MANAGEMENT SCIENCE
Vol. 53, No. 7, July 2007, pp. 1165–1180
ISSN 0025-1909 | EISSN 1526-5501 | 07 | 5307 | 1165
informs®
DOI 10.1287/mnsc.1060.0622

Complexity and the Character of Stock Returns: Empirical Evidence and a Model of Asset Prices Based on Complex Investor Learning

Scott C. Linn
Division of Finance, Michael F. Price College of Business, University of Oklahoma, 205 A Adams Hall, Norman, Oklahoma 73019, slinn@ou.edu

Nicholas S. P. Tay
School of Business and Management, University of San Francisco, 2130 Fulton Street, Malloy Hall, San Francisco, California 94117-1045, tay@usfca.edu

Empirical evidence on the distributional characteristics of common stock returns indicates: (1) A power-law tail index close to three describes the behavior of the positive tail of the survivor function of returns ($\text{pr}(r > x) \sim x^{-\alpha}$), a reflection of fat tails; (2) general linear and nonlinear dependencies exist in the time series of returns; (3) the time-series return process is characterized by short-run dependence (short memory) in both returns as well as their volatility, the latter usually characterized in the form of autoregressive conditional heteroskedasticity; and (4) the time-series return process probably does not exhibit long memory, but the squared returns process does exhibit long memory. We propose a model of complex, self-referential learning and reasoning amongst economic agents that jointly produces security returns consistent with these general observed facts and which are supported here by empirical results presented for a benchmark sample of 50 stocks traded on the New York Stock Exchange. The market we postulate is populated by traders who reason inductively while compressing information into a few fuzzy notions that they can in turn process and analyze with fuzzy logic. We analyze the implications of such behavior for the returns on risky securities within the context of an artificial stock market model. Dynamic simulation experiments of the market are conducted, from which market-clearing prices emerge, allowing us to then compute realized returns. We test the effects of varying values of the parameters of the model on the character of the simulated returns. The results indicate that the model proposed in this paper can jointly account for the presence of a power-law characterization of the positive tail of the survivor function of returns with exponent on the order of three, for autoregressive conditional heteroskedasticity, for long memory in volatility, and for general nonlinear dependencies in returns.

Key words: learning; stock return distribution; power-law; nonlinear dependence; artificial stock market
History: Accepted by Brian Uzzi and Luis Amaral, special issue editors; received August 31, 2004. This paper was with the authors 3 months for 2 revisions.

1. Introduction

Empirical evidence on the distributional characteristics of common stock returns indicates: (1) A power-law tail index close to three describes the behavior of the positive tail of the survivor function of returns ($\text{pr}(r > x) \sim x^{-\alpha}$) (Gopikrishnan et al. 1999, Plerou et al. 1999), a reflection of fat tails; (2) general linear and nonlinear dependencies exist in the time series of returns (Scheinkman and LeBaron 1989, Hsieh 1991, Brock et al. 1991); (3) the time-series return process is characterized by short-run dependence (short memory) in both returns as well as their volatility, the latter usually characterized in the form of autoregressive conditional heteroskedasticity (Bollerslev et al. 1992, Glosten et al. 1993, Engle 2004); and (4) the time-series return process probably does not exhibit long memory (Lo 1991), but the squared returns process does exhibit long memory (Ding et al. 1993, Bollerslev and Wright 2000). Little is known, however, about what behaviors on the part of investors should give rise to jointly observing these phenomena. We propose a model of complex, self-referential learning and reasoning amongst economic agents that jointly produces security returns consistent with these general observed facts. The model features investors who reason inductively through experimentation with new hypotheses while compressing information into a few fuzzy notions that they can in turn process and analyze with fuzzy logic. Our approach is motivated first by the cogent argument set forth by Arthur (1991, 1992, 1994, 1995), Arthur et al. (1997), and LeBaron et al. (1999), who conclude that deductive reasoning must give way to inductive reasoning in complex, ill-defined settings and that real capital markets exhibit a high level of complexity. Second, we follow a stream of thought proposed by Smithson (1987) and Smithson and Oden (1999), amongst others, who conclude that human reasoning can be modeled as if the thought process is described by the application of fuzzy logic. Assuming mental behavior of this sort allows the agent to step outside the rigid confines of more traditional models. We embed this behavior in an artificial stock market model that is utilized as a vehicle for simulating the dynamics of a market from which market-clearing security prices emerge, allowing us to compute realized returns.

The structure of our model extends the important work done in developing the Santa Fe Artificial Stock Market Model studied by LeBaron et al. (1999). While there are important differences between the models, the two machine-learning methods produce similar results. We feel that this is important to understanding how agents learn and reason. One notable implication is that our framework requires, in principle, much less of the agent. We feel that the fact that our model produces results similar to the Santa Fe Institute (SFI) model is both an important statement about how individuals learn and reason, but is also an endorsement of the importance of the SFI model, because, turned around, we are saying that the SFI model produces results similar to a model in which agents learn in a much less structured fashion.

The learning system proposed by LeBaron et al. (1999) makes use of a model in which each agent is assumed to make choices predicated on a large number of rules for the mapping of market conditions into expectations, each with numerous conditions.[1] Agents in their model formulate new rules through the application of behavior emulated by a genetic algorithm. Agents in our model, in contrast, employ only a handful of hypotheses used to generate expectations. These hypotheses are composed of only four rules each. Each rule employs a selection of the information available, which will be used to construct a conditional assignment of values to the parameters of a model of predictions. An agent learns in two ways. First, new hypotheses are generated from existing hypotheses, and low-accuracy hypotheses are replaced, with high probability, with newly generated hypotheses. In this way, agents are selecting the variables to use in the construction of the parameters of their prediction models. Second, values for the parameters of the prediction model, conditional on the hypothesis, are formed by the agent applying fuzzy logic to the observed data. The latter allows us to reduce the number of hypotheses and rules to a quantity that is more like what one might expect versus the long list of rules utilized in the SFI model.

We show that with-dividend returns computed from simulated market-clearing prices for the environment we propose exhibit a tail index in their survivor functions characterized by a power law with exponent on the order of three, exhibit autoregressive conditional heteroskedasticity, and exhibit general nonlinear dependencies and long memory in the volatility process. We also document the presence of similar characteristics for a sample of 50 common stocks traded on the New York Stock Exchange, which act here as a benchmark. The appeal of our results is twofold: First, the behavioral model we propose generates return characteristics for risky securities that are similar to what are observed for actual stock returns; and second, it does so as a product of an environment in which economic agents are endowed with learning and reasoning processes that are close to what many disciplines believe is an accurate depiction of actual behavior.

How investors learn and interact in complex capital market environments is crucial to understanding the nature of financial security return distributions. Complexity demands an alternative approach to the analysis of markets and institutions. The approach we take has its roots in work begun and continuing at the Santa Fe Institute. Examples of such work focusing on the behavior of security prices include Arthur et al. (1997), Brock and Hommes (1998), LeBaron et al. (1999), and Tay and Linn (2001). Tesfatsion (2002) and LeBaron (2000, 2006) provide reviews of this literature.[2]

Our model is motivated by the discrepancy between the idealized well-defined environment that is commonly assumed in neoclassical financial market models and the complex ill-defined markets that are observed in practice. Neoclassical financial market models are generally designed within the context of a well-defined setting so that economic agents are able to logically deduce the expected prices of securities that they in turn employ when setting their demands for those securities.[3] Real stock markets do not, however, typically conform to the severe restrictions required to guarantee such behavior. The actual market environment is usually much more ill defined.

[1] Specifically, LeBaron et al. (1999) assume that each agent possesses 100 rules.

[2] In Tay and Linn (2001), we examine, amongst other things, the time series of prices for a related model and in particular how those prices deviate from rational expectations prices. That study does not address the jointly observed behaviors in returns mentioned at the beginning of this section, and which are the focus of this study. See also Arthur et al. (1997) and Palmer et al. (1994). Leigh Tesfatsion of Iowa State University maintains a comprehensive website devoted to agent-based computational economics and finance (http://www.econ.iastate.edu/tesfatsi/ace.htm). LeBaron (2000, 2006) presents a review of the foundation articles.

[3] An excellent example is the general equilibrium model developed in Brock (1982).

The dilemma is that in an ill-defined environment the ability to exercise deductive reasoning breaks down, making it impossible for individuals to form precise and objective price expectations. This implies that participants would need to rely on some alternative form of reasoning to guide their decision making. We conjecture that individual reasoning in an ill-defined setting can be described as an inductive process. The application of inductive reasoning involves the formulation of tentative hypotheses to fill in the gaps left by incomplete information. These hypotheses are then continually tested in the market and revised as agents seek to improve their understanding of market behavior. Agents in our model generate predictions by the application of an inductive reasoning process in which they rely on fuzzy decision-making rules due to limits on their ability to process and condense information.

Our study is close in spirit to a contemporaneous study by Gabaix et al. (2006), which focuses on a model of large fluctuations in stock returns and is motivated by the presence of a power-law decay in the survivor function of returns as well as trading volume. Our study, however, differs in several important ways from theirs. First, we present a model that jointly produces a power-law decay in the survivor function for returns, autoregressive conditional heteroskedasticity, and general nonlinear dependencies. The central focus of Gabaix et al. (2006) is explaining what gives rise to the power-law decay. Second, our model focuses on the influence of learning and reasoning by agents and the influence of nontraditional aspects of these activities on the distributions of returns. Gabaix et al. (2006) present an insightful model built up from assumptions about the structure of trading and the search for trading partners. We, on the other hand, choose to minimize these influences to highlight how agents learn and reason in an ill-defined environment. In this way, both studies provide important and new insights into what factors potentially give rise to the features of stock returns already mentioned.

This paper is organized as follows. In §§2–5, we begin by presenting empirical results on the existence of a power law in the behavior of the survivor function of common stock returns; on the presence of dependencies in the time series of returns, including autoregressive conditional heteroskedasticity as well as long memory in the volatility process; and on the general existence of nonlinear dependencies in stock returns. Our focus is on a sample of 50 common stocks traded on the New York Stock Exchange. Section 6 goes on to summarize how agents learn in the model and the market environment. Section 7 describes the dynamic simulation experiments of the model. Section 8 presents an analysis of the returns computed from the market-clearing prices generated in the artificial stock market, drawing comparisons with the results presented in §§2–5 for the benchmark sample of stock returns.[4] We also investigate how the characteristics of the returns generated by the model vary with changes in values of key parameters of the model. Section 9 presents our conclusions.

2. The Data and Descriptive Statistics

The data examined herein are comprised of (a) daily with-dividend return series for 50 actively traded NYSE-listed common stocks, and (b) 540 return series computed from simulations of the artificial stock market model. The values of parameters of the model are varied to provide insight into their individual influence on the results of the simulations. Thirty simulations of the market are generated for each set of parameters examined. We defer our analysis of the data from the artificial stock market simulations until §8, following our discussion of the model's structure and the design of the decision-making algorithm ascribed to agents in the model. Instead, we begin by focusing on the characteristics of the stock returns for the 50 NYSE-listed stocks to establish a benchmark for comparison. The source of the stock return data is the Center for Research in Security Prices (CRSP) Daily Return file, and the data included are the 2,780 daily returns ending December 31, 1998.[5]

Table 1 presents summary statistics for the 50 actual stock return series along with results for the simulations, which will be discussed later in the paper. The results for the actual stock returns appear in the column headed "Actual." The sample statistic names are listed in the leftmost column. We present the average values for the statistics and, in parentheses, the standard errors of the point estimates of the statistics computed across the 50 cases. When a test statistic is reported in a table, we present the average value of the test statistic computed across the relevant cases and, in square brackets, the fraction of tests rejecting the null hypothesis.

Notable amongst the descriptive statistics for the actual returns is the high level of kurtosis and the positive skewness. Kurtosis for a normal distribution should equal three, while skewness should equal zero. Both measures for the sample return series deviate from these benchmarks. The Jarque-Bera test (not reported) rejects the null hypothesis of normality for

[4] All computations, including those within the simulations, are performed using the technical computing language MATLAB, a product of MathWorks, Inc.

[5] Returns from the period October 19–21, 1987, are excluded from each series.

Table 1 **Average Values of Descriptive Statistics and Test Statistics**

	Actual	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6	Case 7	Case 8	Case 9	Case 10	Case 11	Case 12	Case 13	Case 14	Case 15	Case 16	Case 17	Case 18
Mean ($\times$100)	0.064	0.013	0.010	0.001	0.012	0.012	0.023	0.000	0.025	0.008	0.024	0.013	0.014	0.012	0.011	0.013	0.012	0.015	0.122
	(0.005)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.005)
Median ($\times$100)	0.000	−0.013	−0.003	−0.003	−0.014	−0.021	−0.035	0.000	−0.020	−0.017	−0.021	−0.010	−0.022	−0.013	−0.014	−0.013	−0.013	0.000	−0.001
	0.000	(0.003)	(0.001)	(0.002)	(0.003)	(0.004)	(0.005)	(0.000)	(0.004)	(0.004)	(0.005)	(0.003)	(0.005)	(0.003)	(0.003)	(0.003)	(0.003)	(0.000)	(0.002)
Std. dev. ($\times$100)	2.120	1.665	1.441	0.241	1.684	1.651	2.174	0.036	2.260	1.479	2.184	1.651	1.656	1.646	1.622	1.686	1.566	1.734	4.943
	(0.132)	(0.032)	(0.052)	(0.069)	(0.046)	(0.045)	(0.027)	(0.008)	(0.081)	(0.038)	(0.047)	(0.051)	(0.043)	(0.034)	(0.041)	(0.035)	(0.050)	(0.035)	(0.097)
Skewness	0.424	1.304	0.948	−0.023	1.561	1.249	2.244	0.080	0.801	1.592	1.515	1.507	1.714	1.546	1.178	1.526	1.348	0.564	0.979
	(0.049)	(0.202)	(0.160)	(0.180)	(0.167)	(0.147)	(0.114)	(0.157)	(0.099)	(0.233)	(0.106)	(0.151)	(0.225)	(0.194)	(0.102)	(0.114)	(0.082)	(0.066)	(0.049)
Kurtosis	10.434	26.967	32.022	129.050	32.747	24.618	29.280	146.490	13.775	34.620	21.992	30.551	32.572	27.906	21.739	26.457	22.844	28.302	13.543
	(0.551)	(4.905)	(2.996)	(13.333)	(3.468)	(3.473)	(2.123)	(22.903)	(2.443)	(4.696)	(1.681)	(4.255)	(5.134)	(3.946)	(2.128)	(2.910)	(2.253)	(1.411)	(0.581)
Hill's tail index	2.550	1.944	1.298	1.419	1.737	1.978	2.102	2.972	2.022	1.872	1.881	1.784	1.927	1.924	1.874	2.009	1.918	0.900	1.779
	(0.088)	(0.082)	(0.094)	(0.075)	(0.090)	(0.093)	(0.028)	(0.122)	(0.076)	(0.087)	(0.041)	(0.089)	(0.079)	(0.072)	(0.085)	(0.082)	(0.083)	(0.017)	(0.026)
Quintos' tail index	2.877	2.546	1.887	1.110	2.295	2.734	2.676	3.794	2.885	2.494	2.544	2.412	2.607	2.595	2.520	2.707	2.620	1.609	2.046
	(0.118)	(0.121)	(0.113)	(0.097)	(0.123)	(0.132)	(0.037)	(0.164)	(0.104)	(0.107)	(0.046)	(0.119)	(0.106)	(0.099)	(0.118)	(0.104)	(0.116)	(0.034)	(0.022)
BDS, $\varepsilon = 1.5\sigma$, $m = 3$	10.45	25.08	25.46	24.13	25.32	24.53	23.93	10.73	26.06	24.03	24.31	25.90	25.68	24.58	24.95	25.43	25.82	24.92	25.06
	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]
BDS, $\varepsilon = 2\sigma$, $m = 3$	9.71	23.74	24.20	23.86	23.87	22.76	22.13	12.67	23.47	22.62	23.10	24.06	23.67	22.82	23.17	23.32	23.87	23.51	21.87
	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]
Ljung-Box Q(5)	31.84	518.81	543.91	604.67	535.74	539.50	424.72	404.45	558.50	524.28	482.17	530.46	502.54	497.90	531.04	498.89	529.13	575.97	424.81
	[0.60]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]
Ljung-Box Q(10)	40.45	530.50	563.86	630.59	556.54	557.04	438.18	412.52	574.57	537.21	495.14	548.65	517.61	509.55	541.84	509.89	544.87	590.42	437.92
	[0.64]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]
ARCH-LM(1)	65.17	233.78	390.03	590.31	257.82	203.27	82.93	467.88	208.64	281.04	136.04	204.66	159.57	138.63	205.36	155.99	159.94	569.88	143.15
	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[0.97]	[1.00]	[1.00]	[1.00]	[0.97]	[1.00]	[0.97]	[1.00]	[1.00]	[1.00]
ARCH-LM(2)	72.61	242.83	416.79	732.15	279.70	216.08	86.88	615.28	220.03	308.69	146.18	219.27	174.78	146.40	217.42	167.47	175.19	614.35	167.64
	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[0.97]	[1.00]	[1.00]	[0.97]	[1.00]	[1.00]	[1.00]	[1.00]	[1.00]	[0.97]	[1.00]	[1.00]	[1.00]
$V(180)$	1.17	1.43	2.15	6.59	1.55	1.45	1.10	5.28	1.28	1.73	1.11	1.56	1.30	1.32	1.55	1.43	1.23	2.05	0.89
	[0.00]	[0.17]	[0.67]	[0.97]	[0.27]	[0.20]	[0.03]	[0.80]	[0.10]	[0.43]	[0.03]	[0.20]	[0.07]	[0.07]	[0.20]	[0.17]	[0.10]	[0.73]	[0.30]
$V(270)$	1.13	1.44	2.10	7.82	1.59	1.48	1.22	6.34	1.28	1.79	1.15	1.60	1.34	1.35	1.60	1.46	1.27	1.93	0.81
	[0.02]	[0.17]	[0.67]	[0.97]	[0.27]	[0.23]	[0.00]	[0.83]	[0.10]	[0.40]	[0.07]	[0.23]	[0.07]	[0.13]	[0.27]	[0.17]	[0.07]	[0.53]	[0.53]
$V_s(180)$	0.761	0.618	0.779	1.026	0.728	0.605	0.513	0.711	0.649	0.730	0.631	0.646	0.697	0.642	0.680	0.648	0.663	0.691	0.616
	[0.62]	[0.77]	[0.60]	[0.07]	[0.73]	[0.80]	[1.00]	[0.50]	[0.73]	[0.63]	[0.87]	[0.80]	[0.77]	[0.77]	[0.70]	[0.83]	[0.70]	[0.77]	[0.87]
$V_s(270)$	0.639	0.535	0.677	0.941	0.639	0.524	0.436	0.657	0.553	0.642	0.540	0.564	0.604	0.554	0.587	0.560	0.576	0.599	0.534
	[0.88]	[0.90]	[0.73]	[0.10]	[0.87]	[0.90]	[1.00]	[0.63]	[0.77]	[0.87]	[0.93]	[0.93]	[0.87]	[0.83]	[0.83]	[0.87]	[0.80]	[0.90]	[0.97]

Notes. The column labeled "Actual" presents results for a sample series of daily with-dividend common stock returns for 50 common stocks traded on the NYSE. Data are from the CRSP archive files. All other columns are for the cases of the artificial stock market listed in Table 3, where each case represents a different selection of values for the parameters of the model. The Hill (1975) and Quintos et al. (2001) estimators of the tail index are described in §3 of the text. The test statistics and the respective null hypotheses being tested for the tests BDS, Ljung-Box Q(q), ARCH-LM(q), $V(q)$, and $V_s(q)$ are described in §4 of the text. Standard errors are in parentheses; fractions of tests rejecting the null hypothesis are in square brackets.

each of the 50 sample series.[6] High kurtosis is consistent with the presence of heavy or fat tails in the distribution, but may also be a manifestation of a peaked distribution.

3. Power-Law Tail Behavior of the Survivor Function

Numerous investigators (see Mandelbrot 1997, Gopikrishnan et al. 1999, Plerou et al. 1999) have presented evidence indicating that the distributions of common stock returns exhibit tails with greater mass than would be predicted if the distributions were Gaussian normal, the conclusion being that these distributions exhibit fat or heavy tails. For a class of distributions characterized by what is referred to as regular variation in the tails, the far-right portion of the tail of the survivor function is a power-law function of the form[7]

$$\mathrm{pr}(r_i > x) \sim \frac{1}{x^{\alpha_i}}, \tag{1}$$

where α_i is the exponent characterizing the power-law tail index for security i and x is a threshold above which the algebraic relation is assumed to be valid. Heavy-tailed models exhibit tails that decay more slowly than the tails of the normal distribution. The condition in essence states that far out in the tail the distribution behaves like a Pareto distribution, in contrast to a normal distribution. As a consequence, the tail probabilities decline according to a power function (Equation (1)). In contrast, the normal distribution decays much faster, in fact, as an exponential function.[8] The principle motivation behind the analysis of what is commonly referred to as the "tail index" $\alpha$is the assessment of the fatness in the tails of the distribution. The index α declines as the tail becomes thicker. A tail index in the range $0 < \alpha < 2$ is consistent with a distribution exhibiting heavy tails and infinite variance.

We compute estimates of α for the panel of stocks using the methods developed in Hill (1975).[9] The Hill estimator is the correct estimator for the tail index if the data are independent and identically distributed (i.i.d.). The empirical literature, however, provides extensive evidence that stock returns follow a process from the class of models characterized by autoregressive conditional heteroskedasticity (Bollerslev et al. 1992, Engle 2004). Quintos et al. (2001) derive a tail index estimator under the condition of GARCH behavior.[10] Therefore, we also present the tail index estimates computed using the Quintos et al. (2001) estimator. It is worth noting that the GARCH(1, 1) process does not exhibit infinite variance. DuMouchel (1983) has shown that the set of observations made up of the top 10% of the ranked-by-value observations of the total sample yield Hill estimates of α that perform well when the overall sample is large.[11] Given that our sample sizes are relatively large, we chose to follow the 10% rule in our calculations. For each stock in the sample, we compute the Hill's estimate of $\hat{\alpha}_i$ using the largest 270 observations of the ranked samples. Table 1 presents the average value of the tail index computed using Hill's estimator and, separately, the Quintos et al. (2001) estimator. The standard error computed using the 50 point estimates is presented below each average. The summary results are: Hill estimator (mean = 2.550, std. err. = 0.088); Quintos et al. (2001) estimator (mean = 2.877, std. err. = 0.118). The average tail index estimate based on the Quintos et al. (2001) estimator is greater than the average based on the traditional Hill estimator. This is perhaps not a surprise, given the general finding that the data exhibit ARCH-type behavior. The standard errors computed using the 50 point estimates indicate that the 50 point estimates are not widely dispersed. These results are in general agreement with results presented by Plerou et al. (1999) and others who conclude that an estimate of α on the order of three describes stock returns for a wide range of return frequencies, including daily data. One conclusion is that

[6] A description of the Jarque-Bera test statistic is provided in the online supplement to this paper, which is provided in the e-companion. (An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.)

[7] Let x be a random variable with cumulative probability distribution function $F_x(x')$. The survivor function $S_x(x')$ is defined as $S_x(x') = 1 - F_x(x') = \mathrm{pr}(x > x')$. The regular variation condition is

$$\lim_{t\to\infty} \frac{S_x(tx')}{S_x(x')} = x^{-\alpha}$$

(Feller 1966, Chap. VIII).

[8] Specifically, the survivor function of the normal distribution is proportional to

$$\frac{1}{\sqrt{2\pi}}\frac{1}{x}\exp\left(-\frac{x^2}{2}\right)$$

(Gourieroux and Jasiak 2001).

[9] Let $x_{(1)} \geq x_{(2)} \geq \cdots \geq x_{(N)}$ represent the ordered values of a sample $\{x\}$. The Hill estimator based on the k largest order statistics is computed as

$$\hat{\alpha} = \left(\frac{1}{k}\sum_{n=1}^{k}\log\frac{x_{(n)}}{x_{(k+1)}}\right)^{-1}$$

(Hill 1975). The Hill estimator is a maximum-likelihood estimator for α of distributions that are characterized by the regular variation condition described in Footnote 7.

[10] Quintos et al. (2001) develop an estimator for the Hill tail index under the condition that the series follows a GARCH(1, 1) process. Quintos et al. show that if a random variable x follows a GARCH(1, 1) process and has a tail index α, then x^2 has tail index $\alpha/2$.

[11] LeBaron and Samanta (2006) have recently presented evidence providing further support for the 10% rule.

while heavy tails are consistent with the data, the typical estimate of α is greater than two, implying a finite variance for the distribution.

4. Temporal Dependence

4.1. General Dependencies: The BDS Test

We begin by presenting a general test for dependence. A test developed by Brock et al. (1987), the BDS test (see also Brock et al. 1996) has the power to reject the null hypothesis of i.i.d. data against several alternative hypotheses including models exhibiting nonlinear dependence.[12]

The BDS statistic tests the propensity of a series to cluster within a distance between points of ε. If the data are i.i.d., the probability of the distance between any pair of points being less than or equal to ε will be constant. In contrast, if the data exhibit nonlinear dependence, then this will not tend to be the case. We set the parameter ε to 1.5 and 2 times the standard deviation of the relevant data series. Kanzler (1999), in an extensive study of the behavior of the BDS statistic, has found that these settings for ε minimize the failure of the BDS statistic to approximate the normal distribution in finite samples under the null hypothesis of i.i.d. data. The BDS test statistic has an asymptotic normal distribution under the null hypothesis of i.i.d. data. Details on the test statistic are presented in the online appendix, including an elaboration on the definition of the "embedding dimension" n.[13] A positive BDS statistic suggests that some form of clustering is present on a too-frequent basis, or in other words, patterns occur too frequently relative to what would be expected if the data exhibited no dependencies. We reject the null hypothesis at the 5% level when $|\text{BDS}_n(\delta, T)| > 1.96$.

Table 1 presents the average BDS test statistics computed for the embedding dimension $n = 3$ and $\varepsilon = 1.5\sigma$ or $\varepsilon = 2\sigma$ for the 50 actual daily time series of returns. The numbers in square brackets are the fraction of tests in which the null hypothesis is rejected at the 5% level. In both cases, the null hypothesis is rejected for every sample series at the 5% level, indicating that some form of dependence is present in each sample series. A range of embedding dimensions was tested and the results always indicated rejection of the null, and so are not presented to conserve space. The BDS test will reject the null hypothesis of i.i.d. for alternative hypotheses in which there exists linear dependence as well as nonlinear dependence. Later we will return to tests in which we explore whether dependence prevails after taking into account both linear and nonlinear dependencies. Before exploring that issue, we turn to the issue of whether dependency is driven by short or long memory in the process generating stock returns.

4.2. Short Memory Tests

In this section, we discuss tests on the presence of short-term memory, while the next section discusses tests of long-term memory. Table 1 presents statistics associated with two tests for short-term memory. The tests are, respectively, the Ljung-Box Q statistic, which tests for the presence of significant autocorrelation in the return series, and the ARCH-LM test of Engle (1982), which is a Lagrange multiplier test of the hypothesis that a series exhibits autoregressive conditional heteroskedasticity.[14] Both statistics are distributed as χ^2_q, where q equals the number of lags over which the statistic is computed. Critical values at the 5% level are 11.07 (five lags) and 18.30 (10 lags). As the numbers in brackets indicate, the Q test results reject the null hypothesis of no autocorrelation in the actual return series for roughly 60% of the sample series. In contrast, the null hypothesis of no short-term ARCH effects is rejected for all of the cases at each lag shown. Additional lags were also tested and the results were relatively invariant to the choice of q. These results indicate that while short-term memory is present in some of the returns series, short-term ARCH effects are present in every series.

4.3. Long Memory Tests

Lo (1991) presents evidence of no long-term memory in stock returns. More recently, other authors have presented evidence consistent with long memory in the squared returns, suggesting long memory in the dynamics of volatility (Ding et al. 1993, Bollerslev and Wright 2000).[15]

Mandelbrot (1972) has suggested using the statistic equal to the range divided by the standard deviation first developed by Hurst (1951) as a method for detecting long-range dependence in a series. Lo (1991) has shown that this statistic is biased in the presence of short-range dependencies and has developed a modified statistic that accounts for the presence of

[12] See also Brock et al. (1991, 1996), LeBaron (1997), and Kanzler (1999) for further discussion and analysis of the BDS test. Several alternative tests have been proposed in the literature, including a test developed by Kaplan (1994), and the test of White (1989). Barnett et al. (1997) show that for the large-sample experiments they study, the BDS test always rejects the null when it should be rejected.

[13] The BDS statistics were computed using the MATLAB routine developed by Blake LeBaron of Brandeis University.

[14] The constructions of the Ljung-Box Q statistic (Ljung and Box 1979) and the ARCH-LM test of Engle (1982) are described in the online supplement to this paper.

[15] The stochastic process of a random variable $\tilde{x}_t$ is characterized as exhibiting long memory (long-range dependence) if there exists a real number H and a constant c such that the autocorrelation function $\rho(j)$ decays as $\rho(j) = cj^{2H-2}$ as $j \to \infty$.

short-range dependence, which has power to detect long memory, but has reasonably behaved size when the null of no long memory is true. The results presented in the prior section suggest that we should account for short-term dependencies in our tests of long-term memory. Details on the test statistic are presented in the online appendix.

We present values for Lo's V statistic for the lags 180 and 270. The 5% two-tailed critical region for the value of the V statistic is equal to $[0.809, 1.862]$ (Lo 1991, Table II). A computed value outside this region would lead us to reject the null hypothesis of no long memory. Table 1 indicates that in most of the cases the V test does not reject the null hypothesis of *no* long-term memory in the actual return series. The average V statistics for the actual returns are always within the critical region. The rejection rates are virtually 0%. We conclude that the actual series do not exhibit long memory. These results are consistent with the results presented by Lo (1991) for daily returns on stock indices.

Long-term memory in the squared series would be consistent with the presence of long memory in the time-varying volatility process associated with the series. The V test statistic, labeled V_s, is used to test the null hypothesis of no long-term memory in the squared returns. The results indicate that the average V_s statistics are always outside the critical region. Further, the rejection rate is relatively high for each lag. We conclude from these results that the volatility processes associated with many of these series tend to exhibit long memory, although this is not universal across all the sample series.

5. Temporal Dependence Accounting for Short-Term Memory

5.1. ARMA-TARCH Effects

We present test results based on the residual series computed from models that account for short memory in both the level of the return series as well as ARCH-type effects (Bollerslev et al. 1992, Engle 2004).[16] We focus our attention on the TARCH, or threshold autoregressive conditional heteroskedasticity, class of models (Glosten et al. 1993). These models account for the possibility that the market reacts in an asymmetric fashion to good and bad news, and have been shown to have good explanatory power. The general structure of the ARMA(m, n)-TARCH(p, q) model for a return series $r_{i,t}$ is composed of an equation for the mean and an equation for the conditional variance:

$$r_t = a_0 + \sum_{\tau=1}^{m} \phi_i r_{t-\tau} + \varepsilon_t - \sum_{\omega=1}^{n} \theta_i \varepsilon_{t-\omega}, \qquad (2)$$

$$\sigma_t^2 = \omega + \sum_{\eta=1}^{q} \alpha_\eta \varepsilon_{t-\eta}^2 + \gamma \varepsilon_{t-1}^2 d_{t-1} + \sum_{\varpi=1}^{p} \beta_\varpi \sigma_{t-\varpi}^2, \qquad (3)$$

where we suppress the i subscript, $d_t = 1$ if $\varepsilon_t < 0$ and 0 otherwise, $\varepsilon_t = \sigma_t z_t$ represents a random error where z_t is a normally distributed random variable with zero mean and variance one, and σ_t^2 is the conditional variance. The indicator variable d_t captures the asymmetric reaction of the market to good and bad news. If the ARMA-TARCH model is correctly specified, then the standardized residuals of the estimated model $\hat{z}_t = \hat{\varepsilon}_t / \hat{\sigma}_t$ will be i.i.d.

We fit ARMA-TARCH models to each of the data series using maximum-likelihood methods.[17] Table 2 presents Q statistics and ARCH-LM test statistics and rejection rates for the standardized residuals of the 50 models for the lags five and 10. The Q statistics and the ARCH-LM test statistics indicate that we cannot, respectively, reject the null hypotheses of zero autocorrelation in the standardized residuals and no autoregressive conditional heteroskedasticity in the standardized residuals. The table also presents BDS statistics for the computed standardized residuals. These tests indicate that in roughly half the cases, dependence still remains after accounting for short-term memory in the series and for TARCH effects.

We recomputed the long-memory test statistics V and V_s described earlier for each of the standardized residual series. The results are similar to those shown in Table 1. The V tests do not reject the null hypothesis of no long memory in the level of each series of standardized residuals; however, the V_s tests do reject the null for the squared residual series.[18]

5.2. Conclusions

We conclude from these results that the sample of common stocks jointly exhibit a power-law decay in the positive tail of the survivor function of returns with a tail index of roughly three, consistent with a distribution exhibiting heavy tails and finite variance. The data also exhibit autoregressive conditional heteroskedasticity and exhibit dependencies after both linear filtering and filtering out the influence of conditional heteroskedasticity. Finally, we find that the squared residuals of models fitted to account for linear effects as well as conditional heteroskedasticity continue to exhibit long memory. We now turn to a discussion of our model of investor behavior and then

[16] See Engle (1982) for the original development of the ARCH model.

[17] Estimation is performed using the MATLAB GARCH Toolbox.

[18] The results are available from the authors upon request.

Table 2 **Average Values of Test Statistics: Tests of Hypotheses About Features of the Standardized Residuals of ARMA-TARCH Models Fit to Actual and Simulated Returns**

	Actual	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6	Case 7	Case 8	Case 9	Case 10	Case 11	Case 12	Case 13	Case 14	Case 15	Case 16	Case 17	Case 18
Ljung-Box Q(5)	1.48	4.25	6.24	5.82	5.86	4.20	5.87	5.02	3.95	4.19	5.02	3.85	4.22	4.99	6.11	5.31	4.83	5.35	5.83
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
Ljung-Box Q(10)	2.41	8.22	11.80	11.55	10.58	9.25	10.42	11.29	6.66	9.19	11.18	9.73	10.02	10.15	10.45	9.59	8.48	10.66	10.41
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
ARCH-LM(5)	4.69	0.94	1.93	0.58	1.78	2.01	0.39	2.47	2.82	1.76	0.99	0.92	1.47	0.97	1.30	1.26	1.37	0.61	1.62
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
ARCH-LM(10)	8.55	1.53	3.55	1.11	3.15	3.23	1.17	3.79	4.43	3.01	1.86	2.05	2.39	2.59	2.16	2.32	1.82	1.43	2.99
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
BDS, $\varepsilon = 1.5\sigma$, $m = 3$	2.71	3.87	3.78	4.61	4.31	4.67	2.06	0.58	2.94	5.93	3.39	3.43	4.20	3.74	4.18	3.37	4.78	4.59	2.07
	[0.66]	[0.67]	[0.77]	[0.80]	[0.70]	[0.70]	[0.47]	[0.30]	[0.50]	[0.67]	[0.60]	[0.60]	[0.67]	[0.53]	[0.63]	[0.73]	[0.80]	[0.87]	[0.50]
BDS, $\varepsilon = 2\sigma$, $m = 3$	2.08	2.55	2.24	2.02	2.78	3.22	1.21	1.70	1.53	4.31	1.93	2.10	3.04	2.16	2.54	2.27	3.27	2.55	1.66
	[0.48]	[0.63]	[0.50]	[0.43]	[0.53]	[0.53]	[0.40]	[0.37]	[0.40]	[0.53]	[0.37]	[0.43]	[0.53]	[0.43]	[0.50]	[0.53]	[0.67]	[0.60]	[0.33]
$V(180)$	1.31	1.23	1.25	1.32	1.16	1.21	0.96	1.35	1.13	1.36	1.05	1.14	1.08	1.10	1.23	1.19	1.12	1.28	1.14
	[0.00]	[0.00]	[0.00]	[0.07]	[0.00]	[0.00]	[0.07]	[0.07]	[0.03]	[0.00]	[0.10]	[0.00]	[0.03]	[0.00]	[0.00]	[0.03]	[0.00]	[0.00]	[0.00]
$V(270)$	1.32	1.21	1.22	1.31	1.16	1.21	1.15	1.48	1.15	1.39	1.09	1.15	1.09	1.08	1.23	1.17	1.11	1.29	1.16
	[0.00]	[0.00]	[0.03]	[0.13]	[0.00]	[0.00]	[0.00]	[0.07]	[0.03]	[0.00]	[0.00]	[0.00]	[0.13]	[0.00]	[0.00]	[0.03]	[0.00]	[0.00]	[0.00]
$V_s(180)$	0.30	0.51	0.58	0.98	0.54	0.48	0.45	0.48	0.43	0.53	0.43	0.47	0.51	0.48	0.52	0.52	0.55	0.59	0.34
	[0.92]	[0.90]	[0.83]	[0.10]	[0.80]	[0.73]	[0.90]	[0.53]	[0.87]	[0.77]	[0.90]	[0.77]	[0.80]	[0.77]	[0.70]	[0.83]	[0.83]	[0.87]	[0.93]
$V_s(270)$	0.25	0.44	0.49	0.87	0.46	0.41	0.38	0.44	0.36	0.46	0.37	0.40	0.44	0.41	0.45	0.44	0.47	0.51	0.28
	[0.92]	[0.90]	[0.83]	[0.27]	[0.83]	[0.80]	[0.90]	[0.53]	[0.93]	[0.77]	[0.90]	[0.80]	[0.80]	[0.80]	[0.73]	[0.90]	[0.83]	[0.90]	[0.93]

Notes. The column labeled "Actual" presents results for a sample series of daily with-dividend common stock returns for 50 common stocks traded on the NYSE. Data are from the CRSP archive files. All other columns are for the cases of the artificial stock market listed in Table 3, where each case represents a different selection of values for the parameters of the model. The test statistics and the respective null hypotheses being tested for the tests BDS, Ljung-Box Q(q), ARCH-LM(q), $V(q)$, and $V_s(q)$ are described in §4 of the text. The ARMA-TARCH model is described in §5 of the text. Fractions of tests rejecting the null hypothesis are in square brackets.

to the empirical experiments, followed by a discussion of how those results compare to those found for the actual stock returns.

6. The Model: Learning, Price Predictions, and Demands

The model begins with a dividend, d_t, announced publicly at time t, and which follows an autoregressive process of the form

$$d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + \upsilon_t, \quad (4)$$

where υ_t is normally distributed with zero mean and constant variance σ_υ^2 (Gaussian noise). Equation (4) is, of course, an abstraction because dividends are not paid continuously. One justification for the specification used is that while dividends are not paid continuously, information that can influence beliefs about firms does arrive on a more frequent basis. We regard the dividend arrival as a proxy for information arrival. A dividend (information) shock may lead to an over- or underreaction, which is then corrected via the regression to the mean process built into the model. Under the partial adjustment condition $0 < \rho < 1$, model (4) exhibits negative feedback.

Agents observe past prices and the dividend. The forecasting model employed by any agent k is assumed to be linear in past prices and dividends and is defined by

$$\hat{E}_{k,t}[p_{t+1} + d_{t+1}] = a_{k,j}(p_t + d_t) + b_{k,j}, \quad (5)$$

where $a_{k,j}$ and $b_{k,j}$ are parameters selected by agent k based on one of the hypotheses j in his set of hypotheses. Agents in the model hold multiple hypotheses for computing the parameter values of Equation (5). The hypothesis selected as a basis for the choice of the parameters is the most accurate hypothesis held by the agent based on the performance of the hypothesis in predicting the price plus dividend during the immediate past. Forecast precision is measured by the inverse of

$$e_{k,j,t}^2 = (1-\theta)e_{k,j,t-1}^2 + \theta[(p_t + d_t) - E_{k,j,t-1}(p_t + d_t)]^2, \quad (6)$$

where k indexes the individual agent and j indexes a hypothesis held by individual k.[19] We set $\theta = 0.02$. The variable $e_{k,j,t}$ is the deviation of the actual price plus dividend from the price plus dividend prediction made by individual k using hypothesis j for period t.

Hypotheses are modified, rejected, or carried forward based on what will be referred to as their fitness in a fashion akin to induction.[20] The fitness of each hypothesis is calculated as $f_{k,j,t} = -e_{k,j,t}^2 - \beta s$, where β is a constant and s is the number of information bits utilized by the hypothesis.[21] The fitness measure imposes higher costs on hypotheses that produce larger squared forecast errors and that employ a greater amount of information.

Agents actively engage in the generation of new hypotheses through what we will call *combination experiments* and *individual experiments*. Combination experiments involve combining elements of pairs of accurate hypotheses (high fitness) in an attempt to discover even more accurate hypotheses. Such activity is assumed to occur with a fixed probability π, the probability of a combination experiment. A new hypothesis so generated then replaces (with high probability) a low-accuracy hypothesis amongst the set of hypotheses held by the agent. Individual experiments involve the modification of already existing hypotheses by alteration of the parameters associated with the hypothesis. Individual experiments occur with probability $(1 - \pi)$. Therefore, individual experiments only occur when combination experiments do not. Agents revise their hypotheses every τ periods.[22] We define τ as the learning frequency. Hypotheses are selected for change based on their fitness values. Stochastic universal sampling is used when selecting hypotheses to combine because this method allows the probability of selection to depend directly on the fitness value. Suppose that a combination experiment is to be implemented. Once selection of two hypotheses has occurred, the method of uniform crossover is employed to arrive at a new hypothesis that combines the two old hypotheses. The selection of the hypothesis to be replaced by the new hypothesis is also implemented by invoking stochastic universal sampling, in which low fitness value hypotheses will be replaced with high probability. The new, untested hypotheses that are created will not, in and of themselves, cause disruptions because they will be acted on only if they prove to be accurate. This avoids brittleness and provides what machine-learning theorists call "gracefulness" in the learning process. A detailed outline of

[19] We set $\theta = 0.02$ in the simulations. Similar specifications are employed by LeBaron et al. (1999) and Tay and Linn (2001).

[20] The argument that individuals will form their expectations by induction in ill-defined environments has been suggested as an alternative to the deductive model usually invoked (for instance, Arthur 1991, 1992, 1994, 1995; Arthur et al. 1997; Blume and Easley 1995; LeBaron et al. 1999; Rescher 1980). Induction is a means of finding the best available answers to questions that transcend the information at hand.

[21] We set $\beta = 0.000005$. Similar specifications are employed by LeBaron et al. (1999) and Tay and Linn (2001).

[22] Learning by induction is modeled using a genetic algorithm (Holland and Reitman 1978, Holland et al. 1986, Goldberg 1989). The processes of recombining and experimentation are referred to as *crossover* and *mutation* in the literature on genetic algorithms.

the process invoked in combination and individual experiments is provided in the online supplement to this paper.

At each date an agent assigns values to the parameters $(a_{k,j}, b_{k,j})$, $j = 1, \ldots, J$, where j denotes the hypothesis j held by the individual k. The values for each $(a_{k,j}, b_{k,j})$ pair are determined through the application of fuzzy logic.[23]

Let X denote a universe of objects and x denote an individual element in that universe. Let A be a subset of X. Define a membership function

$$\mu_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{otherwise}. \end{cases}$$

The value assigned by the function $\mu_A(x)$ for a given x represents the degree of membership of x in A. Classical set theory assigns the discrete values one or zero to $\mu_A(x)$ depending on whether the first or second conditions are met.[24] Fuzzy set theory expands the domain of the membership function to include all numbers in the real interval $[0, 1]$. Intuitively, the larger the value of $\mu_A(x)$, the more $x \in A$. Fuzzy logic deals with propositions of the form $x \in A$, where A is a set and the degree of truth of the proposition is given by $\mu_A(x)$.

The agent observes a set of market descriptors, which we call *fuzzy input variables* and will refer to as *information bits*. Five market descriptors ($p * r/d$, $p/MA(5)$, $p/MA(10)$, $p/MA(100)$, and $p/MA(500)$) are computed. The variables r, p, and d are the interest rate, price, and dividend, respectively. The variable $MA(n)$ denotes an n-period moving average of prices. Thus, the first information bit reflects the current price in relation to the current dividend and is a "fundamental" bit. The remaining four bits are "technical" bits indicating whether the price history exhibits a trend or similar characteristic. The fuzzy output variables are the parameters $(a_{k,j}, b_{k,j})$.

Each agent holds five hypotheses ($J = 5$). Each of these hypotheses contains four rules. A rule maps the information from the inputs to the outputs via a logical statement of the form "if ⟨fuzzy proposition⟩, then ⟨fuzzy proposition⟩," for example, If *fund* is x_1 and *tech*1 is x_2 and *tech*2 is x_3 and *tech*3 is x_4 and *tech*4 is x_5, then a is y_1 and b is y_2. The symbols x_i and y_i are linguistic statements of the form "low, moderately low, moderately high, high." We define $x_1, x_2, x_3, x_4, x_5 \in \{0, 1, 2, 3, 4\}$ and $y_1, y_2 \in \{1, 2, 3, 4\}$, where the codes 1, 2, 3, 4 correspond to the characteristics (low, moderately low, moderately high, high) and the code 0 implies that the information bit is not used in the evaluation of the rule.

The agent begins by resolving all fuzzy statements in the antecedent of each rule of a hypothesis to a degree of membership between zero and one, where the membership functions are described by specific functional forms.[25] For instance, the rule might say "if x is high, then y is high." The first step is to assign a number to the degree of membership that x belongs to the fuzzy set "high." If x is one and the lowest bound of the set "high" is five, then a value for the degree of membership of x in the high set would be zero. If there is only one part to the antecedent, this is the degree of support for the rule. If the rule is of the form "if x is 'high' and z is 'moderate' then y is 'high,'" then there are two antecedents, x and z, and we must establish the degree of membership of each in the respective sets ("high" for x and "moderate" for z). Once the membership function values are established, they must be aggregated into a measure of overall support. The rules in our model are specified like the illustration with the logical connective "and." We therefore seek the intersection of the support levels of each of the antecedents. We thus apply the intersection operator in such situations (Zadeh 1965). In our example with two antecedents, suppose that the membership value for x in "high" was 0.5 and for z in "moderate" was 0.3. The intersection operator would conclude that the overall support for the rule was 0.3. This produces a number that represents the degree of support for the antecedent of the rule as a whole.

The agent then uses the degree of support for the antecedents of the entire rule, a number between zero and one, to determine a membership value for the output variable. In the above example "if x is 'high' and z is 'moderate' then y is 'high,'" the degree of support for the rule was 0.3. This defines an area of the membership function "high" for y. This in turn produces a "mass," as can be seen in the illustration provided in Figure 1 under either the columns labeled a or b, which are the output variables. The masses assigned to y by each rule of a hypothesis are then aggregated using what is referred to as the *centroid method*.

[23] Zadeh (1962, 1965) developed the original exposition on fuzzy logic. Smithson (1987) and Smithson and Oden (1999), amongst others, present evidence on reasoning, and the human thought process that suggests the assertion that individuals reason as if by the axioms of fuzzy logic is supported.

[24] A paradox, however, arises within classical set theory regarding the assignment of membership. Consider the following: *A pile of sand containing one particle is small. Adding one additional particle the pile remains small. Therefore every pile is small (by induction).* This illustration clearly poses a dilemma because our natural reaction would be that not every pile is small.

[25] In the simulations, the input variables are associated with triangular membership functions for the sets moderately low and moderately high, and with trapezoidal-shaped membership functions for the sets high and low. The output variables are associated with Gaussian-shaped membership functions for each of the sets low, moderately low, moderately high, and high.

Figure 1 **Illustration of the Process by Which an Agent Maps Information into Conclusions About the Values of the Parameters *a* and *b* Using a Fuzzy Logic Reasoning System**

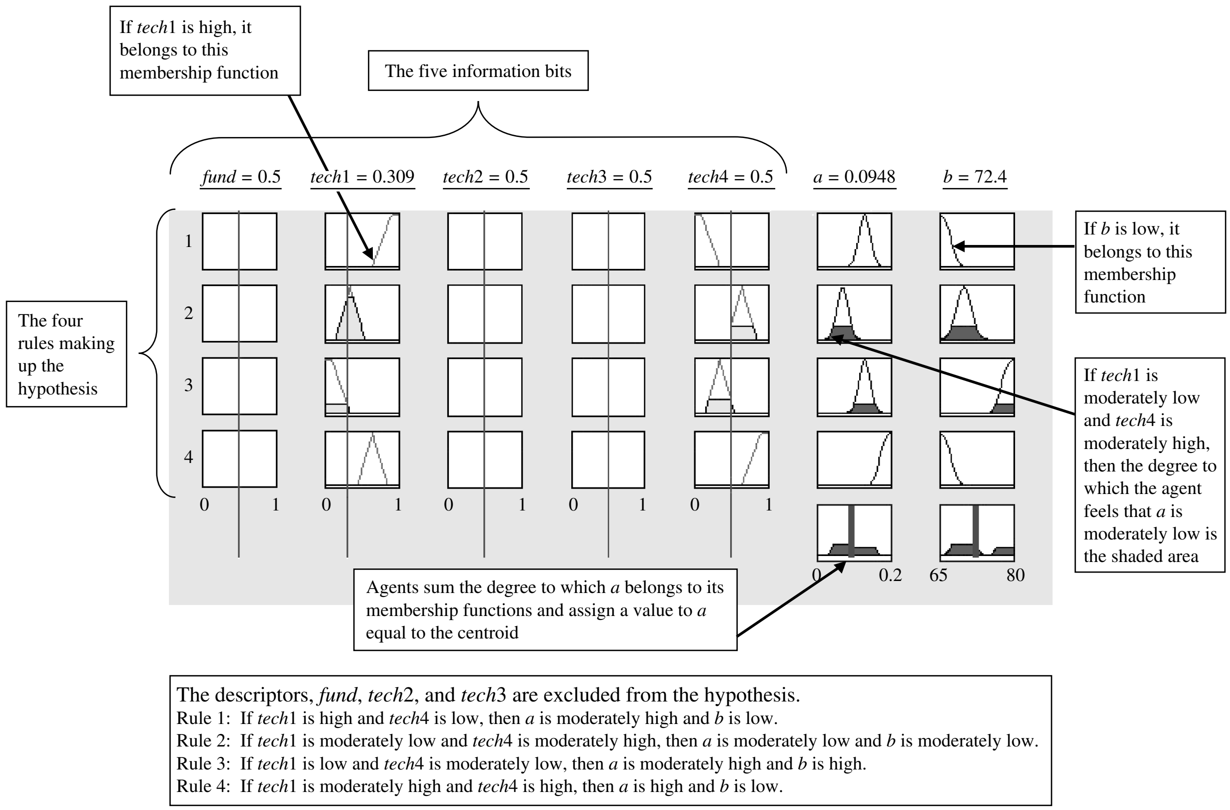

Each agent forecasts the next period's price and dividend ($\hat{E}_{k,t}[p_{t+1}+d_{t+1}]$) using the forecast parameters from the rule base (hypothesis) in her set that has proven to be the most recently accurate. The hypothesis that has performed best in terms of the moving-average squared forecast error described earlier is the one selected as the basis for making a prediction about the next period's price + dividend.

The share demand by agent k at time t equals

$$x_{k,t} = \frac{\hat{E}_{k,j,t}[p_{t+1}+d_{t+1}] - p_t(1+r)}{\lambda \hat{\sigma}^2_{k,j,t,p+d}}, \tag{7}$$

where p_t is the price of the risky asset at time t, λ is the degree of risk aversion, r is the relevant risk-free interest rate for the time horizon, and the expectation (prediction) and variance estimate are conditional on hypothesis j, the most recently accurate of the hypotheses held by the agent.[26]

Agents in the model know that Equation (7) will hold in a homogeneous rational expectations equilibrium when the degree of risk aversion is constant across individuals. However, the fact that they must use induction to form and modify hypotheses and that they use fuzzy rules when forming expectations means that they never know if the market is actually in equilibrium. We assume that agents choose to use (7) when setting their demands, knowing that sometimes the market will be in equilibrium and that sometimes it will not. Each agent is endowed with one share, and hence the market-clearing condition is

$$\sum_{k=1}^{N} x_{k,t} = N. \tag{8}$$

The market-clearing price, p_t, is found by summing Equation (7) over all agents and then setting the sum equal to N, the number of shares available. Once the market clears, the price and dividend at time t are revealed and the accuracies of the rule bases are updated.

[26] The optimal demand function is derived from the first-order condition of expected utility maximization of agents with exponential utility of consumption under the condition that the random variable of interest is normally distributed. However, when the distribution of stock prices is non-Gaussian, the above connection to the maximization of expected utility under an exponential utility function no longer exists, so in those cases we simply take the demand function in (7) as given. The moving average of the squared forecast error for the hypothesis selected (Equation (6)) serves as the estimate of the variance.

Table 3 Parameter Values of the Artificial Stock Market Model

	Learning frequency	$\bar{d}$	ρ	σ_v^2	Risk-free interest rate	λ	Pr(Comb) π	Pr(Indiv) $(1-\pi)$	N	Number of hypotheses (rule bases)
Case 1	30	0.0137	0.5	0.0005	Actual	0.50	0.2	0.8	25	5
Case 2	30	0.0137	0.5	0.0005	Actual	0.50	0.5	0.5	25	5
Case 3	30	0.0137	0.5	0.0005	Actual	0.50	0.8	0.2	25	5
Case 4	30	0.0137	0.1	0.0005	Actual	0.50	0.2	0.8	25	5
Case 5	30	0.0137	0.9	0.0005	Actual	0.50	0.2	0.8	25	5
Case 6	10	0.0137	0.5	0.0005	Actual	0.50	0.2	0.8	25	5
Case 7	1,000	0.0137	0.5	0.0005	Actual	0.50	0.2	0.8	25	5
Case 8	30	0.0137	0.5	0.0005	Actual	0.50	0.2	0.8	10	5
Case 9	30	0.0137	0.5	0.0005	Actual	0.50	0.2	0.8	50	5
Case 10	30	0.0137	0.5	0.0005	Actual	0.50	0.2	0.8	25	3
Case 11	30	0.0068	0.5	0.0005	Actual	0.50	0.2	0.8	25	5
Case 12	30	0.0041	0.5	0.0005	Actual	0.50	0.2	0.8	25	5
Case 13	30	0.0137	0.5	0.0003	Actual	0.50	0.2	0.8	25	5
Case 14	30	0.0137	0.5	0.0007	Actual	0.50	0.2	0.8	25	5
Case 15	30	0.0137	0.5	0.0005	+10% shift	0.50	0.2	0.8	25	5
Case 16	30	0.0137	0.5	0.0005	+10% shift	0.50	0.2	0.8	25	5
Case 17	30	0.0137	0.5	0.0005	Actual	0.10	0.2	0.8	25	5
Case 18	30	0.0137	0.5	0.0005	Actual	0.90	0.2	0.8	25	5

Notes. The table lists the cases investigated in the simulation of the artificial stock market and the parameters varied across the cases. Learning frequency τ: the number of periods between the dates on which any agent updates his hypotheses; $\bar{d}$: the mean dividend (Equation (4)); ρ: adjustment factor for the dividend in the dividend-generating process (Equation (4)); σ_v^2: variance of error in the dividend-generating process (Equation (4)); λ: coefficient of risk aversion for each agent; Pr(Comb) $= \pi$: probability of combination experimentation; Pr(Indiv): probability of individual experimentation ($=1-$ Pr(Comb) $= (1-\pi)$); N: number of agents; number of rule bases: number of hypotheses about the future course of the (price + dividend) by each agent. Each rule base contains four rules used in the construction of the two parameters needed for predicting the next period's (price + dividend) from information on five market-determined variables (five information bits) observed by all agents. Agents represented in the model employ induction and reason as if by fuzzy logic when forming their expectations. The process is modeled as a genetic-fuzzy classifier system. We use the daily one-year T-bill rates in the secondary market for the 5,000 days ending September 5, 2001 when computing demands using Equation (7). In the simulation, we divide this interest rates series by 365 to obtain the approximate daily interest rates. The probability of combination experimentation is the probability that elements of two hypotheses will be split and combined. The probability of individual experimentation is the probability that an agent will have one of his rule bases subjected to random change. When a particular rule base is selected for experimentation, the probability that any individual information bit is changed equals 0.5.

7. The Market Experiments

We examine the implications of our conjecture about the learning and reasoning processes for the underlying hidden structure of the risky security's returns. We present an empirical comparative statics analysis emphasizing how the results change as a consequence of changing the parameter values of the model as well as how the results compare with those documented in Tables 1 and 2 for the 50 sample series.[27]

We examine 18 different combinations of parameter values, each of which we refer to as a *case*, where we use Case 1 as the base for comparison purposes. The parameters of the model and the values examined for each case are tabulated in Table 3. Specifically, we vary the following parameters of the model: learning frequency τ, the mean of the dividend process $\bar{d}$, the adjustment coefficient in the dividend process ρ, the variance of the dividend process σ_v^2, the level of the risk-free interest rate r, the risk-aversion parameter λ, the probability of a combination experiment π and, by implication, the probability of individual experimentation $1-\pi$, the number of agents in the model N, and the number of rule bases (hypotheses) that an agent can hold.

We use the time series of equivalent daily risk-free interest rates computed from the series of one-year U.S. T-bill rates for the 5,000 days ending September 5, 2001 as our proxy for the risk-free interest rate series used in computing demands. Tests of the sensitivity of the results to changes in the risk-free rate involve either a 10% shift up or a 10% shift down in the entire series of daily rates. We keep some parameters fixed throughout, including the number of information bits (market data items) available, set at five, the number of forecast parameters, set at two, and the number of rules in any rule base, set at four. When a particular rule base is selected for experimentation, the probability that any individual information bit is changed equals 0.5.

Thirty simulations are conducted for each case represented in Table 3. The simulations for any case are independent only in the sense that the seeds of the algorithms used to generate pseudorandom variables differ across the simulations for each case. This provides us with 30 separate return series for each case, allowing us to compute estimates of various descriptive statistics as well as test statistics for each of the

[27] The model was coded and simulated using MATLAB, a product of MathWorks, Inc.

30 trials. We believe that 30 observations should provide a good indication of the typical behavior of the model under each assumed set of parameter values.

We began each simulation with a random initial configuration of rules. We then simulated the market for 2,500 periods to allow any asymptotic behavior to emerge. Subsequently, starting with the configuration attained at $t = 2{,}500$, we simulated an additional 2,500 periods to generate price and dividend data for the statistical analyses discussed in the next section. The returns examined in the next section utilize the latter 2,500 price and dividend observations and are computed as

$$r_t = \frac{p_t + d_t}{p_{t-1}} - 1. \tag{9}$$

8. Simulation Results

The columns of Tables 1 and 2 present summary results for the simulation cases described in Table 3. The first column of the table presents the results for the 50 sample series of actual stock returns. The format of the summary data for each simulation case is the same as the format for the summary statistics of the actual stock returns discussed earlier. Case 1 serves as a benchmark. The remaining cases vary the values of one of the parameters of Case 1. Inspection of the results for Case 1 indicates general agreement with the results shown for the actual returns in Column 1. Case 1 produces a larger kurtosis and slightly more positive skewness. The ARCH-LM test rejects the null hypothesis for each of the Case 1 series. Kurtosis may be due to either fat tails, peakedness, or a combination of both. The average estimate of the tail index using the estimator of Quintos et al. (2001) is of the same order of magnitude as the average for the actual returns with a standard error that is also of the same magnitude. Like the actual series, the BDS test rejects the null hypothesis of i.i.d. in all 30 cases. The long memory tests for the Case 1 series tend toward not rejecting the null hypothesis of no long memory (17% of the tests reject the null). This is still greater than what is observed for the actual data. However, the null hypothesis of no long memory in the squared series is rejected roughly the same fraction of times for the Case 1 data as for the actual data. Turning to Table 2, we see that the Q tests and the ARCH-LM tests computed using the standardized residuals of the ARMA-TARCH models fit to the data never reject the null hypotheses for the Case 1 series. As is true, however, for the actual return series, we see that the BDS tests reject the null of i.i.d. roughly two-thirds of the time. Finally, the V test of Lo (1991) tends to reject no long-term memory in the squared standardized residual series, indicative of long memory in the volatility process. We conclude that Case 1 does a reasonable job of producing characteristics similar to actual daily returns.

8.1. Learning Frequency

Learning frequency refers to the frequency at which hypothesis revisions occur. The learning frequency parameter is set as the number of periods between learning events. A low number implies frequent learning (short periods of time between learning events), and that agents will revise their rule bases (hypotheses) more often. When learning is frequent, hypotheses are more likely to be influenced by transient behavior in the time series of market variables. In contrast, when learning is infrequent, agents will have more time between revising their rule bases to test their hypotheses. Furthermore, hypotheses will also tend to be based on longer horizon features in the time series of market variables when learning is infrequent. We test three settings for the learning parameter, fast learning (10 periods between learning events, Case 6), less frequent learning (30 periods, Case 1), and slow learning (1,000 periods, Case 7). All other parameter values remain the same across the three cases.

The results for Case 6 are not vastly different than those for Case 1. However, Case 7, in which learning occurs infrequently, is associated with substantially smaller means for standard deviation and skewness, much higher kurtosis, and a Quintos' tail index average of roughly 3.8. Further, Case 7 is associated with long memory in both the level of returns as well as the volatility of returns, in contrast to the findings for the actual data as well as Cases 1 and 6. Slow learning of the magnitude tested does not appear to be a reasonable assumption.

8.2. The Dividend Process

Three factors impact the dividend in any period, $d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + v_t$: the mean dividend $\bar{d}$, the partial adjustment factor ρ, and the variance of the error term σ_v^2 (through realizations of the error). We test two settings each for these parameters, not including the settings for Case 1. Cases 11 and 12 test two values of the mean dividend, which are larger than the base case used in Case 1. The results, as compared to the actual stock returns and to Case 1, are similar along most characteristics and tests. Kurtosis is slightly more elevated, but the tail index average is in the vicinity of the results for the actual data and for Case 1. Cases 4 and 5 test a lower and higher value for the parameter ρ, 0.1 and 0.9, where the value selected for Case 1 was 0.5. The results again are similar to those for Case 1, suggesting a partial adjustment parameter bracketed by 0.1 and 0.9 is a reasonable assumption. Cases 13 and 14 test two alternative values for the variance of the error in the dividend process, bracketing the assumption used in Case 1. On an annualized basis, the bracket spans roughly 11% to 25.5%, not vastly dissimilar to what is generally observed on an annual basis. Cases 13 and 14 also show little variation from what is observed for Case 1. We conclude that values within the ranges described generate plausible results.

8.3. The Risk-Free Interest Rate and Risk Aversion

The risk-free interest rate and the coefficient of risk aversion influence the demand for the risky asset by any agent (Equation (7)). The structure of the model assumes that the values of these parameters are common across agents. We use the actual daily interest rate in the simulations, so this variable, while common to all agents, varies across dates. The coefficient of risk aversion is assumed to be fixed across agents and across dates. Cases 15 and 16 reflect a 10% shift up and down, respectively, of the entire series of interest rates. Again, we see little variation between these results and those for Case 1. Cases 17 and 18 vary the risk-aversion parameter, 0.1 (Case 17) and 0.9 (Case 18). Case 18 stands out. The standard deviation increases significantly, while kurtosis is smaller relative to Case 1 as well as Case 17. Further, the average tail index for both cases falls. The test results for Cases 17 and 18 also suggest the presence of long memory. However, as Table 2 shows, once ARMA-TARCH models are fit to the series associated with these cases, the tests do not support long memory in the level of returns, but do indicate long memory in the squared residuals.

8.4. Probability of Experimentation

Experimentation in the model is the process by which agents form new hypotheses about the (price + dividend). On those dates when learning occurs and agents revise their hypotheses, new hypotheses are formed by combining existing hypotheses and by altering individual existing hypotheses. The decision about which method for revising hypotheses will be invoked is driven by the probability that a combination experiment will be used, $\Pr(\text{Comb}) = \pi$. The probability of an individual experiment equals $1 - \Pr(\text{Comb})$. Cases 2 and 3 test two values, 0.5 and 0.8, for Pr(Comb). Cases 2 and 3 generate results that are quite different from either the actual data or Case 1. The results presented in Table 1 show that both cases are associated with much smaller tail index averages and evidence of long memory in the level of returns. The results presented in Table 2 show that Case 2 exhibits evidence of long memory in squared returns but not in the level of returns—results that are consistent with what we observe for the actual data and for Case 1. On the contrary, however, the results for Case 3 exhibit little evidence of long memory in the squared returns, but there is evidence of long memory in the level of returns. Hence, between these two cases, Case 2 with Pr(Comb) of 0.5 yields results that more closely resemble the actual data and Case 1. The suggestion is, therefore, that values for Pr(Comb) in the range 0.2 to 0.5 yield results generally consistent with actual data.

8.5. The Number of Agents

The number of agents in the model influences the market-clearing price through aggregate net demand. Case 8 tests a model in which there are only 10 agents, while Case 9 tests a model with 50 agents. In Case 1, we assume that there are 25. Generally, the results for Case 8 more closely resemble the results for the actual returns than do the results for Case 9. In particular, Case 8 is associated with a measure of kurtosis and an average tail index that are more like the measures associated with actual stock returns than with the Case 1 returns. Evidence of long memory in raw returns is more pronounced for the Case 9 returns—inconsistent with the results for the actual stock return data. We conclude that a parameter setting for the number of agents bracketed by 10 and 25 produces results generally consistent with actual data. These results suggest that although participants in real markets are heterogeneous, they may best be described as belonging to groups, the number of which may be reasonably small.

8.6. The Number of Rule Bases

Each agent holds a set of hypotheses about the parameters of the prediction model. These hypotheses are in the form of rule bases that are used to translate market data (information bits) into values for the parameters of the prediction model. Case 10 tests a model in which each agent is allowed only three rule bases, in contrast to the five rule bases allowed in Case 1. The differences between the results for Case 10 and those for Case 1 are inconsequential. Hence, we conjecture that models based on three to five rule bases will produce reasonably similar results.

8.7. ARMA-TARCH Models

We conclude with a brief review of Table 2. Table 2 reports results based on the standardized residuals of ARMA-TARCH models fit to the series associated with the cases listed in Table 3. First, note that in every case the Ljung-Box Q tests, as well as the ARCH-LM tests, never reject the null hypotheses associated with those tests. Essentially, the models remove all short-run autocorrelation in the level of returns, as well as filtering out any ARCH-type behavior. Nevertheless, the BDS statistics reveal that in at least 50% of the trials, the BDS test rejects the null hypothesis of i.i.d. This suggests that some other form of dependence is still present. The tests for long memory in the level of returns overwhelmingly do not reject the null hypothesis of no long memory. However, the tests for long memory in the squares of the standardized residuals suggest that in a large fraction of the cases, long-memory dependence is present.

8.8. Summary

Overall, the results reported in Tables 1 and 2 reveal that only four of the 10 parameters investigated are really critical in the model. These parameters are (a) learning frequency, (b) risk aversion, (c) probability of combination, and (d) number of agents. Conditional on a low probability of combination experimentation in the creation of hypotheses, the results for Case 7 suggest that the simulated market results are not consistent with infrequent revision of hypotheses (a slow learning frequency). Rather, as Cases 1 and 6 suggest, fast to moderate learning is more consistent, all else equal, with the actual data. Cases 17 and 18 involve low and high risk aversion, respectively, both scenarios producing long memory, smaller kurtosis, and smaller tail index estimates than the actual data. Case 3, which involves a high probability of combination experimentation in the formulation of new hypotheses, yields results that are generally not consistent with actual data. Everything else being the same, the conclusion that a low rate of combination experimentation produces results more like the actual data is consistent with the notion that agents are less inclined to create new hypotheses from crossbreeding their existing hypotheses. Case 9 involves many agents and produces long memory in returns, suggesting that the greater the number of agents that are active in the market, the longer it takes for adjustments to be manifested.

9. Conclusions

This paper begins by confirming several characteristics of actual daily with-dividend stock returns: (1) A power-law tail index close to three describes the behavior of the positive tail of the survivor function of returns ($\mathrm{pr}(r > x) \sim x^{-\alpha}$) (Gopikrishnan et al. 1999, Plerou et al. 1999), a reflection of fat tails. (2) General linear and nonlinear dependencies exist in the time series of returns (Scheinkman and LeBaron 1989, Hsieh 1991, Brock et al. 1991). (3) The time-series return process is characterized by short-run dependence (short memory) in both returns as well as their volatility, the latter usually characterized in the form of autoregressive conditional heteroskedasticity (Bollerslev et al. 1992, Glosten et al. 1993, Engle 2004). (4) The time-series return process probably does not exhibit long memory (Lo 1991), but the squared returns process does exhibit long memory (Ding et al. 1993, Bollerslev and Wright 2000). We then go on to present an alternative model of learning and reasoning behavior in capital markets. The model environment in which investors operate is complex and ill defined. Agents learn by induction and the application of fuzzy logic. We assert that models endowing agents with such learning and reasoning processes may account for some of the documented empirical characteristics found in the actual stock returns of our benchmark sample. As such, we embed the learning and reasoning process in an artificial stock market model and conduct dynamic simulation experiments to generate market-clearing prices for a risky security. We compute implied returns using the prices and dividends generated by the experiments and go on to analyze these returns using the same methods applied to our benchmark sample. We find that the characteristics of the returns from our experimental market conform to those for the benchmark sample and that the results are relatively insensitive to variation in the parameter values of the model.

The framework of the model offers an alternative perspective on what generates the behavior of financial security returns that extends beyond the traditional paradigms. A useful extension of our work would be to marry the learning and reasoning process we propose with a model of the structure of trading similar to that proposed in Gabaix et al. (2006). We leave that endeavor for future research.

10. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Acknowledgments

An earlier version of this paper circulated under the title "Fuzzy Inductive Reasoning and Nonlinear Dependence in Security Returns: Results from an Artificial Stock Market Environment." The authors are grateful to numerous colleagues who have provided comments on this and their related work, and especially to the three reviewers for this journal for their many insightful and useful comments.

References

Arthur, W. B. 1991. Designing economic agents that act like human agents: A behavioral approach to bounded rationality. *Amer. Econom. Rev.* **81** 353–359.

Arthur, W. B. 1992. On learning and adaptation in the economy. Working Paper 92-07-038, Santa Fe Institute, Santa Fe, NM.

Arthur, W. B. 1994. Inductive reasoning and bounded rationality. *Amer. Econom. Rev.* **84** 406–411.

Arthur, W. B. 1995. Complexity in economic and financial markets. *Complexity* **1** 20–25.

Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer, P. Tayler. 1997. Asset pricing under endogenous expectations in an artificial stock market. W. B. Arthur, S. Durlauf, D. Lane, eds. *The Economy as an Evolving Complex System II.* Addison-Wesley, Reading, MA, 15–44.

Barnett, W. A., A. R. Gallant, M. J. Hinich, J. A. Jungeilges, D. T. Kaplan, M. J. Jensen. 1997. A single blind controlled competition among tests for nonlinearity and chaos. *J. Econometrics* **82** 157–192.

Blume, L. E., D. Easley. 1995. What has the rational learning literature taught us? A. Kirman, M. Salmon, eds. *Learning and Rationality in Economics*. Basil Blackwell, Cambridge, MA, 12–39.

Bollerslev, T., J. H. Wright. 2000. Semoparametric estimation of long memory volatility dependencies: The role of high frequency data. *J. Econometrics* **98** 81–106.

Bollerslev, T., R. Y. Chou, K. F. Kroner. 1992. ARCH modeling in finance: A review of the theory and empirical evidence. *J. Econometrics* **52** 5–59.

Brock, W. A. 1982. Asset prices in a production economy. J. McCall, ed. *The Economics of Information and Uncertainty*. Conference Report 32, National Bureau of Economic Research, University of Chicago Press, Chicago, IL, 1–43.

Brock, W. A., C. H. Hommes. 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J. Econom. Dynam. Control* **22** 1235–1274.

Brock, W. A., W. D. Dechert, J. Scheinkman. 1987. A test for independence based upon the correlation dimension. Working paper, University of Wisconsin, Madison, WI.

Brock, W. A., D. A. Hsieh, B. LeBaron. 1991. *Nonlinear Dynamics, Chaos and Instability: Statistical Theory and Economic Evidence*. MIT Press, Cambridge, MA.

Brock, W. A., W. D. Dechert, B. LeBaron, J. Scheinkman. 1996. A test for independence based upon the correlation dimension. *Econometric Rev.* **15** 197–235.

Campbell, J. Y., A. W. Lo, A. C. Mackinlay. 1997. *The Econometrics of Financial Markets.* Princeton University Press, Princeton, NJ.

Ding, Z., C. W. J. Granger, R. F. Engle. 1993. A long memory property of stock market returns and a new model. *J. Empirical Finance* **1** 83–106.

DuMouchel, W. H. 1983. Estimating the stable index α in order to measure tail thickness: A critique. *Ann. Statist.* **11** 1019–1031.

Engle, R. F. 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **50** 987–1007.

Engle, R. F. 2004. Risk and volatility: Econometric models and financial practice. *Amer. Econom. Rev.* **94** 405–420.

Feller, W. 1966. *An Introduction to Probability Theory and Its Applications*, Vol. 2. Wiley, New York.

Gabaix, X., P. Gopikrishnan, V. Plerou, H. E. Stanley. 2006. Institutional investors and stock market volatility. *Quart. J. Econom.* **121** 461–504.

Glosten, L. R., R. Jagannathan, D. Runkle. 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *J. Finance* **48** 1779–1801.

Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA.

Gopikrishnan, P., V. Plerou, L. A. N. Amaral, M. Meyer, H. E. Stanley. 1999. Scaling of the distribution of fluctuations of financial market indices. *Phys. Rev. E* **60** 5305–5316.

Gourieroux, C., J. Jasiak. 2001. *Financial Econometrics, Problems, Models and Methods*. Princeton University Press, Princeton, NJ.

Grassberger, P., I. Procaccia. 1983a. Characterization of strange attractors. *Phys. Rev. Lett.* **50** 346–349.

Grassberger, P., I. Procaccia. 1983b. Measuring the strangeness of strange attractors. *Physical Rev.* **9D** 189–208.

Hill, B. M. 1975. A simple general approach to inference about the tail of a distribution. *Ann. Statist.* **3** 1163–1174.

Holland, J. H., J. Reitman. 1978. Cognitive systems based on adaptive algorithms. *Pattern-Directed Inference Systems.* Academic Press, New York.

Holland, J. H., K. J. Holyoak, R. E. Nisbett, P. R. Thagard. 1986. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, MA.

Hsieh, D. A. 1991. Chaos and nonlinear dynamics: Application to financial markets. *J. Finance* **46** 1839–1877.

Hurst, H. E. 1951. Long term storage capacities of reservoirs. *Trans. Amer. Soc. Civil Engineers* **116** 776–808.

Jarque, C. M., A. K. Bera. 1987. A test for normality of observations and regression residuals. *Internat. Statist. Rev.* **55** 163–172.

Kanzler, L. 1999. Very fast and correctly sized estimation of the BDS statistic. Working paper, Department of Economics, University of Oxford, Christ Church, Oxford, UK.

Kaplan, D. T. 1994. Exceptional events as evidence for determinism. *Physica D* **73** 38–48.

LeBaron, B. 2000. Agent based computational finance: Suggested readings and early research. *J. Econom. Dynam. Control* **24** 679–702.

LeBaron, B. 1997. A fast algorithm for the BDS statistic. *Stud. Nonlinear Dynam. Econometrics* **2** 53–59.

LeBaron, B. 2006. Agent-based computational finance. K. L. Judd, L. Tesfatsion, eds. *The Handbook of Computational Economics, II.* Forthcoming.

LeBaron, B., R. Samanta. 2004. Extreme value theory and fat tails in equity markets. Working paper, Brandeis University, Waltham, MA.

LeBaron, B., W. B. Arthur, R. Palmer. 1999. Time series properties of an artificial stock market. *J. Econom. Dynam. Control* **23** 1487–1516.

Ljung, G., G. Box. 1979. On a measure of lack of fit in time series models. *Biometrika* **66** 265–270.

Lo, A. 1991. Long-term memory in stock market prices. *Econometrica* **59** 1279–1313.

Mandelbrot, B. B. 1997. *Fractals and Scaling in Finance: Discontinuity, Concentration and Risk, Selecta Volume E*. Springer-Verlag, New York.

Palmer, R. G., W. B. Arthur, J. H. Holland, B. LeBaron, P. Tayler. 1994. Artificial economic life: A simple model of a stockmarket. *Physica D* **75** 264–274.

Plerou, V., P. Gopikrishnan, L. A. N. Amaral, M. Meyer, H. E. Stanley. 1999. Scaling of the distribution of price fluctuations of individual companies. *Phys. Rev. E* **60** 6519–6529.

Quintos, C., Z. Fan, P. C. B. Phillips. 2001. Structural change tests in tail behavior and the Asian crisis. *Rev Econom. Stud.* **68** 633–663.

Rescher, N. 1980. *Induction: An Essay on the Justification of Inductive Reasoning.* University of Pittsburgh Press, Pittsburgh, PA.

Scheinkman, J. A., B. LeBaron. 1989. Nonlinear dynamics and stock returns. *J. Bus.* **62** 311–337.

Smithson, M. J. 1987. *Fuzzy Set Analysis for Behavioral and Social Sciences.* Springer-Verlag, New York.

Smithson, M. J., G. C. Oden. 1999. Fuzzy set theory and applications in psychology. D. Dubois, H. Prade, eds. *International Handbook of Fuzzy Sets and Possibility Theory: Applications*, Vol. 5. Kluwer, Amsterdam, The Netherlands, 557–585.

Tay, N. S. P., S. C. Linn. 2001. Fuzzy inductive reasoning, expectation formation and the behavior of security prices. *J. Econom. Dynam. Control* **25** 321–361.

Tesfatsion, L. 2002. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life* **8** 55–82.

White, H. 1989. An additional hidden unit test for neglected nonlinearity in multilayer feedforward networks. *Proc. Internat. Joint Conf. Neural Networks*. IEEE Press, New York, 451–455.

Zadeh, L. A. 1962. From circuit theory to system theory. *Proc. Inst. Radio Engineers* **50** 856–865.

Zadeh, L. A. 1965. Fuzzy sets. *Inform. Control* **8** 338–353.