fixed probability distributions of future events and an optimal allocation of resources, but by a stationary process on the (rather large) state space of the strategies used by each of the players.

There are many directions for future work. The tasks involved in analysis are challenging, even for the simple systems discussed here. However, it would be desirable to rigorously establish some attributes for both the stationary distribution of the players' strategies and the resulting time series. There is also the appealing prospect of establishing diffusion (continuous time) limits of the time series (e.g., as the number of players diverge). There are also a multitude of directions to explore in simulations by expanding the predictor set and the players' strategies for using them. On a purely statistical note, it would be interesting to consider how well common time series techniques such as ARCH or GARCH models fit the signals generated by these adaptive systems. Perhaps the most practical endeavor will be to see just how useful these systems are in modeling real financial data.

Acknowledgements

The authors would like to thank W. B. Arthur, J. Carlson, B. LeBaron, M. Leibig, and S. Pepke for useful comments and discussions. In addition, P. Feldman, R. Feldman, S. Rachev, and D. Steigerwald provided helpful information on simulation and time series techniques. The work of ERG was supported by ONR grant N000014-91-J-1502 and by Los Alamos National Laboratories. The work of GHS was supported by NSF grant DMS-9305904.

References

[1] V. Akgiray, "Conditional Heteroskedasticity in Time Series of Stock Returns: Evidence and Forecasts," *Journal of Business*, **62** (1989) 55–80.

[2] G. W. Schwert, "Why Does Stock Market Volatility Change over Time," *Journal of Finance*, **44** (1989) 1115–1153.

[3] A. A. Weiss, "ARMA Models with ARCH Errors," *Journal of Time Series Analysis*, **5** (1984) 129–143.

[4] D. A. Hsieh, "Modeling Heteroskedasticity in Daily Foreign Exchange Rate Changes," *Journal of Business and Economic Statistics* **7** (1989) 307–317.

[5] T. Bollerslev, R. Chou and K. Kroner, "ARCH Modeling in Finance," *Journal of Econometrics*, **52** (1992) 5–59.

[6] W. B. Arthur, "Inductive Reasoning and Bounded Rationality," *American Economic Review*, **84** (1994) 406–411.

[7] W. B. Arthur, "On Learning and Adaptation in the Economy," Santa Fe Institute working paper 92-07-038 (1992).

[8] T. J. Sargent, *Bounded Rationality in Macroeconomics* (Oxford: Oxford University Press, 1993).

[9] W. B. Arthur, J. H. Holland, B. LeBaron, R. G. Palmer, and P. J. Tayler, "An Artificial Stock Market" (in preparation).

[10] R. F. Engle, "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, **50** (1992) 987–1008.