Rational expectations, psychology and inductive learning via moving thresholds

H. Lamba

Department of Mathematical Sciences, George Mason University,

4400 University Drive, Fairfax, VA 22030 USA

T. Seaman

School of Computational Sciences, George Mason University,

4400 University Drive, Fairfax, VA 22030 USA

Abstract

This paper modifies a previously introduced class of heterogeneous agent models in a way that allows for the inclusion of different types of agent motivations and behaviours in a consistent manner. The agents operate within a highly simplified environment where they are only able to be long or short one unit of the asset. The price of the asset is influenced by both an external information stream and the demand of the agents. The current strategy of each agent is defined by a pair of moving thresholds straddling the current price. When the price crosses either of the thresholds for a particular agent, that agent switches position and a new pair of thresholds is generated.

The threshold dynamics can mimic different sources of investor motivation, running the gamut from purely rational information-processing, through rational (but often undesirable) behaviour induced by perverse incentives and moral hazards, to purely psychological effects. The simplest model of this kind precisely conforms to the Efficient Market Hypothesis (EMH) and this allows causal relationships to be established between actions at the agent level and violations of EMH price statistics at the global level. In particular, the effects of herding behaviour and perverse incentives will be examined.

PACS numbers: 89.65.Gh 89.75.Da

I. INTRODUCTION

The Efficient Market Hypothesis (EMH) [8] still has enormous influence in economic theory as well as in the day-to-day operations of financial markets. This is in spite of numerous statistical and experimental studies that invalidate both the underlying assumptions of the EMH [1, 7, 10] (eg. the rationality of economic participants) and the output of mathematical models that are derived from them (eg. stock prices that are described by a geometric Brownian motion resulting in a Gaussian distribution of the changes in the log-price returns) [3, 13]. The near-universal statistical properties of real markets, most of which deviate from those predicted by the EMH, are often referred to as the 'stylized facts'.

The inappropriate use of EMH models and Gaussian statistics by market participants can cause great damage, not just to the participants themselves but to entire economies. The extrapolation from the past into the future becomes less justifiable in an environment where the decreasing probability of large fluctuations follows a power-law rather than an exponential one (i.e. where 'fat-tails' are present) and significant coupling between market participants can develop.

To take a very recent example, at the time of writing a sudden rise in the delinquency rates of so-called sub-prime mortgages in the US has precipitated a credit-crunch in financial markets around the world. The growth of sub-prime lending was stimulated by a period of low global interest rates, inadequate accounting and credit-rating standards and the development of novel financial derivatives such as Collateralized Debt Obligations (CDOs) and Credit Default Swaps (CDSs).

However it is the proximal causes that are of more interest to us here. The herding phenomenon and the resulting 'madness of crowds' that lies at the heart of most, if not all, financial bubbles was present in the general public and their increasingly irrational belief that house prices would continue to rapidly climb. This caused a cycle of additional house-buying and home-equity withdrawal that fuelled prices further. However, on the lending side of the equation, there were significant perverse incentives and moral hazards at play. Realtors, property appraisers and mortgage brokers all received their commissions at the time of the transaction and had no incentive to question the quality of the transaction or lower its amount. The loan originators were similarly able to distance themselves from the default risk on the loans via the use of CDOs and CDSs. This continued up the financial food chain, helped by the fact that there was no open market for these derivatives and their prices could be easily manipulated via creative accounting practices. Finally, when the sub-prime housing market burst, the liquidity assumption inherent in efficient market models — that any desired transaction can be carried out at the current price — failed to hold as buyers of these derivatives vanished from the marketplace.

In short, it is hard to see how models based solely upon the notion of efficient markets can begin to adequately predict or quantify situations such as the one described above. The construction of financial models that can incorporate both rational and irrational agent behaviour, as well as the rational-but-perverse consequences of market defects, is an important undertaking. The simple, yet plausible, class of models introduced in this paper provides a possible framework within which the consequences of different EMH violations can be systematically studied.

The paper is organized as follows. Section II is a summary of previous results obtained for a similar class of models. The reader is directed towards these papers for more detailed arguments concerning the modeling assumptions and for further references. In Section III the proposed modification to these earlier models is described. Briefly, the earlier models used multiple fixed-thresholds for a single agent while the new approach uses a single moving threshold. This new approach appears to be both more elegant and more powerful. Section IV contains some numerical results for the new model showing that the essential qualitative and quantitative features of the earlier models are preserved (the results of a more detailed numerical investigation will be reported elsewhere). Finally in Section V the main conclusions are stated and possible directions for future work are outlined.

II. PREVIOUS RESULTS

In [4–6, 11] a class of models using thresholds was introduced. The models consist of two parts, one defining the price updates and the other determining when agents switch trading positions. The price changes are influenced by both external factors, namely a Gaussian uncorrelated information stream, and changes in the internal supply/demand due to agents changing positions.

The thresholds are involved in the modeling of the agents themselves. In the simplest case, whenever an agent switches position, a pair of *fixed* price thresholds is generated (from some predetermined distribution) that straddles the current price. Then, when the price crosses either of the thresholds the agent switches (taking either a profit or a loss, depending upon which threshold was crossed) and a new pair of thresholds is generated. If one interprets the price thresholds as representing the agent's rational future expectations then the model is, both practically and philosophically, an EMH model. Differing expectations cause trading to occur but the lack of coupling between agents ensures the 'correct' pricing outcome.

A propensity towards herding, either due to subconscious psychological pressures or conscious momentum-trading strategies, was then introduced via an additional threshold mechanism for each agent. When agents are in the minority their herding pressure increases by some amount until the new threshold is exceeded and they switch to join the majority (unless the majority position changes first, or they switch due to the price thresholds being violated). Numerical simulations then demonstrated that, at least within this modeling scenario, herding causes fat-tails (and excess kurtosis) in the price return data but not volatility clustering. However, long-term correlations in the volatility could be induced by additionally allowing the volatility to depend upon the market sentiment (and/or allowing correlations in the information stream). A detailed statistical analysis [11] showed a) the existence of power-law exponents (in both the distribution of the log-price-returns and the decay of the volatility autocorrelation) consistent with those measured from real market data, and b) additional stylized facts were also reproduced.

Other heterogeneous agent models (HAMs) have also successfully replicated the major stylized facts but the approach taken here has significant advantages. Firstly, it is conceptually very simple and has a small number of parameters. Secondly, it would appear to more accurately reflect the human decision-making process — a gradual accumulation of information and/or emotion resulting in a sudden action that is readily described by some type of threshold mechanism. This threshold approach should be contrasted with the more common one taken in the literature of Markovian switching between investment positions or strategies (see [12] amongst others). Here the threshold values act as 'hidden variables' that are history-dependent. Thirdly, the fact that the model can reproduce EMH statistics in certain parameter limits allows us to draw causal inferences regarding EMH-violations that are not possible in more complex HAMs that do not share this property.

However, the above fixed-threshold models can be significantly improved by allowing thresholds to continuously vary with time, rather than being fixed between transactions.

From a modeling perspective, it is far more reasonable to suppose that agents' thresholds are not frozen at the moment of the last trade but instead continuously vary as the market does. Once this change is made, then different motivations and strategies can be simulated by modifying the dynamics of *a single pair of thresholds* as described below. This modeling change is also very natural since the different effects are all relevant to the same investment decision and their contributions to the investment decision can now be considered as cumulative.

III. MOVING-THRESHOLD MODELS

The mathematical moving-threshold model is now described in full but more detailed economic justifications for the modeling assumptions can be found in [5, 6, 11]. The system is incremented in timesteps of length h and each of M agents can be either short or long the asset over each time interval. The position of the i^{th} investor over the n^{th} time interval is represented by $s_i(n) = \pm 1$ ($+1$ long, -1 short). The price of the asset at the end of the n^{th} time interval is $p(n)$ and for simplicity the system is drift-free so that $p(n)$ corresponds to the return relative to the risk-free interest rate plus equity-risk premium or the expected rate of return. An important variable is the *sentiment* defined as the average of the states of all of the M investors

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

and $\Delta\sigma(n) = \sigma(n) - \sigma(n-1)$.

The pricing formula is given by

$$p(n+1) = p(n) \exp\left(\left(\sqrt{h}\eta(n) - h/2\right) f(\sigma) + \kappa \Delta\sigma(n)\right) \tag{2}$$

where $\kappa > 0$ and $\sqrt{h}\eta(n) \sim \mathcal{N}(0, h)$ represents the exogenous information stream. The function f allows the effect of new information on the marketplace to vary with sentiment and is discussed further below. Note that when $\kappa = 0$ and $f \equiv 1$ the price follows a geometric Brownian motion.

Suppose that at time n the i^{th} investor has just switched and the current price is P. Then a pair of numbers $X_L, X_U > 0$ are generated from some specified distribution and the lower and upper thresholds for that agent are set to be $L_i(n) = P/(1+X_L)$ and $U_i(n) = (1+X_U)P$ respectively. Defining the evolution of these thresholds now corresponds to defining a *strategy* for the i^{th} investor. Such strategies can in principle be made arbitrarily complicated and may be partly rational (and conscious) and partly irrational (and subconscious). They can also be constructed to include perverse incentives or inductive learning strategies as desired. It should be understood that there is no requirement for an agent to be aware of the values of their thresholds (or even their existence). This would certainly be the case for an investor driven primarily by psychological effects.

The price interval $[P/(1+X_L),(1+X_U)P]$ defined by the thresholds has multiple interpretations. For example, it incorporates the presence of transaction costs by ensuring that agents do not switch arbitrarily often (if the L_i and U_i are bounded away from 0). In addition it replicates the psychological aspect of human behaviour known as ‘anchoring’ [1, 9] whereby experiments have shown that subjective evaluations of a reasonable price for an asset are often based around the most recent transaction price P. Other seemingly-universal effects such as the predilection for investors to realize profits too quickly and hold onto losing positions for too long can easily be incorporated by modifying the initial distributions of L_i and U_i according to the asset position being taken.

Three simple examples are of particular interest. Firstly, as noted above, the case where the thresholds are fixed (until the agent switches again) completely decouples the agents' behaviour and gives EMH pricing. Secondly, we can mimic the herding effect by causing the thresholds to move together (increasing L_i and decreasing U_i) whenever the i^{th} agent is in the minority. Thus agents in the minority have a higher tendency to switch into the majority than vice versa. Thirdly, we can suppose that a simple, unspecified, perverse incentive is in place, causing agents to prefer one of the states over the other, say $+1$ over -1. This can be recreated by moving the thresholds closer together whenever $s_i(n) = -1$.

IV. NUMERICAL RESULTS

It is not our intention to compare the model against specific data sets. Rather we wish to compare the output of the model against the stylized facts that seem to be consistent across the vast majority of markets and asset classes. The parameter values chosen below are identical to those used in the earlier fixed-threshold models. This is not because the conclusions are sensitive to the values. In fact the opposite is true and all the results reported are robust even in the presence of large parameter changes.

We first select parameters for the model to simulate daily price variations. The time variable h is defined in terms of the variance of the external information stream. A daily variance in price returns of 0.6–0.7% implies that h of 0.00004 should correspond to approximately 1 trading day. The system properties are independent of the number of agents, for large enough M, and $M = 100$ appears to be sufficient. The simulations are run for 10000 timesteps which corresponds to approximately 40 years of trading.

The parameter κ is a measure of the market depth and the relative importance of external noise versus changes in internal supply/demand on the asset price. It is more difficult to estimate a priori than the other parameters, but $\kappa = 0.2$ generates price output that is correlated with, but noticeably distinct from, the EMH price defined by $\eta(n)$ alone.

The initial thresholds X_L, X_U are chosen from the uniform distribution on the interval $[0.1, 0.3]$, corresponding to price moves in the range 10–30%. Simulations have indicated that the models are robust to changes in the exact form of the distribution and so a simple one was chosen. Let us consider first the case where the thresholds are fixed. If $f(\sigma) \equiv 1$ and the initial states are sufficiently mixed so that $\sigma(0) \approx 0$, then as explained in Section III, the model conforms to the EMH and the price is simply

$$p(n+1) \approx p(n) \exp\left(\sqrt{h}\eta(n) - h/2\right). \tag{3}$$

It was argued in [2, 5, 11] that the effect of noise traders (who are not explicitly included since they operate on very short time-scales) is unlikely to be constant over time. We posit that their number, and therefore also their effect, is greater at times of high sentiment, both positive and negative. This is simulated by increasing the price-changing effect of the external information stream at such times by the use of the function $f(\sigma) = 1 + 2|\sigma|$ in (2).

Now let us introduce herding by supposing that for agents in the minority position

$$L_i(n+1) = L_i(n) + C_i h|\sigma(n)|, \quad U_i(n+1) = U_i(n) - C_i h|\sigma(n)|.$$

The thresholds are fixed for agents in the majority position. Note that the change in the position of the thresholds is proportional to the length of the timestep and the magnitude of the sentiment. The constant of proportionality C_i is different, but fixed, for each agent and chosen from the uniform distribution on $[0, 100]$. This range of parameters corresponds to a herding tendency that operates over a timescale of several months or longer. The results of a simulation are shown in Figure 1 and are very similar to those obtained in [5, 11] where multiple, fixed thresholds were used.

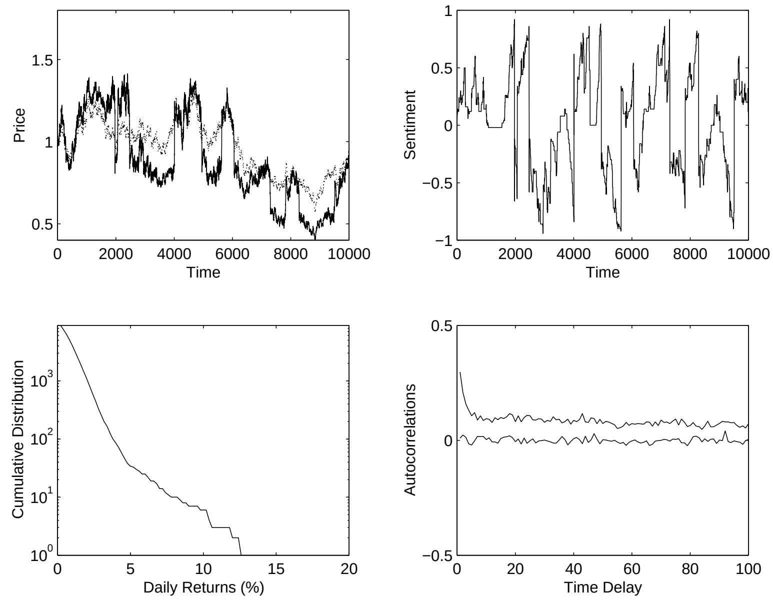

FIG. 1: Results of a simulation over 10000 timesteps including herding.

The top left plot shows both the output price $p(n)$ (more volatile) and the EMH pricing obtained from (3) (less volatile). As can be seen there are significant periods of price mismatching. The top right figure plots the sentiment against time. Periods of bullish and bearish sentiments lasting several years can be observed. The bottom left picture plots the number of days on which the absolute value of the relative price change exceeded a given percentage. It also shows clear evidence of a fat-tailed distribution in the price returns and this is also confirmed by measures of the excess kurtosis which range from approximately 10–30. Finally, the two curves in the bottom right figure are the autocorrelation decays for both the price returns and their absolute values (the volatility). There is no correlation observable in the price returns, even for lags of a single day, while the volatility autocorrelation decays slowly over several months — evidence of volatility clustering. Further numerical testing revealed very similar conclusions to those drawn from the previous fixed-threshold models. These were that, in these models, herding causes fat tails but not volatility clustering since the decay in the volatility autocorrelation vanishes after just a few days when $f(\sigma) = 1$. Also, measurements of power-law exponents similar to those carried out in [11] provided estimates close to those observed in analyses of price data from real markets. These will be reported in detail separately.

We now suppose that, in addition to the herding and sentiment-dependent volatility, a

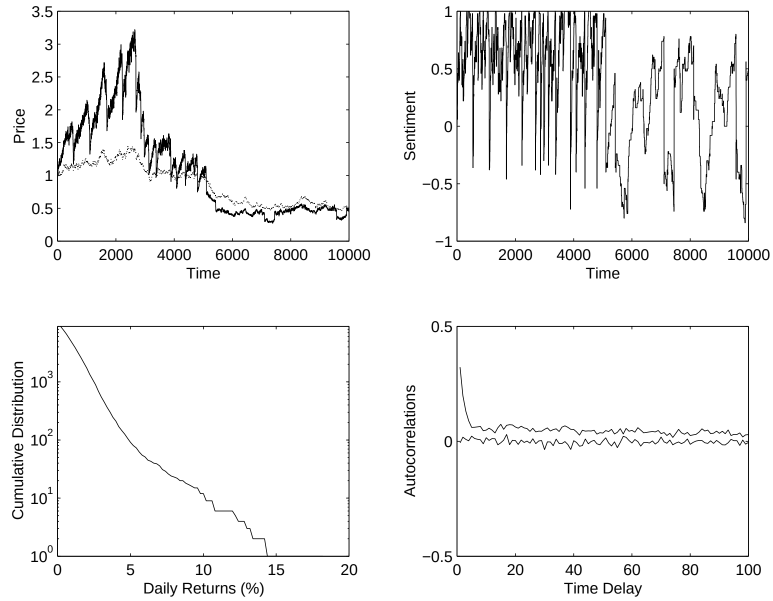

FIG. 2: Results of a simulation over 10000 timesteps including both herding and a perverse incentive. The incentive is removed after 5000 timesteps.

perverse incentive is in place that makes it advantageous to be in the state +1. This is achieved by including the following extra rule when an agent is in the state $s_i(n) = -1$,

$$L_i(n+1) = L_i(n) + Rh, \quad U_i(n+1) = U_i(n) - Rh.$$

The value $R = 100$ was chosen to make the incentive the same order of magnitude as the herding propensity. The incentive is then switched off after 5000 timesteps. Results for a typical simulation are shown in Figure 2.

The effect of the incentive can be seen in both the price and sentiment variables, where both the average value and the volatility of the variables are increased. This simulation merely serves as a prototype for the kinds of studies that can be performed and more systematic investigations are clearly required in order to draw reliable inferences.

V. CONCLUSIONS

In this paper an agent's *strategy* (in the loosest sense of the word) is defined to be a combination of all the rational, inductive learning, psychological and rational-but-perverse factors influencing their decision to switch positions or stay put. We have shown that such strategies can be represented by moving thresholds and can be as simple or as complicated as desired. Indeed the most serious modeling restrictions in the caricature systems simulated above are not in the assumptions behind the moving threshold approach, but in the over-simplified nature of the marketplace itself.

The results presented here, in conjunction with the previous work using multiple fixed thresholds, demonstrate a direct causal relationship between the presence of herding at the agent level and the emergence of fat-tailed power-law returns at the global level. Furthermore, this phenomenon appears to be extremely robust — not just with respect to large changes in the system parameters but also to significant changes in the model itself. The fact that herding by itself does not induce volatility clustering in these models (the autocorrelations die away within a few days) is another very interesting observation that may help clarify the behaviour of real markets.

More elaborate, and realistic, threshold models can be derived by either complicating the trading environment (for example, giving agents the option to exit the market or a choice of different assets) or by using more sophisticated strategies (incorporating inductive learning, say). However, even the simplest moving threshold models can help to shed light upon the EMH assumptions and the consequences of their absence. They may also prove to be more amenable to rigorous mathematical analysis than other HAMs that have been developed thus far.

---

[1] N. Barberis and R. Thaler. A survey of behavioral finance. In G.M. Constantinidos, M. Harris, and R. Stultz, editors, *Handbook of Economics and Finance*, pages 1053–1123. Elsevier Science, 2003.

[2] G.W. Brown. Volatility, sentiment and noise traders. *Financial Analysts Journal*, pages 82–90, 1999.

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

[4] R. Cross, M. Grinfeld, and H. Lamba. A mean-field model of investor behaviour. *J. Phys. Conf. Ser.*, 55:55–62, 2006.

[5] R. Cross, M. Grinfeld, H. Lamba, and T. Seaman. A threshold model of investor psychology. *Phys. A*, 354:463–478, 2005.

[6] R. Cross, M. Grinfeld, H. Lamba, and T. Seaman. Stylized facts from a threshold-based heterogeneous agent model. *Eur. J. Phys. B*, 57:213–218, 2007.

[7] R.H. Day and V.L. Smith. *Experiments in decision, organization and exchange.* North-Holland, Amsterdam, 1993.

[8] E.F. Fama. Efficient capital markets: A review of theory and empirical work. *J. Finance*, 25:383–417, 1970.

[9] J.H. Kagel and A.E. Roth (eds.). *The handbook of experimental economics.* Princeton University Press, 1995.

[10] D. Kahneman and A. Tversky. Judgement under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.

[11] H. Lamba and T. Seaman. Market statistics of a psychology-based heterogeneous agent model. Preprint, Econophysics forum.

[12] 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.

[13] R. Mantegna and H. Stanley. *An Introduction to Econophysics.* CUP, 2000.