Microstructure Dynamics and Agent-Based Financial Markets

Shu-Heng Chen[1], Michael Kampouridis[2], and Edward Tsang[2]

[1] AI-ECON Research Center, Department of Economics, National Chengchi University, Taiwan
chen.shuheng@gmail.com
[2] School of Computer Science and Electronic Engineering, University of Essex, UK

Abstract. One of the essential features of the agent-based financial models is to show how price dynamics is affected by the evolving microstructure. Empirical work on this microstructure dynamics is, however, built upon highly simplified and unrealistic behavioral models of financial agents. Using genetic programming as a rule-inference engine and self-organizing maps as a clustering machine, we are able to reconstruct the possible underlying microstructure dynamics corresponding to the underlying asset. In light of the agent-based financial models, we further examine the microstructure both in terms of its short-term dynamics and long-term distribution. The time series of the TAIEX is employed as an illustration of the implementation of the idea.

1 Introduction and Main Ideas

It comes as no surprise to economists that there is no single strategy which can persistently dominate all other strategies in the market. The idea of the best strategy is simply inconsistent with the intuitive notion of the efficient market hypothesis. While this feature is well expected among economists, the result shown by [4], generally known as the *overreaction hypothesis*, is still very appealing. They have found that successive portfolios formed by the previous five years' 50 most extreme winners considerably underperform the market average, while portfolios of the previous five years' 50 worst losers perform better than the market average.

Recently, a similar phenomenon has been rigorously analyzed and replicated in the agent-based finance literature, in particular, in the H-type model. In this literature, markets at any point in time are composed of different clusters (types) of agents. Agents who follow similar rules are considered to be in the same cluster. Each cluster is defined by the associated behavioral rules. The market microstructure is characterized by the *fractions* (distribution) of individuals over different clusters. Different distributions (microstructure) over the clusters may have different impacts on the aggregates, and both the microstructure and the aggregates are evolving with feedbacks to each other.

Complex dynamic analysis of these models indicates two interesting properties. First, in the short run, it is *likely* that the market fractions are constantly changing. In particular, for each cluster, the market fraction can swing from very low to very high, i.e., switching between the majority and the minority. Second, in the long run, no single strategy can dominate the other, i.e., the market fraction converges to $1/H$ for each cluster. These two properties provide us with a basis to study the complex dynamics of microstructure, which we refer to together as the *market fraction hypothesis*, or as an abbreviation, the MFH. In fact, a number of empirical studies have already attempted to estimate the parameters associated with the MFH.

This paper, however, differs from the H-type models in two regards. First, we do not assume any prefixed behavioral rule (functional form) for any cluster (type) of agents; second, we do not assume that agents of the same type are homogeneous, while they can be *similar*. We consider that this departure will lead us to a more general and *realistic implication* of the MFH. Consider the three-type model as an example. In the fundamentalist-chartist-contrarian model, traders of the same type at any point in time behave in *exactly the same way*, and their functional forms of behavioral rules, in this case, their forecasts of the price in the next period, $\{E_{f\ t}(p_{t+1})\}$, $\{E_{c\ t}(p_{t+1})\}$ and $\{E_{co\ t}(p_{t+1})\}$, are all known. Equations (1) to (3) are typical examples.

$$E_{f\ t}[p_{t+1}] = p_t + \alpha_f(p_t^f - p_t), \quad 0 \leq \alpha_f \leq 1., \tag{1}$$

$$E_{c\ t}(p_{t+1}) = p_t + \alpha_c(p_t - p_{t-1}), \quad 0 \leq \alpha_c. \tag{2}$$

$$E_{co\ t}(p_{t+1}) = p_t + \alpha_{co}(p_t - p_{t-1}), \quad \alpha_{co} \leq 0. \tag{3}$$

Nevertheless, in the real world, the behavioral rules of each trader are expected to be heterogeneous, and even if they can be clustered into types, the representative behavior of each type is normally unknown.[3]

1.1 Genetic Programming as a Rule-Inference Engine

In this paper, we assume that traders' behavior, including price expectations and trading strategies, is either not observable or not available. Instead, their behavioral rules have to be *estimated* by the observable market price. Using macro data to estimate micro behavior is not new as many H-type empirical agent-based models have already performed such estimations [3]. However, as mentioned above, such estimations are based on very strict assumptions upon which a formal econometric model can be built. Since we no longer keep these assumptions, an alternative must be developed, and in this paper we recommend *genetic programming* (GP).

The use of GP as an alternative is motivated by considering the market as an evolutionary and selective process.[4] In this process, traders with different behavioral rules participate to the markets. Those behavioral rules which help traders

[3] While the ideas of fundamentalists and chartists are the results of field work, abstracting the general observed behavior into a very specific mathematical model is a big leap.

[4] See [11] for his eloquent presentation of the *adaptive market hypothesis*.

gain lucrative profits will attract more traders to *imitate*, and rules which result in losses will attract fewer traders. This evolutionary argument in fact is, intuitively, the same as the evolution process considered by the H-type agent-based financial models. For example, their use of the Gibbs-Boltzman distribution is a formalization of this process. Genetic programming is another formalization which, unlike the former, does not rest upon any pre-specified class of behavioral rules. Instead, in GP, a population of behavioral rules is randomly initiated, and the survival-of-the-fittest principle drives the entire population to become fitter and fitter in relation to the environment. In other words, given the non-trivial financial incentive from trading, traders are aggressively searching for the most profitable trading rules. Therefore, the rules that are outperformed will be replaced, and only those very competitive rules will be sustained in this highly competitive search process.[5]

Hence, even though we are not informed of the behavioral rules followed by traders at any specific time horizon, GP can help us infer what these rules are *approximately* by simulating the evolution of the microstructure of the market. Without imposing tight restrictions on the inferred behavioral rules, GP enables us to go beyond the simple but also unrealistic behavioral rules used in the N-type agent-based financial models. Traders can then be clustered based on more realistic, and possibly more complex behavioral rules.[6]

1.2 Self-Organizing Maps as a Clustering Machine

Once a population of rules is inferred from GP, it is desirable to cluster them based on a chosen similarity criterion so as to provide a concise representation of the microstructure. The similarity criterion which we choose is based on the *observed trading behavior*. Based on this criterion, two rules are similar if they are *observationally equivalent* or *similar*, or, alternatively put, they are similar if they generate the same or similar market timing behavior.

Given the criterion above, the behavior of each trading rule can be represented by its series of market timing decisions over the entire trading horizon, for example, 6 months. Therefore, if we denote the decision "enter the market" by "1" and "leave the market" by "0", then the behavior of each rule is a binary string or a binary vector. The length of these strings or the dimensionality of the vectors is then determined by the length of the trading horizon. For example, if the trading horizon is 125 days long, then the dimension of the market timing vector is 125. Once each trading rule is concretized into its market timing vector, we can then easily cluster these rules by applying Kohonen's *self-organizing maps* (SOMs) [9] to the associated clusters.

[5] It does not necessarily mean that the types of traders surviving must be smart and sophisticated. They can be dumb, naive, randomly behaved or zero-intelligent. Obviously, the notion of rationality or bounded rationality applying here is *ecological* [12, 6].

[6] [5] provides the first illustration of using genetic programming to infer the behavioral rules of human agents in the context of ultimatum game experiments. Similarly, [7] uses genetic algorithms to infer behavioral rules of agents from market data.

The main advantage of SOMs over other clustering techniques such as K-means is that the former can present the result in a *visualizable* manner so that we can not only identify these types of traders but also locate their 2-dimensional position on a map, i.e., a distribution of traders over a map. Furthermore, if we suppose that we do not have dramatic crustal plate movement so that the map is fixed over time, then the distribution of traders over the map can, in effect, be comparable over time. This provides us with a rather convenient grasp of the dynamics of the microstructure directly as if we were watching the population density on a map over time.

However, the assumption of crustal stability does not hold in general; therefore, *maps over time are not directly comparable*. To make them comparable, some adjustments are needed. The idea of adjustment is also very intuitive. If the dominant strategy remains unchanged from period A to period B, then when we apply the dominant trading strategy derived from period A to another period B, the strategies should behave in a way that is similar to the dominant strategy derived from period B, if it is not exactly the same. This motivates us to *emigrate* all trading strategies from one map (the home map) to the other (the host map) in such a way that each emigrant shall find its new cluster on the host map based on the same similarity metric. In this manner, we can reconstruct a time-invariant version of the map, and comparison can be made upon this reconstruction.

The rest of the paper is organized as follows. Section 2 provides a brief description of the version of genetic programming used in this paper. Section 3 demonstrates the self-organizing map constructed based on the description in Section 1.2. A time series of these maps is constructed accordingly and the maps are then analyzed both in their short-term dynamic behavior (Section 3.1) and long-term distribution behavior (Section 3.2). The analysis is further consolidated with the results from multiple runs (Section 3.3). Section 4 examines the short-term dynamics and long-term distribution behavior of a rather small self-organizing map. In Section 5, we present our concluding remarks.

2 Genetic Programming

In this paper, we use the financial GP system introduced by Edward Tsang at University of Essex, known as Eddie. Eddie, standing for Evolutionary Dynamic Data Investment Evaluator, applies genetic programming to evolve a population of artificial financial advisors or, alternatively, a population of market-timing strategies, which guide investors on when to buy, to hold, or to sell. These artificial financial agents (market timing strategies) are formulated as decision trees in Eddie, which, when combined with the use of GP, are referred to as *Genetic Decision Trees* (GDTs).

Each of these market-timing strategies (GDTs) is syntactically (grammatically) produced by the Backus Normal Form (BNF) [2]. Figure 1 presents the Backus Normal Form (BNF) of the GP. As we can see, the root of the tree is an If-Then-Else statement. Then the first branch is a boolean (testing whether a technical indicator is greater than/less than/equal to a value). The 'Then' and 'Else' branches can be a new Genetic Decision Tree (GDT), or a decision, to buy or not-to-buy (denoted by 1 and 0).

```
<Tree> ::= If-then-else <Condition> <Tree> <Tree> | Decision
<Condition> ::= <Condition> "And" <Condition> |
                <Condition> "Or" <Condition> |
                "Not" <Condition> |
                VarConstructor <RelationOperation> Threshold
<Variable> ::= MA_12 | MA_50 | TBR_12 | TBR_50 | FLR_12 |
                    FLR_50 | Vol_12 | Vol_50 | Mom_12 | Mom_50 |
                    MomMA_12 | MomMA_50
<RelationOperation> ::= ">" | "<" | "="
Decision is an integer, Positive or Negative implemented
Threshold is a real number
```

Fig. 1. The Backus Normal Form of EDDIE

Given a set of historical data and the fitness function, GP is then applied to evolve these market-timing strategies in a standard way. After evolving a number of generations, what stands (survives) at the end (the last generation) is, presumably, a population of financial agents whose market-timing strategies are financially rather successful. For the details, see [13] and [10].

3 An Illustration from the Taiwan Stock Market

Figure 2 gives a concrete illustration of the idea presented above (Sections 1.1 and 1.2). Here, 500 artificial traders are grouped into *nine* clusters. The parameter value '500' refers to the *population size* used in genetic programming, i.e., the rule-inference stage, whereas the parameter value '9' is due to a 3 × 3 two-dimensional SOM employed in the rule clustering stage. In a sense, this could be perceived as a snapshot of a nine-type agent-based financial market dynamics. Traders of the same type indicate that their market timing behavior is very similar. The market fraction or the size of each cluster can be seen from the number of traders belonging to that cluster. Not surprisingly, they are not evenly distributed. Figure 2 shows that the largest cluster has a market share of 37.6% (188/500), whereas the smallest cluster has a market share of only 0.4% (2/500).

Once we can have a snapshot of the market fraction, we can go further over a series of snapshots so as to have a picture of the dynamics of the market fraction or the dynamics of the market microstructure. However, as we mentioned before, the SOMs constructed from different periods are not directly comparable; therefore, to make them all comparable, we have to first choose a base period and fix the map, i.e., to take the centroid of each cluster as given. In this particular example, we choose the second half of the year 2007 as the base. Once the centroids are given, all points (vectors) in other maps shall *immigrate* into this

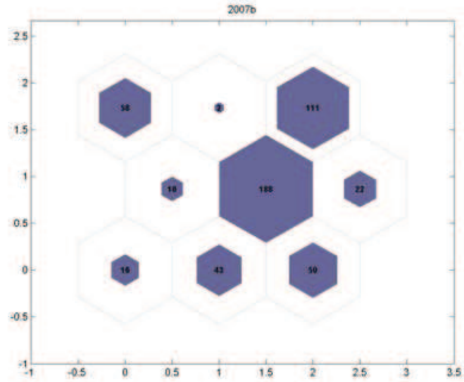

Fig. 2. 3 × 3 Self-Organizing Feature Map

The SOM is constructed based on the 500 financial decision trees generated by GP using the daily data of the TAIEX from July 2007 to December 2007

fixed map, and they are re-clustered based on their similarity to these fixed centroids. Figure 3 shows the reconstruction of these maps in this manner.

This figure has the *market fraction maps* from the year 2006 to the year 2007, crossing 4 different periods. These maps were constructed by using the second half of 2007 as the base period. This figure gives a clear picture of what we mean by *market fraction dynamics.* First of all, we notice that the distribution over the clusters is uneven over time. In each period of time, some clusters obviously dominate others, but that dominance changes over time. This can be seen from the constant renewing of the major blocks. This eye-browsing inspection motivates us to formulate two hypotheses which we already experienced from the dynamics of H-type agent-based financial models.

3.1 Short-Term Dynamics

The first hypothesis regards the *short-run dynamics of market fraction.* Each type of trader can be a dominant group (majority) for some of the time, but the duration of its dominance can only be temporal. The quick turnover of the dominant cluster or its short duration is consistent with the impression of the *swinging dynamics* as we saw in the 2-type agent-based financial models, e.g., [8]. However, in addition to eye-browsing the swing, it is desirable to have an objective measure of *how persistent a dominant cluster can be.* To do so, we need an operational meaning of dominance. Even though there is no unique way of doing this, we find the following threshold to be quite general and useful.

$$\bar{q} = \frac{1+p}{H+p}, \tag{4}$$

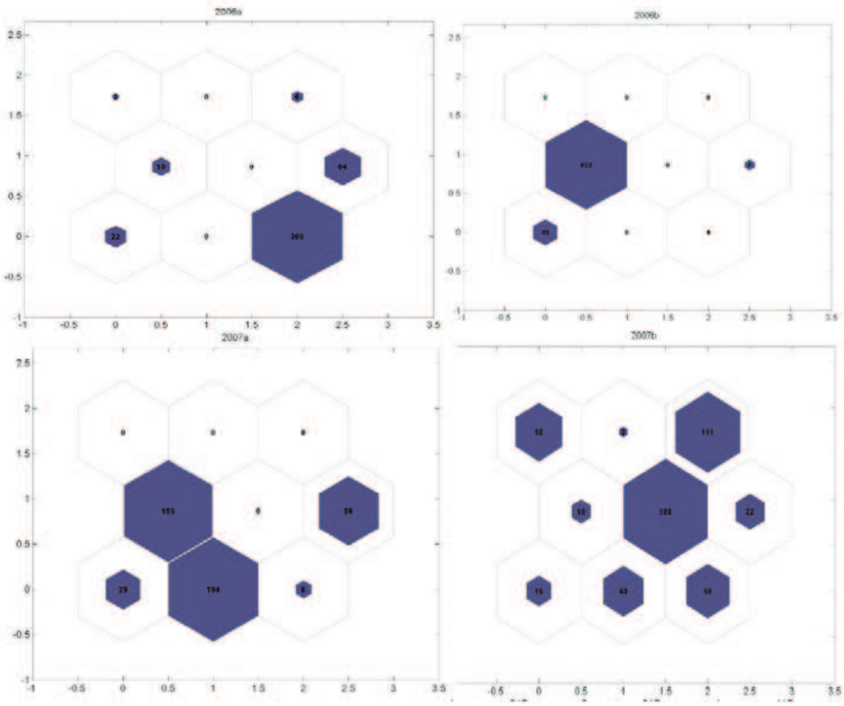

Fig. 3. Market Fraction Dynamics: Map Dynamics

The four SOMs above are constructed using the daily data of the TAIEX from 2006 to 2007. From the top-left panel to the bottom-right panel, they correspond to the first half and seconf half of year 2006 (2006a, b) and the first half and second half of the year 2007 (2007a, b). Except for the last one, 2007b, the other three are reconstructed by using 2007b as the base (see Section 1.2).

where H is the number of clusters, and p, a non-negative integer, is a control parameter for the *degree of dominance*. Hence, a cluster is dominant if its market fraction exceeds this threshold. By varying the parameter p, one can therefore have an operational meaning that is consistent with our intuition regarding dominance. For example, if $H = 2$ (a two-type model) and $p = 2$, a cluster can be dominant only if its market fraction is greater than a $\bar{q}$ of 75%, a standard much higher than just breaking the tie (one half). Of course, the higher the p, the higher the threshold.

Figure 4 presents the *dominance-duration statistics* of each type of trader. Basically, we keep track of the persistent time of each dominance. Once after a type of trader become dominant, we count how many periods in a row that it can remain the dominant cluster. Figure 4 gives three statistics regarding du-

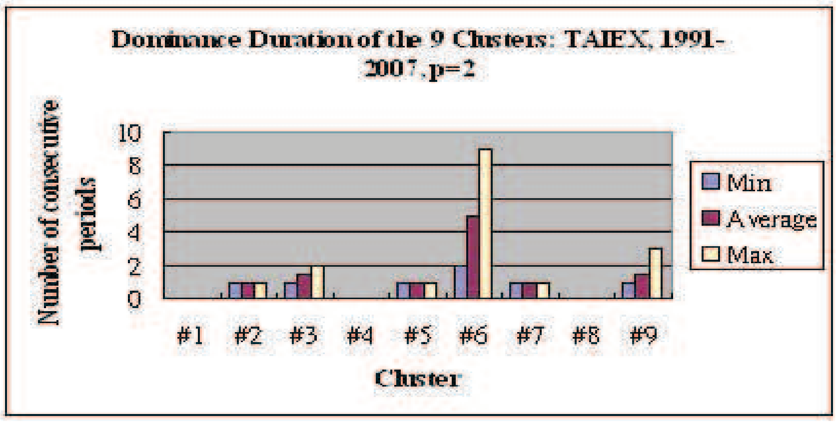

Fig. 4. Duration of Dominance (p=2)

ration, namely, minimum, average and maximum. For example, for Cluster Six, these three statistics are 1, 3 and 9, respectively. In other words, the maximum duration of dominance for Cluster Six is about nine periods, i.e., four and a half years. For other clusters, the longest duration is no more than three periods, i.e, one and half years. So, for most of the time, dominant clusters can hardly continue for long. Hence, we reach the conclusion that, regardless of the types of traders, we can rarely see the consecutive dominance. In this sense, our data lend support to the market fraction hypothesis in a weak sense.

3.2 Long-Term Distribution

The second hypothesis which we can form regarding the market fraction behavior is its *long-term distribution*. Many H-type agent-based financial models can show us that, under some proper parameter values, the long-term market fraction is *even*. In other words, if we have H types of traders, their long-term frequency of appearance should be close to $\frac{1}{H}$. Let $Card_{i\ t}$ be the number (cardinality) of traders in Cluster i in time period t.

$$\sum_{i=1}^{H} Card_{i\ t} = N, \quad \forall t. \tag{5}$$

In our current setting, N, the total number of traders, is 500. The long-term histogram can be derived by simply summing the number of traders over all periods and dividing it by a total of $N \times T$ (# of periods),

$$w_i = \frac{\sum_{t=1}^{T} Card_{i\ t}}{N \times T}. \tag{6}$$

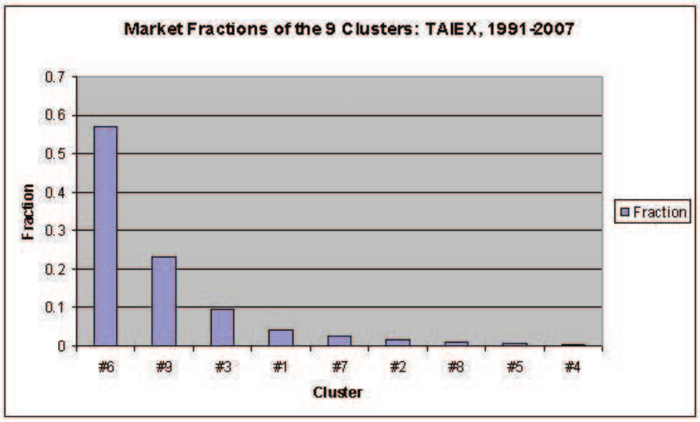

Fig. 5. Long-Term Histogram

Figure 5 gives the long-term histogram of these clusters, $\{w_i\}$. Obviously, they are not equal so that we present them in descending order from the left to the right. Cluster Six has the largest market fraction up to almost 60%, whereas Cluster 4 has the smallest market fraction, which is not even up to 1%.

Of course, this distribution is very different from the uniform one. In order to give a measure of how far it is from the uniform one, we use the familiar *entropy* as a metric. Let us denote the empirical distribution presented in Figure 5 as f_X, and the uniform distribution as f_Y. By definition, $f_Y = \frac{1}{H}$, where H is the number of clusters, which in this case is 9. In order to measure how close f_X is to the uniform distribution f_Y, we calculate the entropy of both distributions. For the discrete random variable, the entropy is defined as

$$Entropy = -\sum_{i=1}^{H} p_i \ln p_i, \tag{7}$$

where p_i is the fraction of each cluster. It is well known that for the uniform distribution $Entropy(Y) = \ln H$. When H=9, it is $\ln 9 \approx 2.2$. The closer $Entropy(X)$ is to 2.2, the closer X is to the uniform distribution. After calculating X's entropy, we find it equal to 1.3, which is only 41% of the entropy of the uniform distribution.

Summary As we have seen in this section and the previous one, both the short-run and the long-run version of the market fraction hypothesis are not well supported. The short-run dynamics indicates the appearance of a long-lasting dominant cluster (up to a maximum of 9 periods). On the other hand, the long-run histogram is very far away from the uniform distribution.

Table 1. Summary Results over 10 runs, for a 3×3 SOM.

	Short-Run		Long-Run
	Mean	Max	E-Ratio
TAIEX (9 Clusters)	2.02	8.25	0.55
TAIEX (3 Clusters)	4.05	8.14	0.80

3.3 Results from Multiple Runs

However, so far we have only presented the results of a single run. To consolidate our results, we further replicate the experiments for an additional nine runs, and Table 1 gives the results for ten runs together.

The first two numeric columns are related to the short-run dynamics and present the averages over the 10 runs for both the average duration and the maximum duration of the 9 clusters. This result is not much different from our earlier single-run results. The mean dominance duration over these ten runs is just about 2 (one year). Nevertheless, the existence of few long-lasting dominant clusters is very evident with the mean maximum duration reaching as high as 8.25 periods (more than 4 years). Hence, the short-run version of the market fraction hypothesis is only weakly supported. The next column presents the ratio of the average realized entropy (over the 10 runs) relative to the base entropy under the null of the uniform distribution, which is 55%, and still quite far away from one. Therefore, the long-term version of the market fraction hypothesis is not well supported.

4 Does the Number of Types Matter?

The illustration presented above is based on a 3 by 3 SOM, which automatically generates nine clusters. This analysis has its limitations mainly because we do not know how many types of agents are really there in the market. In a rather theoretical analysis, [1] showed that it would be enough to characterize the market behavior by a few types, say two to three. Others are rather marginal. Therefore, it would be interesting to investigate the microstructure dynamics based on a smaller SOM corresponding to the few-type agent-based financial models.[7]

In this section, we therefore repeat the above experiments by using a rather small 3 × 1 SOM. We then examine both its short-term dynamics and the long-term histogram. As before, we have 10 multiple runs. The results are shown in Table 1. In terms of duration behavior, we can see that there is no significant difference in the maximum duration between the 9-cluster case and the 3-cluster

[7] Based on [3], the 2-type or the 3-type agent-based financial models are still the most popularly-used classes in the literature.

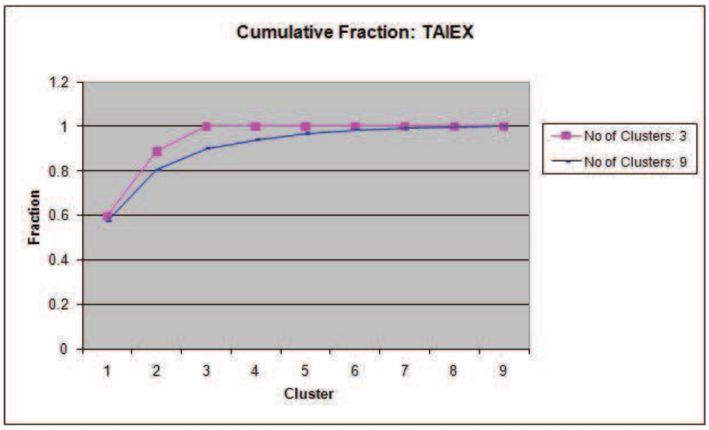

Fig. 6. Cumulative Fractions of 3 Clusters and 9 Clusters

case, for here the mean maximum duration is consistently a little above eight (four years). However, a significant difference in mean duration does exist. What we find here is that when the number of clusters decreases, the mean duration increases from the original 2.02 periods (one year) to 4.05 periods (two years). Therefore, it seems that a smaller number of clusters really drives the short-run dynamics further away from the expectations of the market fraction hypothesis.

On the other hand, if we look at the long-term distribution behavior, we find that a smaller number of clusters does help the distribution (histogram) get closer to the uniform distribution. As shown in Table 1, the realized entropy ratio now increases up to 80% from the original 55%. Hence, the market fraction hypothesis is better supported from a long-term point of view.

Putting them together, what we have observed here is that, when the number of clusters gets smaller, the dominant cluster maintains its position longer, but a different cluster does take the lead in turn, and so, in the long run, they are equally competitive. This observation, of course, is interesting and requires further studies using agent-based financial models.

If the number of clusters does matter for the microstructure dynamics, then it is imperative to know how many clusters we need. To answer this question, Figure 6 presents the cumulative fraction sum from the largest cluster to the smallest cluster. For the 3-cluster case, when the number of clusters (the x axis) gets to 3, the cumulative fraction becomes one, and similarly for the 9-cluster case when the number of clusters gets to 9. However, what we can see here is that when coming to the first five clusters, there is already an accumulation of 96% of the market share. In fact, if we care only about 90% of the market fraction, then 3 clusters are sufficient.

5 Concluding Remarks

After a decade of development, the literature on agent-based financial models has successfully demonstrated the connection between microstructure dynamics and asset price dynamics. The next research agenda would be to gain more understanding of the empirical properties of this microstructure dynamics. In this paper we have shown that the number of types (clusters) of agents may be limited, but the durations of dominant groups are larger than what we may expect from, say, the adaptive market hypothesis [11]. The next step is to explore other financial markets and to see whether this is a universal phenomenon.

References

1. Aoki M (2002), *Modeling Aggregate Behavior and Fluctuation in Economics: Stochastic Views of Interacting Agents*, Cambridge University Press.
2. Backus J (1959) The syntax and semantics of the proposed international algebraic language of Zurich. ICIP, Paris.
3. Chen S.-H, Chang C.-L, Du Y.-R (2010) Agent-based economic models and econometrics. Knowledge Engineering Review, forthcoming.
4. De Bondt W. Thaler R (1985) Does the stock market overreact? Journal of Finance 40(3):793-808.
5. Duffy J, Engle-Warnick J (2002) Using symbolic regression to infer strategies from experimental data, in Chen S (ed.), Evolutionary Computation in Economics and Finance, Springer, 61-82.
6. Gigerenzer G, Todd P (1999) Fast and Frugal Heuristics: The Adaptive Toolbox. In: Gigerenzer G, Todd P and the ABC Research Group, Simple Heuristics That Make Us Smart, 3–34. New York: Oxford University Press.
7. Izumi K, Ueda K (2002) Using an artificial market approach to analyze exchange rate scenarios, in Chen S (ed.), Evolutionary Computation in Economics and Finance, Springer, 135-157.
8. Kirman A (1993) Ants, rationality, and recruitment. Quarterly Journal of Economics, 108: 137-156.
9. Kohonen T (1982) Self-organized foundation of topologically correct feature maps. Biological Cybernetics 43: 59–69.
10. Li J (2001) FGP: A genetic programming based financial forecasting tool. PhD thesis, Department of Computer Science, University of Essex.
11. Lo A (2005) Reconciling efficient markets with behavioral finance: The adaptive market hypothesis. Journal of Investment Consulting 7(2): 21-44.
12. Simon H (1956) Rational choice and the structure of environments. Psychological Review 63:129-138.
13. Tsang E, Li J, Butler J (1998) Eddie beats the bookies. International Journal of Software, Practice & Experience 28(10): 1033-1044.