ON THE NATURE OF THE STOCK MARKET: SIMULATIONS AND EXPERIMENTS

by

Hendrik J. Blok

B.Sc., University of British Columbia, 1993
M.Sc., University of British Columbia, 1995

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE STUDIES

(Department of Physics and Astronomy)

We accept this dissertation as conforming to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

November 2000

Chapter 7

Concluding remarks

7.1 Review

Traditional economic theory interprets stock markets as equilibrium systems driven by exogenous events. But this theory is incapable of explaining some peculiarities—such as the prevalence of large fluctuations—which are observed to be universal across all markets. Instead, these phenomena are traditionally attributed to the exogenous driving factors. The goal of this thesis was to discover whether these anomalies may arise directly from simple interactions between a large number of investors, and not depend on extraordinary external influences.

7.1.1 Anomalous market properties

Some of the peculiarities observed in the markets and not explainable by traditional economic theory follow:

Scaling

Firstly, the distribution of returns (be they price returns for a particular stock or index returns) contain too many outliers to be adequately described by the default Gaussian distribution. In Chapter 5 some alternatives were presented which properly capture the extra "weight" contained in the distribution tails. Empirical analysis suggests the distribution is best described by a Lévy distribution with exponent $\alpha \approx 1.40$ [10] which is truncated for large returns by either a decaying exponential or a power law with an exponent near three.

Clustered volatility

Secondly, although the price series has no (significant) memory—supporting the hypothesis that markets are efficient, containing no arbitrage opportunities—the same cannot be said for market volatility. Volatility, which describes the degree of excitation or uncertainty in the market and is quantified most simply by the absolute value of the price returns, exhibits extremely long temporal correlations. High volatility tends to follow high and low follows low, resulting in clusters of activity. This conflicts with traditional economic theory which states that fluctuations should be regular and uncorrelated.

To test the hypothesis that these properties may emerge spontaneously from the interactions of many simple investors, two market simulations, the Centralized and Decentralized Stock Exchange Models (CSEM and DSEM) were constructed in Chapters 2 and 3, respectively.

7.1.2 Centralized stock exchange model

CSEM was a traditional simulation, building on similar models developed over the last few years. Its main features include centralized trading (all traders deal with a single *market maker*), synchronous updating and forecasting of returns. Each forecast was deliberately nudged by a normally-distributed amount with standard deviation σ_ϵ, the forecast error. It was discovered that as the forecast error was reduced the system underwent a second-order (critical) phase transition near $\sigma_c \approx 0.08$, below which the price diverged (or would have if it wasn't artificially bounded).

When the distribution of the price returns was computed it was discovered that CSEM was only able to produce an overabundance of outliers (compared with the Gaussian) below the critical point, precisely in the regime where the price series is known to be unrealistic. Above the critical point the distribution fit very well to a Gaussian. Thus, CSEM is unable to capture the anomalous "fat tails" phenomenon observed empirically. Since it failed this first test, it was not tested for any of the other properties mentioned above. Instead, focus was shifted to the decentralized model. (In retrospect, CSEM may have been abandoned too rashly. By allowing multiple values of the control parameter, as in DSEM, more realistic dynamics may be realizable. This hypothesis will be tested.)

7.1.3 Decentralized stock exchange model

DSEM arose from dissatisfaction with the structure of CSEM: synchronous, centralized trading was replaced with asynchronous, decentralized trading directly between market participants and the need for forecasting was eliminated with a simple *fixed* *investment* strategy in which agents trade in order to maintain a balance between stock and cash. To drive the dynamics the ideal investment fraction was allowed to be affected by exogenous news events (modeled as a discrete Brownian process) and endogenously by price movements.

The dynamics were observed to have three phases of existence, depending on the strength of the agents' response to price movements: when the price response was in the region $r_1 > r_p > r_2$ autocorrelations in the price series were relatively weak but as the price response passed the critical point $r_1 = 1$ very strong positive correlations emerged and the price diverged rapidly. The third phase was found when the price response dropped below $r_2 \approx -0.33$, revealing a first-order phase transition. Below this point the price series was strongly anticorrelated.

When all agents were forced to share the same price response scaling in the price return distribution could not be induced except in the phases which exhibited unrealistic memory effects. However, if the price response was sampled from a two-point distribution, scaling (with a realistic truncation for large returns) was found for a number of simulations, the best predictor for scaling being that the upper price response exceeded one, $r_{hi} > 1$. For those runs which did exhibit scaling the exponent was found to be $\alpha = 1.64 \pm 0.25$, which compares favourably with the best known empirical quantity 1.40 ± 0.05 [10].

Having found that DSEM could capture this anomalous property of empirical markets it was also tested for memory effects, again using the two-point price response distribution. It was found that the price series did not have a significant memory provided that the lower bound of the price response was in the region $0.5 \le r_{lo} < 1$. Similarly, volatility clustering was observed when the upper limit exceeded $r_{hi} > 1.25$ or when the lower limit was below $r_{lo} < -0.5$.

All three requirements were met when $0.5 \le r_{lo} < 1$ and $r_{hi} > 1.25$. What this means for real markets will be discussed below. But first we review the remainder of the thesis.

7.1.4 Fixed investment strategy

DSEM was constructed on the principle of the fixed investment strategy (FIS) which states that one should adjust one's portfolio in order to maintain a balance between the capital invested in a risky stock and the capital held in a safe(r) asset. In Chapter 6 the results of an experiment intended to test the credibility of the FIS in a "real-world" situation (with trading costs, etc.) were reported.

It was discovered that the FIS actually underperformed when compared with a simple "Buy-and-hold" strategy, at least over this particular realization. This is probably due to the strong trend observed in the portfolio over the course of the experiment, in which the capital nearly doubled. The FIS is designed to take advantage of fluctuations in the price series and is sub-optimal in the presence of a long-term trend.

However, the experiment did reveal an interesting (and possibly advantageous) feature of the FIS: it minimized the risk in the sense that it reduced the frequency of large events (both up and down) as measured by the excess kurtosis. By applying the FIS large fluctuations were scaled down bringing them in line with the Gaussian distribution which is typically assumed. Of course, it should be remembered that these conclusions are less than rigorous, being the result of a single brief experiment with a particular portfolio.

7.1.5 Log-periodic precursors

While the FIS experiment was running the hypothesis that market crashes are heralded by log-periodic precursors was also tested. The theory derives from discrete scale invariance and suggests that, in some cases, systems approaching a critical event may exhibit accelerating oscillations in the power law describing the critical point.

It has been suggested that detecting these oscillations may improve predictions of the critical event time and recent work in seismology is promising. But the financial data from the FIS experiment indicate that, even if log-periodic precursors do exist, technical optimization difficulties prevent any accurate forecasts of large fluctuations therefrom.

7.2 Conclusions to be drawn from this research

The main point the reader should draw from this thesis is that it is possible to replicate realistic market dynamics with a many-agent model with simple driving forces. DSEM was driven by a simple (discrete) Brownian motion without fat tails and having no memory, but through the interactions of the agents both fat tails and long memories (in the volatility) emerged. Similarly, these properties may arise endogenously in real economic systems, and appeals to anomalous external events to explain them may be unwarranted.

Interestingly, the most realistic simulations were observed when the price response (control parameter) was centered around a critical point at $r_p = r_1 = 1$. If DSEM is assumed to properly capture the essence of real markets the question is naturally raised: "Why are the markets tuned to this region of parameter space?" The fact that this region encompasses a critical point is suggestive of a concept called *self-organized criticality* (SOC) which claims that many dynamical systems spontaneously evolve towards a critical point [11, 12, 88]. The problem with this description is that it adds nothing to our knowledge: it does not tells us how or why the market self-organizes.

In a simple economic model involving producers and consumers it was discovered that the system self-organizes to the critical state in order to maximize efficiency [95, Ch. 11]. On one side of the critical point the supply outweighs demand and on the other the reverse is true. In this example it is easy to see why the market would self-organize. To test whether a similar process could drive DSEM to the critical state DSEM has been extended to allow the agents to adjust their *preferences* (news response and price response parameters) when their current choices are performing poorly. This is discussed further below.

Another interesting consequence of the observation that the price response is centered around $r_p = 1$ is that—if DSEM is at all meaningful—real investors do watch (and base decisions on) trends in stock prices. In DSEM, to get realistic behaviour, even the least responsive agents had to have $r_{lo} \geq 0.5$ which can roughly be interpreted as the perceived autocorrelation between successive returns. DSEM suggests that there do not exist any (pure) *fundamentalist* traders (who respond only to fundamental information about the company and are unconcerned with the stock's price movements) in real markets. Unfortunately, while an interesting hypothesis, it is not clear how this assertion could be tested empirically.

7.3 Relation to other work in the field

Quite a few market models have been developed over the last few years. In this section some of these models are contrasted with CSEM and DSEM.

We begin by comparing how the price is chosen in the models. Recall that in CSEM the price is set by an auctioneer in order to balance supply and demand. In DSEM, however, the price is simply the most recent trading price. In most of the models reviewed the price was set by an external *market maker* as in CSEM [17, 27, 28, 30–36, 96–99] the only exceptions being reaction-diffusion models [47, 100] in which buyers and sellers diffuse in price space and a trade is executed when they meet. DSEM provides a new mechanism for allowing the price to emerge directly from the agents' decisions.

Another major difference between the CSEM and DSEM is in how they are updated: in CSEM trades are executed synchronously, once per day while DSEM allows trading in real time, with agents choosing their own activation times. On this front it appears that asynchronous updating is becoming more prevalent [17, 32, 34, 98] with more of the older sources choosing parallel updating [27–30, 33, 35, 47]. This is fortunate because a mounting volume of evidence suggests that parallel updating may introduce spurious artifacts into simulation dynamics [14, 49–52].

The preferred litmus test for each of these models is whether they can reproduce fat tails in the price return distribution and many of them can [30, 31, 33, 34, 96–99, 101].

The Cont-Bouchaud percolation model [31] has received a great deal of attention lately [34,96–99]. It is characterized by a network of information which produces herding effects. The advantage of the model is that analytic results exist [31, 101] which predict that the price return distribution should have a (truncated) power law distribution (with a scaling exponent $\alpha = 3/2$). It has also been demonstrated to exhibit clustered volatility [97, 99]. DSEM provides an alternative explanation which does not require herding. However, it would be interesting to know what the consequences of herding would be, which brings us to directions for future research.

7.4 Avenues for further work

I conclude this thesis with some thoughts on how DSEM may be extended to produce new insights and on further statistical properties which could be tested:

As discussed above, one of the most pressing issues is whether scaling and clustered volatility can emerge spontaneously without requiring tuning of the price response parameters. This can be tested by allowing the agents to choose their preferences (response parameters) as they see fit. To do so, a *meta-strategy* is required which controls when an agent adjusts its preferences and by how much. An arbitrary but reasonable choice is to allow preference adjustments when the agent's performance is demonstrably poor: for instance, if the agent sells shares at a price below the average price it bought them for. When this occurs the agent randomly shifts its preferences by some amount. This has been recently coded into DSEM and research is ongoing.

Another interesting direction to explore is the extension of DSEM to support multiple stocks. This idea was inspired by Bak et al. [47] in which they described adding a new stock as adding a new dimension in price space. It is well known that the dimensionality is one of the few factors which can impact the character of a critical point [61] so it would be interesting to see how the critical point in DSEM would be affected.

On the surface CSEM and DSEM are quite different. However, it should be possible to modify DSEM such that all trades are handled by a centralized control or market maker. The agents could respond to orders called out by an auctioneer in similar manner to CSEM. Discovering whether scaling and clustered volatility are robust to these changes would be very informative.

On the experimentation side, there are a number of statistical properties which could be tested for. One of these is an asymmetry between up- and down-movements in the price series. Roehner and Sornette [102] found that peaks tend to be sharp but troughs (lows) tend to be flat. Since DSEM is symmetric in its response to up- and down-moves it would be surprising if this asymmetry could be replicated.

Another interesting property which is currently being tested (but did not make it into this thesis) is Pareto's law for the distribution of incomes which states that the richest segment of the population have incomes in excess of that predicted by the log-normal distribution (which fits the majority of the population). This is thought to be an amplification effect whereby the richest individuals are able to leverage their wealth to increase their income faster than others [83]. Data are being collected to test for this effect in DSEM.

Beyond that, the price series contains more information than just the distribution of returns. For instance, the intra-trade interval and bid-offer spread are also interesting with testable distributions [48].

Finally, evidence is mounting that the distribution of empirical returns is truncated by an inverse cubic power law [6,7,69] rather than the exponential assumed in Section 5.1. It would be useful to determine which hypothesis DSEM obeys. To do so, much larger datasets are required in order to determine the distribution of very large returns (since it is difficult to distinguish the two hypotheses on scales studied in this dissertation). Alternatively, the moments of the distribution could be explored: if the exponential truncation holds then all moments should be finite but the inverse cubic implies that the k-th moment should diverge as the index k increases to three. Either way, it would be valuable to determine if the distribution of returns in DSEM is truncated by an inverse cubic as appears to be the case for empirical data.

In short, many exciting possibilities remain for future research into DSEM.