Random time in Agent-Based Market Models

Gilles Daniel

Computer Science Department

University of Manchester, UK

gilles@cs.man.ac.uk

Abstract

Agent-based models of financial markets traditionally adopt a discrete-time approach to represent the interactions between agents, mainly because financial time series commonly used by Economists and practitioners are available on a daily basis only. Nevertheless, one cannot discard the intraday activity, more difficult to observe but probably of importance to explain the global dynamics of the markets: new information arrives, traders update their beliefs, and prices move constantly. To address this issue, Boitout and Delahaut recently built upon an existing popular model and introduced the notion of random duration between asynchronous events. Their elegant approach, mixing discrete and continuous time, gives them a definite advantage when modeling clustered volatility in periods of intense intraday activity *and* allows them to compare their artificial time series with real ones.

Keywords: Agent-Based Market Models, discrete-time, continuous-time, discrete-events, volatility clustering, long range dependence of volatility

1 Introduction

In the last few years, Econophysicists have been trying to understand the dynamics underlying the price fluctuations of financial assets (mainly stocks, currencies and interest rates) by considering markets as complex adaptive systems made of interactive agents with bounded rationality [AAP88, BG94, ADL97, Wei99, MS99]. The aim of these so-called Agent-Based Market Models (ABMMs) is to propose a microscopic structure for the traders — how they form expectations, take decisions, learn, communicate, etc. — that could replicate, through extensive simulation, the macroscopic stylized facts emerging from their interactions and universally exhibited by real financial time series [Con01]. Such a model would then be considered as a plausible candidate to explain the inner dynamics of real financial markets.

A striking phenomenon prevalent in the ABMMs literature is the drastic simplification made about time. Whereas time is a crucial parameter appearing both at the microscopic and macroscopic level (in the modeling of agents but also in the statistical properties of the emerging time series), it is the rule rather than the exception to consider discrete-time models in which the time increment between two states of the system is arbitrarily set to one trading day [PAH$^+$94, CZ97, Art99, JHHJ00, GB03]. The obvious reason is to be found in the importance and popularity of daily data in financial markets, most stylized facts being observed and documented at this frequency or at its multiples (weeks, months, etc.). Nevertheless, choosing this particular value for the evolution of the agents themselves might represent a limitation of the models, since in reality the traders' beliefs, together with their strategies, wealth, etc., might evolve between two consecutive observations, as a result of their intraday activity.

To address this issue, Boitout and Delahaut [BD03] have recently proposed to mix continuous-time and discrete-time: as an improvement to a popular existing ABMM (Lux model, [Lux98, LM99, LM00]), they allow the trading time between two events to evolve randomly, but record the observable market variables (price, volume, etc.) at discrete calendar intervals only — typically at the end of each day. Thus, they can model the intraday activity of the market *and* compare their artificial time series with empirical ones. We investigate in this paper the supposed advantages introduced by this mixed approach, proposed for the first time in the context of ABMMs.

2 Daily stylized facts of financial time series

Strangely enough, financial time series exhibit some statistical properties universally shared by individual shares, stock indexes, foreign exchange rates, etc. [Pag96]. Needless to say, once discovered, those common signatures were under huge scrutiny by Economists and Investment Bankers, the latter seeing a good opportunity to make money by predicting the unpredictable: indeed, a prevailing theory since the 60's, the Efficient Market Hypothesis [F$^+$69], states that the price of financial assets immediately reflects all the information available, preventing anybody from constantly beating the market since the arrival of information is by definition unpredictable itself. As a matter of fact, price changes of usual assets are reported to exhibit very little autocorrelation over horizons longer than 20 minutes, and the hypothesis that the price would follow a random walk (presence of a unit root) can never be rejected. This was quite disappointing for practitioners, since it invalidated Technical Analysis (in absence of correlation between price changes, the use of historical data is useless to predict tomorrow's move) and simply asserted that the only way to raise ones expectations in the long term was to successfully forecast the arrival of new information. Nevertheless, other stylized facts are more promising; in particular, some long range dependence in the volatility, computed as the square or absolute value of the price returns, has been observed, together with a tendency for moments of high volatility to cluster in time. In other words, if the prediction of direct price returns appears to be a myth, investigating their absolute value might be worthwhile.

Those stylised facts are consistently studied on a daily, weekly, monthly or yearly basis, for the simple reason that data are usually available at these frequencies only. One just needs to open the central pages of the Financial Times to realise that information about intraday variables is almost nonexistent. Nevertheless, this should not prevent Econophysicists from modeling the behaviour and interactions of agents at higher frequencies, since in reality traders might update their beliefs or change strategy many times a day, new information arrives constantly and prices vary every minute or so. The model proposed by Lux was among the first to adopt this approach.

3 Discrete-time intraday activity

Our aim is not to explain extensively the Lux model, but to outline the differences between this model and its continuous-time version proposed by Boitout. As a consequence, we will expose only its main characteristics and direct the reader to [Lux98, LM99, LM00] for more details.

In this models, n traders are initially randomly divided into three categories: fundamentalists (n_f), who constantly expect a mean reversion of the price p toward its fundamental value p_f, optimistic chartists (n_+) — buyers influenced by the opinion of other traders and by the short term trend of the market, and pessimistic chartists (n_-) — chartists sellers. The key idea is that the main stylized facts would be generated by the intraday switch of traders between these three categories. In this context, an event is defined as one of the following:

1. a pessimistic trader becomes optimistic

2. an optimistic trader becomes pessimistic
3. a fundamentalist trader becomes optimistic
4. an optimistic trader becomes fundamentalist
5. a fundamentalist trader becomes pessimistic
6. a pessimistic trader becomes fundamentalist
7. the price moves up
8. the price moves down
9. the fundamental value moves up
10. the fundamental value moves down

Each of these ten events is associated with a time-varying hazard rate β_j, $1 \leq j \leq 10$, which depends on the current state of the market (i.e. the value of the market variables $n_+(t)$, $n_-(t)$, $p(t)$ and $p_f(t)$) and the recent price trend over a short time horizon. The time line is linearly divided into unit time steps t, representing trading days and used to record the market variables periodically. Finally, in order to account for the intraday activity of traders, each unit time step is divided again into 100 micro-intervals Δt, during which events occur asynchronously (one event only per micro-interval). The probability for an event j to occur during Δt equals $\beta_j \Delta t$, but only the quickest is executed.

The problem with this discrete-time approach of the intraday events is that in case of volatility burst, i.e. when the price varies brutally for a while, the maximum price change per day is artificially restricted by the number of micro-intervals. This limitation comes from the fixed number of events per day, and can be overcome only by momentarily augmenting this number, i.e. reducing the length of the micro-intervals. Hence, Lux reports that the time increment had to be manually reduced from $\Delta t = 0.01$ to $\Delta t = 0.002$ during periods of high volatility.

To avoid this rather *ad hoc* disruption in the simulation, and to account for extremely frequent events during volatility bursts, Boitout and Delahaut replaced the constant micro-interval Δt by the duration of the events themselves, enabling the number of events per day to fluctuate according to the activity, like in real-life. To put it in a nutshell, in order to introduce more flexibility in the way the intraday events are handled, they switched from a discrete-time to a discrete-events (or continuous-time) approach.

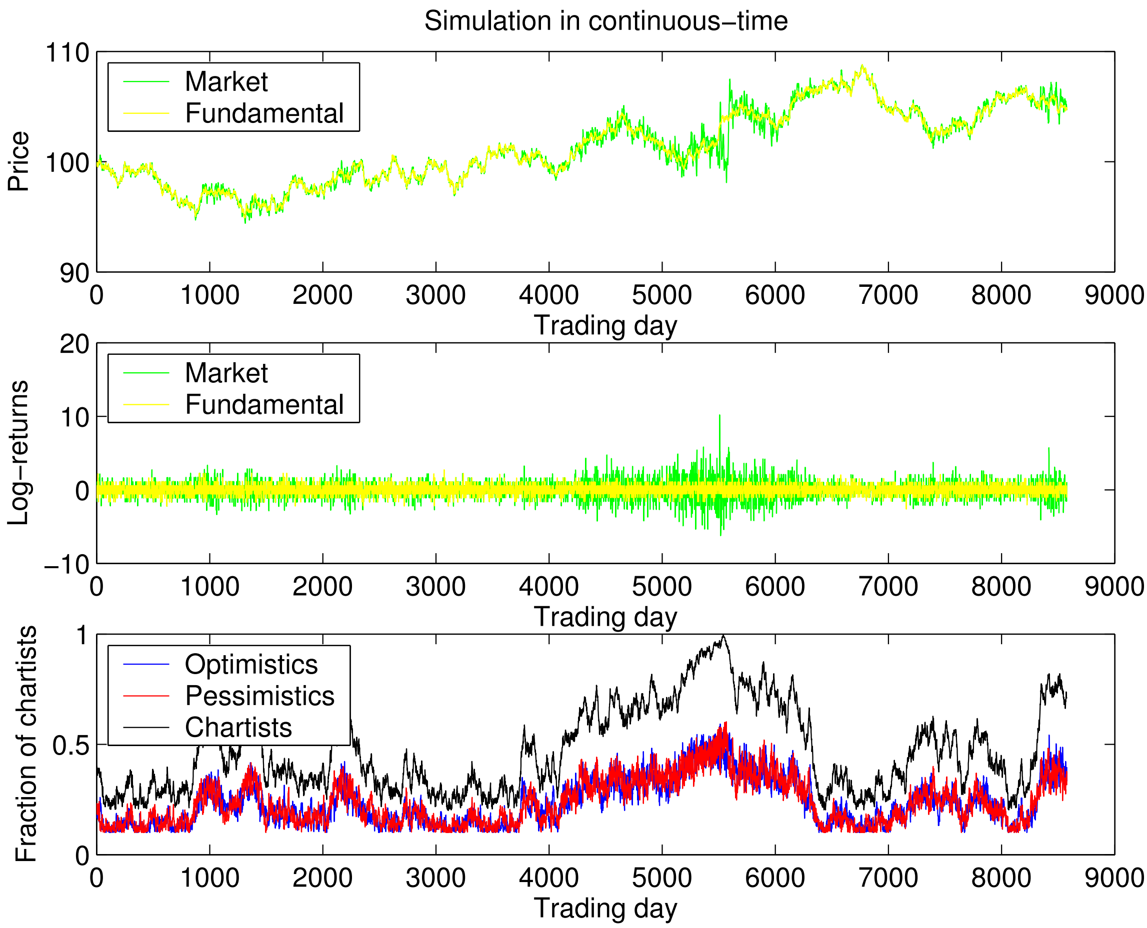

Figure 1: Simulated market with 100000 events, representing 8576 trading days. One can observe periods of calm, succeeded by bursts of volatility when chartists become prominent.

4 Continuous-time approach

We run a simulation of the Boitout model with the set of parameters used in [BD03], with 100000 events, and obtain a first time series $Y = [\frac{n_+(i)}{n}\ \frac{n_-(i)}{n}\ p(i)\ p_f(i)]$, $1 \leq i \leq 100000$ describing the market variables after each event, coupled with $E = [eventNumber(i)\ duration(i)]$, the number of the event executed and its duration, for every iteration. Using E, we can perform time aggregation and post-process Y to get the value of the market variables at the end of each day (a day has a constant length, in calendar time). The resulting time series, $calendarY = [\frac{n_+(t)}{n}\ \frac{n_-(t)}{n}\ p(t)\ p_f(t)]$, $1 \leq t \leq 8576$, can then be directly compared with daily data

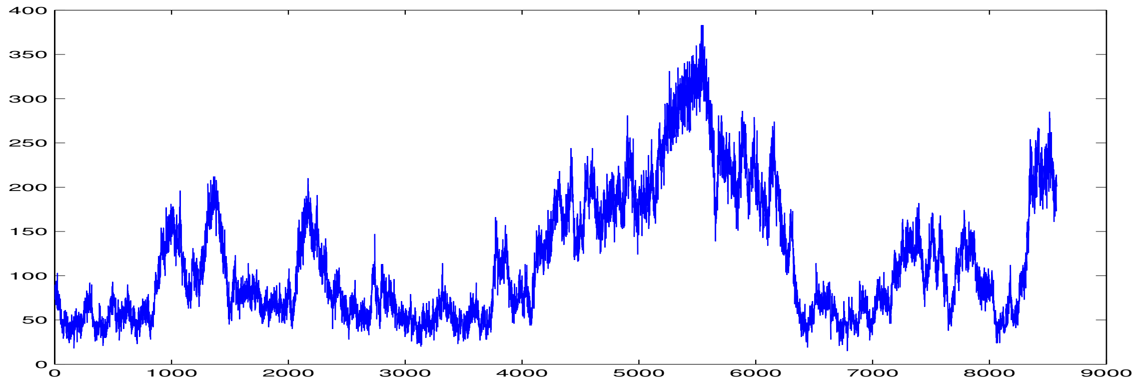

Figure 2: Number of events per day. In period of high volatility, the duration of events scales down and we observe a burst of number of events per day.

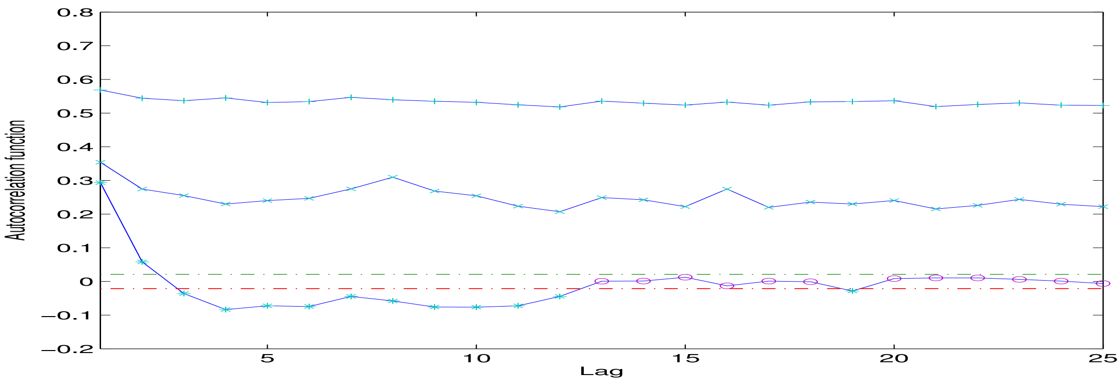

Figure 3: Autocorrelations of raw, squared and absolute log-returns (bottom to top).

extracted from empirical financial time series.

We present in Figure 1 an example of realisation. One can observe that in general, the market price fluctuates regularly around the fundamental value, due to the majority of fundamentalists. Moreover, even if the variations of the fundamental value are Gaussian by construction, the market log-returns clearly exhibit departures from normality. In particular, the market experiences from time to time some periods of instability ($1000 \leq t \leq 1500$, $2000 \leq t \leq 2200$, $4100 \leq t \leq 6500$, $8400 \leq t \leq 8550$) characterised by a high volatility and a dramatic increase in the proportion of chartists. Those periods of clustered volatility lead to extreme price variations, indicating that the continuous-time approach to handle intraday events successfully scales up and down. To verify this, we show in Figure 2 the evolution of the number of events per day, where volatility bursts are clearly identifiable. Finally, we plot in Figure 3 the sample autocorrelation at different lags for raw, squared and absolute log-returns; although raw returns exhibit a strongly high predictability for a one day lag, the results are qualitatively the one expected: we observe a lack of dependence for raw returns, but a slow decay in squared and absolute log-returns, denoting a long range dependence in the volatility.

On plots, the simulated price and log-return time series look very similar to financial time series, and a very complete package of statistical tests performed by Boitout and Delahaut show that they can successfully reproduce the lack of linear or non-linear dependences of returns, the long range dependence of the volatility and its tendency to cluster in time, all of which are phenomena universally observed in financial markets.

5 Conclusion

Agent-based market models have to face an important contradiction when choosing an internal representation of time: on one hand, financial data are most commonly available on a daily basis, which explains why the Econophysics literature traditionally focuses on stylized facts empirically observed at this time horizon. On the other hand, traders themselves interact constantly with each other, new information arrives continuously and prices keep changing. The first ABMMs [PAH$^+$94, CZ97, Art99] adopted a common discrete-time approach for both the market variables and the trader's activity (their state was modelled at the end of each time period only), resulting in a loss of information about the event history within the time period. Lux then introduced a smaller time scale for intraday activity; unfortunately, this framework needed to be artificially tweaked during periods of high activity due to the fixed number of events per day. Finally, Boitout and Delahaut recently extended this model to allow the duration between events to fluctuate according to the market activity, while still observing their market variables on a daily basis. Such a mixed approach, which reconciliates the continuous intraday activity of the market with the daily horizon of recorded variables, looks very promising, and it would be interesting to see if it can be adjusted to a pure ABMM such as the one proposed by Giardina and Bouchaud [GB03], which actually models the evolution of individual traders through time, and not only their repartition.

References

[AAP88] P. W. Anderson, K. Arrow, and D. Pines. *The Economy as an Evolving Complex System*. Redwood City, Addison-Wesley Co., 1988.

[ADL97] W. B. Arthur, S. Durlauff, and D. Lane. *The Economy as an evolving complex system II*. Addison-Wesley, Reading MA, 1997.

[Art99] W. B. Arthur. Complexity and the economy. *Science*, 284:107–109, 1999.

[BD03] N. Boitout and T. Delahaut. Long range dependence in a real trading time agent-based model. 2003. http://www.univ-orleans.fr/DEG/LEO/activitescientifique/seminaires2002.html.

[BG94] E. Bonabeau and G.Theraulaz. *Intelligence collective*. Hermès, Paris, 1994.

[Con01] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(3):223–236, 2001.

[CZ97] D. Challet and Y.-C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407, 1997.

[F$^+$69] E. F. Fama et al. The adjustment of stock prices to new information. *International Economic Review*, 10:1–21, 1969.

[GB03] I. Giardina and J. P. Bouchaud. Bubbles, crashes and intermittency in agent based market models. *Science & Finance*, 2003. To appear.

[JHHJ00] P. Jefferies, M. L. Hart, P. M. Hui, and N.F. Johnson. From market games to real-world markets. 2000. cond-mat/0008387.

[LM99] T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397:498–500, 1999.

[LM00] T. Lux and M. Marchesi. Volatility clustering in financial markets: a microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance*, 3:675–702, 2000.

[Lux98] T. Lux. The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. *J. Economic Behavior and Organization*, 33:143–165, 1998.

[MS99] R. Mantegna and H. E. Stanley. *Introduction to Econophysics*. Cambridge University Press, 1999.

[Pag96] A. Pagan. The econometrics of financial markets. *Journal of Empirical Finance*, 3:15–102, 1996.

[PAH+94] R. G. Palmer, W. B. Arthur, J. H. Holland, B. LeBaron, and P. Tayler. Artificial economic life: a simple model of a stock market. *Physica D*, 75:264–274, 1994.

[Wei99] G. Weiss. *Multiagent Systems, A Modern Approach to Distributed Artificial Intelligence*. MIT Press, 1999.