

A Threshold Model of Investor Psychology

ROD CROSS¹
MICHAEL GRINFELD²
HARBIR LAMBA³
TIM SEAMAN⁴

Abstract

We introduce a class of agent-based market models founded upon simple descriptions of investor psychology. Agents are subject to various psychological tensions induced by market conditions, and endowed with a minimal ‘personality’. This personality consists of a threshold level for each of the tensions being modeled, and the agent reacts whenever a tension threshold is reached. This paper considers an elementary model including just two such tensions. The first is ‘cowardice’, which is the stress caused by remaining in a minority position with respect to overall market sentiment, and leads to herding-type behaviour. The second is ‘inaction’, which is the increasing desire to act or re-evaluate one’s investment position. There is no inductive learning by agents, and they are only coupled via the global market price and overall market sentiment. Even incorporating just these two psychological tensions, important stylized facts of real market data, including fat-tails, excess kurtosis, uncorrelated price returns and clustered volatility over the timescale of a few days, are reproduced. By then introducing an additional parameter that amplifies the effect of externally generated market noise during times of extreme market sentiment, long-time volatility correlations can also be recovered.

Keywords: investor psychology, volatility clustering, kurtosis, herding

1 Introduction

For many years it has been apparent that models of financial markets using sets of standard assumptions, often known as *efficient market hypotheses* (EMH), are not capable of reproducing important features of observed market behavior. This manifests itself most clearly in the real-world phenomena of non-Gaussian market statistics such as fat-tails, excess kurtosis and volatility clustering (and the corresponding market bubbles and crashes). The EMH allow for closed-form mathematical solutions to certain fundamental

¹Department of Economics, University of Strathclyde, Sir William Duncan Building, 130 Rottenrow Glasgow G4 0GE, Scotland, UK

²Department of Mathematics, University of Strathclyde, Livingstone Tower, 26 Richmond Street, Glasgow G1 1XH, Scotland, UK

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

⁴Department of Mathematical Sciences, George Mason University, MS 3F2, 4400 University Drive, Fairfax, VA 22030 USA

problems, such as the pricing of plain vanilla European options [4]. Increasingly however, attention has shifted towards trying to model more realistic market behavior.

There are two basic approaches to this modeling problem. The ‘top-down’ approach is to try and suitably perturb efficient-market models and/or their solutions, based upon observed historical deviations from the predicted market behavior. However the universality of such non-Gaussian statistics across differing financial products, cultures and trading rules suggests that basic human psychology is an underlying cause. Thus the ‘bottom-up’ approach is to try and mimic the ways in which the EMH are violated by individual market participants. There are many such agent-based market models in the literature ([26, 42, 3, 24, 28, 29, 30, 31, 16, 8, 38, 20, 5, 25, 18, 19, 27, 41] provides a representative selection that can be mined for yet further references). The problem has, in particular, captured the attention of statistical physicists because of the strong similarities to spin models of magnetization, such as the Ising model. Our models are also agent-based and we start by briefly justifying our overall approach before providing details of the models and numerical simulations in the remainder of the paper.

There are various versions of the EMH (see [13, 14] and standard texts on mathematical finance) which are often classified as being either strong, semi-strong or weak. In the strong version, market prices reflect all the information relevant to the determination of the asset price, including insider-information. The semi-strong form includes just publicly-available information in the information set. In the weak form, the information set incorporates only information on the past behaviour of asset prices. At the individual level, the key assumption is that market participants are ‘rational’ in the sense that they maximize ‘well-behaved’ expected utility functions: when they receive new information they update their beliefs correctly according to Bayes’ law. At the aggregate level, the key assumption is that of *Markovian Markets*:

All new information is acted upon instantaneously so that the current market price reflects all previously known data, implying that future prices are independent of the market history.

In what follows, we provide a class of mathematical models that violate the EMH postulates in a self-consistent and psychologically-plausible manner. This is achieved by endowing our hypothetical agents with minimal ‘personalities’ that determine their reaction to such sources of emotion as investing performance (relative to the market average or investor expectations, say), inaction (the increasing urge to modify or re-evaluate one’s position) and ‘cowardice’ (the discomfort caused by being in the minority at any given time, often referred to in the literature as the ‘herding tendency’ or imitative behavior [2]). Each agent reacts by re-evaluating their market position whenever one of these emotional levels breaches their individual tolerance level. Thus we have replaced perfect rationality by ‘bounded rationality’ (in the sense of Simon [37]). An immediate consequence of this approach is that the history-dependent nature of agents’ emotional level(s) leads to non-Markovian market processes.

A major motivation for this work was to develop models that are as simple as possible, but which are still capable of reproducing the non-Gaussian market statistics observed in

real markets. One major simplification is the absence of any ‘spatial’ relationships between agents and interactions between ‘neighboring’ individual agents. Instead agents are only coupled via the overall market sentiment and the global market price, which itself is partly determined by the market sentiment via an invocation of the law of supply and demand. This simplifies the models considerably since no spatial structure (and all the associated assumptions, parameters etc) needs to be introduced. A further simplification is the absence of any inductive learning by the agents. Inductive learning is a common feature of such heterogeneous agent models, especially in the modeling of minority games (see for example [1, 6, 7] and the econophysics forum at <http://www.unifr.ch/econophysics/>). A key observation of the work presented here is that, even without modeling these effects, excess kurtosis, fat-tails and volatility clustering are observed to a significant degree, in broad agreement with real market data.

It should be emphasized that it is not our intention here to simulate actual historical data or to make specific predictions. Nor do we contend that all participants in any given market at all times can be modeled in the manner we propose. In real markets some groups of agents may indeed be better modeled by including local spatial coupling and/or inductive learning strategies. However, reacting to certain fundamental psychological pressures is a basic human trait and, as such, the effects of this behaviour need to be understood and appreciated. Thus the approach taken here is one of experimentation and analysis on models that are as simple as possible, consistent with underlying reasonable assumptions on human psychology. This allows us to isolate the effects of different forms of human ‘bounded rationality’ and their consequences for market statistics and aggregate behaviour. We regard the ‘caricature’ nature of these models (and the very small number of critical parameters) as an advantage rather than a drawback. Isolating phenomena within a minimalist framework has often proved to be the key to understanding the same events in more realistic environments.

The paper is organized as follows. In Section 2 we introduce a base model that incorporates just two sources of tension. The first is ‘cowardice’ (the increasing stress caused by being in the minority position) and the second is ‘inaction’ which is the increasing urge to act or re-evaluate one’s position since the last decision. The inaction tension also introduces a form of hysteresis into the model, and this is briefly discussed. We then describe various ways in which the base model can be modified. In Section 3 we present numerical simulations which forcefully demonstrate the non-Gaussian nature of the resulting market statistics, in particular, the generation of fat-tails and short-time volatility correlations. An interesting observation is that if the cowardice tension is excluded then the volatility clustering and excess kurtosis vanish. Thus the herding tendency can be directly linked to the existence of these phenomena, at least within this class of models. Finally, we show that heteroskedasticity, and the corresponding long-time volatility correlations, can be introduced simply by allowing the effect of new market information to be a function of market sentiment. The resulting model thus reproduces the most important observed ‘stylized facts’ of market data.

2 The Model

We now describe an initial model, incorporating tension thresholds, that has been chosen for its simplicity. The system is evolved in discrete timesteps of length h and there are M investors, all of equal weight in the market, who can be either long or short in the market (represented by $+1$ and -1 respectively) over each time interval. Whenever a particular tension reaches an agent's pre-defined tolerance level, then that agent will be forced to switch position in order to eliminate that tension (a closely-related model, applied to the classic El Farol bar problem [1], can be found at [9]).

We incorporate just two sources of psychological tension and describe both of them in detail before providing the necessary mathematical formulation. The first tension, which we are referring to as 'cowardice', is the increasing stress that an agent feels while they are in the minority position. We assume that over each time interval, the cowardice tension level of an agent who is in the minority is increased by an amount proportional to the magnitude of the difference in the overall long/short positions. When this agent's tolerance threshold is exceeded, the agent will switch positions to (attempt to) join the majority. At that time the cowardice tension is reset to zero since that participant has now extinguished that particular source of tension. Agents with a lower cowardice tolerance can, alternatively, be interpreted as 'momentum traders' who attempt to ride waves of market exuberance or gloom.

The second included tension is caused by 'inaction' of the participant. There are several ways of incorporating such a tension into the model. The most obvious is simply to measure the length of time since the investor last acted and then force a re-evaluation when that interval becomes too large. However this does not take into account fluctuations in the market price. An investor may be perfectly happy to stay put for a long time in a flat market. It is more reasonable to suppose that the investor will switch due to inaction whenever the price moves by a certain percentage, up or down, from the price at which they last switched. This percentage can then be regarded as an inaction threshold. Specifically, if investor i last made a switch when the market price was P_i then she will next be forced to switch as soon as the price leaves the interval $[P_i(1 - H_i), P_i(1 + H_i)]$ where $100 \times H_i$ represents her tolerance level expressed as a percentage. Since we shall be using a geometric pricing model, we instead use the interval $[P_i/(1 + H_i), P_i(1 + H_i)]$ (the difference between the two is negligible for small H_i). Note that we are not taking into account whether the investor made or lost money. In other words this switching represents either 'cutting one's losses' or 'taking one's profits', depending upon the situation, and is symmetric with respect to both. Experimental and other evidence [12, 22, 32] has shown that this is not the case. However, we consider this observed asymmetry between investors' reactions to profits and losses to be a minor effect at this stage of the modeling process. It can be easily incorporated into a more complex model by, for example, shifting the intervals $[P_i/(1 + H_i), P_i(1 + H_i)]$ so that they are no longer log-symmetric about the last decision price.

The above modeling of the inaction tension can be regarded as a form of hysteresis

[10]. Resetting the price of the i^{th} investor has the effect of introducing a price range within which that investor is comfortable with their current position (ignoring, of course, the cowardice tension). This range then moves, shadowing the market price, as the system evolves and the investor switches positions. Furthermore, the inaction tension, and the existence of the above dynamic price ranges, is closely related to other economic and psychological factors. For example, the inclusion of transaction costs (that prevent agents from switching positions arbitrarily often) will result in very similar hysteresis intervals [39][Volume I]. Also, the inaction tension can be linked to the notion of ‘anchoring’ [23] whereby investors’ subsequent behaviour is dependent upon some, perhaps arbitrary, previous value.

Both the above tensions are recognizably ‘human’. However, it might be argued that the above psychological considerations apply most naturally to non-professional investors. Whether or not this is the case, it is also interesting to consider some other motivations of professional investors and fund-managers. Such professionals, who are rewarded (or fired) according to the performance of the funds under their control, can, if they are operating a managed fund, want to be seen to be doing something, even if they have no particular justification for doing it. Or, in the case of a mainly tracker-fund, they may not want to draw attention to themselves by standing apart from the crowd for too long.

In order to isolate the effects of the cowardice and inaction propensities we abstract from the institutional detail of trading arrangements. Typically trades are conducted via market orders, whereby trades are conducted more or less immediately at the best available price; or via limit orders, where a worst allowable price is specified, and the orders are stored in a queue or limit order book if they fail to trigger an immediate transaction. As demonstrated in [11], non-trivial properties emerge from such a more realistic specification of how prices evolve in an order-based market. It would be interesting to see how investor psychology interacts with market and limit orders in the price change process, but this is beyond the purpose and scope of the present paper.

We now introduce some notation and define a model based upon the above considerations. The state (long or short) of the i^{th} investor over the n^{th} time interval is denoted by $s_i(n) = \pm 1$ and

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

is a measure of the ratio of long to short investors. A value of σ closer to -1 thus denotes a bearish market while a value closer to +1 denotes a bullish one. The market price at the end of the n^{th} time interval is denoted by $p(n)$ and updated via

$$p(n+1) = p(n) \exp \left(\sqrt{h} \Delta W(n) + \kappa \Delta \sigma(n) \right) \quad (2.1)$$

where $\Delta W(n)$ is a standard Gaussian random variable that represents the creation of new, uncorrelated and globally available, information over that time period. Note that the timestep h is measured in units where $h = 1$ corresponds to the interval over which the random variable $\Delta W(n)$ has unit variance. The variable $\Delta \sigma(n) = \sigma(n) - \sigma(n-1)$ is

the most recent change in market sentiment and the constant $\kappa > 0$ determines the effect that a single agent has on the market price — the larger the value of κ , the more the market price is influenced by internal market dynamics as opposed to the generation of new market information. This linear dependence of the log-return rate on change in market sentiment has been used previously in many other agent-based models, even though empirical evidence [33] suggests a concave dependence of market price on supply/demand. In addition, a ‘fundamental price’ $W(n)$ can be determined from the information stream $\Delta W(n)$ by setting $\kappa = 0$ in (2.1) and iterating. It is valuable to compare $W(n)$ and $p(n)$ in simulations as this helps identify extended periods of price deviations due to investor sentiment.

Based on the new price in (2.1), the tension levels and states of all the participants are updated as follows. Let the i^{th} participant have inaction and cowardice thresholds $H_i > 0$ and $C_i > 0$ respectively, assigned at the start from predetermined (possibly joint) probability distributions. Then our rules for updating tensions and inducing switching are:

1. The cowardice level for investor i at time n is defined by $c_i(n)$, and is updated by applying

$$c_i(n+1) = c_i(n) + h|\sigma(n)| \quad \text{if } s_i(n)\sigma(n) < 0 \quad (2.2)$$

$$c_i(n+1) = c_i(n) \quad \text{otherwise.} \quad (2.3)$$

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 and the cowardice tension is also reset to zero.

Note that the factor h in the updating of the cowardice tension (2.2) takes into account the length of time that the participant is in the minority, so that the model is independent of the (somewhat arbitrary) timestep used. Also, the cowardice level of all investors in the majority remains unchanged, while the levels for all minority investors are increased by the same amount $h|\sigma(n)|$ (the differences between investors’ responses are due to their different threshold levels C_i).

We now make some important remarks about the above caricature model, henceforth referred to as the *base model*. These take the form of a list of some of the most significant assumptions that have been made, together with suggestions as to how these could be relaxed in future models.

- Many other psychological tensions could be included, for example ‘contrariness’ which can be thought of as an ‘anti-herding’ tendency caused by an increasing belief that the majority market sentiment is about to switch. This contrariness tension has not been incorporated into the base model because its effect is, to a large degree, replicated by the inaction tension: when an investor is in the majority her cowardice

tension level stops increasing and, unless the market sentiment switches sign first, her next switching (which will be due to the inaction tension) can be interpreted as a belief that the market is about to switch.

- The mechanism for updating the cowardice tensions has, as with much else in the base model, been chosen to be as simple as possible. A more realistic algorithm could perhaps allow the cowardice tension level to subside while an agent is in the majority, or permit some other complex interaction between the two included tension levels.
- In (2.2) it is implicitly assumed that every agent correctly perceives the market sentiment to be $\sigma(n)$. In practice, such knowledge is likely to be imperfect, gleaned as it is from a mix of trading-floor statistics and chatter, and numerous sentiment surveys, indices and reports. This could be more realistically simulated by defining an individual $\sigma_i(n)$ for each agent by simply applying some random quantity to $\sigma(n)$. If a spatial structure is introduced onto the set of agents, then a more complicated model could, for example, adapt the cowardice tension so that agents are affected by the sentiment of their neighbours rather than the market as a whole. However, one of the key motivations for the base model is the absence of any such spatial structure.
- Each agent is forced to switch, with probability 1, whenever a threshold is reached. This can be modified so that there is a non-zero chance of not switching whenever a re-evaluation occurs.
- A market drift (which could reflect either the risk-free interest rate plus the equity-risk premium, or the expected rate of return) is easily incorporated into the model by modifying $\Delta W(n)$ and the price thresholds of the inaction tension so that they are time-dependent and relative to the drift. However this change does not alter the results significantly and so has been ignored for the sake of clarity.
- All agents are assumed to be of equal weight in the market. However, obvious modifications allow for a Pareto (power-law) distribution of agent sizes (see, for example, [34] for models incorporating such wealth distributions).
- We assume that every participant is able to carry out any transaction they have decided upon. In particular, we do not attempt to match up buyers with sellers, specify the trading rules in operation, or incorporate effects due to margin calls and limited cash supplies. This last point is especially noteworthy as it could be modeled by an additional tension that might significantly affect the market dynamics, especially during market corrections. Isolating the effects of such a change would be an interesting application of the modeling approach taken in this paper.
- As mentioned earlier, the base model is perfectly symmetric with respect to profits and losses. In fact, since the profit-and-loss history does not directly affect an agent's decision-making, it need not even be recorded. However, it is a simple matter to develop related models where this symmetry is broken.

3 Numerical simulations

We now choose parameters for a numerical investigation of the base model. The first parameter to be fixed is the timestep h and, for the simulations presented here, we wish this to correspond approximately to one trading day. There are numerous market studies in the literature that calculate standard deviations of price returns over various time periods. However, we prefer not to directly quote any such values for the following two reasons. Firstly, we wish to emphasize that the base model is a caricature based upon fundamental psychological considerations and not upon a particular financial product or market. Secondly, the value assigned to h reflects only the market price fluctuations due to the information stream (i.e. excluding market sentiment) and this number can be bounded from above, but not inferred, from market data alone. For these reasons, we choose the standard deviation of daily returns in the fundamental price $W(n)$ to be approximately 0.6–0.7%. This range seems to be consistent with many different applications and leads us to choose $h = 0.00004$. Our simulations will be over 10000 timesteps which, given that there are approximately 250 trading days in a year, corresponds to a period of about 40 years.

The above choice of timescale now determines the ranges of sensible values for the other parameters in the model. In particular it places limits on the fastest effects that can be modeled. For example, noise traders that operate on the timescale of minutes or hours are now excluded from being represented directly as agents in the model (although their effects can be interpreted as having been incorporated into the forcing term $\Delta W(n)$). We thus choose the inaction threshold for the i^{th} participant from a uniform distribution over the range 10% to 30%. Similarly we assume that the typical participant will switch due to cowardice over timescales ranging from weeks to several months, and so C_i is picked from the uniform distribution over the interval $[0.001, 0.004]$. Notwithstanding our earlier comments regarding managed funds and tracker funds, we suppose that agents with a lower cowardice threshold will also have a lower inaction threshold. Such agents will switch more often and so are more ‘short-term’ than investors with higher thresholds. Therefore, to fix the thresholds for the i^{th} agent, a single random variable z_i is chosen from the uniform distribution on $[0, 1]$ and the thresholds are determined by

$$C_i = 0.001 + 0.003z_i, \quad H_i = 0.1 + 0.2z_i. \quad (3.1)$$

The parameter κ is more difficult to fix a priori. However, trial-and-error revealed that a value of $\kappa = 0.2$ results in market prices that deviate from the fundamental price $W(t)$ to a significant degree, while still remaining highly-correlated with it.

At this point, we note that the model contains two qualitatively different types of parameter. The first are the pure numbers h and κ . The second are the joint probability distributions of the cowardice and inaction thresholds. The model is, predictably, independent of changes in h (provided that h remains small enough). However, at least qualitatively, it appears to be very robust with respect to changes in κ and the probability distributions. In particular, the numerical results shown below are not the result of our choice of correlated threshold distributions. As an example, if the cowardice and

inaction thresholds for the i^{th} agent are chosen from the same distributions as before but using two independent random variables $z_{C,i}$ and $z_{H,i}$, then qualitatively indistinguishable results are obtained. Of course the true situation probably lies somewhere between these two simplified extremes and future work by experimental economists could, in principle, provide better estimates for such joint distributions.

Results from a typical simulation are shown in Figure 1. There are $M = 100$ agents with threshold values generated from the above distribution. Their initial tension levels are randomly assigned, consistent with their threshold levels, as are their initial states ± 1 . Figure 1a) shows the market sentiment as a function of time. It can be seen that there are dramatic shifts in overall sentiment, typically occurring every few years. This is interesting since individual participants may be switching every few weeks or months. Note that there are periods where almost every agent in our model is taking the same position. This is perhaps too extreme, and is due to our modeling every participant in the market in a similar fashion. Nonetheless it seems plausible that in a real market there will be periods during which the vast majority of investors whose actions are dominated by the above sources of tension will be taking the same position in the market.

Figure 1b) shows the daily price returns $\frac{p_{n+1}}{p_n} - 1$ as a function of time. Time intervals containing a few exceptional daily price moves of the order of 10%, can be seen. These correspond to large and dramatic shifts in market sentiment. Figure 1c) plots the fundamental price $W(t)$ (dashed line) versus the market price $p(t)$ (solid line). These (detrended) prices are plotted on a linear scale to accentuate the differences between them, even though a geometric pricing model is being used. Figure 1d) plots the percentage of agents that switched market position at the end of each time period and is thus a measure of the market activity.

Figure 1e) shows a cumulative log-plot of the number of daily returns (positive or negative) that exceed a certain value, from 0–20%. If the daily returns followed a Gaussian distribution then an approximate straight line would be expected. Instead, there is a clear exponential distribution for small daily returns (up to approximately 2%) but then clear evidence of an excess of large daily fluctuations. This is in agreement with market observations. The kurtosis of the daily market returns was also computed over many such runs. It is a highly variable quantity, depending as it does upon the precise number and magnitude of the largest daily fluctuations, but typically ranged between 10 and 100 (for a Gaussian distribution the value should be 3), thus displaying excess kurtosis. Finally, figure 1f) shows the autocorrelations of both the daily market returns and their absolute value (the volatility). The market returns are completely uncorrelated over intervals of a day or longer, as indicated by the lower line, and this is also in accordance with real market time series where these autocorrelations typically vanish over timescales of more than a few minutes. On the other hand, volatility correlations are observed for intervals upto about 4 days – and this is evidence of volatility clustering in the model. However, real market data usually reveal volatility correlations decaying slowly over weeks or months [21, 15, 31]. At the end of this section we shall introduce a simple addition to the base model that introduces volatility clustering over these longer timescales.

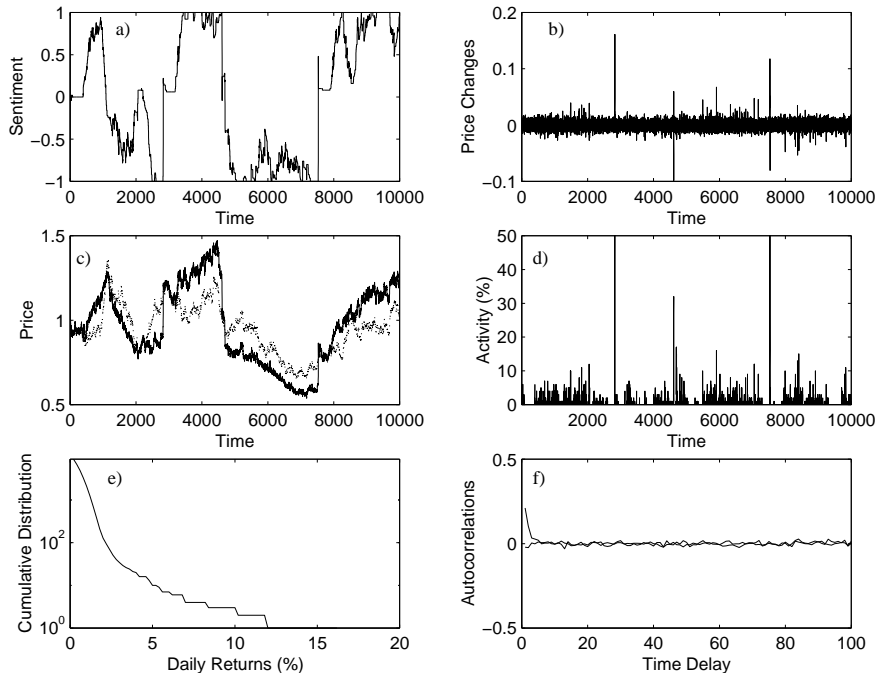


Figure 1: Numerical simulation of the base model.

Before continuing, we note that the dynamics of the model does not depend significantly on the number of agents M . To demonstrate this, the base model was run with 1000 agents, and the results are shown in Figure 2. The two sets of pictures display very similar ranges and timescales for all the important market statistics, and also qualitatively similar behaviour. Thus for the remaining simulations we shall revert to using 100 agents.

In order to ascertain the effect of herding in the model, we now remove the cowardice tension by redefining all the cowardice thresholds $C_i = \infty$. Keeping all other parameter values the same, we observe results typified by Figure 3. The most obvious change from Figure 1 is the complete absence of extreme events in the distribution of daily price returns. Also, although the volume of trading is comparable to before, the vast majority of trades cancel each other and the market sentiment remains very close to zero. In turn, the market return statistics are essentially Gaussian, as displayed by the negligible difference between the market price and fundamental price in Figure 3c), and by the virtually straight line in Figure 3e). The measured value of kurtosis also collapses to 3 ± 0.1 . Thus, at least in the base model, the existence of heavy-tailed distributions and short-term volatility correlations can be attributed entirely to cowardice-induced switching.

The differences between Figures 1 and 3 can be explained as follows. The agents' personalities are defined by the quantities $z_i \in [0, 1]$ that determined their threshold levels. Thus we can picture them as lying along the interval $[0, 1]$. In the absence of the cowardice tension, if agents' initial market positions, long or short, are perfectly mixed along this interval, then they will remain so at all later times. In other words, any change in market price will switch almost equal numbers of long and short investors, which

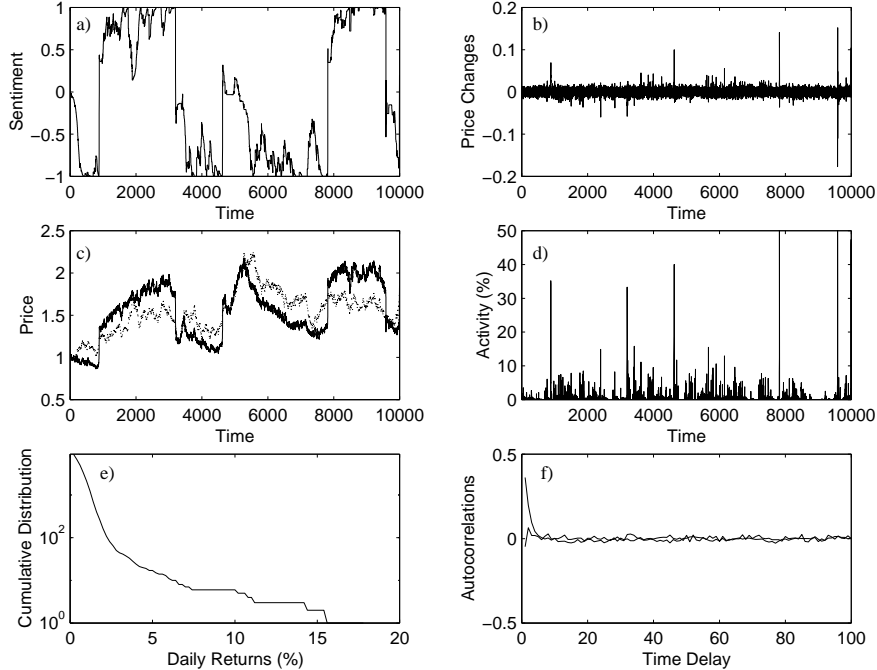


Figure 2: Numerical simulation of the base model with 1000 agents.

then cancel each other out and leaves the sentiment unchanged, apart from discretization effects due to the finite number of traders. The presence of the cowardice tension alters this situation by ‘unmixing’ the market positions along this interval. Interestingly, the number of switches induced by the cowardice tension that are needed to make this occur is relatively small. In other words, in Figure 1, the large majority of agent-switches are caused by the inaction threshold being violated, not the cowardice threshold. A more detailed analysis of the base model, and its market statistics, will appear elsewhere.

We conclude by making a simple modification to the base model that introduces heteroskedasticity. It has been observed (see, for example [35, 36, 40] and references therein) that periods of extreme market volatility often coincide with periods of extreme market sentiment, such as occurred in the 1920s and early 1930s, and the late 1990s. While this relationship is, as yet, neither well-quantified nor well-understood (see [17] for a review and a volatility clustering model based on intermittency and coexistence of attractors), it suggests modifying the base model in such a way that the influence of external shocks on the price $p(t)$ is now a function of market sentiment. Then, since market sentiment displays long-time correlations, this provides a plausible mechanism for volatility clustering over similar timescales. To this end, we assume that the generation of the uncorrelated, randomly occurring market information $\Delta W(n)$ is unchanged, but that its effect on the market is magnified at times of polarized markets, both bullish and bearish (other more complex dependencies are of course possible). Within the context of our base model, we can effect this change most easily by replacing (2.1) with

$$p(n+1) = p(n) \exp \left(\sqrt{h} \Delta W(n) (1 + \alpha |\sigma(n)|) + \kappa \Delta \sigma(n) \right). \quad (3.2)$$




Figure 3: Numerical simulation of the base model with the cowardice tension removed.

One possible interpretation of the new parameter $\alpha > 0$ is that the number and activity of noise (momentum) traders increases during such polarized market conditions. In Figure 4 we choose $\alpha = 2$ (so that news, good or bad, entering a highly-biased market moves the market 3 times further than when market sentiment is completely neutral) and all other parameter values are unchanged from those used to generate Figure 1. As can be seen, the volatility auto-correlation now decreases slowly over the timescale of months while the fat-tails, excess kurtosis and uncorrelated price returns of the base model are retained.

4 Conclusions

A simple framework has been introduced, within which the effects of various aspects of investor psychology can be studied. The inaction tension (and the hysteresis that results) is capable of multiple interpretations, and provides the driving force to ensure that agents never cease trading. By itself, it cannot reproduce non-Gaussian market statistics, but the inclusion of a herding-tendency, as described by the "cowardice" tensions and thresholds, generates fat-tails and short-term volatility correlations while retaining the property of uncorrelated daily price returns. If in addition, the market is modeled as reacting more strongly to new information at times of extreme sentiment, then long-term volatility clustering also appears. Thus, the simultaneous modeling of these three effects reproduces the most significant stylized facts of real-world market statistics.

The models presented simulate a very small number of fundamental phenomena, and

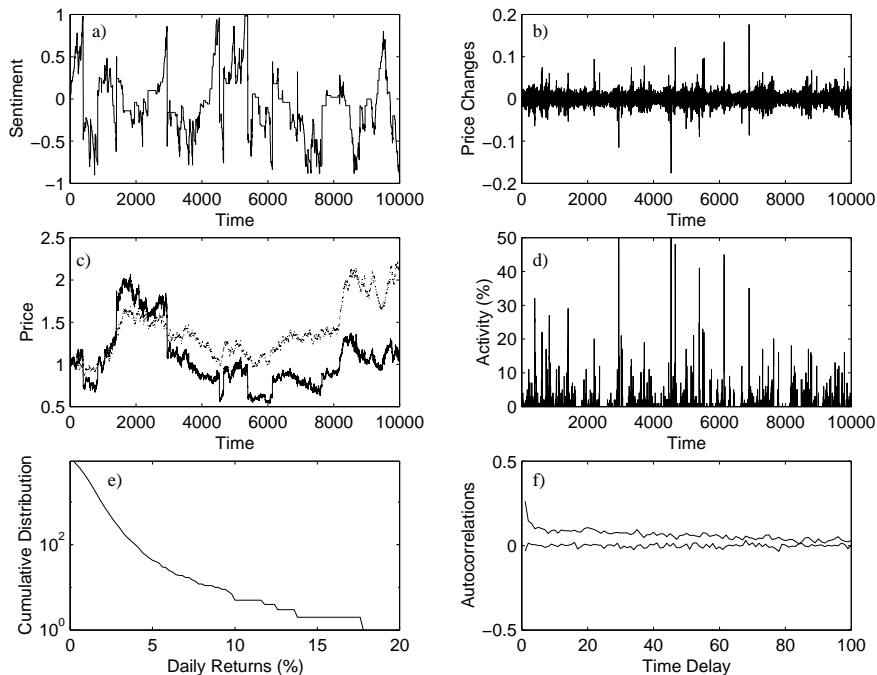


Figure 4: Numerical simulation incorporating parameter $\alpha = 2$.

it is this property that allows us to draw broad qualitative conclusions from the numerical simulations. However, if related models are to be used in a more quantitative manner, it seems likely that some of the modeling assumptions that were made will have to be replaced by experimental or empirical data. It is to be hoped that novel approaches to the heterogeneous agent modeling problem, such as the one taken in this paper, will in turn stimulate new research in both experimental economics and the analysis of market data.

Acknowledgments

H.L. gratefully acknowledges the hospitality and financial support of the Fields Institute for Research in the Mathematical Sciences, Toronto, Canada.

References

- [1] W.B. Arthur. Inductive reasoning and bounded rationality. *Am. Econ. Rev.*, 84:406–412, 1994.
- [2] C. Avery and P. Zemsky. Multidimensional uncertainty and herd behavior in financial markets. *Amer. Econ. Rev.*, 88:724–748, 1998.
- [3] N. Barberis, A. Shleifer, and R. Vishny. A model of investor sentiment. *Journal of financial economics*, 49:304–343, 1998.
- [4] F. Black and M. Scholes. The pricing of options and corporate liabilities. *J. Polit. Econ.*, 81:637–654, 1973.

- [5] S. Bornholdt. Expectation bubbles in a spin model of markets: intermittency from frustration across scales. *Int J. Mod. Phys. C*, 12:667–674, 2001.
- [6] D. Challet and Y.C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407–418, 1997.
- [7] D. Challet and Y.C. Zhang. On the minority game: analytical and numerical studies. *Physica A*, 256:514–532, 1998.
- [8] R. Cont and J. P. Bouchaud. Herd behavior and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*, 4:170–196, 2000.
- [9] 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*. forthcoming 2004. Can also be obtained at math.gmu.edu/~harbir/far0908.pdf.
- [10] R. Cross, A.M. Krasnosel’skii, and A. Pokrovskii. A time-dependent preisach model. *Physica B*, 306:206–210, 2001.
- [11] M.G. Daniels, J.D. Farmer, L. Gillemot, G. Ori, and E. Smith. Quantitative model of price diffusion and market friction based on trading as a mechanistic random process. *Phys. Rev. Lett.*, 90:108102, 2003.
- [12] R.H. Day and V.L. Smith. *Experiments in decision, organization and exchange*. North-Holland, Amsterdam, 1993.
- [13] E.F. Fama. Efficient capital markets: A review of theory and empirical work. *J. Finance*, 25:383–417, 1970.
- [14] E.F. Fama. Efficient capital markets. *J. Finance*, 46:1575–1618, 1991.
- [15] J.D. Farmer. Physicists attempt to scale the ivory towers of finance. *Comput. Sc., Eng*, 1:255–272, 1999.
- [16] 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.
- [17] A. Gaunersdorfer, C.H. Hommes, and F.O.O. Wagener. Bifurcation routes to volatility clustering under evolutionary learning. Technical report, University of Amsterdam, Centre for Nonlinear Dynamics in Economics and Finance, 2003.
- [18] I. Giardina and J.-P. Bouchaud. Bubbles, crashes and intermittency in agent based market models. *Eur. Phys. J. B*, 31:421–437, 2003.
- [19] I. Giardina and J.-P. Bouchaud. Volatility clustering in agent based market models. *Physica A*, 324:6–16, 2003.

- [20] 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.
- [21] J.Y.Campbell, A.W. Lo, and A.C.Mackinlay. *The Econometrics of Financial Markets*. Princeton UP, NJ, 1999.
- [22] J.H. Kagel and A.E. Roth (eds.). *The handbook of experimental economics*. Princeton University Press, 1995.
- [23] D. Kahneman and A. Tversky. Judgement under uncertainty: Heuristics and biases. *Science*, 185:263–291, 1974.
- [24] T. Kaizoji. Speculative bubbles and crashes in stock markets: an interacting-agent model of speculative activity. *Physica A*, 287:493–506, 2000.
- [25] T. Kaizoji, S. Bornholdt, and Y. Fujiwara. Dynamics of price and trading volume in a spin model of stock markets with heterogeneous agents. *Physica A*, 316:441–452, 2002.
- [26] A. Kirman. Ants, rationality and recruitment. *Quart. J. Econ.*, 108:137–156, 1993.
- [27] A. Krawiecki and J.A. Holyst. Stochastic resonance as a model for financial market crashes and bubbles. *Physica A*, 317:597–608, 2003.
- [28] B. LeBaron. Agent-based computational finance: suggested readings and early research. *J. Econ. Dyn. Control*, 24:679–702, 2000.
- [29] T. Lux and M. Marchesi. Scaling and criticality in a stochastic mult-agent model of a financial market. *Nature*, 397:493–500, 1999.
- [30] 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.
- [31] R. Mantegna and H. Stanley. *An introduction to econophysics*. CUP, 2000.
- [32] T. Odean. Are investors reluctant to realize their losses? *J. Finance*, 53:1775–1798, 1998.
- [33] V. Plerou, P. Gopikrishnan, X. Gabaix, and H. Eugene Stanley. Quantifying stock-price response to demand fluctuations. *Phys. Rev. E*, 66, 2002.
- [34] M. Raberto, S. Cincott, S. Focardi, and M. Marchesi. Traders’ long-run wealth in an artificial financial market. *Comp. Econ.*, 22:26–39, 2003.
- [35] G.W. Schwert. Why does stock market volatility change over time? *J. Finance*, 44:1115–1153, 1989.
- [36] G.W. Schwert and P.J. Seguin. Heteroskedasticity in stock returns. *J. Finance*, 45:1129–1155, 1990.

- [37] H.A. Simon. A behavioural model of rational choice. *Quart. J. Econ.*, 69:99–118, 1955.
- [38] D. Sornette. Stock market speculation: spontaneous symmetry breaking of economic valuation. *Physica A*, 284:355–375, 2000.
- [39] H.R. Stoll. *Microstructure: the Organization of Trading and Short-term Price Behavior*. Edward Elgar, 2002.
- [40] P. Veronesi. Stock market overreaction to bad news in good times: A rational expectations equilibrium model. *Rev. Financial Studies*, 12:975–1007, 1999.
- [41] F. Wagner. Volatility cluster and herding. *Physica A*, 322:607–619, 2003.
- [42] M. Youssefmir and B. A. Huberman. Clustered volatility in multiagent dynamics. *Journal of Economic Behavior and Organization*, 32:101–118, 1997.