13. Are Efficient Markets Really Efficient?: Can Financial Econometric Tests Convince Machine-Learning People?

Shu-Heng Chen and Tzu-Wen Kuo

AI-ECON Research Center, Department of Economics National Chengchi University, Taipei, Taiwan 11623
email: chchen@nccu.edu.tw, kuo@aiecon.org

Using Quinlan's Cubist, this paper examines whether there is a consistent interpretation of the efficient market hypothesis between financial econometrics and machine learning. In particular, we ask whether machine learning can be useful only in the case when the market is not efficient. Based on the forecasting performance of Cubist in our artificial returns, some evidences seems to support this consistent interpretation. However, there are a few cases whereby Cubist can beat the random walk even though the series is independent. As a result, we do not consider that the evidence is strong enough to convince one to give up his reliance on machine learning even though the efficient market hypothesis is sustained.

13.1 Introduction

A series of applications of *genetic programming* (**GP**) to model agent-based markets has been conducted in [13.1], [13.3], [13.4] and [13.6]. One of the issues typically examined in these studies is whether the agent-based markets are satisfied with the *efficient market hypothesis* (**EMH**). The conjecture that agent-based markets can generate time series data which are satisfied with the EMH was first established in [13.1].

[13.1] simulated *a cobweb model* with GP-based producers and speculators. From their PSC([13.9]) testing of their artificial financial time series, among the 40 series examined, 38 have no linear processes at all, i.e., they are all identified as $ARMA(0,0)$. [1] This result indicates that *the GP-based artificial market is so efficient that there are hardly any linear signals left.* To some extent, this can be considered as a match for the classical version of the EMH.

[13.1]'s result is further strengthened by [13.3] and [13.6] in the agent-based modeling of stock markets. These two papers do not allow traders to imitate other traders directly. Instead, the only way that traders can learn is through a built-in *business school.* With this setup, they find that the stock

[1] One of the by-products of the PSC filter is to inform us of the linear AutoRegressive-MovingAverage process, i.e., the $ARMA(p,q)$ process, extracted from the original series.

S.-H. Chen et al. (eds.), *Computational Intelligence in Economics and Finance*

return series are not only linearly uncorrelated, but also nonlinear uncorrelated or independent.[2] Therefore, in terms of statistical independence, these two markets are even more efficient than the real financial markets.

Thanks to [13.8], there is a *puzzle* about this result. If the return series is indeed so efficient (no signals left at all), then *what is the incentive for traders to search?* In fact, both [13.3] and [13.6] examine traders' profiles and produce two interesting findings. First, most traders are not believers of the EMH, or technically, they are not martingale believers. Second, during the dynamics of the stock markets, there are many traders that show a "successful" search for new ideas.

[13.2] attempt to solve this puzzle by assuming the existence of *brief signals* in an agent-based stock market. They employ the technique known as the *complexity function* to detect whether there are brief signals in the series, and find that brief signals are very weak as opposed to what we experience from real financial data. Therefore, the puzzle remains unsolved. However, there is another possibility, i.e., while the market is efficient in the *econometric* sense, it may not be efficient in the *machine-learning* sense. Those statistics of market efficiency, such as the BDS statistics, may not be able to tell us how hard one must forecast a series in question. In this paper we shall examine whether machine-learning people can share a similar interpretation as econometricans on the BDS testing results. We can do this by using a promising machine-learning tool. The machine-learning tool chosen in this paper is the well-known Quinlan's **Cubist**.

The rest of the paper is organized as follows. Section 13.2 gives a brief introduction to Quinlan's **Cubist**. Section 13.3 describes the data used in the paper. Experimental results with some discussions are given in Section 13.4. Section 13.5 then leaves the concluding remarks.

13.2 Cubist

The following is a very brief introduction to **Cubist**.

The "**Data Mining with Cubist**" website address is:
http://www.rulequest.com/cubist-info.html, which is also the company page of Professor Quinlan. Tutorials describing and illustrating the use of Cubist can be found on http://www.rulequest.com/cubist-win.html.
The following materials are adapted from the tutorial materials put up on the websites.

Cubist is a tool for generating *rule-based predictive models* from data. The *target value* for prediction in our case is the return for the next period, r_t, and the set of *attributes* are simply the historical returns $\{r_{t-i}\}_{i=1}^{10}$. Cubist's job

[2] In both [13.3] and [13.6], the ARMA(p,q) processes extracted from the return series were all ARMA(0,0). Furthermore, in [13.6], the BDS test failed to reject the null of IIDness in all return series, while it was rejected once in [13.3].

Table 13.1. Data description

High-Frequency Financial Data		
Time Series (Type)	Period	# of Obs. (N_1, N_2)
HSIX (one-minute)	12/1/98-12/31/98	4585 (3475,1100)
HSIX (one-minute)	1/4/99-1/29/99	4695 (3584,1100)
\$ECU/\$US (tick-by-tick)	2/25/99 (8:00)- 2/26/99 (7:59)	11571 (8678, 2893)
\$ECU/\$US (tick-by-tick)	3/1/99 (1:00)- 3/2/99 (0:59)	11206 (8396, 2799)
Artificial Data		
MARKET	High BDS	Low BDS
A	10 (7.02)	1 (4.05)
B	5 (2.69)	1 (-0.18)
C	6 (-1.62)	2 (-0.68)
D	1 (4.92)	4 (3.95)

N_1 refers to the size of in-sample data, and N_2 refers to the size of hold-out sample.

is to find out how to predict r_t from $\{r_{t-i}\}_{i=1}^{10}$. Cubist does this by building a *model* consisting of *rules* (conjunctions of conditions) associated with a *linear regression.* Cubist thus constructs a *piecewise linear model* to predict r_t. See Fig. 13.1 and 13.2 for some illustrations.

Cubist allows those piecewise linear models extracted to be combined with the *nearest neighbor* (**NN**) *models.* For some applications, the predictive accuracy of a rule-based model can be improved by combining it with a nearest-neighbor model. However, Cubist uses an unusual method for combining rule-based and nearest neighbor models.

In this paper we experiment with both versions of Cubist, namely, the one using simple rule-based models (coded as "[Cubist]") and the one using the rule-based models and the NN models (coded as "[Cubist + NN]").

13.3 Data Description

The real financial data considered in this paper are described in Table 13.1. Two types of financial data are employed and they are both high-frequency data. One is the one-minute Hang Seng stock index and the other is the tick-by-tick \$ECU/\$US exchange rate. For both series, we transform the original data into the return series, and then the forecasting is made based on the return series. The time series plot of the return series is given in Fig. 13.3-13.6. The dotted lines in these figures show the cutoff point where the whole sample is divided into the training sample and the hold-out sample.

The artificial data employed in this paper is sampled from the artificial returns of Markets A, B, C, and D. (See [13.7] for a description of the gener-

Market D (High BDS)

Rule 1: (1195 cases)
If Rt_1 <= -5e-006
Then return = 0.00023 + 0.967 Rt_1 - 0.018 Rt_4 + 0.016 Rt_9 - 0.001 Rt_3 - 0.001 Rt_7 + 0.001 Rt_10 - 0.001 Rt_6

Rule 2: (56 cases)
If Rt_1 > -5e-006 and Rt_4 > -5e-006
Then return = 0.00036 - 0.077 Rt_4 - 0.037 Rt_5 + 0.031 Rt_9 + 0.028 Rt_10 - 0.027 Rt_8 - 0.025 Rt_3 - 0.024 Rt_6 - 0.017 Rt_2 + 0.014 Rt_1 - 0.006 Rt_7

Rule 3: (249 cases)
If Rt_1 > -5e-006 and Rt_4 <= -5e-006
Then return = 0.01059 + 0.777 Rt_4 - 0.334 Rt_5 - 0.241 Rt_3 + 0.206 Rt_9 - 0.195 Rt_8 + 0.177 Rt_10 - 0.143 Rt_6 - 0.134 Rt_2 - 0.071 Rt_7 + 0.052 Rt_1

Market D (Low BDS)

Rule 1: (1227 cases)
If Rt_1 <= -4.1e-005
Then return = 0.00043 + 1.002 Rt_1 - 0.001 Rt_4

Rule 2: (209 cases)
If Rt_1 > 0.003977 and Rt_4 > -0.004665
Then return = 0.00668 + 0.33 Rt_9 - 0.193 Rt_6 - 0.177 Rt_2 - 0.164 Rt_4 - 0.117 Rt_1 + 0.073 Rt_10 - 0.039 Rt_5 - 0.025 Rt_3 - 0.001 Rt_7

Rule 3: (42 cases)
If Rt_1 > -4.1e-005 and Rt_1 <= 0.003977
Then return = 0.00192 - 0.032 Rt_4 - 0.03 Rt_5 - 0.025 Rt_2 - 0.017 Rt_3 - 0.016 Rt_6 + 0.009 Rt_9 - 0.002 Rt_10 - 0.002 Rt_8 - 0.001 Rt_7

Rule 4: (22 cases)
If Rt_1 > 0.003977 and Rt_4 <= -0.004665
Then return = 0.01899 + 2.134 Rt_4 - 0.697 Rt_5 - 0.322 Rt_3 - 0.246 Rt_2 - 0.136 Rt_6 - 0.124 Rt_1 - 0.083 Rt_10 + 0.054 Rt_9 + 0.026 Rt_8 - 0.001 Rt_7

Fig. 13.1. Decision trees extracted from the artificial data by using Cubist

HSIX (Dec, 98)
Rule : return = 6e-006

HSIX (Jan, 99)
Rule : return = 4e-006 + 0.005 Rt_4

$EUC/$US (2/ 25/ 99 – 2/ 26/ 99)
Rule : return = 0

$EUC/$US (3/ 1/ 99 – 3/ 2/ 99)
Rule : return = -1e-006

Fig. 13.2. Decision trees extracted from the financial data by using Cubist

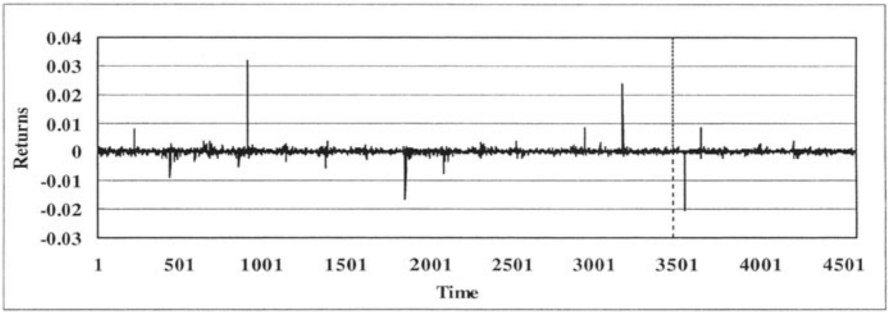

Fig. 13.3. Time series polt of stock returns: one-minute HSIX (12/1/98-12/31/98)

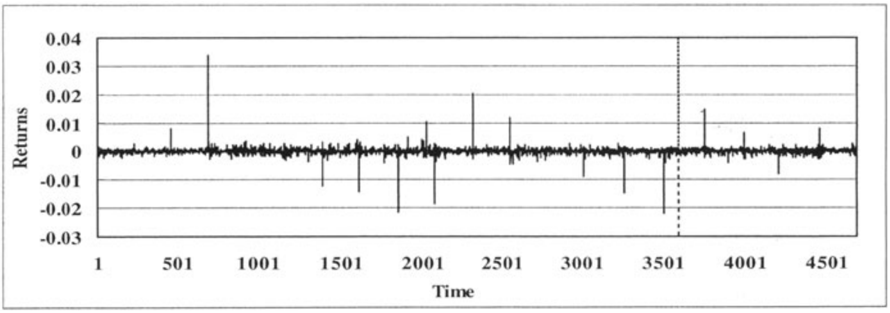

Fig. 13.4. Time series polt of stock returns: one-minute HSIX (1/4/99-1/29/99)

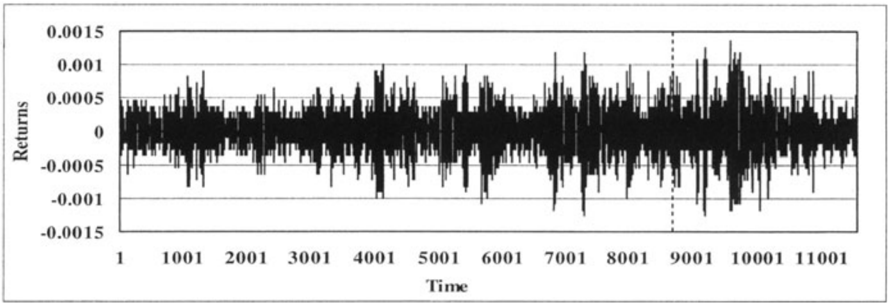

Fig. 13.5. Time series plot of *ECU*/US returns (tick-by-tick, 2/25/99-2/26/99)

Table 13.2. The forecasting performance of Cubist on high-frequency returns

	MSE	MAPE	MAE
×	10^{-4}	10^{0}	10^{-3}
Hang-Shen Index: Dec. 98			
Cubist	0.01	98	0.36
Cubist+NN	0.01	207	0.47
RW	0.01	100	0.36
Hang-Shen Index: Jan. 99			
Cubist	0.01	99	0.54
Cubist+NN	0.01	202	0.59
RW	0.01	100	0.54
\$ECU/\$US: 2/25/99 8:00 - 2/26/99 7:59			
Cubist	0.00	100	0.20
Cubist+NN	0.00	90	0.18
RW	0.00	100	0.20
\$ECU/\$US: 3/1/99 1:00 - 3/2/99 0:59			
Cubist	0.00	100	0.19
Cubist+NN	0.00	92	0.18
RW	0.00	100	0.19

"Cubist" refers to the *simple rule-based model*, whereas "Cubist+NN" refers to the composite of the *simple rule-based model* and the *nearest-neighbor model.*

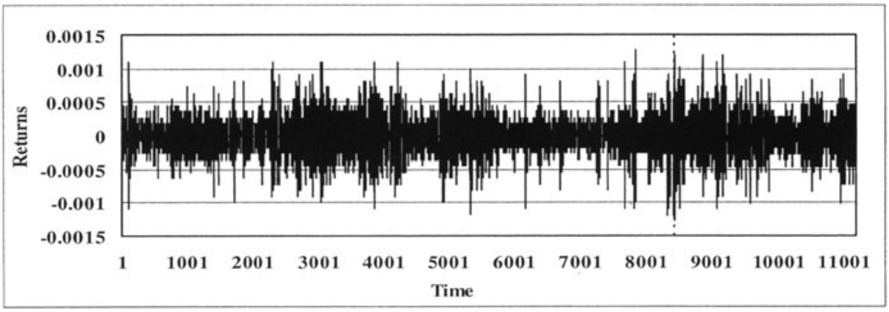

Fig. 13.6. Time series plot of *ECU*/US returns (tick-by-tick, 3/1/99-3/2/99)

ation process of these return series.) However, we do not use the whole series. Instead, based on the BDS value, we only select two subseries, namely, the one with the smallest absolute value and the one with the largest absolute values. (See Table 5 of [13.3].) Therefore, we have two series, each with 1999 observations, from these four markets. The series which are singled out by this criterion and their corresponding BDS values are given in Table 13.1. For example, for Market A the series with the highest absolute value of the BDS statistic is Series 10, which has a BDS value of 7.02. We then take the first 1,499 observations as the training sample and the others (500 observations) as the hold-out sample.

Table 13.3. Experimental results: markets A and B

	MSE	MAPE	MAE
×	10^{-4}	10^{0}	10^{-3}
Market A: High BDS			
Cubist	0.29	103	2.58
Cubist+NN	0.18	91	2.16
RW	1.21	100	7.01
Market A: Low BDS			
Cubist	0.46	115	2.99
Cubist+NN	0.33	136	2.43
RW	1.52	100	7.48
Market B: High BDS			
Cubist	0.75	140	4.36
Cubist+NN	0.74	255	4.87
RW	0.74	100	5.03
Market B: Low BDS			
Cubist	3.07	269	9.60
Cubist+NN	3.18	247	10.79
RW	2.86	100	9.93

"Cubist" refers to the *simple rule-based model*, whereas "Cubist+NN" refers to the composite of the *simple rule-based model* and the *nearest-neighbor model.*

13.4 Experimental Results

The performance criteria to evaluate Cubist are the *mean squared error* (**MSE**), *mean absolute percentage error* (**MAPE**) and *mean absolute error* (**MAE**). The benchmark is the *random walk* (**RW**). For each financial or artificial return series, the two versions of Cubist mentioned in Section 13.2 are applied and the results are given in Tables 13.2, 13.3, and 13.4.

Let us take a look at the financial data first. In all cases and all criteria, [Cubist] performs at least as good as RW. On the other hand, [Cubist+NN] is the best in the tick-by-tick return series of the exchange rate. However, the number of rules discovered by the Cubist is *only one*, and the contents of those rules are also very simple and naive (See Fig. 13.2). In a few cases, the rules simply took "0" as the prediction value. In other words, *what Cubist did was to rediscover the EMH.* Therefore, if someone argues that these series are efficient, then Cubist would not argue over that.

An important question remains: how about the artificial return series generated from the agent-based stock markets? Several interesting results deserve our attention. First, when we apply Cubist in a highly efficient market, it does not mean that it will simply give "0" as the prediction value. For example, consider Market C. This market is shown to be highly efficient in [13.6]. However, the Cubist suggests 17 rules for the series with a high BDS value, and 20 rules for the one with a low BDS value. Moreover, in terms of

Table 13.4. Experimental results: markets C and D

	MSE	MAPE	MAE
×	10^{-4}	10^{0}	10^{-3}
Market C: High BDS			
Cubist	4.65	209	10.55
Cubist+NN	4.58	335	11.87
RW	4.11	100	11.61
Market C: Low BDS			
Cubist	4.53	160	9.56
Cubist+NN	4.39	197	11.78
RW	4.24	100	11.45
Market D: High BDS			
Cubist	1.13	141	4.94
Cubist+NN	1.11	756	5.46
RW	1.58	100	7.21
Market D: Low BDS			
Cubist	1.36	125	4.51
Cubist+NN	1.18	416	4.93
RW	1.74	100	7.19

"Cubist" refers to the *simple rule-based model*, whereas "Cubist+NN" refers to the composite of the *simple rule-based model* and the *nearest-neighbor model*.

the criterion MAE, both series beat the random walk. Therefore, the efficient market hypothesis does not necessarily imply that machine-learning people can do nothing.

Let us consider Market D. In [13.3], this market is shown to be not efficient. Therefore, we anticipate that Cubist can do something on these series, and it in fact does. For the series with a high BDS value, 3 rules are discovered, and for the series with a low BDS value, 4 rules are discovered. However, does it work? The answer is "not quite". If we take MAPE as the criterion, then Cubist clearly is beaten by the random walk in both cases.

13.5 Conclusions

Examples like this are many and can be found easily from Tables 13.3 and 13.4. Due to these negative examples, machine-learning people may think that econometric test results are irrelevant. In agent-based stock markets, we may have many traders like this. Therefore, the puzzle proposed in Section 13.1 may be solved from this direction.

References

13.1 Chen, S. -H., Kuo, T. -W. (1999): Towards an Agent-Based Foundation of Financial Econometrics: An Approach Based on Genetic-Programming Artificial Markets. In: Banzhaf, W., Daida, J., Eiben, A. E., Garzon, M. H., Honavar, V., Jakiela, M., Smith, R. E. (Eds.), GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference, Vol. 2. Morgan Kaufmann, 966–973
13.2 Chen, S. -H., Tan, C. -W. (1999): Brief Signals in the Real and Artificial Stock Markets: An Approach Based on the Complexity Function. In: Arabnia, H. R. (Ed.), Proceedings of the International Conference on Artificial Intelligence, Vol. II. Computer Science Research, Education, and Application Press, 423–429
13.3 Chen, S. -H., Yeh, C. -H. (1999): On the Consequence of "Following the Herd": Evidence from the Artificial Stock Market. In: Arabnia, H. R. (Ed.), Proceedings of the International Conference on Artificial Intelligence, Vol. II. Computer Science Research, Education, and Application Press, 388–394
13.4 Chen, S. -H., Yeh, C. -H. (2001): Evolving Traders and the Business School with Genetic Programming: A New Architecture of the Agent-Based Artificial Stock Market. Journal of Economic Dynamics and Control, 25, 363–393
13.5 Chen, S. -H., Wang, H. -S., Zhang, B. -T. (1999): Forecasting High-Frequency Financial Time Series with Evolutionary Neural Trees: The Case of Hang-Seng Stock Index. In: Arabnia, H. R. (Ed.), Proceedings of the International Conference on Artificial Intelligence, Vol. II. Computer Science Research, Education, and Application Press, 437–443
13.6 Chen, S. -H., Yeh, C. -H. (2002): On the Emergent Properties of Artificial Stock Markets: The Efficient Market Hypothesis and the Rational Expectations Hypothesis. Journal of Economic Behavior and Organization, 49(1), 217–239
13.7 Chen, S. -H., Liao, C. -C. (2002): Testing for Granger Causality in the Stock-Price Volume Relation: A Perspective from the Agent-Based Model of Stock Markets. Information Sciences, forthcoming
13.8 Grossman, S. J., Stiglitz, J. (1980): On the Impossibility of Informationally Efficiency Markets. American Economic Review, 70, 393–408
13.9 Rissanen, J. (1986): Stochastic Complexity and Modeling. Annals of Statistics, 14(3), 1080–1100