The behavioural foundations of extreme market movements

Johan du Plessis (10235086)

9th January 2012

1 Introduction

It is by now common knowledge that market movements do not always correspond to the random walk model imposed in standard economic theory. Market returns exhibit 'fat tails', that is price movements far larger than would be expected under the standard model, and other 'anomalies' (Lux and Marchesi, 2000; Lux, 2008). Of interest is not the technical definition of fat tails as corresponding to distributions that exhibit high kurtosis, but rather the (multitude of) ways in which observed price movements deviate from the standard model, in particular where these moves are more extreme.

Non-normal return distributions can result from the underlying fundamentals of the economy - that is the news that affects investors' beliefs about the value of stocks (Lux and Marchesi, 1999). For example, if factories burn down with some fat-tailed distribution, this will result in similar price movements for a stock relying on these factories. However, evidence of excess volatility of stock prices as compared to fundamental values and overreaction (and underreaction) to news suggest that something else may be at work (Lux, 2008). This coupled with abundant evidence on non-rationality in human decision making suggest that it is worth examining how investor behaviour, including irrational and boundedly rational behaviour, can result in extreme price movements.

There is a growing body of literature that attempts to do just this using the concepts of dynamical systems, which have proven effective in describing phenomena in the natural sciences. These consider the non-linear interaction between a number of agents and the overall properties of the system that can emerge from this (Lux, 2008). Alfarano *et al.* (2005) call the notion that such interactions are the cause of observed stylised facts the *interacting agents hypothesis* (in opposition to the *efficient markets hypothesis*).

The fact that market returns do not seem to be normal (at least when measured at frequencies higher than monthly) is somewhat surprising, given the consequences of the central limit theorem. Even daily returns can be seen as the aggregation of many thousands of intra-daily returns (except in illiquid markets) (Lux, 2008). To explain this one needs one may do away with one or more of the core assumptions that the central limit theorem needs. One is that of finite variance, which naturally leads to the consideration of stable distributions and the generalized central limit theorem. However, these distributions appear to be *too* fat tailed (Lux, 2008). The other assumption is independence of returns. Agent-based models tackle primarily this last assumption by creating dependence between agents.

As the body of available literature is vast, in this essay I will consider primarily a particular class of models developed by Thomas Lux and his collaborators. The models I will examine will rely very much on imitation (that is, herding behaviour) of investors as well as the contrast between two types of investors, usually identified as chartists and fundamentalists. Relatively simple rules followed by each class will be seen to lead to surprisingly rich behaviour of the system as a whole.

The structure of the essay is as follows. In section 2 I will briefly mention some of the stylised facts that appear to be present in almost all financial markets. Section 3 will give an overview of some of the most basic elements that are used in agents based models, including those I will discuss. In section 4 I give a brief overview of the models examined by Lux and his collaborators and the most important results. This is followed by a detailed discussion of the models presented and some suggestions for further research in section 5. The final section concludes.

2 Empirical regularities

From empirical financial data, researchers have found a number of characteristics that seem to be pervasive. I shall briefly mention them. For details please see the relevant literature.

- Standard statistical tests are unable to reject the hypothesis that financial prices follow a random walk. This, at least superficially, confirms the efficient markets hypothesis (EMH) (Lux and Marchesi, 2000). It lends to credence to the observation that most of the time the markets are well behaved. There are, however, a number of statistical discrepancies which show that markets do not in fact follow a random walk.

- The overall distribution of returns exhibit fat tails. In particular, the kurtosis of the distribution is significantly higher than would be expected under the normal distribution, leading to larger price movements than would be expected (Lux and Marchesi, 2000). It appears these price movements have tails that follow a power-law distribution (Lux, 2008). An interesting feature is that this property disappears the larger the time intervals over which returns are measured (for monthly, yearly, etc. returns the distributions are approximately normal).

- Volatility clustering is also observed. The market has periods of relative calm interspersed with periods of highly volatile prices. This is reflected in the autocorrelation of absolute returns and squared returns being significant even with long lags. They also seem to decay with a power law. This effect is particularly pronounced for absolute returns (Lux, 2008).

- Sornette (2003) report another indication of dependence, but in the distribution of drawdowns. Drawdowns are successive periods in which the price moves in the same direction. The distribution of drawdowns is found to be much as expected for small drawdowns, that is price movements appear to be independent most of the time. However, there are more large drawdowns than predicted by the standard model, indicating periods of dependence.

- There is also literature asserting excess volatility of stock prices as compared to fundamental values as well as over- and underreaction to news events (Lux, 2008). Only this last item is strictly incompatible with the EMH. However, it is also the hardest to verify as we do not have access to the news process that feeds into fundamental prices.

3 The main ingredients

3.1 A dynamical systems approach

Lux and Marchesi (2000) suppose that linear models may be insufficient to describe all the above empirical properties of markets. This naturally leads to a consideration of dynamical systems. This considers how the non-linear interaction of individuals based on a set of relatively simple rules can lead to emergent properties which mimic the empirical phenomena we have discussed. These systems may have infrequent but sudden transitions, corresponding to critical points (or singularities) where the system makes a transition from disorder to order. Such properties may explain the relative calmness of markets most of the time, interspersed with periods of volatility. Furthermore dynamical systems naturally lead to power-laws, for which there is considerable empirical evidence in market data (Sornette, 2003; Lux, 2008).

These models are often hard to manage analytically (even despite drastic simplifications). A common solution is to use approximations to derive some analytical results (for instance regarding stability of equilibria), in particular variables are replaced with their mean value to obtain a deterministic system. Simulation is then used to empirically validate the properties of the system(Lux, 1998; Sornette, 2003). Statistical tests can be run on the data to determine if it has the same qualitative features as what is observed in real markets.

3.2 The main players

It is standard in economic literature to examine two types of traders, often identified with the chartists (technical analysts) and fundamental traders seen in real markets. The interaction between these groups can create a rich and surprising set of dynamics (Sornette, 2003).

3.2.1 Noise traders

These are called 'noise traders' as they are meant to be irrational (Lux, 1995). However, in the models we examine, their irrationality could be disputed. They are often identified as chartists, who in real life look at price patterns in order to determine the trend in which the market is moving. They are trend-followers, buying when the the trend is up (and expected to continue) and selling otherwise.

We may consider chartists as being optimistic (tending to buy) or pessimistic (tending to sell). This may depend on the opinions of other traders (imitation), on observed price behaviour (trend following), and on some measure of expected returns of buying vs selling (Lux (1995) and the papers that follow include all three). All of these reflect an aspect of real chartist behaviour and give some tendency to follow the crowd.

3.2.2 Fundamentalists

This group is usually meant to be more rational. They may be modelled as having rational expectations (Sornette (2003) describes models in which this is done). However, we shall consider models where their behaviour is modelled more explicitly.

Fundamentalists believe the market will revert to its fundamental value, often defined as the expected present value of real dividends (Lux and Marchesi, 1999). This is the premise of fundamental traders in real life and they attempt to determine what this fundamental price might be. The simplest way to model this behaviour is merely to assume that the fundamentalists do know the fundamental price (as in (Lux, 1995) and the papers that follow) - in practice it is of course subject to much uncertainty.

Fundamental traders are contrarian. They will buy when the price is below fundamental value and sell when it is above. Unlike noise traders they are supposed to have a stabilising effect on the market price.

3.3 Imitation dynamics

Imitation plays an important role in the models we will consider. Often this process is called 'contagion','infection' or 'herding', which perhaps captures the main idea more emotively. Opinions can pass from trader to trader like a virus. It is an attempt to capture in a simple way trader psychology (Lux, 1995). There is evidence that such psychological factors are at play in the minds of real traders (Sornette, 2003). We will be particularly interested in how optimism or pessimism may spread among the traders. Imitation is self-reinforcing in the sense that the more traders there are in a certain group, the more likely it is for other traders to imitate them (Lux and Marchesi, 2000; Sornette, 2003).

This may be modelled in several (not necessarily mutually exclusive) ways. Traders may consider the opinions of their peers in determining a strategy to follow. In this case each trader has access only to a subset of the whole population (Sornette, 2003). One may also suppose that traders can probe the general mood of the market, i.e. the opinions of all other traders (or of all other traders of the same type), as is done in the models of Lux (1995) and related papers which we will examine. These papers essentially use a mean-field approximation where the effect of all other agents on a single agent is replaced by an average effect.One may also consider that traders imitate each others' strategies, for instance switching from fundamentalist to chartist or from optimistic to pessimistic, based on the realised (or expected) profits of these strategies. Such a situation is encountered in Lux and Marchesi (2000) and Lux (1998).

Additionally, there may be a force that counteracts imitation, namely some private signal received by each trader which may (or may not) contradict the market opinion. This reflects, for example, that traders may examine different data or reach different conclusions given the same data. This is can be modelled explicitly (Sornette, 2003). However, in the models of Lux (1995) and those that follow, this is implicit by considering a probabilistic approach to switching between groups. Switching is more likely to occur to the majority opinion, but there is still a small probability of switching the other way.

As long as there are sizeable numbers of traders in every group, we may say the market is in a state of disorder. Participants' actions tend to counteract each other and the market is stable. When any one group dominates order is created in the market - many traders agree and take the same action. Their actions reinforce each other and in the case of a dominance of optimistic (or pessimistic) noise traders, exaggerate market movements away from the fundamental value (Sornette, 2003).

Whether imitation should be considered irrational behaviour is not entirely clear and depends on the specific situation being modelled as well as the information assumed to be available to traders. For traders who do not have access to other information (or where such information is unreliable), following the opinions of others may well be very rational (Lux, 1995). Similarly where it is important for traders to maintain their reputation, that is for them not to underperform compared to peers, it makes sense to follow the opinion of the crowd as this minimises the risk of underperformance (Sornette, 2003). In any event such imitation qualifies as boundedly rational and seems to reflect something of real human behaviour.

3.4 Price dynamics

One needs a mechanism by which the price responds to the decisions of the different actors in the model. One may assume that the market price reacts instantaneously, equating demand and supply (for instance Lux (1995) do so in order to simplify their final model). However, one may also wish to consider a slower and possibly noisy adjustment of price toward equilibrium, which would more closely reflect the functioning of actual markets. The approach often taken (sometimes only implicitly) is to assume the existence of a market maker who will accept any imbalance of demand and supply. The market maker will adjust the price in the next period (or possibly the current period) according to some, possibly noisy, decision rule. That is it will increase the price if there is excess demand and lower it if there is excess supply.

3.5 Order decisions

Excess demand/supply for each group of traders can be formulated in different ways. A distinction needs to be made between position-based strategies and order-based strategies. In the former an agent chooses their desired net position in the stock based on current information and then places an order for the difference between their current position and the desired position. In order-based strategies, the agent decides on the quantity to order independently of their current holdings (Farmer, 2002). The models considered here use order-based strategies (and ignore holdings), the appropriateness of which is discussed in section 5 as the results can be very different.

3.6 Synchronicity

A distinction also needs to be made between synchronous and asynchronous decision-making. In the former all agents make decisions (such as choosing strategies) at specified time-steps; decisions are made simultaneously. In the latter case the time-step is chosen small enough (in the limit this gives a continuous-time model) such that only one agent can make a decision at a time and so decisions are not simultaneous (Lux and Marchesi, 2000; Farmer, 2002). The models I will consider use asynchronous decisions, which may again have different results from the other case (Lux and Marchesi, 2000).

4 The models and results in brief

Lux (1995) describes a model containing all the basic elements already introduced, which he then simulates and extends in Lux (1998), Lux and Marchesi (1999) and Lux and Marchesi (2000). There is a population of traders consisting of chartists and fundamentalists. Chartists may be optimistic or pessimistic. With some endogenous probability optimists may become pessimists and vice versa. This probability depends on the overall sentiment of the market (the proportion of optimistic versus pessimistic chartists), the price trend, and a measure of the instantaneous returns of buying/selling. The model further allows for imitation between investor types, that is chartists may become fundamentalists and vice versa based on a myopic measure of profits experienced. It is supposed that the fundamental value of the stock follows a standard random walk in the most complicated model (in the others it is held constant).

An interesting result from Lux (1995) and Lux (1998) is that deterministic but chaotic dynamics can generate seeming fat tails. The author constructed a deterministic model to approximate the stochastic model (replacing variables with their means) and simulated the deterministic dynamics. Statistical tests on the data also showed leptokurtosis which decreased with lower frequencies. The dynamics show waves of optimism and pessimism (that is, bubbles). The author conjectures that the following dynamics are at work. Positive price changes will tend to convert agents into optimistic chartists. This reinforces the rising trend. However, at some point the optimistic chartists dominate and the price increase slows down (as demand no longer accelerates). The force to switch to the optimistic strategy reduces. The price will now be far from its fundamental value, which creates large potential profits for fundamentalists. This results in more traders converting to the fundamental strategy, who sell the stock. This will eventually result in the price falling and create a wave of pessimism. There is thus a kind of symmetry in these models.

This deterministic approach should be contrasted with the approach from Lux and Marchesi (2000) which examines stable equilibria where the price equals the fundamental price on average. There is some critical percentage of chartists below which the system is stable and above which it becomes unstable. If the system becomes unstable due to the critical point being reached it leads to large fluctuations as the price is driven away from fundamental value, before the system is (endogenously) driven back to the stable region. This approach leads to periods of calm with transitory volatile behaviour. The fluctuations in the percentage of chartists are caused stochastically (even in a stable market we may consider some fluctuations due to the personal situations of traders). When the number of chartists increases toward the critical level, the volatility increases. This is only temporary as the price is driven back toward fundamental levels by the action of fundamentalists who see trading opportunities.

They find that in the long term, the price does equal the fundamental value on average, so fundamentalists have a stabilising effect. Price distributions appear to follow a power-law (which has high kurtosis - theoretically infinite in fact) and kurtosis reduces the larger the time-periods considered as in empirical data. There is a strong dependence between the magnitude of successive price changes and volatility seems to cluster (periods of calm interspersed with periods of high volatility). Furthermore the standard tests cannot reject the random walk hypothesis. It must be noted that these features occur despite that the simulated data, based on a discrete approximation of the model, will be stationary and have all moments finite. They thus ‘fool’ the statistical tests (Lux and Marchesi, 2000). These features are illustrated in figure 1.

The above models have two main drawbacks. The first is the fact that with increasing number of agents the distributions found converge to a normal distribution. The problem lies mainly in the mean-field approach using the opinion index, whose fluctuation becomes negligible (Lux, 2008). The second drawback is their analytical intractability. Alfarano *et al.* (2008) attempt to overcome both problems in one go. They develop a model with local (as opposed to global) interactions between chartists. Agents are assumed to meet randomly in pairs and one of the agents may convert to the mood (optimistic or pessimistic) of the other. The model is somewhat simplified as compared to the model of Lux and Marchesi (2000) as it does not incorporate explicit trend following, a compar-

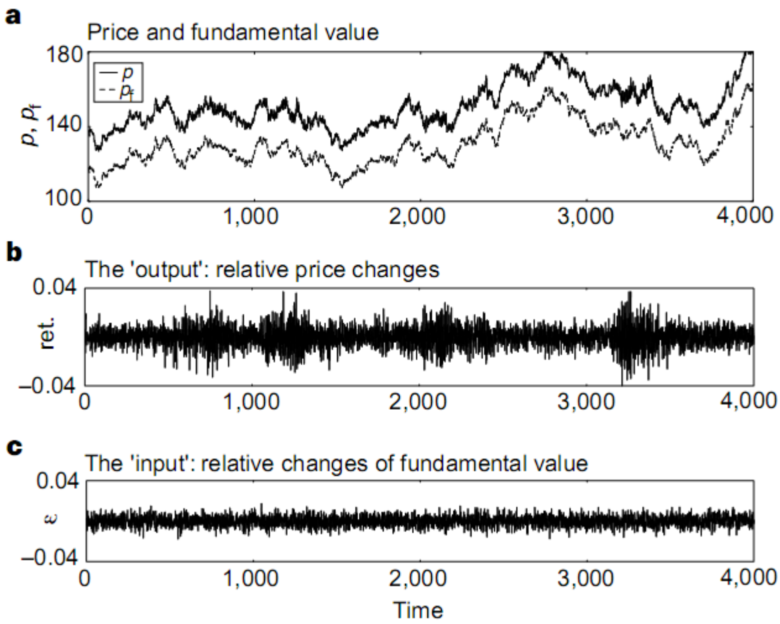

Figure 1: *This figure (taken from Lux and Marchesi (1999)) illustrates visually (for one simulation) the properties of the price dynamics discussed. Panel (a) shows the evolution of the market price and the fundamental value (shifted down for visual clarity). It is clear that the former tracks the latter. Panel (c) shows the the evolution of changes in fundamental value and (b) relative price changes. The former is a white noise process, whereas the latter shows greater volatility and also volatility clusters.*

ison between profits or the possibility of switching between chartist and fundamentalist strategies. Prices are also assumed to adjust instantaneously (i.e. no market maker).

The authors find distributional properties for the new model (but again under an approximation to the system) similar to those already discussed, namely leptokurtosis and clustered volatility. More interestingly they are able to show that the distributions do not produce *actual* power law behaviour but are able to produce data that are hard to distinguish from power-laws using statistical tests. The non-normality of distributions persist if the frequency of pairwise meetings increases linearly with the number of agents.

5 Discussion and suggestions for further research

Perhaps the biggest drawback of these models is that interesting dynamics are essentially a finite size effect (Lux, 2008). As the number of agents increase the distributions converge to normality. Lux (2008) show clear disappointment at this result. Ironically, a staunch rational agent economist may have viewed this as confirmation of the standard model. That many real markets have both a large number of agents and still display the stylized facts suggests that the number of *effective* agents are lower due to strong dependencies between them (a concept which certainly needs a more rigorous definition) (Lux, 2008).

Lux (2008) mention an uneasiness with increasing the frequency of meetings in the Alfarano *et al.* (2008) model, which would overcome the finite size problem. Admittedly it seems like a forced means of obtaining a desired result and one may argue that real people would not interact with significantly more other people in larger markets. However, Alfarano *et al.* (2008) show that this need not be devoid of intuition. For instance, for particles in a fixed area/volume frequency of interaction would increase with the number of particles. This could model increased group pressure or interaction via mass media. However, it might well be that these are better modelled (at least more intuitively) by Lux (1995)'s opinion index, as they may reflect the average population.

Lux (2008) mention research into network structures where agents are only influenced by local peers, which may also do the job. However, the results achieved are not yet satisfactory. I conjecture that another (or complementary) means of doing this would be to consider the size of agents. The models considered all assumed that agents have an equal effect on each other and that they buy in the same volume. Larger agents may have a disproportionate influence in the market in terms of both trade volume and opinion formation. They may be strongly connected to each other. Not only are there few of them, but they may in fact act as (approximately) only one agent. One may consider for instance pension funds, some of whom are very large, and who do seem to act in a somewhat concerted fashion (Sornette, 2003). They may not be influenced very much by the numerous small agents - influence in the other direction may, however, be strong. Sornette (2003) describe a hierarchical network model in which some agents are more highly connected than others. The least connected agents may be used to represent large traders (however, they are only sparsely connected to each other). It may be more efficient to construct an explicit structure that makes the distinction more apparent.

The realism of order-based strategies can be questioned. With such a strategy a value investor who sees a stock is mispriced will continue increasing his position in the stock as long as the mispricing persists. This can lead to unbounded positions before the price again reaches fundamental value. With investors in real life having limited capital and probably having limits on how much they are willing to bet on a trade, this seems like an unrealistic assumption (Farmer, 2002). Farmer (2002) show that analogous position-based strategies may not result in prices being driven to fundamental value. They also discuss a class of position-based strategies that *do* have this property.

It is at least somewhat unsatisfactory that despite having a theoretical model, it is necessary to resort to fitting simulated data to determine whether a power-law distribution (or other effects) is at work. Analytical results are available for simplified models. However, even the more complicated models are already much oversimplified and it appears that making them more realistic will require non-trivial additional complexity. In a sense, mathematical theory is lagging behind the needs of economic modelling (which is possibly a novel situation). This may motivate effort on expanding the mathematical toolkit for analysing dynamical systems.

All the models assert the existence of a fundamental price. That investors should all agree on what the fundamental price is, is questionable and thus noisy perceptions of the fundamental value deserve some attention. Farmer (2002) considers such a case. They find that for linear demand functions it is sufficient to consider a representative agent with average demand. This justifies the approach taken in the models discussed here. The very existence of such a thing as a fundamental price could, of course, be disputed and it may be useful to consider a model of interaction which does not make reference to it. It may be sufficient that there are some investors who believe that such a price exists and attempt to find it based on some information (indeed, the Farmer (2002) model considers such an interpretation).

In recent times the advent of algorithmic trading has resulted in the majority of trading volume on large exchanges being driven by algorithmic traders (Rogow, 2009). Many believe their existence has increased the fragility of the markets (Goldfarb, 2010).

An agent-based approach seems an appropriate framework for studying this effect. Some algorithmic trades (e.g. such as for portfolio insurance) have a definite destabilising effect (Sornette, 2003). Algorithmic traders may thus act as very fast-moving chartists. It is not, however, clear if algorithms can (or should) be modelled with the same behavioural rules as human traders. The *flash crash* and similar events that have been observed in several markets are thought to have been caused by algorithmic traders (Goldfarb, 2010; Cui and Lauricella., 2011). Sornette (2003) specifically analyses bubble and crash dynamics. However, it is not clear whether these flash events would conform to this analysis and an extension would be interesting.

The idea of a passive market maker is a simplifying assumption, which though it is useful, is not very realistic. Market makers may in fact have a non-trivial role to play in market dynamics. Lux (2008) (in a footnote) mention a situation in which it is in fact optimal for the market maker to move the market toward a chaotic state. That market makers may in fact reduce liquidity offered in troubled times may exacerbate market instability. Many market makers are now also algorithmic (Rogow, 2009). This means this analysis could overlap with the previous suggestion. Farmer (2002) mentions (but does not model) some attributes that more realistic market maker models could include.

It is hoped that these and similar models could serve as laboratories for testing the impacts of different market regulation (Lux, 2008). It is somewhat ironic, then, that one of the motivations for their use, is that they can explain regularities that appear to be universal and independent of market structure (Lux, 2008). Nevertheless, they may be able to identify just what kinds of regulations are important for market efficiency.

Not all the stylized facts of financial markets are explained by these models. For instance under- and overreaction to news is not modelled. However, I agree with Lux (2008) that this is not a 'deep' anomaly as the detection of such reactions is any case subjective. However, as these are not compatible with the EMH it would serve as ammunition against it should agent-based models produce these phenomena. The distribution of drawdowns was also not examined in the models and so it is not known whether these are similar to those in real markets. However, it is likely that the distribution of drawdowns gives a better (or at least complementary) impression of the dependence inherent in the data and is worth examining (Sornette, 2003).

One important point that must be made is that these models do not (and cannot) prove that the efficient markets hypothesis (EMH) is incorrect. In fact, except for the last point, none of the stylized facts mentioned contradict the EMH (they do however contradict the far stronger notion that markets follow a random walk) (Lux, 2008). Lux's argument that it is unlikely that the news processes in all markets follow a similar distribution is not convincing (Lux, 2008). Human behaviour, at least, is partly observable (maybe even quantifiable) whereas the supposed news process driving fundamental prices is far more ethereal. If the source of market turmoil is primarily from trader behaviour, then we may hope to curtail it by appropriate policies and education. However, very little can be done if it follows from fundamental processes.

The decision rules adopted by the models considered (and most other papers) are somewhat ad-hoc. It would be nice (but not strictly necessary) to see a justification for particular rules from some form of (possibly myopic) utility maximisation based on demonstrated psychological biases (Lux, 2008). This may also help in identifying more realistic decision rules and possibly refute some common rules, such as the order-based rules examined here.

The main use of the agent-based models discussed is qualitative rather than quantitative. Alfarano *et al.* (2005) assert that the complexity of the models and the lack of analytical solutions make it hard to obtain empirical estimates. Lux and Marchesi (2000) confine themselves to choosing parameters that give numerical results that are within the range observed in markets. This does not, however, allow us to infer much about the (real-life) values of parameters supposed to govern agent behaviour. Alfarano *et al.* (2005) introduce a much simplified model (along the lines of Alfarano *et al.* (2008)) with analytical results which does allow them to perform estimation. In any case we can only observe the *effects* of agent interaction in the form of market variables. Given the simplicity of the models (meaning there is most certainly a model specification error) and the probably very noisy observations we would be wise to distrust the parameter estimates obtained. How useful these numerical values would be is any case unclear, except that they may be able to tell us something about the degree of importance of for instance contagion of opinion versus, say, trend-following. In such cases one must be aware that different effects can be confounded. For instance contagion of opinion is most likely to show up when the trend is rising or falling and thus trend-following is also present.

6 Conclusion

Agent based models based on dynamical systems theory offer a means of explaining how the interactions of agents (traders) translate into emergent market behaviour. The field is hampered by the complexity of the systems needed to describe realistic behaviour and the resistance of these systems to the derivation of analytical results.

However, the ability of the systems to generate data that resemble very closely what is observed in real markets using simple but plausible behavioural rules, is encouraging. The particular models briefly examined in this essay give insight into the importance of human psychology in the form of imitation. They also show how the interaction of chartists and fundamental traders can result in waves of optimism or pessimism with the price diverging from the fundamental value, before being driven back by the actions of fundamental traders.

We have discussed one class of models that possess these properties. However, there are a number of other models, making different assumptions, which possess similar properties. It is not always clear which assumptions are best - consider the case of synchronous or asynchronous decisions. A more general picture of the consequences of different assumptions should hopefully soon emerge.

Besides this a vast unexplored wilderness of further research lays ahead. This should hopefully continue to stimulate work in this area. Of particular interest will be to see if these models can explain the interdependence between agents which results in the central limit theorem apparently not applying for high-frequency data (Lux, 2008). Other work include the setting up of rigorous microfoundations, developing better approximations for the complicated systems or even analytical solutions (although this may lie more in the domain of applied mathematics than economics), empirical validation via the estimation of parameters, the explanation of more stylised facts, and the use of these models to make policy recommendations.

References

Alfarano, S., Lux, T. and Wagner, F. (2005). Estimation of Agent-Based Models: The Case of an Asymmetric Herding Model, *Computational Economics* **26**(1): 19–49.

Alfarano, S., Lux, T. and Wagner, F. (2008). Time variation of higher moments in a financial market with heterogeneous agents: An analytical approach, *Journal of Economic Dynamics and Control* **32**(1): 101–136.

Cui, C. and Lauricella. (2011). Mini 'Crashes' Hit Commodity Trade.
URL: *http://online.wsj.com/article/SB10001424052748704322804576303522623515478.html*

Farmer, J. D. (2002). Market force, ecology and evolution, *Industrial and Corporate Change* **11**(5): 895–953.

Goldfarb, Z. A. (2010). Report examines May's 'flash crash,' expresses concern over high-speed trading.
URL: *http://www.washingtonpost.com/wp-dyn/content/article/2010/10/01/AR2010100103969.html?sid=ST2010100107554*

Lux, T. (1995). Herd behaviour, bubbles and crashes, *The Economic Journal* **105**(431): 881–896.

Lux, T. (1998). The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions, *Journal of Economic Behavior & Organization* **33**(2): 143–165.

Lux, T. (2008). Stochastic behavioral asset pricing models and the stylized facts.

Lux, T. and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market, *Nature* **397**(February): 498–500.

Lux, T. and Marchesi, M. (2000). Volatility clustering in financial markets: A microsimulation of interacting agents, *International Journal of Theoretical and Applied Finance* **3**(4): 675–702.

Rogow, G. (2009). Rise of the (Market) Machines.
URL: *http://blogs.wsj.com/marketbeat/2009/06/19/rise-of-the-market-machines/*

Sornette, D. (2003). *Why stock markets crash*, Princeton University Press, Woodstock.