News and price returns from threshold behaviour and vice-versa: exact solution of an agent-

based market model

News and price returns from threshold behaviour and vice-versa: exact solution of an agent-based market model

Damien Challet

Nomura Centre for Quantitative Finance, Mathematical Institute, Oxford University, 24-29 St Giles', Oxford OX1 3LB, UK

E-mail: challet@maths.ox.ac.uk

Abstract

Starting from an exact relationship between news, threshold and price return distributions in the stationary state, I discuss the ability of the Ghoulmie–Cont–Nadal model of traders to produce fat-tailed price returns. Under normal conditions, this model is not able to transform Gaussian news into fat-tailed price returns. When the variance of the news is so small that only the players with zero threshold can possibly react to news, this model produces Levy-distributed price returns with a -1 exponent. In the special case of super-linear price impact functions, fat-tailed returns are obtained from well-behaved news.

PACS numbers: 89.65.Gh, 89.75.Da

(Some figures in this article are in colour only in the electronic version)

Few agent-based models of financial markets are amenable to mathematical analysis. Notable exceptions include percolation-based imitations models [1], behavioural switching models [2, 3] and minority game-based models [4, 5]. Mathematical solutions are valuable tools for exploring the mechanisms underlying financial market dynamics, particularly with a view to understanding the causes of unpredictability and large fluctuations. Threshold mechanisms, whereby the agents react to events exceeding their personal tolerance levels, provide an indirect way to avoid forcing the agents to trade at each time step, and also implicitly define timescales.

A new simple Markovian model of agents that adjust their thresholds on past volatilities was recently introduced and exhibited promising features [6]. Here, I derive exact relationships between price return, threshold and news distribution in the limit of infinite number of agents, and discuss in detail when this model produces fat-tailed returns from well-behaved news distributions.

The model works as follows [6]: there are N agents with individual threshold $\theta_i, i = 1, \ldots, N$, who monitor a single source of external news. At time t the news intensity and sign is $\epsilon(t)$, drawn from the symmetric distribution $\rho(\epsilon)$; a positive ϵ corresponds to good news and reversely. Only the agents whose thresholds are smaller than the absolute value of the news react to it by placing a order to buy if $\epsilon(t) > 0$ and reversely. Mathematically, if $\phi_i(t)$ denotes the action of agent i at time t,

$$\phi_i(t) = \Theta(\epsilon(t))\Theta(\epsilon_t - \theta_i(t)) - \Theta(-\epsilon(t))\Theta(\theta_i - \epsilon(t)), \tag{1}$$

where $\Theta(x)$ is the Heaviside function; note that $\phi_i(t) \in \{-1, 0, 1\}$. This gives rise to global demand imbalance $\Phi(t) = \sum_i \phi_i(t)$. The price evolves according to

$$\log p(t + 1) = \log p(t) + r(t), \tag{2}$$

where r is called the price return. Its relationship with demand imbalance is assumed to be fully characterized by a constant price impact function g:

$$r(t) = g(\Phi(t)/N). \tag{3}$$

At each time step, agent i may re-adjust his threshold θ_i with probability s to $|r(t)|$. In that way, the agents are given minimal learning capabilities, tracking changes of volatility with a reaction time scale of order $1/s$. It is worth noting that the agents have no influence on the news process. The inventors of this model produced evidence that the agents transform Gaussian news into price return distributions $P(r)$ whose shape and properties depend on the parameters N, s and ρ; some of them seemingly no longer Gaussian [6].

Stationary state

Since $\rho(\epsilon)$ is symmetric, the model can be simplified by replacing ϵ by $\epsilon' = |\epsilon|$ and $\rho(\epsilon')$ by $\rho'(\epsilon') = 2\rho(|\epsilon|)$. Dropping the primes, the dynamics of individual actions reduces to

$$\phi_i(t) = \Theta(\epsilon_t - \theta_i(t)). \tag{4}$$

Because of equation (3), since Φ is discrete for finite N, $r(t)$ and θ_i take discrete values. The master equation for $P(\theta, t)$ is therefore

$$P(\theta, t + 1) = (1 - s)P(\theta, t) + s\delta\left[\theta - g\left(\sum_{\theta' \leqslant \epsilon_t} P(\theta', t)\right)\right]. \tag{5}$$

Letting $\tau = \theta/N$ and $\epsilon' = \epsilon/N$, and dropping the prime, one can transform this problem into a finite continuous interval (r, τ and ϵ belong to [0, 1]), resulting in

$$P(\tau, t + 1) = (1 - s)P(\tau, t) + s\delta\left[\tau - g\left(\int_0^{\epsilon_t} d\tau' P(\tau', t)\right)\right]. \tag{6}$$

By definition, $P(\tau, t)$ becomes $P(\tau)$ in the stationary state; the final step is to average over all possible ϵ, which yields

$$P(\tau) = \int d\epsilon\, \rho(\epsilon)\delta(\tau - g[F(\epsilon)]) = \frac{\rho[F^{-1}(g^{-1}(\tau))]}{P[F^{-1}(g^{-1}(\tau))]g'(g^{-1}(\tau))}, \tag{7}$$

where g' is the first derivative of g.

It should be noted that the disappearance of s in this equation is a sign that s is a time scale; the auto-correlation function of $r(t)$ stays positive for a time scale of order $1/s$, as can be seen in the original paper. In the stationary state, since $r = g(F(\epsilon))$, the pdf of r has exactly the same functional form as that of τ, because the thresholds are adjusted to the last
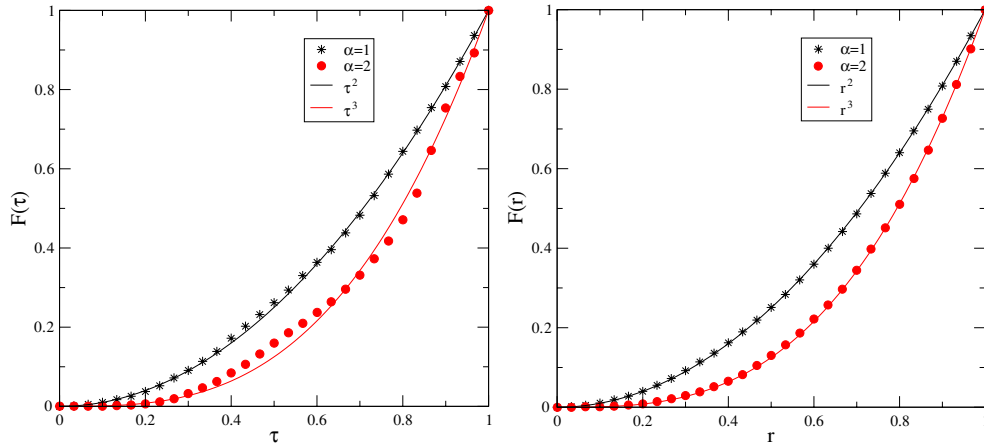
Figure 1. Threshold (left) and price return (right) cumulative distribution function $F(\tau)$, $F(r)$, respectively, obtained with $\rho(\epsilon) = (\alpha + 1)^2 \epsilon^{\alpha(\alpha+2)}$ for $\alpha = 1$ (black symbol and line) and $\alpha = 2$ (red symbol and line). $N = 100\,00$, $s = 0.001$, $100\,000$ iterations; $F(\tau)$ is averaged over 9 equally spaced snapshots in the last half of the iterations.

price return with probability s. The multiple appearances of $P(\tau)$ under various forms on the right-hand side of equation (7) are due to the self-referential nature of the model.

The dynamics produces $P(\tau)$ from $\rho(\epsilon)$. However, equation (7) seems quite difficult to solve; it is much easier to find which $\rho(\epsilon)$ corresponds to a given $P(\tau)$, since

$$\rho(\epsilon) = P[g(F(\epsilon))]g'(F(\epsilon))P(\epsilon) \tag{8}$$

does not depend on itself in an intricate way.

Linear price impact function

The original model reports simulations with $g(r) = \lambda r$, which simplifies the discussion. Without loss of generality, one takes $\lambda = 1$. The original model also assumes Gaussian news intensity, that is, $\rho(\epsilon) = 2\mathcal{N}(0, D)$, $\epsilon \geqslant 0$. While it would be hard to find $P(\tau)$ that leads exactly to such a distribution, $P(\tau) = 2\mathcal{N}(0, D)$ corresponds to

$$\rho(\epsilon) = \frac{2}{D^2\pi} e^{-\frac{\operatorname{erf}[\frac{\epsilon}{\sqrt{2}D}]^2}{2D^2}} e^{-\frac{\epsilon^2}{2D^2}}. \tag{9}$$

In particular, for $\epsilon \geqslant D$, $\rho(\epsilon)$ approximates very well a Gaussian. It is therefore possible to conclude that Gaussian news produces Gaussian-tailed price return distributions. Numerical simulations show that Gaussian news can produce one peak or two symmetric peaks with or without one central Levy distributed peak (in finite size systems) whose origin is explained below.

Nevertheless, one might question the assumption that the news is Gaussian. For instance damages caused by natural disasters have fat tails [7]. The power-law case yields a more direct relationship between news and price returns: $P(r) = (\alpha + 1)\tau^\alpha$ ($\alpha > -1$) corresponds to $\rho(\epsilon) = (\alpha + 1)^2 \epsilon^{\alpha(\alpha+2)} = (\alpha + 1)\epsilon^\beta$ for r and $\epsilon \in [0, 1]$: it is therefore impossible to produce $\alpha < 0$ from $\beta > 0$ with linear price impact functions, hence, well-behaved news cannot trigger fat-tailed price returns; however, fat-tailed news do trigger fat-tailed prices. Figure 1 confirms the validity of the above calculations: numerical and theoretical values of
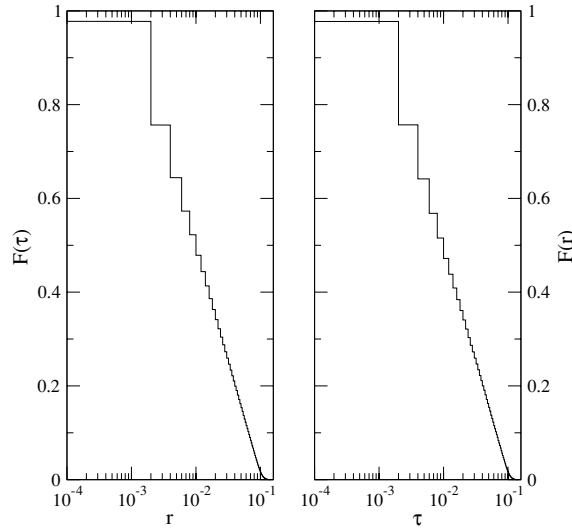
Figure 2. Cumulative distribution of r and τ for $N = 5000$, $D = 0.1/N$, $s = 0.1$, 10^5 iterations.

$P(r)$ match perfectly. The case of $P(\tau)$ is interesting: even after 10^5 time steps, there is a discrepancy between numerical simulations and theoretical predictions; this comes from the temporal correlations of the thresholds over time scales of order $1/s$, which causes the presence of peaks in any snapshot of $P(\tau)$. When one averages $P(\tau)$ over several snapshots, one finds that the theoretical prediction and the numerical simulations match better. Alternatively, one could plot the distribution of the values taken by the threshold of a given agent. As a final example, an exponential distribution of thresholds $P(\tau) = e^{-\tau}$, $\tau \geqslant 0$, corresponds to $\rho(\epsilon) = e^{e^{-\epsilon}} e^{-\epsilon}$.

So far, the discussion has only addressed the model in the limit of infinite systems. However, the behaviour of the model changes completely for finite size if $D \ll 1/N$ and $sN > 1$. In that case, it is easy to convince oneself that the only situation yielding non-zero price returns occurs when $P(0) \neq 0$; in other words, in the case of very small news, the only agents who react to them are infinitely susceptible, as they react in fact to any news amplitude. The mechanism is the following: the smallest non-zero threshold is $1/N > \epsilon_t$ for all t. Assume that $P(\theta = 0, t_0) = 0$, which implies that $r(t_0) = 0$. Then on average sN agents adjust their threshold to $\theta = 0$. At the next time step, all of them react to any news amplitude, which leads to an average price return $r(t_0 + 1) = s$. Because of the threshold adjustment process, the number of agents with a zero threshold follows a multiplicative random walk $r(t + 1) = r(t)\eta(t)$, where $\mathcal{P}(\eta) \simeq \mathcal{N}(1 - s, s(1 - s)/N)$, with a barrier at $1/N$; when $P(0, t') < 1/N$, r jumps back on average to s. Numerical simulations fully confirm this mechanism. They also show that $P(r) \propto r^{-1}$ (and of course $P(\tau) \propto \tau^{-1}$), with an upper cut-off at about s. Multiplicative random walks with lower boundary are well known to give rise to power-law distributed density functions whose exponent γ is given by (see for instance [8])

$$\int d\eta \, P(\eta)\eta^{\gamma - 1} = 1. \tag{10}$$

In the large N limit, $\mathcal{P}(\eta)$ converges to a Dirac function, which leads to $(1 - s)^{\gamma - 1} = 1$, hence $\gamma = 1$. The normalization of $P(r)$ is not problematic, since r is bounded from above.

Intermediate cases occur when $sN \sim 1$: the centre of the distribution is of Levy type with -1 exponent and its tail is truncated.

Nonlinear price impact functions

Average real market price impact functions are reportedly concave functions, either $g(x) = x^\gamma$ with $\gamma < 1$ [9, 10] or $g(x) = \log x$ [11]. It will be enough to discuss the power-law case. Using $P(r) = (\alpha + 1)\tau^\alpha$ ($\alpha > -1$), one finds that

$$\rho(\epsilon) = \gamma(\alpha + 1)^2 \epsilon^{\gamma(\alpha+1)^2 - 1} = \gamma(\alpha + 1)^2 \epsilon^\beta. \tag{11}$$

Equivalently, $\alpha = \sqrt{\frac{\beta+1}{\gamma}} - 1$. The condition for power-law distributed price returns ($\alpha < 0$) is $\beta < \gamma - 1$. Therefore, sublinear impact functions dampen ρ, and fat-tailed returns are only obtained from sufficiently fat-tailed news; indeed, if $\gamma - 1 < \beta < 0$, the news are fat-tailed, but the price returns are not. A super-linear g on the other hand makes it possible to have a well-behaved ρ and fat-tailed price returns; but, as stated above, average price impact functions are not convex functions.

Conclusion

In short, in the $N \to \infty$ limit, the above derivation not only shows the crucial importance of the functional form of the news on the price properties, but also that this model produces very fat-tailed price returns in the limit of infinitesimally small news with exponent 1, and only for finite size systems; the other possibility for fat-tailed price returns arises when the price impact function is super-linear. Therefore, under realistic assumptions, this model cannot transform Gaussian news into fat-tailed price returns, but the news intensity distribution may not be Gaussian in real life.

Acknowledgment

References

[1] Cont R and Bouchaud J-P 2000 *Macroecon. Dyn.* **4** 170
[2] Alfarano S and Lux T 2003 Long Memory in Economics (Berlin: Springer)
[3] Hommes C H 2005 *Handbook of Computational Economics* ed K L Judd and L Tesfatsion (Amsterdam: Elsevier)
[4] Challet D, Marsili M and Zhang Y-C 2005 *Minority Games* (Oxford: Oxford University Press)
[5] Coolen A A C 2005 *The Mathematical Theory of Minority Games* (Oxford: Oxford University Press)
[6] Cont R, Ghoulmie F and Nadal J-P 2005 *J. Phys.: Condens. Matter* 212
[7] Embrechts P, Klüppelberg C and Mikosch T 1997 *Modelling Extremal Events for Insurance and Finance* (Berlin: Springer)
[8] Maslov S, Zhang Y-C and Marsili M 1998 *Physica* A **253** 403
[9] Lillo F, Farmer J D and Mantegna R 2003 *Nature* **421** 129
[10] Plerou V, Gopikrishnan P, Gabaix X and Stanley H E 2004 *Quant. Finance* **4** C11
[11] Bouchaud J-P, Gefen Y, Potters M and Wyart M 2004 *Quant. Finance* **4** 176