

Simulating Financial Markets using MASON Framework

CSS 739 Class Project Team*

Center for Social Complexity, George Mason University, USA
{rax222}@gmu.edu

Key words: Agent-based Modeling, Computational Social Science, Financial Markets

1 Motivation and Objectives

The purpose of this project is to create computational platform that would be used to implement agent-based models of financial markets. In particular, following functionalities are desired:

- Simulation may host heterogeneous agents that follow news processes, learn, obey budget constraints etc.;
- Many different assets traded on markets using different trading rules may be present (based on price-impact heuristics, double auction);
- Flexible GUI and reporting facilities, including ability to summarize batch results from sensitivity analysis.

This paper outlines the prototype prepared during the CSS739 class, presents three sample models and follows the steps taken to ensure correctness of our implementation.

2 Platform Architecture

The implementation took place using MASON simulation framework, see Luke et al. [2005]. The static structure of our platform is outlined on Figure 1, where UML class diagram is presented (minus classes corresponding to orders). Figure 2 presents interactions taking place during a sample iteration between two agents and an orderbook.

3 Implemented Models

3.1 Cont Model

The Cont model (Cont [2006]) is the simplest model implemented. All traders (agents) follow the same behavioral rules. They are heterogeneous in the sense

* Programming Group: Justin Briggs, Francesco Trusiano, Maciej Latek, Randy Marks. Data Analysis Group: Sandler Lin, Cheryl Hansen and Thomas Tzitzouris.

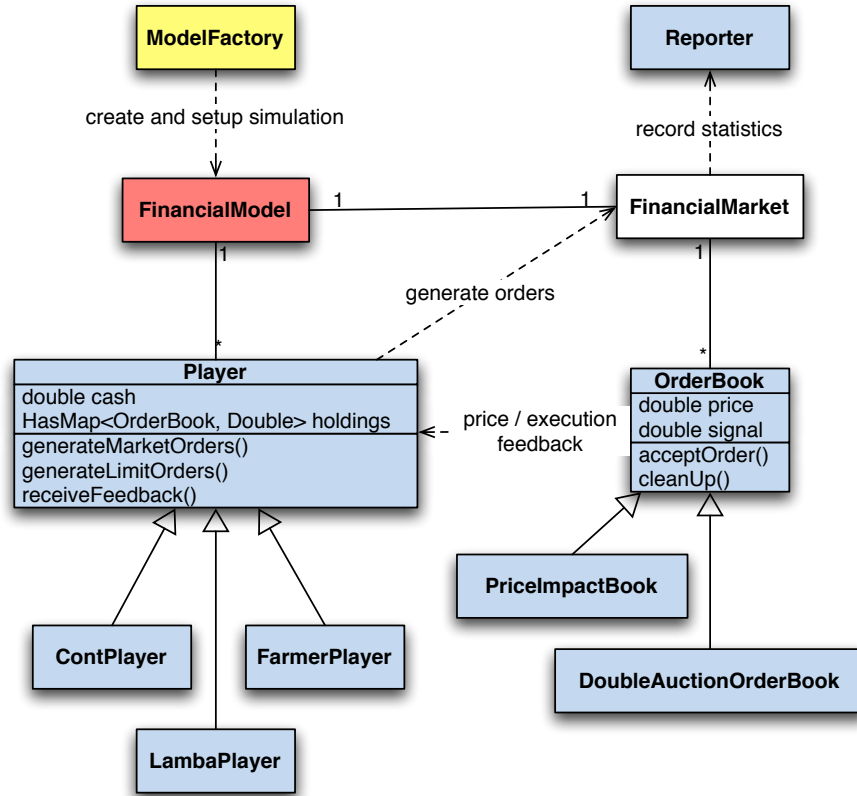


Fig. 1. High-level UML class diagram of the main components and relations in the FinancialMarketModel, including the main attributes of Players and OrderBooks. Agent classes (light blue) inherit from the MASON `Steppable` interface while the master class is implementation of MASON's `SimState`.

that they are given independently assigned (subjective) volatility thresholds $\theta_i(t)$. Each period, all agents receive a common signal, which can be interpreted as public information or "news", in the form of IID Gaussian random variables $\epsilon_t \sim N(0, D^2)$. Each agent i responds to this signal by selling if $\epsilon_t < -\theta_i(t)$, buying if $\epsilon_t > \theta_i(t)$, and otherwise sitting out the period. The market then determines the excess demand and arrives at the market clearing price by means of a market impact function. Lastly, with probability s , agents update their threshold to match the absolute value of the return rate for the current period. Somewhat surprisingly, even this very simple "zero intelligence" model of agent behavior yields fairly realistic market movements, as presented in Figures 3 (for small s) and 4 (for large value of s).

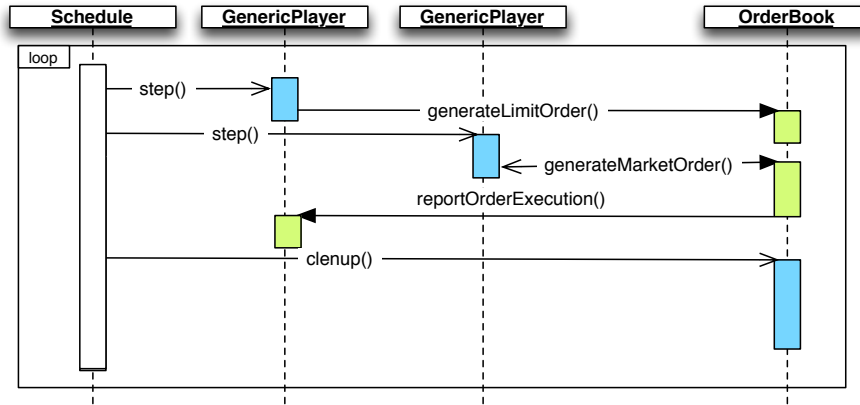



Fig. 2. High-level UML sequence diagram of the interactions between main object of the FinancialMarketModel, with messages send between two agents and a single orderbook during a single tick shown. Agents are activated by MASON simulation engine and may generate different kinds of `MarketOrders`, depending on their internal states and signals from environment (prices, news process). The `MarketOrder` gives an agent instantaneous feedback, whereupon effects of `LimitOrder` may be delayed in time. The `cleanup` signal causes the Orderbook to arbitrage all possible trades between `LimitOrders`, recalculate new prices and returns (if applicable) and remove expired orders.

One down side of the Cont model that is not noted in the papers is that there appears to be a great degree of periodic cyclicity in the volatility of the markets. This is especially apparent with higher probability s of updating thresholds. With $s = 0.1$, this is readily apparent if you look at the autocorrelation of absolute returns over longer time scales than shown in the paper. This effect can be particularly well seen when averaging over many runs, as on Figure 11. Holding s constant at 0.1 and varying D , we find the period and amplitude to be predictable.


3.2 Farmer Model

The Farmer model (Doyme Farmer et al. [2003]) has two types of agents, which trade a single asset by means of a continuous double auction similar to those normally used on stock markets. "Patient" traders place limit orders, which specify both a price and quantity they wish to buy or sell. Limit orders do not execute immediately, and may be canceled at any time. "Impatient" traders place market orders, specifying only a quantity, which execute immediately at the current best price, assuming there is sufficient liquidity provided by outstanding limit orders. This model also has simple behavioral logic for traders. Behavior is regulated by:



(a) Trade volumes.


(b) Price



(c) Raw and absolute log returns.

(d) Distribution of Returns

(e) Distribution of Returns (log scale)



(f) Annualized moving average volatility

(g) Autocorrelation of returns.

(h) Autocorrelation of absolute returns.

Fig. 3. Examples of outputs and statistics from a single run of the Cont FinancialModel simulation with $D = 0.001$, $s = 0.01$.

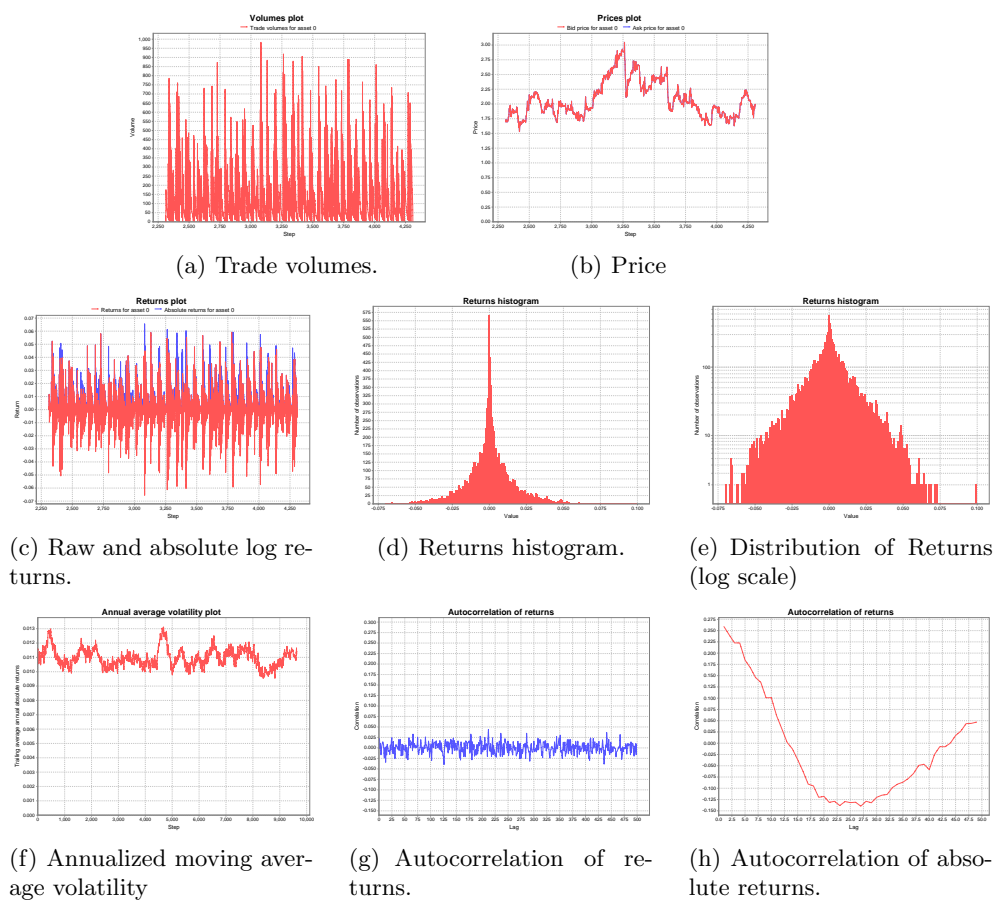


Fig. 4. Examples of outputs and statistics from a single run of the Cont FinancialModel simulation with $D = 0.001$, $s = 0.1$.

- Patient agents place **LimitOrders** of constant size σ at a Poisson rate of α per unit time, and with a randomly generated price. Orders are equally likely to be either buy or sell orders. Buy limit orders are placed uniformly anywhere in the semi-infinite interval below the current ask price ($-\infty < p < a(t)$), where p is the *logarithm* of the price; similarly for sell orders. A uniform distribution in log prices is equivalent to an exponential distribution in prices. In addition, outstanding limit orders are cancelled at a Poisson rate of δ per unit time. All of these processes are independent.
- Impatient agents place **MarketOrders** of constant size σ at a Poisson rate of μ per unit time.
- A double auction order book. This manages the pending limit orders and facilitates trades between market and limit orders. It also provides aggregate market information such as return rate per unit time and bid-ask spreads.

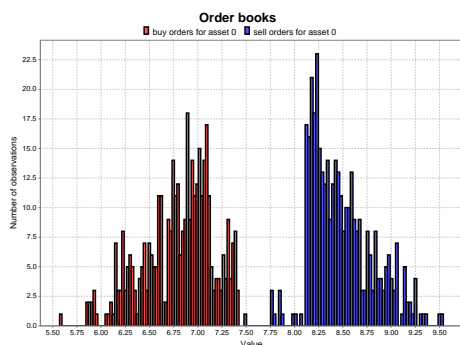
This model explains a large portion of the spread and price diffusion of actual stocks given an order flow rate, which implies that the double auction structure has a significant impact on the nature of market movements. The model also yields market behavior qualitatively close to actual markets in many respects, see Figure 5.

However, as a model of market dynamics, it fails on some fronts. Notably, we do not see realistic volatility clustering. Volumes are fairly consistent, as are average returns. There is no lingering positive autocorrelation of absolute returns. Also notice the amplitude of single-period returns in Figure 5(g). The average per-step absolute returns are around 20%, much higher than empirically plausible, at least if a step is meant to signify short period, such as a day or less. This is clearly seen when we use actual prices in the order book. The same prices which look somewhat reasonable when viewed on a log scale appear quite volatile prices when viewed on a normal scale as in Figure 6.

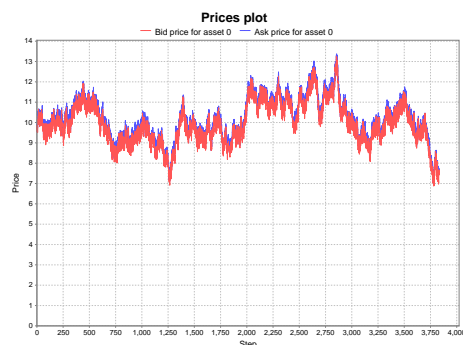
3.3 Farmer-Cont Model

As an experiment, we constructed a model combining some aspects of the Farmer model with some aspects of the Cont model. The result is a model which is still essentially zero intelligence, but which exhibits empirically positive characteristics of both models. In this model there are two types of agents, which interact via the continuous double auction mechanism, as in Farmer. In fact, the Patient traders, which place limit orders and thus provide liquidity, remain the same. Impatient traders are modeled after Cont. They have heterogeneous volatility thresholds, and place market orders depending upon how public information in the form of a normally distributed random variable $\epsilon_t \sim N(0, D^2)$ compares with these thresholds. Some fraction s of agents then adjust their volatility thresholds in light of current volatility by setting it to the trailing period's return rate.

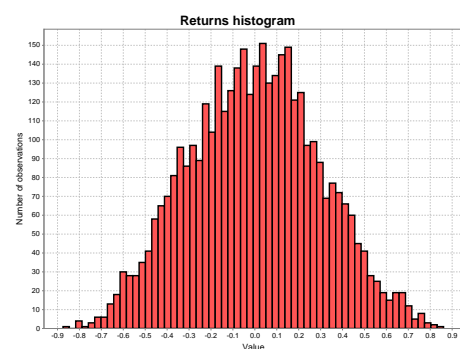
The result is shown in Figure 7. When there are sufficient limit orders being placed by Farmer's patient agents to ensure adequate liquidity for Market orders, we get dynamics that have many positive aspects compared with reality —



(a) Snapshot of Orderbook.



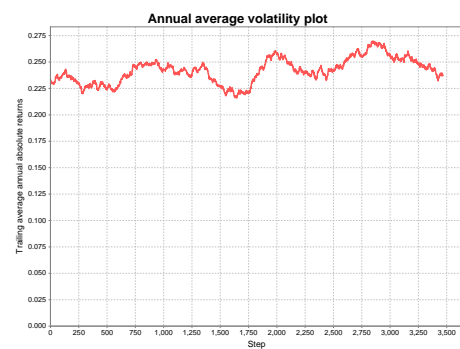
(b) Bid and ask log prices.



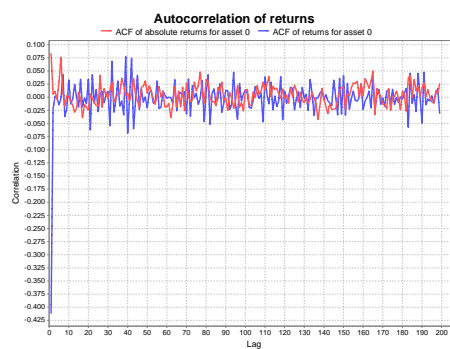
(c) Log returns histogram.



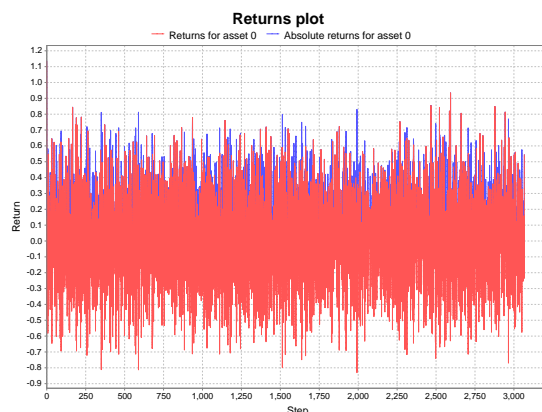
(d) Trade volumes.



(e) Rolling Average Volatility.



(f) Autocorrelation of returns.

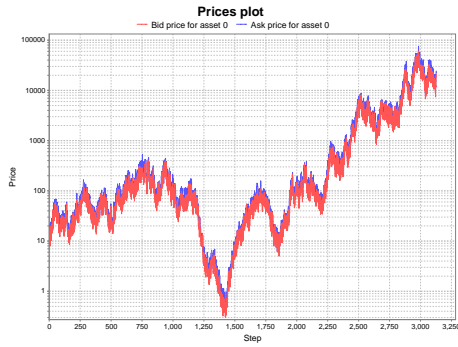


(g) Raw and absolute log returns.

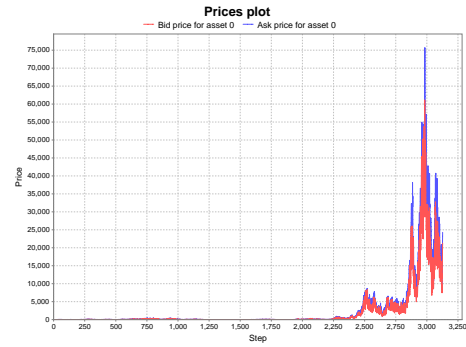
Fig. 5. Examples of outputs and statistics from a single run of the Farmer FinancialModel simulation with $\alpha = .585$, $\delta = .39$, $\mu = .325$, $\sigma = 1$. These parameters are derived from empirical data.



(a) Snapshot of Orderbook with actual prices (log x scale).



(b) Prices (log scale).



(c) Same prices (natural scale).

Fig. 6. Examples of outputs and statistics from a single run of the FinancialModel simulation with $\alpha = .585$, $\delta = .39$, $\mu = .325$, $\sigma = 1$ shown in actual, rather than log prices

approximately power law return distribution, volatility clustering, realistic variability in *actual* prices, volumes, and returns, and essentially no predictability of future returns based upon past returns.

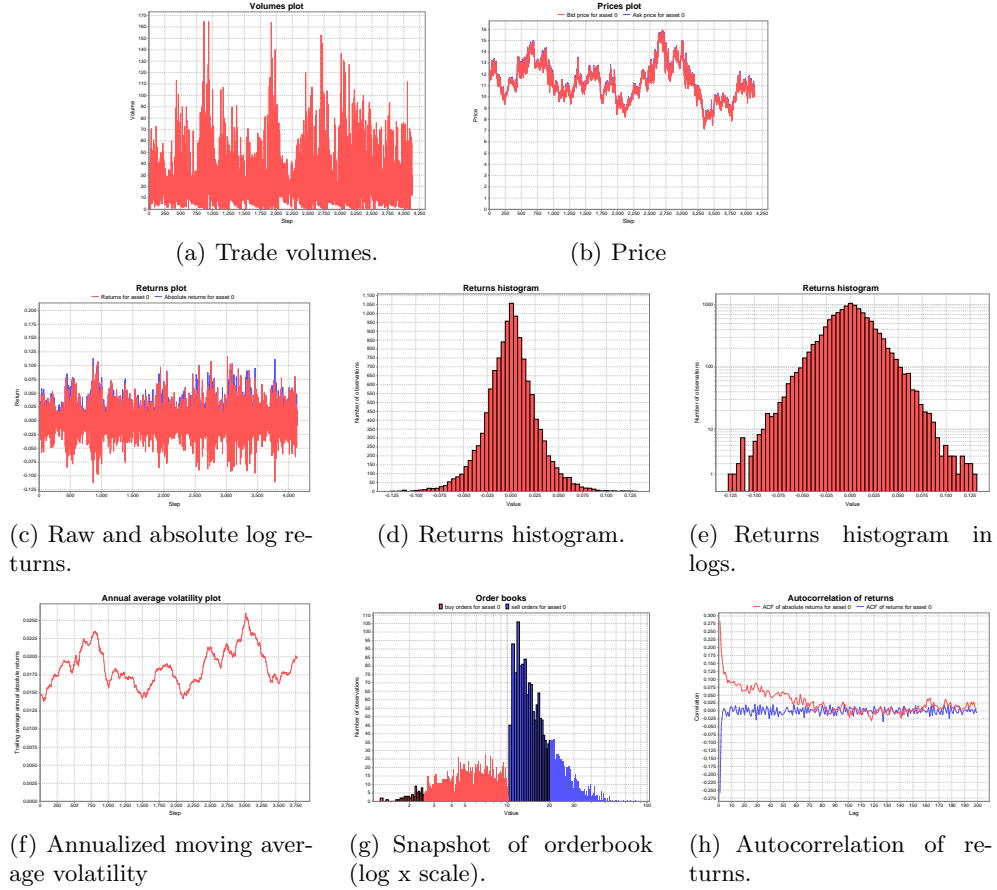


Fig. 7. Examples of outputs and statistics from a single run of the Farmer-Cont hybrid FinancialModel simulation with actual prices, $D = 0.001$, $s = 0.01$, $\alpha = 0.9$, $\delta = 0.2$, and $\mu = 0.4$.

3.4 Westerhoff Model

The Westerhoff model (Westerhoff [2004]) includes a limited number of markets in which two types of agents operate. Fundamentalists specialize in a certain market while chartists may switch between markets. More specifically, chartists

use extrapolative methods to select markets that shows price trends but that are not too misaligned.

The model considers $k = 1, 2 \dots K$ asset markets of equal size: the log fundamental value of asset k in period $t + 1$ is regulated by a news arrival process:

$$F_{t+1}^k = F_t^k + N$$

Market makers are present in all markets and stand ready to absorb imbalances between buyers and sellers: they adjust prices S_t^k according to a loglinear price impact function. Traders submit orders according to their expectations of the price movements. Fundamentalists expect the prices of the assets to return to their fundamental values and their demand is expressed as:

$$\begin{aligned} D_t^{F,k} &= a^F [E_t^F(S_{t+1}^k) - S_t^k], \\ E_t^F &= S_t^k + b^F (F_t^k - S_t^k) \end{aligned}$$

On the other hand, chartists follow simple technical analysis rules:

$$\begin{aligned} D_t^{C,k} &= a^C [E_t^C(S_{t+1}^k) - S_t^k], \\ E_t^C &= S_t^k + b^C (S_t^k - S_{t-1}^k) \end{aligned}$$

Chartists try to identify the attractiveness of markets with a fitness measure that entails the risk of being captured in a bursting bubble. The attractiveness of market k is defined as:

$$A_t^k = \log \frac{1}{1 + f(F_t^k - S_t^k)^2}$$

The relative percentage of chartists choosing market k depends on the attractiveness of that market and is defined as:

$$W_t^k = \frac{\exp(gA_t^k)}{\sum_{k=1}^K \exp(gA_t^k)}$$

As shown in Figure 8, this model is capable of generating complex price dynamics, volatility clustering, fat tails for the distribution of the returns and positive autocorrelation for absolute returns.

4 Verification of Correctness

The supreme objective was to recreate results reported by others rather than to create a standalone model of the market. Therefore, the validation focused on comparing outputs from our simulations to graphs and statistics from other papers. These compare well with data shown in the original paper (shown in Figures 9 and 10).

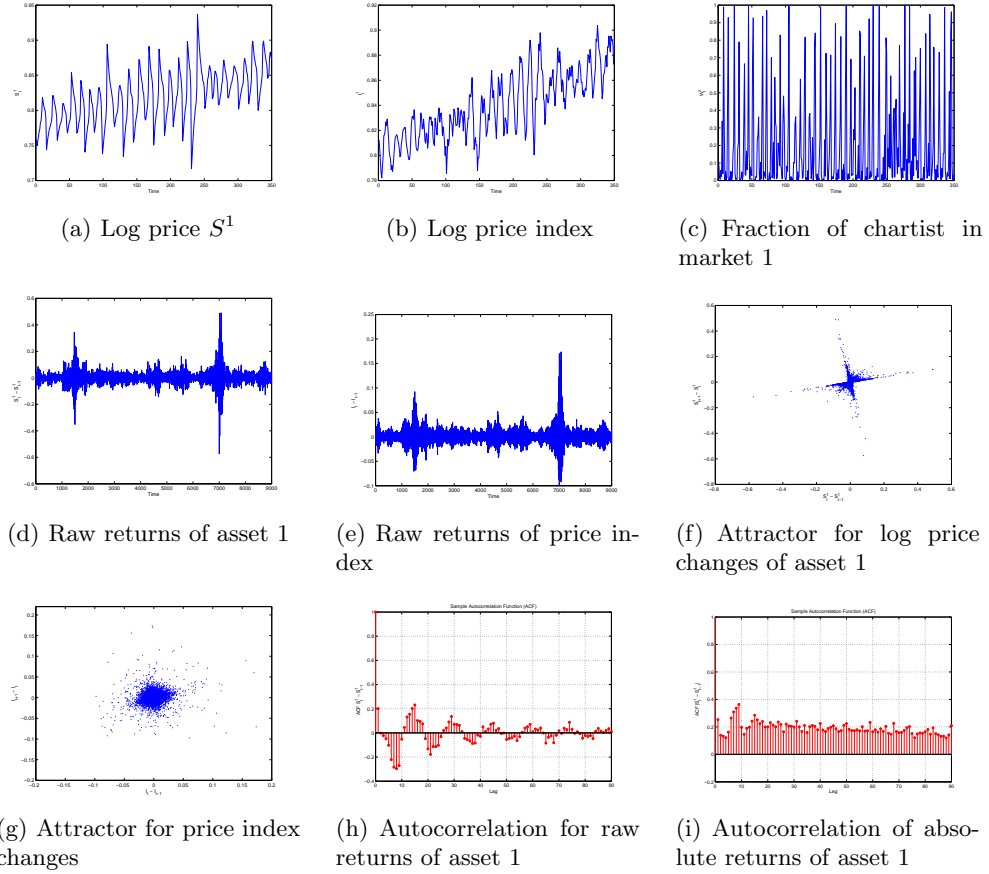


Fig. 8. Examples of outputs and statistics from a single run of the Westerhoff simulation with $K = 5$, $N = 0.0002$, $a^M = 0.1$, $a^c, b^c = 5$, $a^F, b^F = 0.2$, $f = 1,000,000$ and $g = 1.2$.

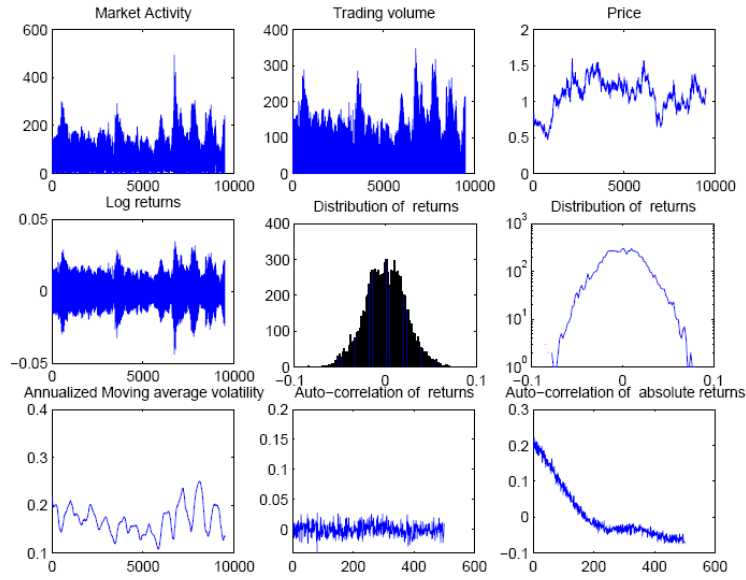


Fig. 9. Outputs from Cont's paper with $D = 0.001$, $s = 0.01$. Compare directly to Figure 3.

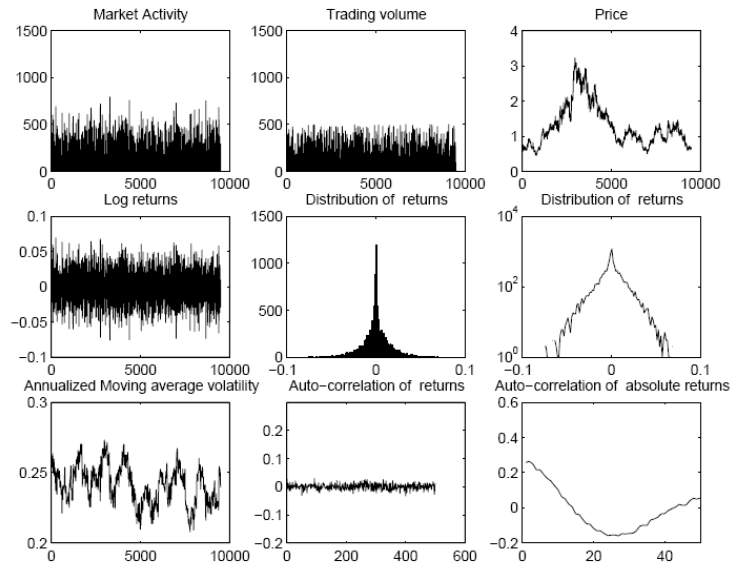


Fig. 10. Outputs from Cont's paper with $D = 0.001$, $s = 0.1$. Compare directly to Figure 4.

5 Batch Mode Instrumentation

FinancialMarketModel includes a set of wrappers that are useful for conducting multi-run batch parameter space exploration or sensitivity analysis. A simple wrapper is included in `experiments` package. Post-run analysis can be performed by using Matlab, sample scripts are included in `scripts` folder and can be used to render shape of autocorrelation curves and histograms of returns as a function of simulation settings. Figure 11 presents a sample prepared by varying the D news process variance for Cont's model.


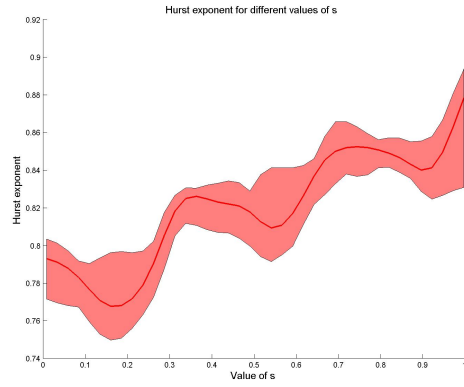
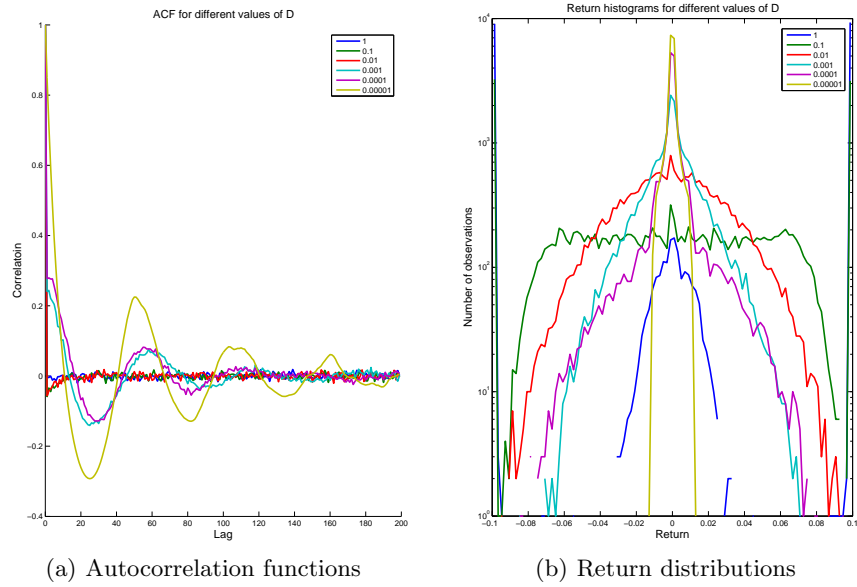


Fig. 11. Autocorrelation and distributions in Cont simulation with $s = 0.1$ and varying D , averaged over many runs. In bottom row, kernel regression with 95% resampled confidence intervals on Hurst exponents of time series generated from Con't model for varying levels of s and D is presented.

Bibliography

- Rama Cont. *Long Memory in Economics*, chapter Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models. Springer Berlin Heidelberg, 2006.
- J. Doyne Farmer, P. Patelli, and I. I. Zovko. The Predictive Power of Zero Intelligence in Financial Markets. *ArXiv Condensed Matter*, September 2003.
- Sean Luke, Claudio Cioffi-Revilla, Liviu Panait, Keith Sullivan, and Gabriel Balan. MASON: A Multiagent Simulation Environment. *Simulation*, 81:517–525, 2005.
- Frank H. Westerhoff. Multiasset Market Dynamics. *Macroeconomic Dynamics*, 8(05):596–616, November 2004. URL http://ideas.repec.org/a/cup/macdyn/v8y2004i05p596-616_04.html.