

Market statistics of a psychology-based heterogeneous agent model

HARBIR LAMBA¹

TIM SEAMAN²

Abstract

We continue an investigation into a class of agent-based market models that are motivated by a psychologically-plausible form of bounded rationality. Some of the agents in an otherwise efficient hypothetical market are endowed with differing tolerances to the tension caused by being in the minority. This herding tendency may be due to purely psychological effects, momentum-trading strategies, or the rational response to perverse marketplace incentives.

The resulting model has the important properties of being both very simple and insensitive to its small number of fundamental parameters. While it is most certainly a caricature market, with only boundedly rational traders and the globally available information stream being modeled directly, other market participants and effects are indirectly replicated. We show that all of the most important ‘stylized facts’ of real market statistics are reproduced by this model.

Another useful aspect of the model is that, for certain parameter values, it reduces to a standard efficient-market system. This allows us to isolate and observe the effects of particular kinds of non-rationality. To this end, we consider the effects of different asymmetries in agent behaviour and show that one in particular leads to skew statistics consistent with those seen in some real financial markets.

1 Introduction

The standard models of mathematical finance have provided a rigorous intellectual framework within which to analyze and quantify one of the most complex aspects of modern society. The key assumptions underlying much of this work are collectively known as the

¹(Corresponding author) Department of Mathematical Sciences, George Mason University, MS 3F2, 4400 University Drive, Fairfax, VA 22030 USA

²Department of Mathematical Sciences, George Mason University, MS 3F2, 4400 University Drive, Fairfax, VA 22030 USA

Efficient Market Hypothesis (EMH)[15, 16] and they have enormous influence in both economic theory and the day-to-day operations of financial markets.

However, in recent years two things have become apparent. Firstly, work by experimental economists and psychologists has shown that human beings often behave irrationally even in very simple economic games — in direct conflict with the EMH. This phenomenon is not confined to ‘unsophisticated’ people and it is not random — there would appear to be fundamental psychological pressures that result in statistically predictable and experimentally reproducible deviations from rationality. Secondly, numerous analyses of data from financial markets show some agreement with EMH predictions (such as uncorrelated price returns except over the very shortest time-scales) but also significant, and surprisingly consistent, anomalies such as non-Gaussian log-price returns³. These core observations have become known as the *stylized facts* of real markets and appear to be largely independent of geography, trading systems, culture and so on [9]. They also do not appear to have diminished significantly over time even though, with the advent of computerized trading and vastly superior information flow/processing speeds, one might feel that potential violations of the EMH ought be much reduced from earlier decades or centuries.

A key element of the EMH, over and above the existence of heterogeneous rational agents who are capable of maximising their individual utility functions, is the concept of *rational expectations*. This states that although the agents will have differing predictive models and future expectations, their average behaviour is correct and therefore so are the price changes caused by the arrival of new globally available information. One likely explanation for the existence of the stylized facts is the breakdown of the rational expectations assumption due to significant coupling developing, albeit temporarily, between agents as the market evolves in time. In order to investigate such possibilities various heterogeneous agent models (HAMs), requiring direct simulation of the agents rather than just the calculation of averaged properties, have been developed (see for instance [10, 17, 20, 21, 27, 31, 33, 39, 40, 55]).

In this paper we study a particularly simple class of such models, previously introduced

³It should be noted that non-Gaussian price returns are not logically inconsistent with efficient, memory-free markets, since the possibility of infinite variances in the independent price increments would invalidate the Central Limit Theorem, but they are certainly suggestive of the possibility that efficient pricing is not occurring.

in [12, 50]. The original motivation for this research was to develop a methodology whereby psychological (and irrational) effects could be systematically introduced into an otherwise efficient market and then to compare the output of the model against the stylized facts. This is in contrast with most extant HAMS where it is difficult, if not impossible, to disentangle the chain of cause-and-effect between particular agent properties and the stylized facts. Our primary mechanism for providing coupling between agents will be a *herding tendency*. While an individual/organization is holding a minority opinion/position they feel an increasing pressure to conform that eventually becomes unbearable, at which point they will change to join the majority (unless enough of the agents with majority positions switch first). This ‘threshold’ approach is psychologically plausible in that it mimics the human decision-making process — a gradual accumulation of pressure and/or new information results in a sudden action. Different agents will have different threshold levels that are, to some extent, a reflection of their personality or trading philosophy (at the extreme are self-confessed momentum traders).

However, the herding phenomenon is not just applicable to (amateur) individual investors and although it is tempting to describe such behaviour as ‘irrational’, or ‘boundedly rational’ in the sense of Simon [53], this is not a correct characterization for many types of agent. Some would lose their job/bonus/investment capital if they significantly underperformed the average market return for even a short time — when herding, they are exhibiting behaviour that is no more irrational than animals herding when surrounded by predators. The use of complex, possibly automated, trading strategies may help to mitigate the herding pressure but it is unlikely to completely remove it. For example, buy/sell signals from automated trading programs may unwittingly be reacting to the herding effects of other agents.

The models that we propose are constructed in such a way that, for certain values of the system parameters, an efficient pricing model is recovered. This is a valuable property when it comes to investigating the effects of agent-coupling in markets. It allows us to assume the existence of an underlying, hypothetical, efficient market whose statistical properties are well understood. The mechanisms by which prices are set, orders are carried out, information is propagated and so on do not need to be specified. We merely presuppose that they are in place and constant, and then change parameter(s) to introduce particular types of agent-coupling into the model. This will allow us to directly relate specific agent-behaviour modifications to the emergence of stylized facts, at least within

our modeling paradigm. It is important to note that we shall not require all agents to be subject to the herding tendency. Many agents may either be immune or be trading solely for reasons of pure hedging, say. We do not simulate such agents directly – instead we assume that the rational expectations assumption still applies to them and focus the model on those agents where significant coupling is likely to occur.

We now outline the paper. In Section 2 the complete market model is introduced. Then in Section 3 we recall the stylized facts of market behaviour and compare these with the output of our model. This statistical analysis is far more detailed than that of [12] and we show that all of the important details are replicated, over a wide range of parameter values. In Section 4 we show how the model can be used as a laboratory to investigate the effects of boundedly rational behaviour by focusing upon observed asymmetries in stock market returns. Since the models studied in Section 3 and in [12] are completely symmetric with respect to rising and falling prices, we investigate various ways of breaking this symmetry that result in similarly skewed market data. In Section 5 we offer some concluding remarks and directions for future research.

2 The Model

The system is evolved in discrete timesteps of length h which will be chosen to correspond to one trading day for the simulations in this paper. There are M investors, all of equal size, who can be either long or short in the market over the n^{th} time interval. The market price at the end of the n^{th} time interval is $p(n)$. For clarity $p(0) = 1$ and we assume that the expected price is drift-free so that, in reality, $p(n)$ corresponds to, say, the price corrected for the risk-free interest rate plus equity-risk premium or the expected rate of return.

There are two, essentially separable, components to the model. The first is the market pricing model which describes how $p(t)$ depends upon both new extraneous information and the internal market dynamics. The second is the set of rules by which each agent determines if and when to switch positions.

We start by postulating the existence of a globally available, uncorrelated information stream that is external to the market. This takes the form of Gaussian increments $\sqrt{h}\eta_n$ with variance h that represent all new information generated and disseminated over the

n^{th} time period. This allows us to define, for each time n , a fundamental or ‘correct’ price $p_F(n)$ via a standard geometric pricing model

$$p_F(n+1) = p_F(n) \exp\left(\sqrt{h}\eta_n\right). \quad (2.1)$$

Such a market pricing model (after the introduction of an Itô drift term that has been omitted here for clarity of exposition) is in perfect agreement with the strong EMH and leads directly to the standard theory of financial mathematics. Thus our first objective is to modify (2.1) by introducing additional sources of information and pricing pressure that are generated by the market itself.

Denote the position of the i^{th} investor over the n^{th} time interval by $s_i(n) = \pm 1$ (+1 long, -1 short), and the sentiment of the market by the average of the states of all of the M investors

$$\sigma(n) = \frac{1}{M} \sum_{i=1}^M s_i(n). \quad (2.2)$$

The change in market sentiment from the previous time interval is defined by $\Delta\sigma(n) = \sigma(n) - \sigma(n-1)$ and the resulting pricing pressure is incorporated into the model via

$$p(n+1) = p(n) \exp\left(\sqrt{h}\eta_n + \kappa\Delta\sigma(n)\right),$$

where κ , the market depth, is a parameter that defines the relative importance of market sentiment versus new information on the asset price. This linear relationship between price and sentiment/buy-bid pressure has been often used in the literature, although there is evidence that the relationship is better described by a concave function [49].

Our second modification is to include a parameter that reflects the excess influence of new information $\sqrt{h}\eta$ on market prices during times of extreme sentiment (i.e. when $|\sigma|$ is close to 1). Such a relationship between market volatility and sentiment is very natural and can be explained both by the increase in very short-term speculators (e.g. day traders who are not directly simulated in our model) and the increased attention that is paid to such information by all traders during highly polarized markets. This leads us to introduce the following price update formula

$$p(n+1) = p(n) \exp\left(\sqrt{h}\eta_n(1 + \alpha|\sigma(n)|) + \kappa\Delta\sigma(n)\right).$$

It was noted in [12] that any non-trivial causal relationship between market sentiment and volatility will result in volatility clustering since the market sentiment typically displays

long-time correlations. To date however there seems to have been very little, if any, research quantifying the correlation between the two.

In order to close the model we must now specify how the states of the individual agents are determined, i.e. how the i^{th} agent decides when to switch. This of course is where the individual agents' bounded rationality enters the model and we first consider the herding tendency caused by being in the minority. We imbue each agent with a fixed threshold $C_i > 0$ which determines their tolerance to being in the minority. At time n , her herding tendency level is denoted by $c_i(n)$. This level is incremented via $c_i(n+1) = c_i(n) + h|\sigma(n)|$ (i.e. increased by an amount proportional to the length of the time interval and the severity of the inconsistency) whenever $s_i(n)\sigma(n) < 0$. Otherwise, the investor's herding level remains unchanged from one timestep to the next and $c_i(n+1) = c_i(n)$. As soon as $c_i(n)$ exceeds C_i , the investor switches market position and c_i is reset to zero.

Of course, now some mechanism is required to ensure that at least some agents eventually abandon the majority. This is achieved by introducing an 'inaction' tension — as the current price moves further away from the one at which agent i last traded, she feels an increasing desire to switch positions. Note that although we are introducing this effect as a psychological tension, it does have other interpretations (see below) and indeed is the very core of the model. Let P_i be the price at which the i^{th} investor last switched positions and let $H_i > 0$ be an inaction threshold, chosen randomly from an interval $[H_L, H_U]$. Then, as long as the current price $p(n)$ stays within the geometric interval $[P_i/(1 + H_i), P_i(1 + H_i)]$, the investor maintains her position. But if the current price $p(n)$ leaves this interval, the investor switches.

The model is completely defined as follows:

$$p(n+1) = p(n) \exp \left(\sqrt{h}\eta_n(1 + \alpha|\sigma(n)|) + \kappa\Delta\sigma(n) \right). \quad (2.3)$$

1. The herding tension level for investor i at time n , $c_i(n)$, is updated via

$$c_i(n+1) = c_i(n) + h|\sigma(n)| \quad \text{if } s_i(n)\sigma(n) < 0 \quad (2.4)$$

$$c_i(n+1) = c_i(n) \quad \text{otherwise.} \quad (2.5)$$

2. Let P_i be the price at which the i^{th} investor last switched. If either $p(n+1) \notin [P_i/(1 + H_i), P_i(1 + H_i)]$ or $c_i > C_i$ (or both) then the long/short position of the investor is reversed, P_i is reset to the current price, a new H_i is chosen from the uniform distribution on $[H_L, H_U]$ and c_i is also reset to zero.

The only difference between the model presented above and the final model of [12] is that the inaction thresholds H_i are reassigned after each switching rather than being fixed throughout the simulation. Thus many of the simplifications and assumptions that have been implicitly made in the above model have already been discussed there. For example, each agent may have a different value of the *perceived* overall sentiment, close to but different from the true value. Details such as these can be incorporated into the model very easily but appear to make very little difference, if any, to the qualitative statistical behaviour of the model and its ability to replicate the stylized facts. Another important feature of the model is that it appears to be extremely robust — no fine-tuning of the parameters or the way in which the inaction and herding tensions are updated/interact is required. Furthermore, most of the parameters can easily be ‘ball-parked’ and, in principle, could be refined further via suitable statistical analyses of market behaviour and psychological experimentation. It is to be hoped that new models, such as presented here, will help stimulate such investigations.

At this point it is worthwhile to step back and compare the complete model with a real financial market, especially with regard to the different jobs and roles of the people and institutions involved. This is because the model does not attempt to directly simulate all the market participants, merely those whose bounded rationality is most significant over the timescales of interest. We reiterate that our modeling philosophy is to systematically examine the effects of agent-coupling on an otherwise efficient market, not to create as realistic a model as possible (or at least not yet).

Firstly, the information stream $\sqrt{h}\eta_n$ is created by companies themselves, myriad government departments, accounting firms, research institutions and analysts whose job is to collect, examine or publish data on particular firms or the economy as a whole. While $\sqrt{h}\eta_n$ cannot be observed directly (in the presence of either $\kappa > 0, \alpha > 0$) we are assuming implicitly that all this information/research is correct/unbiased and instantaneously available to all the agents (via the financial press, for example). In recent years these assumptions have become increasingly suspect, especially with regard to research that has become compromised and with regard to dubious accounting practices. Nevertheless, it is not our intention to relax those parts of the EMH here.

In the standard financial model (2.1), which corresponds to $\kappa = 0, \alpha = 0$ in (2.3), arbitrageurs and market makers will determine the ‘correct’ price $p_F(n)$. We suppose that they are still present, and operating in the same manner, but do not consider them

directly. In particular, we do not specify the particular mechanism used by the market maker to determine the price. The stylized facts are prevalent in almost all financial markets, independent of whether they are, for example, limit-order or market-order driven.

The M agents themselves are those herding-influenced traders whose timescale is longer than h but shorter than the duration of the simulation. Agents with longer timescales need not be considered (since they will not trade much, if at all); those with shorter timescales, such as day traders, are incorporated into the term $1 + \alpha|\sigma|$; and those with no herding tendency are collectively assumed to be covered by the rational expectations assumption.

Finally, the thresholds H_i can be interpreted as modeling several different effects simultaneously: a predetermined price range that reflects the agent’s predictive model and expectations; as a limit to the price changes that can be tolerated by the agent before needing to either cut losses or take profits; or the point at which the agent, if she is in the majority, believes that sentiment is about to switch. The first interpretation shows how such analysts (and inductive learning) are indirectly included in the model. Of course, in reality the quantity H_i may be a number that changes in real-time, rather than being fixed a priori at the time of the last switch in position, but simulations suggest that this unnecessarily complicates the model. Furthermore, as noted in [12], the presence of the inaction thresholds leads to hysteresis in the agents’ individual market responses [11] and is capable of directly replicating other effects such as transaction costs and the phenomenon of ‘anchoring.’ This occurs when the value placed upon an asset or estimate by an individual is influenced by her recent exposure to other numerical quantities (in this case, the price at which she last traded) and has been well-documented by behavioural economists and psychologists [30, 52].

There are two interesting limits to consider with respect to the main model parameters. Firstly, as stated above, in the limit $\kappa \rightarrow 0, \alpha \rightarrow 0$, the pricing formula decouples from the trading model, and the fundamental price p_F generates Gaussian log-price statistics. Nonzero values for κ, α provide the linkage between the pure information stream $p_F(n)$ and the overall market sentiment $\sigma(n)$. The market then becomes non-Markovian since the current and future behaviour of agents is strongly dependent upon their trading history and previous values of σ and the price p .

There is however a second, more illuminating, limit via which we can recover the efficient-market price $p_F(n)$. If we let $\alpha = 0, M \rightarrow \infty, C_i \rightarrow \infty$ then herding effects

vanish and all trading is due to inaction thresholds being violated. Assuming that the initial sentiments of all the participants are sufficiently mixed, $\sigma(n)$ will remain arbitrarily close to 0 for all n , implying that the price $p(n)$ is arbitrarily close to the fundamental price $p_F(n)$ in (2.1). Therefore, in this limit, we have created a market in which agents have a rationale for trading (interpreting the interval $[P_i/(1 + H_i), P_i(1 + H_i)]$ as being based upon the i^{th} agent’s market analysis) and which is always in a quasi-equilibrium with long and short agents being balanced. This is philosophically very close to the hypothetical markets of standard mathematical finance operating under the strong EMH and rational expectations — trading occurs as a consequence of differing expected future performance and market analysis of individual traders, yet coupling does not occur and the market price remains correct.

3 Stylized facts and numerical simulations

For the numerical simulation we shall use the same values as in the initial investigation [12]. There is no characteristic timescale in the model — this is introduced via the parameter h which corresponds to the variance of the information stream over one timestep. If this is to reflect one trading day, and we suppose that the standard deviation of daily price moves is of the order of 0.6–0.7%, then this motivates a choice of $h = 0.00004$. All simulations will be for 10000 timesteps, i.e. for a 40 year period. As demonstrated in [12], the number of agents $M = 100$ appears to be sufficient — unlike many other HAMs in the literature, the dynamics are not dependent upon the size of M provided it is sufficiently large.

Once h has been fixed, we suppose that the C_i are chosen from the uniform distribution on $[0.001, 0.004]$, as this leads to herd-induced switching on the timescale of weeks and months for those agents in the minority. The price ranges for the inaction tension are chosen randomly after every switching from the uniform distribution on the interval 10%–30%, i.e. $[H_L, H_U] = [0.1, 0.3]$. Day traders would of course have much smaller values but our choice of h means that we cannot attempt to directly model changes occurring over such short timescales. The value of the market-depth κ is, in our model, a measure of the relative ability of extraneous information versus internal dynamics to influence the market price. Simulations using the above parameters suggest that a value of $\kappa = 0.2$ results in dynamics that are strongly influenced by both, being correlated to $p_F(n)$ but

differing significantly during periods of extreme market sentiment. Finally, we choose $\alpha = 2$ so that the effect of extraneous news, good or bad, is to move the market three times as much during a completely polarized market than during a neutral one.

As mentioned earlier, this model is the same as that introduced in [12], except that the inaction thresholds for individual agents are randomized after every switching. This allows the model to mimic the agents' own market analysis and research (possibly including other effects such as inductive learning). The agents' herding propensities, as measured by the thresholds C_i , are kept constant throughout the simulation as we consider this to be reflective of an agents' personality or trading philosophy and therefore less likely to change with time. The model is surprisingly robust with respect to both major changes in the switching rules and the parameters. No fine-tuning was required (or attempted) and all of the statistical and qualitative observations that follow appear to arise very naturally from the fundamental premises upon which the model was founded. Figure 1 shows results for a typical run. Figure 1a) shows both the price $p(n)$ (solid line) and

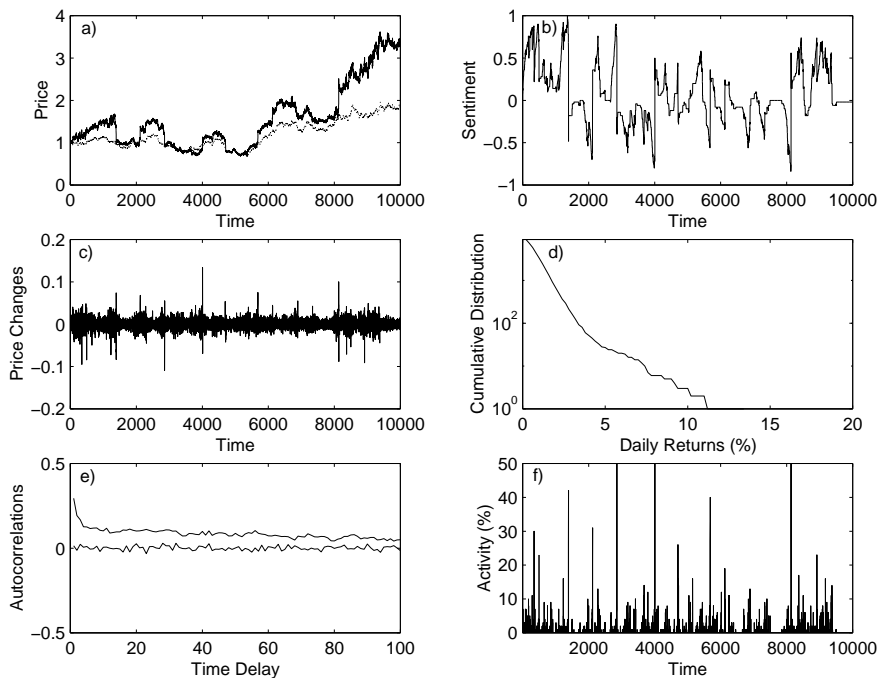


Figure 1: Numerical simulation over a 40-year period. See the text for full details.

fundamental price $p_F(n)$ (dashed line). There are significant long-lasting deviations from p_F over several years, often ended by sudden and dramatic trend reversals. This is also reflected in Figure 1b), which plots the market sentiment $\sigma(n)$ versus time. Figure 1c)

shows the daily log price returns $r(n) = \log\left(\frac{p(n)}{p(n-1)}\right)$. Note that there are several days over the simulation period during which the price change is of the order of 10%. The kurtosis of $r(n)$ was ≈ 37 , as opposed to 3 for a Gaussian dataset (i.e. an excess kurtosis of 34), but this quantity fluctuates to a large degree between runs, typically ranging between 10 and 100.

There is also clear visual evidence of both an excess of extreme events and volatility clustering. This is confirmed by Figures 1d) and 1e). Figure 1d) shows the cumulative absolute returns (i.e. the number of absolute daily price changes larger than a given percentage). If the price returns were Gaussian then one would expect an approximate straight line on this semilog scale. Instead there is a clear excess of days on which high price moves ($> 3\%$) occur. Figure 1e) shows the autocorrelations between daily returns (the almost flat curve) and the absolute daily returns (the price volatility). Thus we see that there is negligible autocorrelation in the price returns even with just a one-day timelag, but that the volatility autocorrelation decays very slowly over a timescale of months – clear evidence of volatility clustering.

The above observations form the basis for the more detailed numerical and statistical investigations that follow. In particular, we visit each of the following ‘stylized facts’ in turn and compare the output of our model against previous studies of real market data:

- uncorrelated price returns
- excess kurtosis in the price return data
- fat-tails in the price return data
- volatility clustering
- Hurst exponent $\approx \frac{1}{2}$ for price returns
- Hurst exponent $\neq \frac{1}{2}$ for volatility
- Slow convergence to Gaussian returns over long time-scales.

Other non-Gaussian phenomena have been noted (such as the price return asymmetry to be explored in Section 4) but the above list forms a core set of properties that a model should be capable of reproducing.

Uncorrelated price returns

There are two very closely related measures of price returns that are both used in the literature, the daily log price returns $r(n)$ defined above and the daily percentage price change $R(n) = \frac{p(n)}{p(n-1)} - 1$. None of the simulations showed any significant difference between the two and, in particular, there is no evidence of correlations between either $r(n)$ or $R(n)$ over periods of a day or longer. This is in agreement with the very many studies of real market data, see for instance [1, 13, 15, 23, 45].

Excess kurtosis

Our observed kurtosis range of 10–100 is consistent with many previous statistical studies. For instance, Longin [36] studied an index of the most traded stocks on the New York Stock Exchange and reported that the distribution of the logarithm daily percentage had excess kurtosis of 22. Pagan [45] reported excess kurtosis ranging from approximately two to approximately 200 when measuring stock returns, changes in bond rates and changes in exchange rates. Cont and Bouchard [10] reported that the distribution of returns of almost all financial stocks exhibits excess kurtosis, which ranges between 2 and 50 for daily returns. The high variability of this statistic is due to the fact that it is so sensitive to the number and magnitude of the the largest fluctuations in the time-series.

Fat-tails

In order to quantify the fat-tails in the price-return data we examine the cumulative distribution of the absolute returns $|r(n)|$ normalized by dividing by their standard deviation. This is plotted on a log-log scale in Figure 2 together with two similar plots for the positive returns and the negative returns taken separately. (The reason for displaying negative and positive returns separately is to emphasize that the model, as it stands, is completely symmetric with respect to price changes in either direction. Modifying this behaviour is the subject of Section 4).

As can be seen, all three data sets are well-approximated by a power-law decay with very similar exponents. This exponent was estimated to be in the range [2.8, 3.2] for every simulation tested. This is in strong agreement with empirical measurements of real market data. For instance, Gopikrishnan *et al* [22] studied a database of every trade for all stocks in three major US stock markets from January 1994 through December 1995 and reported an asymptotic power-law behavior for the cumulative distribution with an exponent of approximately 3. Similar values were reported in [35]. In another study, Plerou *et al* [48] studied the stock price fluctuations of individual companies and found that the tails

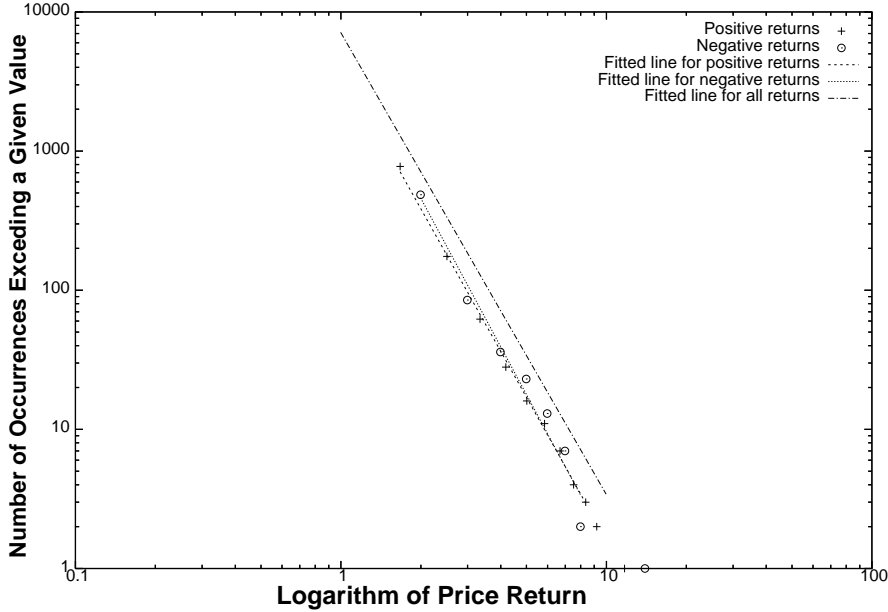


Figure 2: The cumulative distributions of the log price returns (positive only, negative only and absolute value) along with the line of best fit.

of the distributions could be well described by power-law decay, characterized by an exponent $2.5 < \alpha < 4$, for time-scales from five minutes up to approximately 16 days. Dacorogna *et al* [14] studied foreign exchange rate returns and found that all estimates of the tail exponent were greater than three but less than five. See also, for example, [23, 36, 37, 38, 43].

Finally we note that the existence of fat-tails can be directly linked to the herding effect. For if the herding tendency is removed by setting $C_i = \infty \forall i$ (but keeping α unchanged), then the fat-tails vanish.

Volatility clustering

Another extremely important stylized fact is the phenomenon of volatility clustering. Whilst the price returns are uncorrelated over all but the very shortest time-scales, their absolute value (the volatility) shows a slow autocorrelation decay over periods on the order of months. Figure 3 shows this autocorrelation plotted on a log-log scale from 1—100 days. An approximate power-law decay is consistent with many previous empirical studies [41]. Using linear regression to estimate the decay exponents resulted in values typically between -0.3 and -0.6 for all simulations. This range agrees well with, for example, the value of -0.37 obtained by Cont in his recent study of S&P Index future

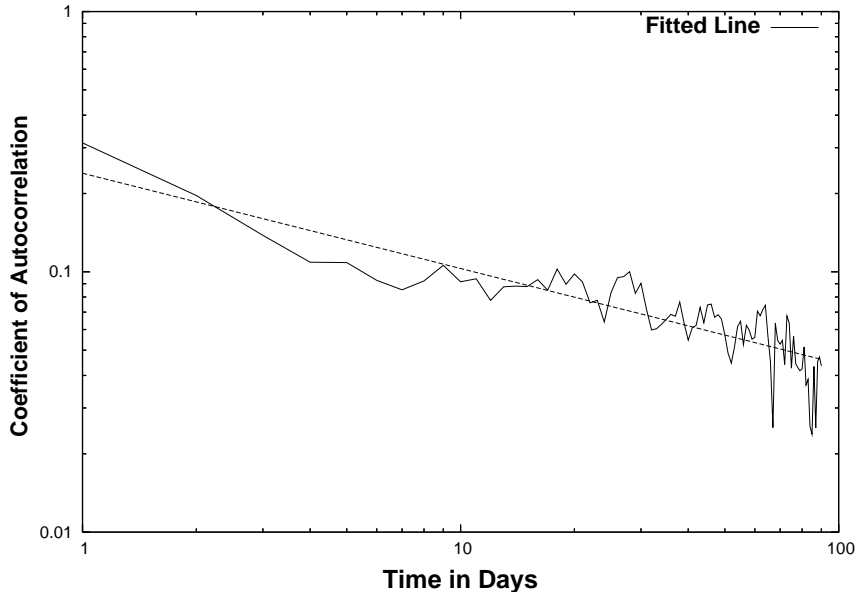


Figure 3: Decay of correlation of volatility and approximation with power law.

prices [8] and the value of -0.3 reported in [35] for S&P 500 price data.

It turns out that the volatility clustering can be directly related to the presence of $\alpha \neq 0$ in the model. When α is reduced to 0, the fat-tailed price returns and kurtosis remain but volatility correlations disappear except over time-scales of 3–5 days. This phenomenon was previously reported in [12].

Hurst exponent for price returns

The Hurst exponent H of a time-series is closely related to the existence of autocorrelations and the decay rate of the power-law spectrum [57]. In particular, for a pure uncorrelated Brownian motion $H = \frac{1}{2}$. Data with $H > \frac{1}{2}$ are said to demonstrate ‘persistence’ and show positive correlation effects while $H < \frac{1}{2}$ is referred to as ‘anti-persistence.’

We estimated the Hurst exponent of time-series generated in the simulated market using the Detrended Fluctuation Analysis (DFA) method, full details of which can be found in [47]. Prior to using this implementation of the DFA method to estimate the Hurst exponent of time-series data generated by the simulated market, we compared the DFA method to the well-known R/S method which goes back to Hurst himself [42]. We also tested both methods for bias by running each method through three tests: both the DFA method and the R/S method were tested on true Brownian motion, $H = 0.5$, and on fractional Brownian motions with known Hurst exponents equal to 0.25 and 0.75

True fBM	R/S	DFA
0.25	0.315	0.256
0.50	0.537	0.491
0.75	0.744	0.748

Table 1: Computed Hurst exponent by each method compared to the Hurst exponent of the fractional Brownian motion (fBM) used as input.

(all time-series were of length $N = 10,000$). The Brownian motions, both true and fractional, were generated from an algorithm that creates a time-series of known Hurst exponent [46]. Both methods detected the correct Hurst exponent reasonably well, but our results showed that the DFA method was the more accurate. The results of these tests are summarized in Table 1. Using DFA, the Hurst exponents of both $r(n)$ and $R(n)$ were to be $H = 0.49 \pm 0.01$, close enough to $H = 0.5$ for the time-series to display no ‘memory’ and mimic an efficient market. This value of $H = 0.49$ for the simulated market is consistent with measures of the Hurst exponent in many real, well-developed financial markets. Carbone, Castelli and Stanley [6] reported on measuring a time-dependent Hurst exponent in financial time-series of the German stock market index DAX and a government bond market BOBL. Log returns for the DAX stock index produced a Hurst exponent of $H = 0.490$ with a standard deviation of $\Delta H(\%) = 4.16$. Over this same time period, the Hurst exponent for the log returns of the BOBL averaged $H = 0.486$ with a standard deviation of $H(\%) = 3.23$, although H ranged from approximately 0.46 to approximately 0.54. Lipka and Los looked for persistence in the time-series of prices, returns and logarithms of price returns of indices of eight European stock markets using seven different methods [34]. The eight Hurst exponents reported were in the range $[0.41, 0.55]$.

However it is interesting to note that some estimated values of H for real markets have varied significantly from $\frac{1}{2}$. In their study of the evolution of market efficiency in Latin American and Asian markets, Cajueiro and Tabak [5] measured Hurst exponents and found them to range widely between $H = 0.45$ and $H = 0.75$. Also, in their recent study [57], Kyaw, Los and Zong report higher values of Hurst exponents, indicating long-term dependence of financial time-series in several Latin American markets — especially those markets that are comparatively small, less liquid and generally considered as ‘emerging.’ It would be very interesting to see whether suitable changes to our model and the agents’

Time unit	Kurtosis
day	16.69
week	19.47
month	13.88
year	3.03

Table 2: Mean kurtosis for the normalized distribution of relative price changes computed for selected time intervals. The kurtosis for a Gaussian distribution is 3.

behaviour could reproduce this significant difference between newer and more established markets.

Hurst exponent for volatility data

Over ten runs, the average Hurst exponent was $H = 0.86$ for both $|r(n)|$ and $|R(n)|$. Such a large Hurst exponent indicates highly-significant long-term dependencies in the time-series and is in strong agreement with the value $H \approx 0.9$ for the volatility of the S&P 500 as reported in [7].

Slow convergence to Gaussian returns over long time-scales

Price returns often exhibit slow convergence to a Gaussian distribution as the time interval increases, i.e. yearly price returns display far less excess kurtosis compared to monthly, weekly and especially daily returns. Previous work on actual market data can be found in [23, 48].

To investigate this phenomenon, we ran our model 100 times and computed the mean kurtosis using time-scales reflecting daily, weekly, monthly and yearly price changes. The numbers reported in Table 2 are in broad agreement with the very wide range of empirically determined kurtosis reported in the literature and over yearly time intervals the excess kurtosis is negligible (see, for example, [25, 45, 26]).

4 Introducing asymmetries into agent behaviour

The model introduced in Section 2 is completely symmetric with respect to falling or rising prices. However many aspects of agent behaviour, as repeatedly verified by experimental economists and psychologists, are asymmetric [2, 4, 3, 19, 24, 30, 44, 51, 56]. It has also been noted that for many financial assets there are more large drawdowns in prices but not

equally large upward movements (gain/loss asymmetry) [9, 28, 29]. In order to examine possible causal relationships between asymmetric agent behaviour and market statistics, we now consider some asymmetric modifications of the inaction and herding tensions.

People, on average, are far quicker to realize a gain than a loss, often trading too soon in the first case and too late in the second [44, 51]. In order to replicate this effect, the inaction interval for the i^{th} trader is shifted so that it is asymmetric with respect to the last trading price P_i . To be explicit, for an agent who has just switched from short to long, it is redefined as

$$[P_i/(1 + 3H_i/4), P_i(1 + H_i/4)]$$

while for an agent who has just switched to being short it is

$$[P_i/(1 + H_i/4), P_i(1 + 3H_i/4)].$$

However, no resulting asymmetries were noted in the price returns. On reflection, this is not surprising since the modification that was introduced is asymmetric with respect to the position held by an individual agent, but not with respect to the market price. Variations on the above approach that shifted the interval consistently in one direction were also attempted. These represented other possible effects, such as the fact that short-selling is more difficult than owning a stock (and has potentially unlimited liabilities) or selling induced by margin calls. However none of them introduced any noticeable asymmetries into the price-return data. This is not to say that such effects are unimportant or do not result in asymmetries, merely that their inclusion via shifting the inaction thresholds does not result in asymmetric data.

Next, an asymmetry was introduced into the herding threshold. The basic premise is that the herding pressure is significantly stronger in a bullish market than a bearish one. In other words, one is made to feel more foolish, more quickly, being contrarian in a bullish environment. To induce this effect in our model we simply keep the thresholds C_i unchanged but increase the rate at which the tensions c_i increase when $\sigma > 0$:

$$\begin{aligned} \text{when } \sigma > 0 : \quad & c_i(n+1) = c_i(n) + 2h|\sigma(n)| \quad \text{if } s_i(n)\sigma(n) < 0 \\ \text{when } \sigma < 0 : \quad & c_i(n+1) = c_i(n) + \frac{1}{2}h|\sigma(n)| \quad \text{if } s_i(n)\sigma(n) < 0. \end{aligned}$$

With the above change in place, the price returns now display marked asymmetries. In Figure 4 (to be compared with Figure 2) the (normalized) positive and negative log

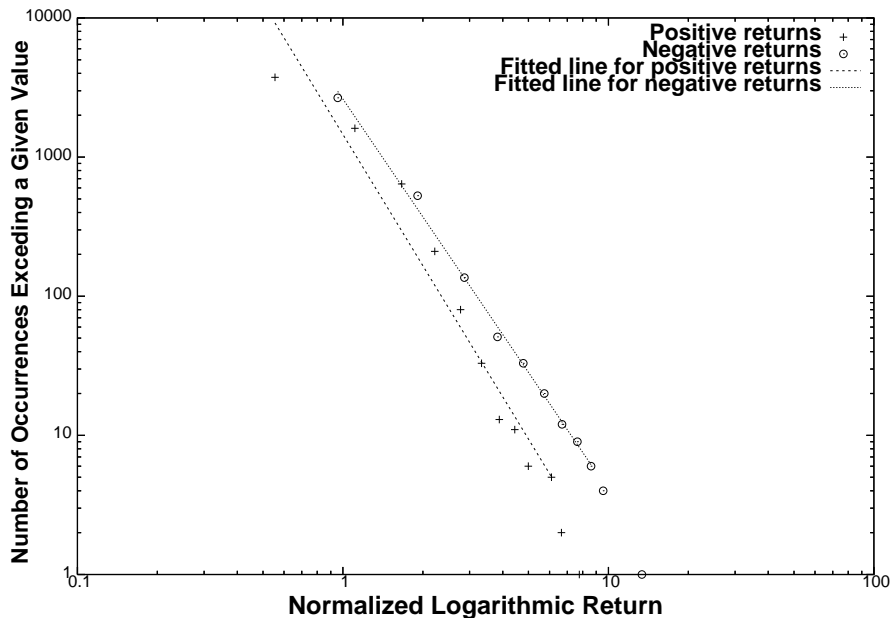


Figure 4: Distribution of normalized logarithmic returns separated into increases (positive) and decreases (negative).

returns are plotted separately and there is a marked majority of decreases over increases that exceed a given (large) magnitude.

In fact, this modification has another immediate consequence, namely that price returns should display noticeably different behaviours in bull markets as opposed to bear markets. Recent work by Kaizoji [32] on the Japanese Nikkei from 1975 to 2002 provides support for this assertion. In that paper, the Nikkei data was split into two parts — the bullish period before December 1989 and the bearish one after. Kaizoji showed that the pre-December 1999 returns were well-replicated by a power-law distribution, and the post-December 1999 by an exponential one.

In order to attempt a comparison with the results of [32], we split our time-series into ‘bullish’, ‘bearish’ and ‘sideways’ markets according to whether $\sigma > \frac{1}{3}$, $\sigma < -\frac{1}{3}$ or $\sigma \in [-\frac{1}{3}, \frac{1}{3}]$ respectively. The log price return cumulative distributions were calculated over the bullish and bearish periods and plotted in Figures 5 and 6, respectively. The tail of the bull market returns is well-fitted by a power-law decay, as before, but the bearish periods are better-fitted by an exponential (note the linear x -scale). Intuitively, this difference can be explained by the fact that as the importance of herding effects is reduced, the market is closer to a Gaussian one in which exponential tails are to be

expected.

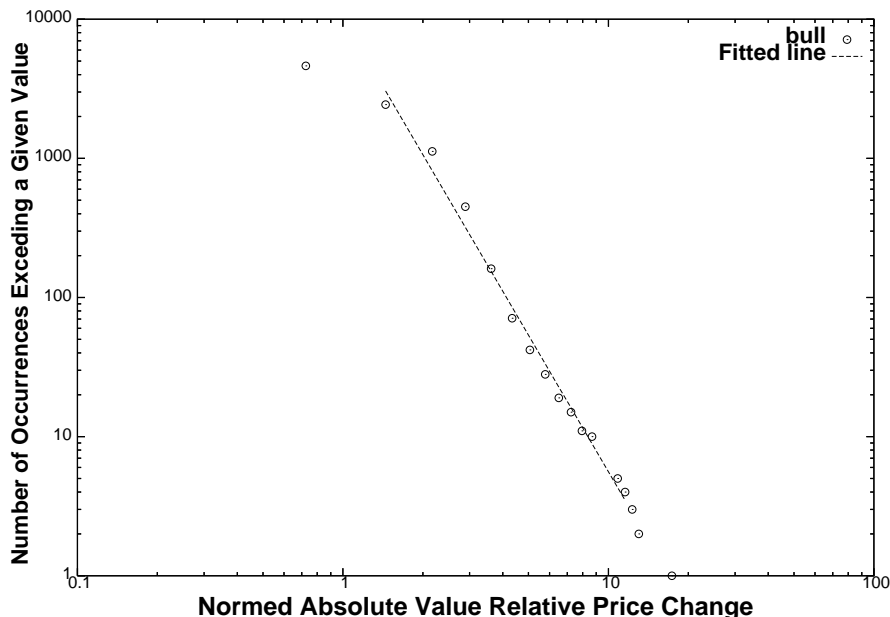


Figure 5: Cumulative distribution of bull market normed absolute value of relative price changes.

5 Conclusions

In this paper we have showed that a very simple psychology-based heterogeneous agent model can reproduce all of the most important stylized facts. Given the robustness of the emergent statistics with respect to the small number of parameters, this closeness of fit to real market data is especially noteworthy.

To be more specific, in [12] it was demonstrated that the introduction of a threshold-type herding tendency into a pure EMH model results directly in excess kurtosis and fat-tailed returns. Here we have demonstrated that the rate of decay of the tails is consistent with a power-law exponent of ≈ 3 , closely matching observations of real data. Similarly, in [12] the inclusion of the term $\alpha \neq 0$ into (2.3) induces a long-term volatility correlation and here it was shown that again it has the correct functional form and decay exponent. Finally, in Section 4 we showed that skewed statistics can be achieved by certain asymmetries in agent behaviour related to the herding tendency (but not by more obvious changes to the inaction thresholds).

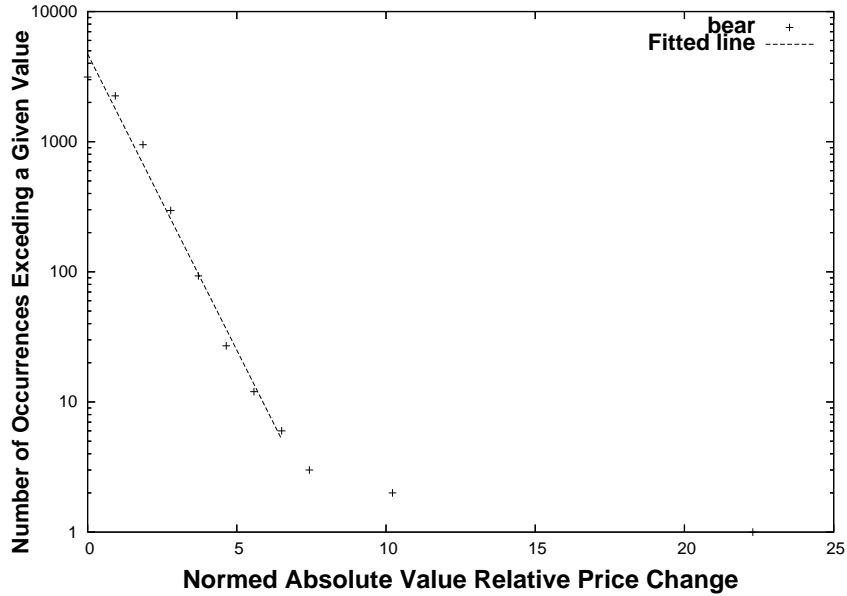


Figure 6: Cumulative distribution of bear market normed absolute value of relative price changes.

Much of the simplicity (and appeal) of the model comes from the fact that it contains no spatial information (such as neighbour-neighbour interactions) or inductive learning. That it replicates the stylized facts so well does not mean that these extra complications are unimportant in real marketplaces, but it does suggest that they may not be necessary for a deep understanding of the causes of the stylized facts.

Finally, we remark that the mere existence of power-laws in such a simple system is extremely interesting. Some previous explanations for the existence of inverse-cubic price-return decays [54, 18] are predicated upon the fact that the distribution of agent sizes follows a Pareto law with known exponent. However, the above results show that the same exponent can be observed in a model with all agents of equal size. A deeper theoretical understanding of how power-law statistics can be induced by threshold descriptions of herding may yield important insights.

References

- [1] A. Arneodo, J-P. Bouchard, R. Cont, J-F. Muzy, M. Potters, and D. Sornette. Comment on “Turbulent cascades in foreign exchange markets”. *cond-mat. 9607120*,

1996.

- [2] B.M. Barber and T. Odean. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The Journal of Finance*, 55(2):773–806, 2000.
- [3] B.M. Barber and T. Odean. Boys will be boys: Gender, overconfidence, and common stock investment. *The Quarterly Journal of Economics*, 141:261–292, 2001.
- [4] N. Barberis, M. Huang, and T. Santos. Prospect theory and asset prices. *The Quarterly Journal of Economics*, 116:1–53, 2001.
- [5] D.O. Cajueiro and B.M. Tabak. The Hurst exponent over time: testing the assertion that emerging markets are becoming more efficient. *Phys. A*, 336:521–537, 2004.
- [6] A. Carbone, G. Castelli, and H.E. Stanley. Time-dependent Hurst exponent in financial time series. *Phys. A*, 344:267–271, 2004.
- [7] P. Cizeau, Y. Liu, M. Meyer, C.-K. Peng, and H.E. Stanley. Volatility distribution in the S&P500 stock index. *Phys. A*, 245:441–445, 1997.
- [8] R. Cont. Scaling and correlation in financial data. *cond-mat. 9705075*, 1997.
- [9] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1:223–236, 2001.
- [10] R. Cont and J. P. Bouchaud. Herd behavior and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*, 4:170–196, 2000.
- [11] R. Cross, M. Grinfeld, H. Lamba, and A. Pittock. Frustration minimization, hysteresis and the El Farol problem. In M.P. Mortell, R.E. O’Malley Jr., A. Pokrovskii, and V. Sobolev, editors, *Relaxation Oscillations and Hysteresis*, pages 61–72. SIAM, 2005.
- [12] R. Cross, M. Grinfeld, H. Lamba, and T. Seaman. A threshold model of investor psychology. *Phys. A*, 354:463–478, 2005.
- [13] M.M. Dacorogna, U.A. Müller, R.J. Nagler, R.B. Olsen, and O.V. Pictet. A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. *Journal of International Money and Finance*, 12(4):413–438, 1993.

- [14] M.M. Dacorogna, U.A. Müller, O.V. Pictet, and C.G. de Vries. The distribution of extremal foreign exchange rate returns in extremely large data sets. Technical report, Olson and Associates Research Group, 1995.
- [15] E.F. Fama. Efficient capital markets: A review of theory and empirical work. *J. Finance*, 25:383–417, 1970.
- [16] E.F. Fama. Efficient capital markets. *J. Finance*, 46:1575–1618, 1991.
- [17] J.D. Farmer. A simple model for the nonequilibrium dynamics and evolution of a financial market. *Int. J. Theor. Appl. Fin.*, 3:451–468, 2000.
- [18] X. Gabaix, P. Gopikrishnan, V. Plerou, and H.E. Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 2003.
- [19] S. Gervais and T. Odean. Learning to be overconfident. *The Review of Financial Studies*, 14(1):1–27, 2001.
- [20] I. Giardina and J.-P. Bouchaud. Bubbles, crashes and intermittency in agent based market models. *Eur. Phys. J. B*, 31:421–437, 2003.
- [21] I. Giardina and J.-P. Bouchaud. Volatility clustering in agent based market models. *Phys. A*, 324:6–16, 2003.
- [22] P. Gopikrishnan, M. Meyer, L.A. Nunes Amaral, and H.E. Stanley. Inverse cubic law for the distribution of stock price variations. *European Physical Journal B*, 3:139–140, 1998.
- [23] P. Gopikrishnan, V. Plerou, L.A. Amaral Nunes, M. Meyer, and H.E. Stanley. Scaling of the distribution of fluctuations of financial market indices. *Phys. Rev. E*, 60:5305–5316, 1999.
- [24] M. Grinblatt and M. Keloharju. What makes investors trade? *The Journal of Finance*, 55(2):589–616, 2001.
- [25] D.M. Guillaume, M.M. Dacorogna, R.R. Davé, U.A. Müller, R.B. Olsen, and O.V. Picket. From the bird’s eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and Stochastics*, 1:95–129, 1997.

- [26] G.W.Schwert. Indices of U. S. stock prices from 1802 to 1987. *The Journal of Business*, 63(3):399–426, 1990.
- [27] M. Hart, P. Jeffries, P.M. Hui, and N.F. Johnson. Crowd-anticrowd theory of multi-agent market games. *Eur. Phys. J. B*, 20:547–550, 2001.
- [28] M.H. Jensen, A. Johansen, and I. Simonsen. Inverse statistics in economics: The gain-loss asymmetry. *cond-mat/0211039*, 1, Nov 2002.
- [29] A. Johansen and D. Sornette. Large stock market price drawdowns are outliers. *Journal of Risk*, 4(2):69–110, 2001/2002.
- [30] D. Kahneman and A. Tversky. Judgement under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- [31] T. Kaizoji. Speculative bubbles and crashes in stock markets: an interacting-agent model of speculative activity. *Phys. A*, 287:493–506, 2000.
- [32] T. Kaizoji. Inflation and deflation in financial markets. *Phys. A*, 343:662–668, 2004.
- [33] T. Kaizoji, S. Bornholdt, and Y. Fujiwara. Dynamics of price and trading volume in a spin model of stock markets with heterogeneous agents. *Phys. A*, 316:441–452, 2002.
- [34] J.M. Lipka and C.A. Los. Long-term dependence characteristics of European stock indices. Technical report, Kent State University, 2003.
- [35] Y. Liu, P. Gopikrishnan, P. Cizeau, M. Meyer, C-K. Peng, and H.E. Stanley. Statistical properties of the volatility of price fluctuations. *Physics Review E*, 60:1390–1400, 1999.
- [36] F.M. Longin. The asymptotic distribution of extreme stock market returns. *Journal of Business*, 69(3):383–408, 1996.
- [37] T. Lux. The stable paretian hypothesis and the frequency of large returns: An examination of major German stocks. *Applied Financial Economics*, 6(6):463–475, 1996. Available at <http://ideas.repec.org/a/taf/apfiec/v6y1996i6p463-75.html>.
- [38] T. Lux. The limiting extremal behaviour of speculative returns: an analysis of intra-daily data from the Frankfurt Stock Exchange. *Applied Financial Economics*, 11(3):299–315, 2001.

- [39] T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397:493–500, 1999.
- [40] T. Lux and M. Marchesi. Volatility clustering in financial markets: A micro-simulation of interacting agents. *Int. J. Theor. Appl. Finance*, 3:675–702, 2000.
- [41] R. Mantegna and H. Stanley. *An introduction to econophysics*. CUP, 2000.
- [42] T. Di Matteo, T. Aste, and M. M. Dacorogna. Scaling behaviors in differently developed markets. *Phys. A*, 324:183–188, 2003.
- [43] U.A. Müller, M.M. Dacorogna, R.B. Olsen, O.V. Pictet, M. Schwarz, and C. Morgenegg. Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. *Journal of Banking and Finance*, 14(6):1189–1208, 1990.
- [44] T. Odean. Are investors reluctant to realize their losses? *J. Finance*, 53:1775–1798, 1998.
- [45] A. Pagan. The econometrics of financial markets. *Journal of Empirical Finance*, 3:15–102, 1996.
- [46] V. Paxson. Fast, approximate synthesis of fractional Gaussian noise for generating self-similar network traffic. *Computer Communication Review*, 27(5):5–18, October 1997.
- [47] C-K. Peng, S.V. Buldyrev, S. Havlin, M. Simons, H.E. Stanley, and A.L. Goldberger. Mosaic organization of DNA nucleotides. *Phys. Review E*, 49:1685–1689, 1994.
- [48] V. Plerou, P. Gopikrishnan, L.A. Nunes Amaral, M. Meyer, and H.E. Stanley. Scaling of the distribution of price fluctuations of individual companies. *Physics Review E*, 60(6):6519–6529, 1999.
- [49] V. Plerou, P. Gopikrishnan, X. Gabaix, and H.E. Stanley. Quantifying stock-price response to demand fluctuations. *Phys. Rev. E*, 66, 2002.
- [50] T. Seaman. *Models of selected problems in mathematical finance and numerical methods for stochastic differential equations*. PhD thesis, George Mason University, 2006.

- [51] H. Shefrin and M. Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance*, 40(3):777–790, 1985.
- [52] R. J. Shiller. *Irrational Exuberance*. Princeton University Press, 2000.
- [53] H.A. Simon. A behavioural model of rational choice. *Quart. J. Econ.*, 69:99–118, 1955.
- [54] S. Solomon and P. Richmond. Power laws of wealth, market order volumes and market returns. *Physica A*, 2001.
- [55] F. Wagner. Volatility cluster and herding. *Phys. A*, 322:607–619, 2003.
- [56] C. Zhou. Informational asymmetry and market imperfections: Another solution to the equity premium puzzle. *The Journal of Financial and Quantitative Analysis*, 34(4):445–464, 1999.
- [57] S. Zong, C.A. Los, and N.N. Kyaw. Persistence characteristics of Latin American financial markets. Technical Report 0411013, Economics Working Paper Archive at WUSTL, November 2004. Available at <http://ideas.repec.org/p/wpa/wuwpfi/0411013.html>.