Physica Scripta. Vol. T106, 41–47, 2003

Statistical Mechanics of Conventional Traders May Lead to Non-Conventional Market Behavior

Lev Muchnik* and Sorin Solomon**

Racah Institute of Physics, The Hebrew University, Jerusalem 91904, Israel

Received April 2, 2003; accepted April 7, 2003

PACS Ref: 89.65.Gh, 87.23.Ge, 89.20.–a, 89.20.Ff, 89.75.Kd

Abstract

We describe the main idea and the conceptual architecture of a platform for simulating a large number of asynchronously interacting agents in continuous time. We show how the generic capabilities of the platform apply to the simulation of realistic stock market interactions. A particular example of a very dramatic market event that took place in Financial Times Stock Exchange (FTSE) on September 20, 2002 is used to uncover the parameters characterizing the classical investor types within the market. The simple microscopic rules governing the individual agents behavior are shown to result in a collective market behavior similar to the one of a damped harmonic oscillator. Specifically, the aggregated influence of the fundamentalist traders is formally related to Hooke's law while the behavior of the trend followers corresponds to inertia and viscous friction forces.

1. Introduction and background

Over the last few years the emergence of complex phenomena has been studied in a wide range of systems. A generic approach is the use of computer simulations. The development of computer power and computational methods allow realistic, fully controlled virtual ("computer gedanken") experiments.

In the present paper we report the application of this general approach to stock market dynamics. The stock market is from a physicist's point of view the largest, most well tuned, efficient and well maintained emergence laboratory in the world. Measurements in this laboratory are performed, recorded, transmitted, stored and documented with perfect precision. The challenge is to develop the tools for simulating, analyzing and predicting the outcome of experiments in this extraordinary laboratory.

A variety of simplified microscopic models of the stock market have been introduced over the last decade [1–27]. Most of these models focus on specific aspects of the problem: basic features of the trader's behavior or of the stock exchange procedures. They show that even a small set of simple assumptions can explain the set of "stylized" experimental facts characterizing generically the market [28]: power (Pareto–Zipf) laws, fat tails (and/or Levy-stable distributions), (multi-) fractal dynamics (Hurst exponents), long range correlations (clustered volatility), criticality (scaling exponents).

It is obvious that in order to go beyond these generic "stylized" facts, one has to consider more realistic features [29]: detailed stock market procedures, individual trader behavior, communication lags, external events (news arrival, economic fundamentals, etc.) [8,13]. Understanding complex behavior usually boils down to finding a limited set of relevant simple microscopic features. However, what are the relevant microscopic details is obviously not known in advance. Consequently, the simulation has to introduce and control an—in principle—arbitrary number of such details and test their respective relevance.

To achieve this we developed a platform that simulates an arbitrary number of traders that interact with an arbitrary range of behaviors. On the resulting virtual stock market one can emulate the behavior of the real markets and confront it on a trade by trade basis with data from on-line measurements.

The platform can be exploited to achieve several purposes:

- Experiment with the effect on the market behavior of various features and events.
- Compare the efficiency of different trading strategies.
- Isolate the influence of (groups of) traders' strategies on the market.
- Study the co-evolution of traders' behavior.
- Find ways to improve market efficiency and stability.

In the present paper we show how the platform can be used on a particular interesting case study.

We start with a description of the non-trivial architecture of the simulation platform.

2. Simulation platform

The main difficulty in building an efficient and realistic simulation environment turned out not to be the proliferation of traders' parameters. Rather, one encounters a more fundamental obstacle that is characteristic generically to computer-based simulations and that to our knowledge was not satisfactorily addressed until now: the spatially-distributed, time-parallel structure of reality. We will describe below the conceptual challenge and our solution to this problem.

2.1. *Computer simulation of continuous time asynchronous systems*

Turing conceived computers as sequential discrete machines: one single operation per computing cycle updating one single quantity. The structure of the real world is totally different: at any time, any of the basic elements can change (and initiate a sequence of further changes unto the others). This introduces two problems:

* e-mail: lev@topspin.co.il
** e-mail: sorin@cc.huji.ac.il

- One does not have an *a priori* unit time step that would map the real time onto a discrete simulation step: events can take place at any time.
- One does not know in advance which is the next player that will act and therefore one cannot have a pre-ordered list of future operations (in fact the action of each player is likely to affect the timing and the identity of the element/agent performing the next operation).

Traditionally, the approximations assumed that changes can take place only at fixed (pre-appointed) times. The problem of independent agents discretely and asynchronously interacting in continuous time was not solved until the present work. We feel that beyond its relevance for the stock market, this problem is crucial for the reliable computer simulation of a very wide class of emergent systems [38].

2.2. *The Gedanken football match*

To make the problem clear and to liberate the ideas from the clutter of finance market technicalities, we describe in this section the main concept in a context that is universally well known: football. We hope that the apparent "lightness" (game-like) nature of the present example will not divert the reader from the seriousness of the problem and the value of its solution.[1]

Let us list schematically the features characterizing a football match from the conceptual point of view: One has 25 independent agents (the players and the referees—we neglect for the moment the otherwise crucial feedback from the public) interacting in the following (feedback-driven way):

1. Each agent has a personal perception of the current game situation (that might include history elements).
2. Each agent can revise at discrete times his perception of the situation (the times at which these revisions take place depend on the individual agent and the situation).
3. Each agent decides on responses to his new perception of the situation. The time it takes the agent to reach a decision depends on the agent and the situation.
4. After taking a decision, the agent acts upon it. The resulting actions are, of course, not necessarily the ones decided/intended: they are influenced by the laws of reality (physical or otherwise). Also in the time interval between the internal decision and the action, the situation might change (the other agents may have acted in the meantime).
5. The actions of the agents, together with external factors (physical environment, field and ball condition, time passing between the steps 2–4), change of the objective situation in the game.

The element closing the dynamical loop is the fact that the changes (or non-changes) in the objective situation (step 5) can lead to revisions in the perceived situation by each of the agents (step 2).

[1]According to Professor Stauffer—the honored person of this conference—real men drink beer, watch football and write Fortran. Which definitely connects the football game to the Turing machines (especially after a few beers).

How would one faithfully simulate this model on the computer?

One possibility, analogue to the molecular dynamics procedure in physics and chemistry is to divide the time in a very fine mesh. For each moment of this fine mesh one checks for changes in the objective situation. In this case, one has to poll each agent and check for changes in his perceptual-decision-actuation internal state. The time unit in molecular dynamics is decided by the ratio between the spatio-temporal features of the interaction potential and the typical speed of the molecule. In the present case, the ball features are of the order of 10^{-1} m and the speed of the order of 10 m/s. Therefore, the time step has to be significantly lower than 10^{-2} s. Even a delay of 10^{-2} s in accessing the ball can result in missing it (goaaaal!).

With such a small time step, one cannot afford to poll too many players at each time step. The situation becomes unbearable if one introduces the supporters in the system. If one has 10^5 fans, each with a probability rate of acting of $O(10\,s^{-1})$ then, in the interval between two player actions, one can have up to $O(10^5)$ interventions. To take it faithfully into account one is forced to a time step finer then $O(10^{-6})$.

To make things worse, the molecular dynamics technique accumulates errors as the number of steps increase. Therefore, with a microscopic step-size, the simulation becomes unreliable after a very short real time.

2.3. *What is our solution to this tantalizing situation?*

Before getting into technicalities, let us describe the main idea:

Every player (or referee, or fan in the tribune) maintains his perception of the situation and updates it from time to time. The times at which the perception of a player is updated are not necessarily known in advance, they are triggered in one of two ways:

- spontaneously by some (possibly random, and/or situation-dependent) internal clock of the player,
- by an important enough change in the perceived external situation: the ball being sent to him or menacing to reach the gate (goal). The reaction threshold depends on the current sensitivity profile of the individual player, and the situation, and may involve stochastic elements.

As a perception update is triggered, the player may respond with a decision for action. The decision then reaches (after a realistic time depending on the player and the conditions) the effector muscles and the physical objects (own body, ball, other players) generating new events (that will then generate new reactions by players). Note that the program is not actively simulating the time intervals between events: it jumps directly from one event to the subsequent one advancing appropriately the simulated time (which has nothing to do with the ticks of the machine clock).

External conditions (e.g., the coach calling on one of the players to go to the bench, rain drops falling on the players) can be treated as agents too.

Note that the crucial idea is that instead of polling each agent at each time unit, one leaves in the charge of each agent to "wake up" according to its own criteria (time passage,

certain (possibly unexpected) events taking place, etc). Therefore, no simulation effort is performed at times when nothing happens and no polling of momentarily inactive players is necessary.

The order of the agents' interventions is not pre-defined. In fact, it can change in every moment: an unexpected event may wake up an agent promptly that otherwise planned to be quiet for the next 90 minutes. Yet the order is not random, it respects in detail the causality rules that nature and the players' behavior dictates to the system.

Let us now explain how we realize this promise.

3. Description of the algorithm; the World Manager

Let us now describe the asynchronous continuous time algorithm in a generic context.

The description above is correct at the conceptual level: the simulation relies on an event-driven mechanism (the evolution of the system is the sequence of events generated by the dynamical agent behavioral rules). However, on a sequential machine it is not realistic to expect the process representing each player to send interrupts to the CPU. Moreover, those interrupts' real timings would have nothing to do in general with the timings of the simulated events.

A way to avoid this problem is to introduce the "World Manager". The World Manager is the core of the platform and has the following tasks and features:

- The World Manager is the repository of the full information of the current system state.
- The World Manager is in charge of updating the current system state by acting changes according to the decisions that the agents passed to him.
- In our simulation, as in reality, action is not applied immediately but is delayed (immediate action is technically allowed but, of course, not realistic).
- The World Manager is the repository of an ever-changing events schedule. At any moment, an event that used to be the first in line for execution might be "scooped" by a faster, newly introduced event.
- The World Manager is aware of all the decisions and conditions for performing an action, therefore, he has all the information necessary to decide which is the next event to take place and the time at which it will take place (even though the waiting list may have been reshuffled by the latest implemented event).
- Note that as soon as the latest event is completed, the simulation advances the time to the next event on the list. Therefore in the most rigorous way, and without any approximation, the World Manager can advance the time of the simulation (without spending any additional CPU time) to the time of the next event.
- The World Manager receives, together with every decision from an agent, the agent's current threshold criteria, i.e., the conditions necessary steer the agent into new decision making (the "steer-up" criteria may include sheer time passing from the last decision attempt).
- This gives the World Manager the full necessary information to decide whether to "steer up" the agent for new decisions after each change in the system.
- Consequently, every time that a change in the current system state is executed by the World Manager, the World Manager sends "decision steering" signals to the agents whose thresholds are exceeded. The other agents are left alone (and no CPU is spent on them).
- Each of the informed (steered) agents can choose to either neglect or respond to the event. The usual response is a decision for some action. The actual action/change is to be implemented on the system at the appropriate (delayed) time by the World Manager. Also modifications in the "steer-up" thresholds may be decided at this stage.

One should contrast the present approach with most of the other simulation schemas (like molecular dynamics or Monte Carlo family simulations) which sample the system evolution at discrete times causing systematic inaccuracy. In our simulation the events take place at any real-valued timing required by the realistic continuum-time character of our dynamics. In our formalism, the time scale can be viewed as continuous, and not composed of small adjustment time slices. The simulation supports discrete events, which happen at precisely the moment they are supposed to happen—not at the closest time tick.

As in continuum time reality, there is no problem here with events taking place exactly at the same time (even if several chains of events take place during overlapping time periods, the realistic delays in the decision actuation imply that they are taken care of in the most rigorous and efficient way).

It is important to clarify that the asynchronous event-driven simulation rate does not depend on the number of agents but on the number of actions (events). This allows extremely effective implementation for systems with a huge number of components.

The fact that agents are normally supposed to submit their request for action and wait until the time arrives does not make them inaccessible during this time. They can still respond to any event they receive.

4. Oscillating stock market model example

Once the simulation platform is constructed it can be used for the study of the stock market through computer experiments paralleling real market measurements.

It is a common approach in experimental physics to study the forces governing the system behavior by throwing it out of equilibrium and watching its return to the equilibrium state: one has to disrupt a pendulum in order to measure its mass, elastic constant and friction. Accepting the same approach to the stock market study led us to the following line of thought: Since in most cases, the system is in equilibrium (or at least very close to it depending on which version of the market efficient hypothesis one favors) the forces acting on it cancel effectively and are not fully revealed. In order to see those forces in action and to study them, one has to watch for some abnormal event that dragged the system far from equilibrium.

It is usually very costly to bring a large stock market out of equilibrium so as to undo the efficient market hypothesis even for a short time. One has rather to look for anomalous events capable of kicking the market far from its steady state (further than the usual fluctuations size). With enough luck one may be able to identify a colossal mistake such as an order erroneous by a factor of 100. Events like this are rare (someone has to pay for them, and usually the same person does not get the opportunity to make them twice). But quite a number of them were reported over the last years. Interestingly enough most of the errors were due to the use of computers (e.g., keystroke error [32]) which were introduced in the first place to avoid them. In one case, an extra zero was accidentally added to a sell order and one executed a £$300 \cdot 10^6$ sell order instead of the intended £$30 \cdot 10^6$ [32]. The order went through just before the market was closed, causing a 3.5% Financial Times Stock Exchange 100 Share Index (FTSE 100) crash. The mistake did not have long-time influence on the market: it recovered next morning.

The event which we wish to discuss below took place on September 20, 2002 at the London Stock Exchange [30–35]. The event's total duration was around 20 min. and lead within minutes to losses of the order of £$100 \cdot 10^6$ [33]. The entire trade volume (£$3.2 \cdot 10^9$) of the event was larger than the volume of an average trading day. The trade activity (trades/second) was so high that some computer systems failed to cope and prices of shares were delayed. The event began at 10.10 am and within 5 minutes the FTSE 100 index rose from 3,860 to 4,060. Within another few minutes, the index fell to 3,755. After some oscillations, the index returned to the initial value at the end of the 20 min.

It is believed that the event was triggered by a huge order (£$1.2 \cdot 10^9$) mistakenly submitted twice (or even three times) causing all trade participants to try to exploit the opportunity.

The price evolution throughout the event is sketched in Fig. 1. What drew our (and others) attention were the actual oscillations before the return to equilibrium. This graph was likened to the vibrations of a muffled guitar string [30,31]. We suggested a theoretical explanation of it. We claim that once the initial mistake is postulated, one can explain the rest of the market behavior solely on the

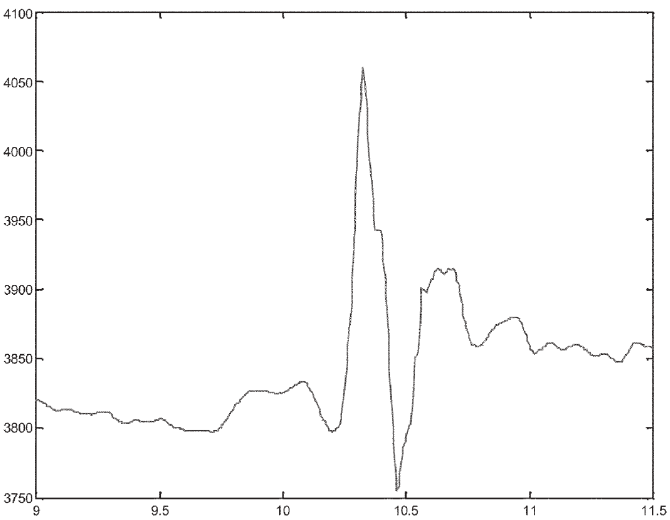

Fig. 1. FTSE 100 fluctuations as a result of a trading mistake on September 20, 2002.

basis of the usual forces known to act within the stock market. The fact that indeed the trigger was a mistake is confirmed by the eventual return (within 20 min) of the market to the initial price.

It turns out that the usual trader profiles that are traditionally invoked to explain the usual market behavior are sufficient to explain the oscillations. The only necessary additional ingredient is a dramatic event to take place (or the perception of such an event). The event is needed to take the market out of equilibrium without affecting its fundamental dynamical parameters. In the present particular case, this was more evident since the actual fundamental value was not affected by the mistaken order. In the generic case, where the event (war, loss/gain warnings, executive resignations) affects the future expectations, the oscillations take place on the background of a long term change in the average price and the oscillating effect is somewhat masked (though still discernable). Rare cases as the present one constitute important occasions to measure the relevant market constants parametrizing respectively:

1. The inertia of the market (corresponding to the mass of the oscillator);
2. The reversion to the mean (corresponding to the oscillator elastic constant); and
3. The friction (damping force).

In fact each of these parameters could be associated with some classical investor types:

- **Fundamentalists** believe in the market efficiency and that even if the current value of the price P_{market} is far from the fundamental one $P_{fundamental}$, the market will eventually revert to the fundamental price. However, they realize that this can take some time and therefore they have to discount (K) for it. The greater the difference between the P_{market} and $P_{fundamental}$, the greater will be the difference between P_{market} and the price P_{order} offered by the fundamentalists.

  Hence, the price of the order they suggest will be determined by:

$$P_{order} = P_{market} - K * \left(P_{market} - P_{fundamental}\right). \quad (1)$$

- **Chartists**, unlike fundamentalists, never care about the real stock value. They make their money by following the market trend. Chartists believe that any trend they discover can be exploited to predict the future stock value by extrapolating its future value $P_{order}(t_+)$ based on the recent past $P_{order}(t_-)$ and the present $P_{order}(t_0)$. The simplest choice for short times is a linear extrapolation (though this does not do justice to the sophisticated voodoo that they are actually practicing; the appropriate name would rather be "trend-followers"):

$$P_{order}(t_+) = P_{market}(t_0) + (1 - C)(t_+ - t_0) \cdot \frac{P_{market}(t_0) - P_{market}(t_-)}{t_0 - t_-}$$

The constant C expresses the expected departure from the present trend. For $C = 0$ the traders assume that the present trend will continue unchanged. For $C > 0$, they assume the trend might reverse in the next $\Delta t = t_+ - t_0$ interval. For $C < 0$, they assume that the trend might strengthen. In fact a panic state, where the trader is afraid that the market is going through a very significant change may have a $C \ll 0$ which is bound to lead to the amplification of the present trend. In this case, we call C the "panic parameter".

- The third kind of trader is the **noise-trader**. The noise-traders are making "random" offers at a price P_{order} randomly distributed around the current price P_{market}.

  $P_{order} = P_{market} + s \cdot RANDOM(-1, 1)$.

  In general, they have an overall effect of damping the oscillations (because they dilute the signal generated by trends), but in the presence of panicky traders, the random bids of the noise-traders may trigger an avalanche (as described in the next section).

To simplify the experiment and limit the number of free parameters we made the following generic assumptions and applied them to all traders:
All three types of traders submit limit orders only.

1. All order volumes are equal.
2. The only way to increase the chance of their order to be satisfied, traders must lower the price of the sell order or raise the price of the buy order.
3. In order to maximize their revenues, trader activity boosts when they feel the market is not at its optimum.
4. Traders act independently from each other.

During normal times, the market is calm: the fundamentalists believe that the fundamental price is somewhere in the neighborhood of the market price, the trend followers have nothing to follow and the noise traders are just inducing limited, incoherent fluctuations. Consequently, in this model, the market is boring and predictable until some erroneous order (limit order with a price high above the fundamental one or—similar to what happened on September 20, 2002—a huge market order) is produced. This single order triggers a series of responses from the participating traders that cause the entire market to oscillate.

It is important to point out that there is absolutely no direct interaction between traders. In this model every microscopic element acts autonomously. The only interaction they experience is produced through the traders' collective action on the market.

Before presenting the results of our realistic simulations of the model above let us estimate approximately, in terms of the various trader types the average (without noise) price evolution.

Let us assume that the fundamental price does not change significantly during the brief event that we are studying:

$$P_{fundamental}(t) = P_0.$$

Moreover we assume for simplicity that there are in average M orders per time interval dt. By considering Eqs. (1) and (2), one obtains the following increment in the market price evolution:

$$x(t + \mathrm{d}\ t) - x(t) = -K \cdot M \cdot x(t) + (1 - C) \cdot M \cdot [x(t) - x(t - \mathrm{d}\ t)]. \quad (3)$$

Then, the passage from discrete to continuous equations is obtained by defining:

$$x(t) = \frac{[x(t + \mathrm{d}t) - x(t - \mathrm{d}t)]}{2 \cdot \mathrm{d}t}$$

and

$$x(t) = \frac{[x(t + \mathrm{d}t) - 2 \cdot x(t) + x(t - \mathrm{d}t)]}{\mathrm{d}t^2}.$$

By solving these 2 linear equations with respect to $x(t + \mathrm{d}t)$ and $x(t - \mathrm{d}t)$ one can express $x(t + \mathrm{d}t)$ and $x(t - \mathrm{d}t)$ in terms of $x(t)$, $x(t)$ and $x(t)$ to obtain:

$$x(t + \mathrm{d}t) = \tfrac{1}{2} x(t) \cdot \mathrm{d}t^2 + x(t) + x(t) \cdot \mathrm{d}\ t$$

and

$$x(t - \mathrm{d}t) = \tfrac{1}{2} x(t) \cdot \mathrm{d}t^2 + x(t) - x(t) \cdot \mathrm{d}t$$

Substituting these values in Eq. (3), and taking $\mathrm{d}t \to 0$ (and rescaling appropriately the constants M, C and K), one gets for the market price the differential equation:

$$mx(t) + c \cdot x(t) + k \cdot x(t) = 0.$$

This is the damped harmonic oscillator and its parameters (originating in the individual traders behavior) determine the global market behavior: coarsely speaking, the period of the market oscillations is expected to be $4\pi m/\sqrt{4km - c^2}$ and the amplitude of the oscillations is expected to decay exponentially as: $e^{\frac{-c}{2m}t}$.

For the appropriate parameter values, one obtains oscillations similar to the actual event of September 20, 2002. In fact Fig. 2 shows the output of the simulation platform for a realistic choice of the individual trader parameters. Except for the fluctuations induced by the

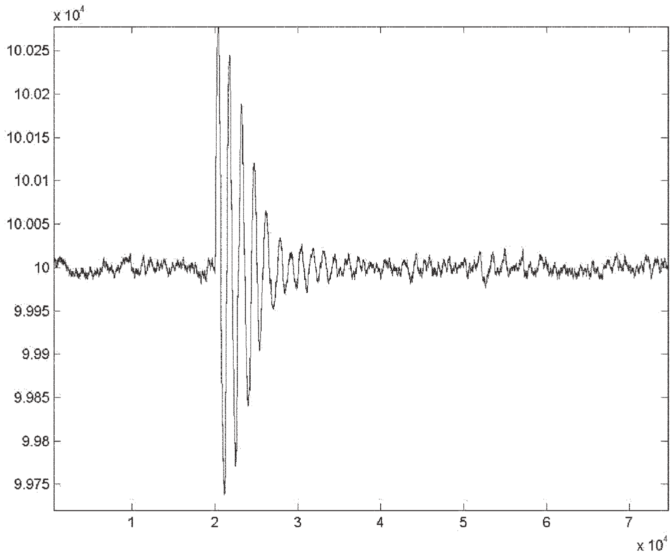

Fig. 2. Externally triggered oscillator-like simulated market price evolution.

noise traders, the simulation presents quite faithfully the oscillations predicted by the theory.

For panic periods, the viscous friction c is negative and the amplitude increases. This is exploited by the computer experiment described in the next section.

4.1. *Spontaneous price oscillations*

In this section we alter the model defined in the section above to correct the two obvious problems it possesses:

- The model requires an unnaturally large fluctuation to initialize the fluctuation pattern. While in the FTSE event of September 20, 2002 this seems to have been the case, in general large endogenous fluctuations might appear even without being triggered by a bid error.
- The fluctuations are only decreasing, while in some realistic events, there is an initial phase in which the fluctuations amplitude increases.

The only change one has to introduce to have these effects is to connect the extrapolation coefficient of the trend followers C to the actual market dynamics $P_{order} - P_{market}$.

More precisely, we introduce a panic factor C that depends on the actual speed and size of the present trend. It is well known [36,37] that the perception of the market changes is affected by these factors: C may take large negative values $C(t)$ once the investors are scared by a fortuitous random fluctuation or by scary (but not necessarily fundamental-value-changing news). If this happens, the market will spontaneously enter a period of increasing oscillations, until either some of the investors get out of the market (producing a crash), or the market calms down gradually (by external intervention or by itself).

Therefore, in such cases, one does not require a special event to displace the market price far from its normal values. One can imagine that nervous traders (the ones that do not know what to expect and are ready to follow any trend) interpret some larger random fluctuation as a tendency and amplify it. On the other hand, it does not take them long to realize that the price deviation is not going to cause any long-lasting effects and to "bring" it back to the fundamental price. As a result, the simulation on the market platform (and in reality) displays a market price evolution of the type shown on Fig. 3.

One should keep in mind that the above simple applications are only the beginning of the experimental program that the market simulation platform will support.

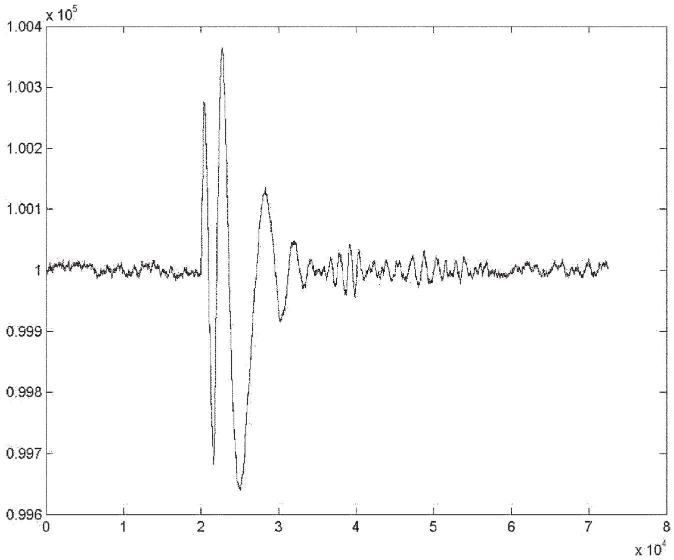

Fig. 3. Spontaneous price oscillations in a simulation with trend-dependent panic parameter C.

5. Conclusions

In this paper we present a general method for simulating the spatially distributed continuous asynchronous dynamics of real systems faithfully and efficiently in a sequential discrete state machine. We describe the main idea and its application in the context of stock markets. As an application we show that the usual types of simple trader behavior are sufficient to explain the colossal oscillations that took place in the morning of September 20, 2002 at FTSE. Moreover we indicate how similar events can help estimating the parameters governing the market behavior for generic market instances.

We hope to be able to present in the future more systematic and exhaustive market studies and expose more complex emergent phenomena using our generic simulation platform.

Acknowledgement

Research supported in part by the Israeli Science Foundation.

References

1. Muchnik, L., Slanina, F. and Solomon, S., The Interacting Gaps Model: Reconciling Theoretical and Numerical Approaches to Limit-Order Models. To be published in Physica A.
2. Bak, P., Paczuski, M. and Shubik, M., "Price Variations in a Stock Market with Many Agents," Physica A, Dec (1997) cond-mat/9609144.
3. Stauffer, D., Int. J. Modern Phys. C **11**, 1081 (2000).
4. Caldarelli, G., Marsili, M. and Zhang, Y.-C., Europhys. Lett. **40**, 479 (1997).
5. Maslov, S., Physica A, **278**, 571 (2000), cond-mat/9910502.
6. Slanina, F., Phys. Rev. E **64**, 056136 (2001) cond-mat/0104547.
7. Levy, M., Levy, H. and Solomon, S., Economics Lett. **45**, 103 (1994).
8. Levy, M., Levy, H. and Solomon, S., "Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena." ISBN 0124458904, hardback, 251 pages (Academic Press, Published August 2000) http://www.elsevier-international.com/catalogue/title.cfm, ISBN=0124458904.
9. Levy, M., Levy, H. and Solomon, S., J. Physique **I**, 1087 (1995).
10. Solomon, S., "Generalized Lotka Volterra (GLV) Models of Stock Markets," p. 301 in "Applications of Simulation to Social Sciences," (edited by G. Ballot and G. Weisbuch), (Hermes Science Publications 2000) cond-mat/9901250.
11. Solomon, S., "Pioneers on a new continent on physics and economics," Quantitative Finance. February (2003).
12. Palmer, R. G., Arthur, W. B., Holland, J. H., LeBaron, B. and Tayler, P., Physica D **75**, 264. (1994).
13. Moss de Oliveira, S., de Oliveira, P. M. C. and Stauffer, D., Evolution, Money, War, and Computers-Non-Traditional Applications of Computational Statistical Physics, (Teubner, Stuttgart–Leipzig 1999), ISBN 3-519-00279-5.
14. Galluccio, S., Caldarelli, G., Marsili, M. and Zhang, Y.-C., Physica A **245**, 423 (1997).
15. Mantegna, R. and Stanley, H. E., "An Introduction to Econophysics," (Cambridge University Press, 1999) ISBN: 0521620082.

16. Maslov, S., Mills, M., Physica A **299**, 234 (2001), cond-mat/0102518.
17. Challet, D. and Stinchcombe, R., Physica A **300**, 285 (2001), cond-mat/0106114.
18. Luckock, H., "A Statistical Model of a Limit Order Market," Sidney University, preprint (28 September, 2001).
19. Willmann, R.D., Schuetz, G. M. and Challet, D., Physica A **316**, 430 (2002), abs/cond-mat/0206446.
20. Beja, A. and Goldman, M. B., J. Finance **35**, 235 (1980).
21. Bouchaud, J.-P. and Cont, R., Euro. J. Phys. B **6**, 543 (1998). http://www.cmap.polytechnique.fr/~rama/papers/langevin.pdf.
22. Kempf, A. and Korn, O., J. Financial Markets **2**, 29 (1999), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=36092.
23. Zhang, Y.C., Physica A **30**, 269, (1999) http://arxiv.org/PS_cache/cond-mat/pdf/9910/9910072.pdf.
24. Biais, B., Hilton, P. and Spatt, C., J. Finance **50**, 1655 (1995), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=5586.
25. Biais, B., Hilton, P. and Spatt, C., J. Finance **50**, 50 (1995), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=6970.
26. Bouchaud, J.-P., Mézard, M. and Potters, M., Quantitative Finance **2** 251 (2002) cond-mat/0203511.
27. Lux, T. and Marchesi, M., Nature **397**, 498 (1999), http://www.ge.infm.it/~ecph/papers/lux/lux-marchesi.ps.gz.
28. Lux, T., Heitger, F. and Takayasu, H., "Micro-Simulations of Financial Markets and the Stylized Facts. The Empirical Science of Financial Fluctuations: Econophysics Approach on the Horizon," (Berlin: Springer, in press).
29. Solomon, S., Computer Phys. Commun. **121** 161 (1999), http://xxx.lanl.gov/abs/adap-org/9901003.
30. Ball, P., Nature 01/10/(2002), http://www.nature.com/nsu/020923/020923–18.html.
31. Frauenfelder, M., How Wall Street Is Like a Guitar String, Business 2.0, March (2003), http://www.business2.com/articles/mag/0,1640,47145,00.html.
32. UK shares reverse "blip", 15 May, 2001, BBC, http://news.bbc.co.uk/1/hi/business/1331167.stm.
33. English, S., Footsie plunges on single trade 15/05/(2001), Telegraph, http://www.telegraph.co.uk/news/main.jhtml?xml=/news/2001/05/15/nfste15.xml.
34. Smith, C., "The Brave New World of Screen Trading, Commodity Trading Consumer Research," http://www.ctcr.investors.net/guide/ptscreen.htm.
35. Treanor, J. and Hume, N., "Bank loses J100 m in two-minute frenzy," 21 September, (2002). The Guardian, http://www.guardian.co.uk/business/story/0,3604,796284,00.html.
36. Kahneman, D. and Tversky, A. (Eds.), "Choices, Values and Frames," (New York: Cambridge University Press and the Russell Sage Foundation, 2000).
37. Thaler, R., "Quasi Rational Economics," (Russell Sage Foundation, New York, 1994).
38. The importance of being discrete: Life always wins on the surface, Shnerb, N. M., Louzoun, Y., Bettelheim, E. and Solomon, S. Proc. Natl. Acad. Sci. USA, **97**, 10322 (2000) http://xxx.lanl.gov/abs/adap-org/9912005.

