System Dynamics and stock markets

Martin Timothy Rafferty

Abstract

This paper presents an analysis of the behaviour of a stock market; the London stock exchange main market as expressed in the FTSE 100 index. The paper examines the main features of the literature relating to the academic and practitioner views on market behaviour. One of the main pillars of current academic understanding of stock markets, the efficient market hypothesis, is examined and tested. A novel variation on a known flaw in the efficient markets hypothesis is examined; the sub-Monday variation on the Monday effect. Using actual data this variation is tested and found to be in violation of the efficient markets hypothesis. The paper describes two index/market designs; one with actual data and one with hypothetical data. Limitations of system dynamics models in data rich environments are illustrated. The models presented here are put forward as prototypes for the development of further stock market simulations. This paper presents a proof of concept that cyclical behaviours exist in stock markets and that stock markets are therefore amenable to analysis using the system dynamics paradigm.

Keywords

Stockmarket, “Stock market”, “FTSE 100”, “Efficient markets hypothesis”, “EMH”, “Integration error”, “sub-Monday effect”, “Monday effect”

Introduction[1]

The behaviour of financial markets occupies much intellectual effort from both the academic and practitioner communities where a fairly sharp difference of opinion is in operation. However, it is also a field where the academic and the practical approaches occasionally unite to create crossover systems (Black and Scholes, 1973).

From the academic point of view the field has been dominated for many decades by the idea that markets behave in a near random fashion. A milestone in the development of this paradigm came when Eugene F. Fama gave full voice to the efficient markets hypothesis (Fama, 1970). Fama acknowledged that the hypothesis was and remains a work in progress or as he concludes “we certainly do not want to leave the impression that all issues are closed” (page 416). To supply some context to the analyses that follow the three categories that Fama (1970) provided as degrees of the efficiency of the market with some basic implications for practitioners are outlined.

Strong form efficiency would mean that market participants have access to and act on all relevant and knowable information regarding price formation. Any market that was efficient at this level would preclude earning of any excess returns other than by chance. No analysis of the market would consistently provide excess returns. **Semi-strong form efficiency** means that market participants have access to and act on all publicly available information. No analysis of the market, that did not involve private information, would yield better than average returns other than by chance. **Weak form efficiency** means that market participants have access to and act on historic price information. No technical analysis of past price performance would yield better than returns produced by chance.

Underpinning all of these degrees of efficiency is the idea that the current price of an asset reflects, to one of the degrees noted above or some shade thereof, the available information about that asset. Hence, as new information arises prices will move to reflect that new information. As new information arises in a near random fashion, causes are numerous and opaque, asset price behaviour will reflect this near random information flow. This is the random walk (Malkiel, 2003, page 3). In recognition of the observation that markets tend to rise over time the random walk idea is generally supplemented with a component for this gentle rise upwards. The basic model of market behaviour then becomes the random walk with drift.

[1] This is not a paper about system dynamics per se it is a paper about the value of system dynamics in modelling stock markets. This research makes use of the system dynamics paradigm to describe an empirical phenomenon. Hence this paper is to be read as a proof of concept rather than a definitive exercise in market modelling.

If the efficient markets hypothesis holds true the implications are profound. There is little incentive for anyone to carry out research on stock markets and little room for system dynamics type simulations to uncover any feedback processes. Any attempt to model a market which was strong form efficient would be a simple random process with a degree of drift. Any attempt to model a market which was semi-strong form efficient could model only the strong version with additional feedbacks limited to the effects of private information. Any attempt to model a market that was weak form efficient could only model the semi-strong form with additional feedbacks limited to the effects of publicly available information that was not already incorporated in the share price.

Adhering to this philosophy would mean that models of market behaviour can still be valuable but are likely to be limited to special situations or to testing the degree of efficiency that is in operation. Examples of these types of analysis are provided by Dubow and Monteiro (2006) and Monteiro et al (2007) which both focus on insider trading.

However, the efficient markets hypothesis is not universally accepted and even in those areas where it is accepted this acceptance can be conditional on so called exceptions or anomalies. Malkiel (2003) identifies several calendar based market anomalies including the January effect, Monday effect, turn of the month and holiday effects. In addition to asserting that these effects are not dependable Malkiel further dismisses them as follows "these nonrandom effects (even if they were dependable) are very small relative to the transaction costs involved in trying to exploit them" (page, 64). The results of his analysis are a kind of acceptance, albeit grudging, that the markets Malkiel examined are not entirely efficient at any level since all of the effects mentioned above can be isolated by examination of historic price series. If a market is not efficient at the lowest level it cannot be efficient in any higher or stricter sense.

In addition to the admission that markets are less than efficient from an academic viewpoint the practitioner viewpoint can differ radically. A reasonably typical, if there is such a thing amongst such a diverse population, practitioner view is expressed by George Soros (Soros, 1994). When discussing the idea of a random walk in stock price behaviour, a foundation of the efficient markets hypothesis, he asserts that "the theory is manifestly false" with the justification that "I have disproved it by consistently outperforming the averages over a period of twelve years" (page 47). In a wider sense Soros' approach can be summed up as follows: investor behaviour and stock prices as well as behaviours and prices in other asset classes, are interlinked in such a way that patterns are occasionally observable due to the existence of a feedback process. Soros refers to this feedback process as reflexivity. It is the ability to discern and take action on these patterns that illuminates the feedback mechanism which produces further actionable patterns or investment opportunities.

Soros' approach is at the more philosophical end of the spectrum of processes that rule in the world of practitioners. Nonetheless the central core of Soros' thinking appears to be that there are patterns in the data which are exploitable by those with sufficient insight to see them; this is the general thrust of much of the literature and other material, broadcast media for example, on practical investing (Rafferty, 2012).

There is therefore some uncertainty about whether a system dynamics based investigation of stock market behaviour or the efficient markets hypothesis would be valuable. Further, there is no point in modelling the supply and demand relations of the participants if prices are apparently determined by supply and demand though in reality are completely determined by a near random information flow.

To date, explicit attempts to model stock market dynamics in the system dynamics literature are sparse. Getmansky (2003) implicitly examined market efficiency/inefficiency when examining arbitrage opportunities available to hedge funds. Chiarella and Gao (2002) modelled differences in the dynamic behaviour revealed by comparing noise with fundamental trends on the Standard and Poor's 500 index. They also supply a critique of attempts to model the market pricing mechanism using traditional regression models. Provenzano (2002) created a market model with the focus on behaviours that arose as part of the trading style of the participants. Askar et al (2007) modelled the Egyptian stock market with the focus on the implications for a stock index of differing investor behaviours. Sterman (2000) discusses the business cycle and economic long wave which hint at a link with stock market behaviour. Most of these analyses focus on supply and demand with delays and feedbacks produced by the participants as price determining mechanisms. As noted above this is unlikely to result in a meaningful analysis if prices, via supply and demand, are fundamentally determined by the near random information flow.

From a system dynamics viewpoint we can postulate a link, much as Soros does, between the information flow and the actions of the market participants; each influencing the other. At a fundamental level demand is expressed as a positive price reinforcing loop and supply acting as a balancing loop to counter the positive thrust of the demand. We could frame such a scenario from the viewpoint of a simple price responder "prices are rising so buy, prices are falling so sell" which also gives the basis for unwarranted boom and bust cycles. We could further characterise boom or bust behaviour as just a significant divergence of current valuation from true value; though there is no universally accepted method for determining true value and therefore no certain way to tell if boom or bust behaviour is occurring other than with hindsight. In this model value and price are detached. The up phase of this type of cycle was characterised by Alan Greenspan as "irrational exuberance" (Greenspan, 1996) and we could venture to characterise the down phase as irrational pessimism.

That there are feedback effects within stock markets is undeniable. The price yield figure below shows the relationship between share price and yield on the FTSE 100 index (yield has been multiplied by 700 for ease of comparison and is an average of 3.76% of price over the period shown).

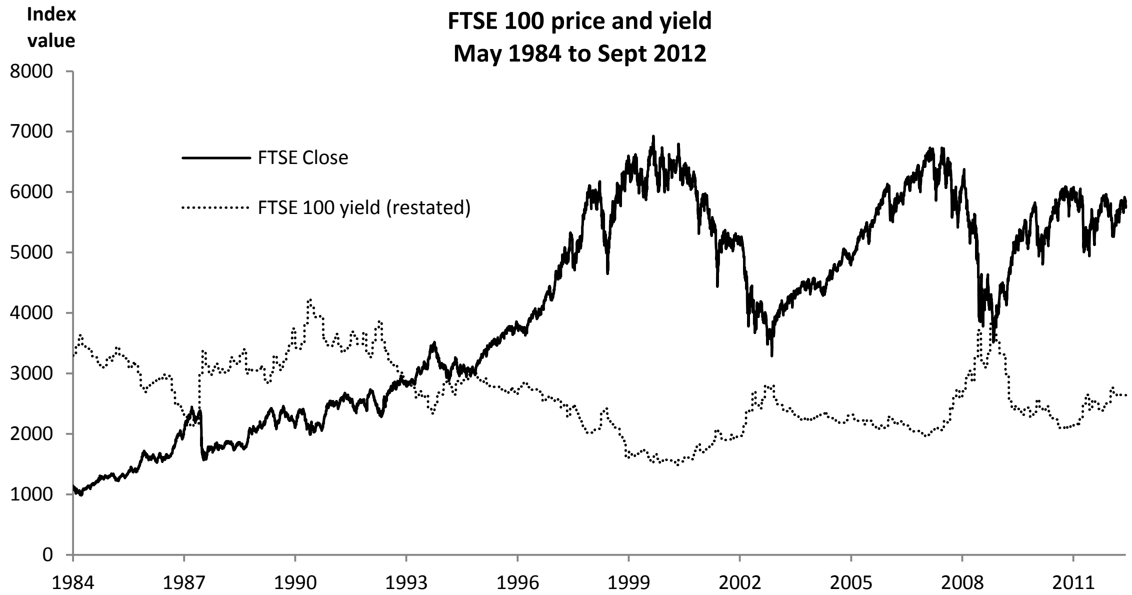

Figure 1[2]

The key finding from the initial analysis is that, according to the efficient markets hypothesis the part played by market participants as buyers and sellers is minimal in terms of determining initial prices then subsequently in determining price changes. Rather starkly, this perspective views buyers and sellers of stocks as mere conduits for the implementation of the information flow which is the essential price setting mechanism. In contrast the practitioner view is largely the opposite of the academic view; that the markets are inefficient and hence there is scope for uncovering exploitable opportunities. In pursuance of this theme the analysis presented here focuses on investigating the random nature of the information flow as revealed and expressed in market prices. A consequence of revealing any significant non-random behaviour would be that the system dynamics paradigm could be meaningfully employed in modelling stock market behaviour.

Whilst the apparently free actions of the market participants are in this view anything but free such a view is not unknown throughout the breadth of academic thought. Consider, for example, Richard Dawkins view that humans and all other creatures that pass on their inheritance genetically are

[2] Data sourced from Rafferty (2012, page 30).

unwittingly acting as vessels for the selfish gene that is the true determinant of their behaviour (Dawkins, 2006).

Focus of this analysis

If the information flow were random with positive drift the first day of any trading period after a weekend break would theoretically yield the greatest returns as it has two or more days of information to incorporate in the share price; this is the Calendar time hypothesis (French, 1980). If any other result were systematically obtained then the information flow could not be random given that in the EMH information flows are the key price determinants. Previous research has suggested that the returns on Mondays may be inconsistent with a random flow of information and the calendar time hypothesis (French, 1980). This phenomenon is known as the Monday effect and it generally states that in certain markets share prices are lower on Mondays than on any other day of the week.

For this paper it is a variant on the Monday effect[3] that will form the focus of attention. The variation that is used here is to invest, via the index, on a Monday when that Monday is below the price of the previous Monday – when there was a Monday in the previous five trading days. This variant, together with others, is introduced and described in Rafferty (2012, pages 70-71).

French (1980) investigated the basic Monday effect in some detail over the period 1953 to 1977 on the Standard and Poor's composite index. French concludes both that the "persistently negative returns for Monday appear to be evidence of market inefficiency" and that "an active trading strategy based on the negative expected returns would not have been profitable because of transaction costs" (page 68). In other words the effect had/has statistical significance but was/is of limited practical value. French goes on to suggest that the greatest value that the effect had was to allow an investor who was going to invest anyway to delay those purchases until Monday when prices were relatively depressed or reschedule sales to sometime other than Monday. By induction this would lead to outperformance and a weakening of the status of the efficient markets hypothesis. In large part this investigation focuses on French's observation that the value of this anomaly lies solely in applying it to a passive rather than an active trading strategy. It is noted that the passive investing technique is not the property of the efficient markets hypothesis it is merely the most efficient way of accessing the underlying paradigm. Those who do not accept the efficient markets hypothesis as valid may still make use of passive investing.

[3] The Monday effect is also known as the weekend effect or the (Monday) day of the week effect and the variant is described in Rafferty (2012) as the sub-Monday effect.

Draper and Krishna (2002) analyse the anomalous behaviour of Monday returns in the UK with reference to a set of explanatory variables. They confirm the existence of the Monday effect and note that it can be largely explained with reference to the explanatory variables. They have sought to add insight into cause to assist in explaining the effect. The addition of cause to effect would not reduce the significance of the Monday effect or the implications of this effect for other explanations of more general market behaviour including the efficient markets hypothesis. A significant element of the explanatory power they describe is attributed to their observation that UK dividends tended to be paid on Mondays and in particular they suggest that nearly 94% of ex-dividend days in their sample are Mondays. Dividends paid by FTSE 100 constituents are now normally paid on Wednesdays and have been since 2001; prior to 2001 ex-dividend days on the London Stock exchange were Mondays (McStravick, 2000a). A timetable is produced by the exchange each year detailing the acceptable ex-dividend dates for the forthcoming year. Any index constituent is not bound to follow the timetable but there are additional responsibilities for those that do not. A cursory examination of adherence to the timetable by FTSE 100 companies was undertaken via an analysis of ex-dividend days for the six months to 28th September 2012; Mondays were significantly underrepresented as ex-dividend days. Announcement dates, record dates and payment dates in the same period were also examined and again Mondays were significantly underrepresented. The move to Wednesdays occurred on Monday 5th February 2001 and is briefly documented and confirmed in (McStravick, 2000b). At the same time as the change from Monday to Wednesday ex-dividend days was introduced, changes to the settlement period, shortened from 5 to 3 days, were introduced; record dates remain as Fridays.

Although the majority of academic writing concerning the Monday effect finds that it persists Connolly (1991) provides an abstract statistical analysis of the Monday effect which finds that it is weak or easily explained away. This is, Connolly suggests, due to the perceived unreliability of abstract statistical analysis who states that “inferences about the DOW and weekend effects may reflect a deficiency of classical statistical methods rather than systematic (and anomalous) return behavior”, (page 94) the analysis provided here is as straightforward as possible; thereby avoiding the criticism levelled by Connolly. Support for Connolly’s general position on the Monday effect is provided by Schwert writing in Constantinides et al (2002) who offers the observation that “the weekend effect seems to have disappeared, or at least substantially attenuated, since it was first documented in 1980” (page 945).

Model design[4]

Figure 2 below shows one interpretation of the EMH in graphical format.

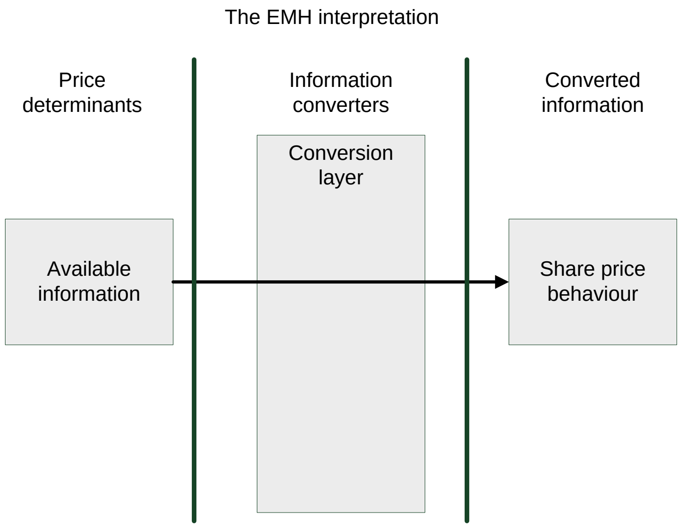

Figure 2

Figure 2 presents an interpretation of the prevailing academic view. Figure 2 shows the transformation of information into share price behaviour. The available information includes all information of all types that may affect share prices. The conversion layer contains buyers and sellers, brokers, market makers, stock exchanges and so forth that convert the information into a form that can be implemented as share price behaviour. The conversion layer is a kind of clear lens. In the EMH this lens is highly efficient and transparent in its effects – information is rapidly and accurately transformed into share price behaviour.

We may now turn to the practitioner view. A version of the practitioner view is expressed graphically in figure 3 as the non-EMH interpretation.

[4] As this research is primarily aimed at providing a proof of concept the models provided with the paper are primarily illustrative only.

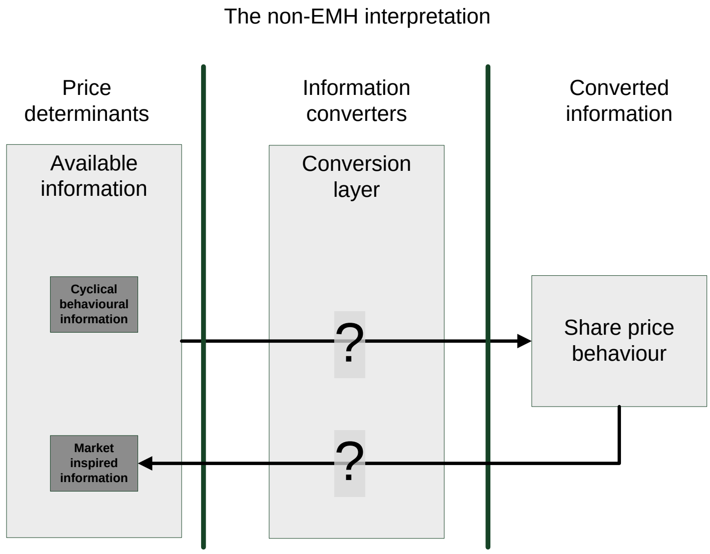

Figure 3

Figure 3 shows the transformation of information into share price behaviour in a less efficient manner. As before, the available information includes all information of all types that may affect share prices. Explicitly including market inspired information and behavioural investing patterns that are to some extent predictable; the documented Monday and sub-Monday effect for example. In this interpretation the conversion layer is a kind of translucent lens. In the EMH the conversion layer is highly efficient and transparent in its effects – information is rapidly and accurately transformed into share price behaviour. In this version the conversion layer is less efficient, containing frictions (bid-ask spreads, information asymmetries...etc) and is somewhat opaque in its effects. The track of the information through the conversion layer is neither straightforward nor exact. Figure 3 differs substantively from figure 2 in that it includes a looped structure from the share price behaviour to the available information. This subset of available information has been labelled *market inspired information* and includes, for example, the dividend yield information/relationship described in figure 1. For the purposes of this analysis, differentiating market inspired information from other information is a process that is complex to the point of impossibility due to the number of links and feedback processes in operation within the set of all information.

Examples of the subset of all available information labelled *cyclical behavioural information* are that such as mood/weather related information (Garrett et al, 2004) and temperature information (Cao and Wei, 2005) which have been postulated as partial determinants of price behaviour and are cyclical in nature with the progression of the seasons.[5]

[5] Note that these two pieces of research, and similar, are noted as unpersuasive for a number of reasons primarily due to the data being tied to physical market location which may differ radically from market

Two index designs are presented for consideration. The first uses actual closing price data for the FTSE 100 for each trading day in the period expressed in index points. This design covers the period 1st November 2007 to 31st October 2012. Data has been sourced from commercial suppliers. The second design uses artificial data constructed from theory and calibrated with observed values.

In addition to the closing price data an indication of whether the closing price belongs to a sub-Monday or not is required. This takes the form of a binary signal which reads 1 if this is a sub-Monday and a 0 otherwise.

The data series contains 1262 trading days of which 235 (18.62%) are Mondays and 104 (8.24%) are sub-Monday variants. Figure 4 illustrates the overall shape of the data.

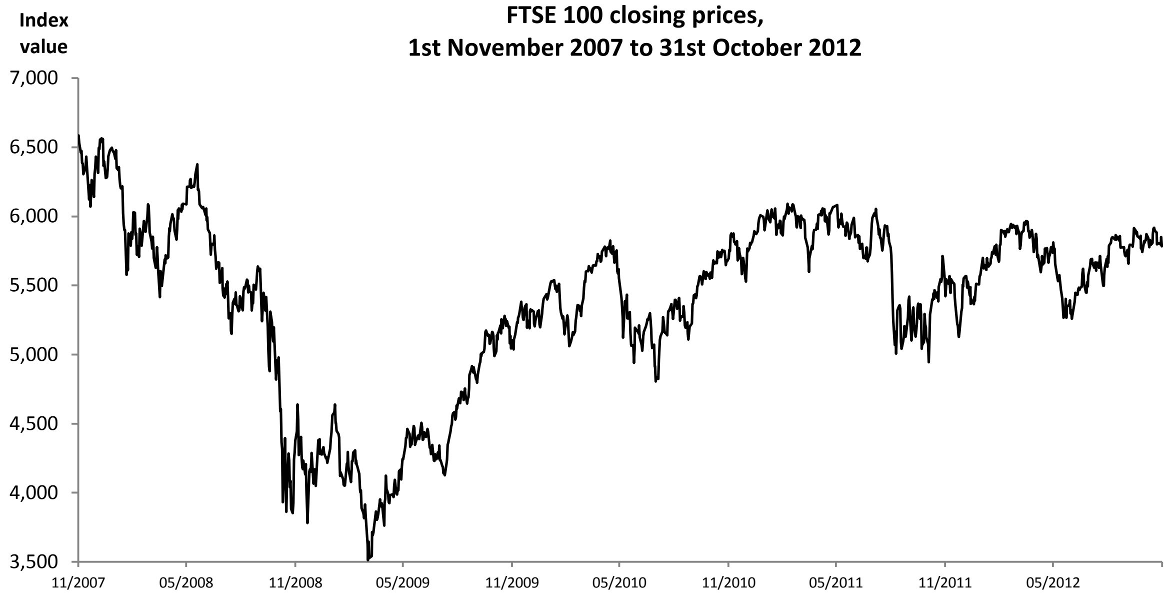

Figure 4

Conceptually the model adopts a measure that is designed to give an indication of normal market behaviour; a benchmark, the behaviour of any variant can then be measured against this benchmark. In this case the benchmark is the 'buy all rule' described in Rafferty (2012, pages 57-60). The buy all rule is an implementation of the three questions (a) what did it cost? (b) What is it worth now? And (c) what is the difference? The three questions are applied to a process that just buys into the index every available trading day and can be expressed as the equation below:

participant location. It is acknowledged that some attempt has been made to address this issue "[the effect] remains strong even after controlling for the geographical dispersion of investors relative to the city where the stock exchange resides" (Cao and Wei 2004, page 1561). Other issues are with inter-year data comparisons (are returns increased in very cold years?) or between large scale differences in ambient environment (are returns lessened in areas where light levels are generally lower?). Overall, even where these questions are considered insufficient detail on these is provided to reach a conclusion of the validity of the research.

$$Benchmark\ return = \left(n * P_t - \sum_{i}^{t} P\right) / \sum_{i}^{t} P * 100$$

Where; *n* is number of transactions (all buys), *P* is the index price, *i* is the initial time and *t* is the current time. The result of this formulation for any given day is the percentage of profit that the benchmark is returning. Capital purchase costs are included in the formulation but transaction costs are not.[6] The benchmark figure is equivalent to the excess of the current value of all purchases over the costs of buying the index every available trading day over the time period chosen.

Constructing the return figure for buying on a sub-Monday is the same as the benchmark except that all purchases are conditional on the day being a sub-Monday. The calculation of returns from buying on a sub-Monday compared to the benchmark is shown in the equation below. Symbols are as per the benchmark return except that P' is the price on a sub-Monday, RR is the net running return and BR is the benchmark return specified above.

$$RR = \left(n * P_t - \sum_{i}^{t} P'\right) / \sum_{i}^{t} P' * 100 - BR$$

The formulation shown above underestimates the discrete return on the anomaly due to the inclusion of the anomalous returns in the benchmark return.

Findings

First we discuss the simulation based on actual data.

Initially the model was run with a time step of 1 to calibrate the model exactly to the number of days that comprise the data series. The output of the model as expressed by the net running return figure is illustrated below.

[6] Transaction costs are assumed to be the same for a passive investing strategy using the efficient markets hypothesis or the timed, investing on sub-Mondays, passive investing strategy illustrated here.

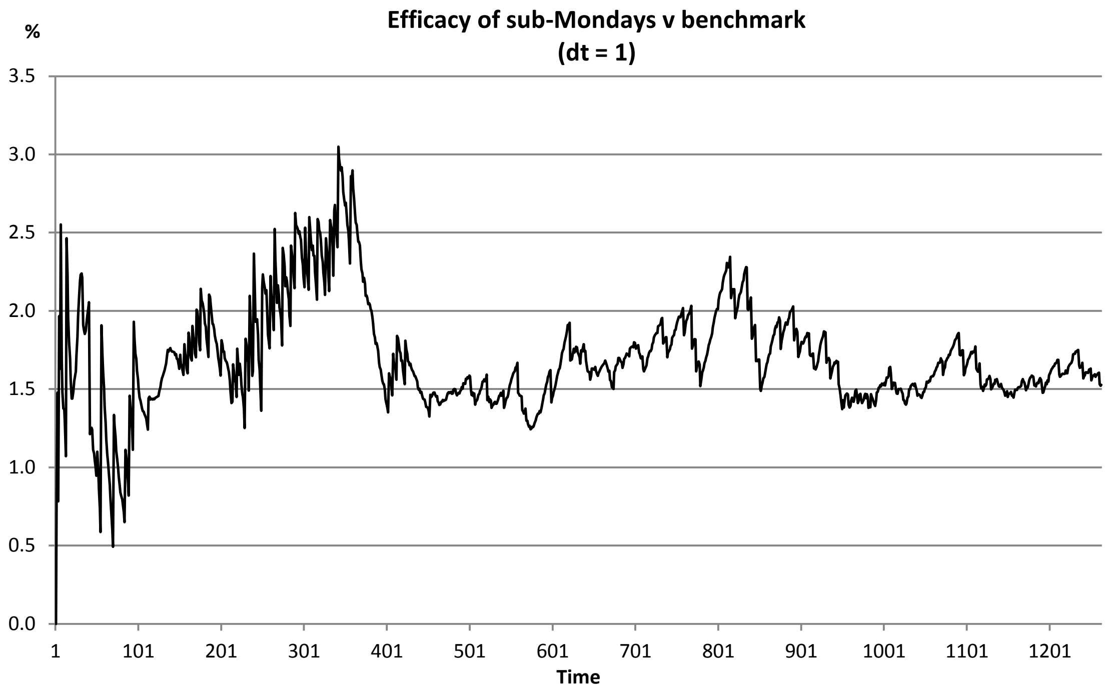

Figure 5

The model output shown in figure 5 accords exactly with the expected performance of the differing buying strategies, that is, the buying on sub-Mondays option outperforms the benchmark consistently. In this case the sub-Monday option outperforms the benchmark 100% of the time. Technically it outperforms 100% - 1 observation as on the initial pass both figures are zero.

The model is then reset to have a time step of 0.25 and is run again. The outputs are shown in figure 6.

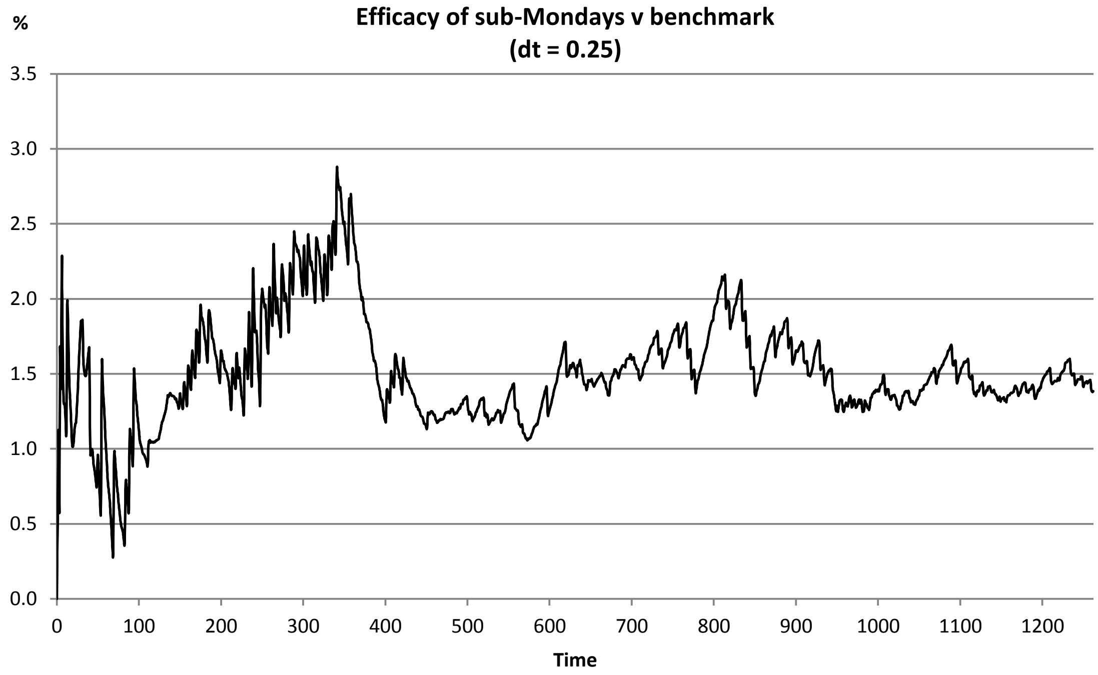

Figure 6

The outcome of this run is that the overall behaviour remains unchanged, sub-Monday outperforms buy all, though the detail of the series is different. To highlight these differences between the two series the values from figure 6 have been subtracted from those underlying figure 5 and are displayed as figure 7.

One specific aspect of behaviour that arises from this simulation is an inconsistency regarding the different buy signals. With timestep set to less than 1 any change in signal will result in a gradual change from 1 to 0, and vice versa, according to the size of the timestep. When mutually exclusive buying regimes are tested this can then result in buying on signals which are in reality mutually exclusive processes. The differences that arise in this fashion do not change the overall shape of the data here but they do present a small, though fundamental, inconsistency in the ability of the model to mimic reality. Such differences are in the nature of continuous simulations and are exacerbated with the particular model design selected here. To accept this type of simulation we must be willing to write this off as good enough for simulated behaviour. It would not be acceptable for modelling actual behaviour.

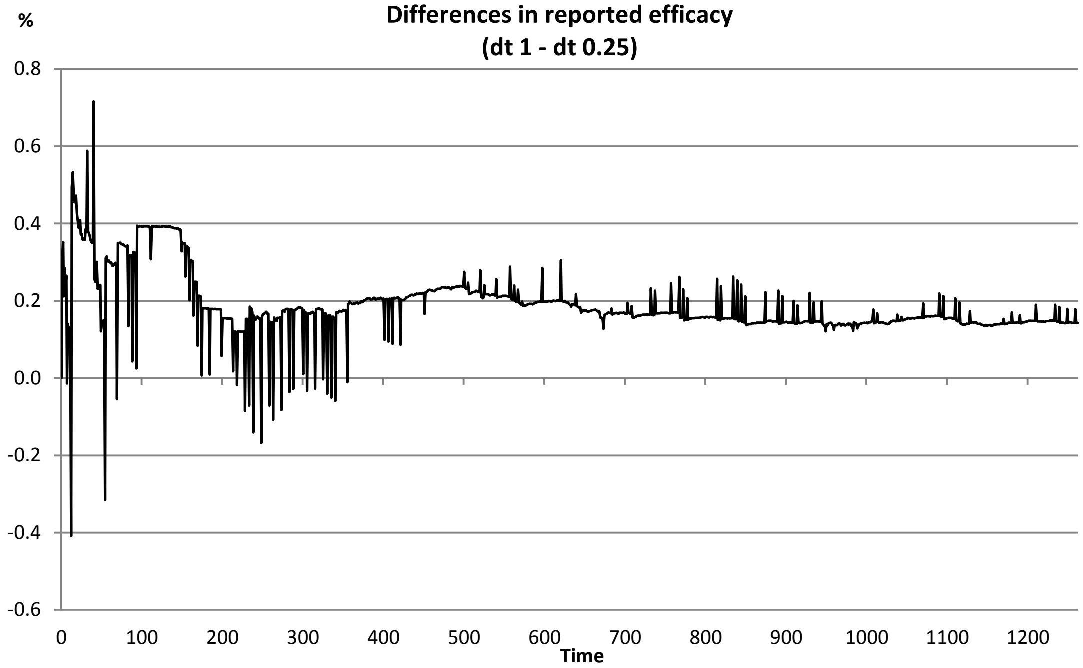

Figure 7[7]

This section discusses a simulation with hypothetical index data calibrated from real values.

Part of the strength of system dynamics is that it presents the opportunity to distil the essence of systemic behaviour and use that to model structural factors that give an insight into similar but different systems. For example, in this case it would be appropriate to see this similar but different system as the behaviour of the FTSE 100 over the coming 5 years or some other index over the same period.

To create a hypothetical simulation for any stock market we must look to the existing theory and data to guide us in replacing the actual index data. As we have discussed above the EMH asserts that the market is a random process with some degree of drift. We can therefore look to this basic paradigm to provide the framework on which to construct the model.

Choosing a random number from 0 or 1 would provide the basic model of the market where the switch was the determinant of the direction of market movements; say 1 for a move upwards and 0 for a move downwards. Looking to the data would provide us with some refinement on this basic structure. In particular if we look at the number of moves upward/downward in our calibration period these reveal the drift in terms of how many moves have been observed in each direction. If we then extend this analysis to uncover the size of the moves in either direction we can further refine the model.

The method outlined above is not the only way to model the market. The figure below shows the distribution of the data for the period 1st November 2007 to 31st October 2012. With so much data available we are spoilt for choice regarding how to calibrate our model(s).

[7] As the two series are of different lengths, there are four times as many data points for the series with dt = 0.25, the values at the end of each time period, corresponding to one day, only are included in the figure.

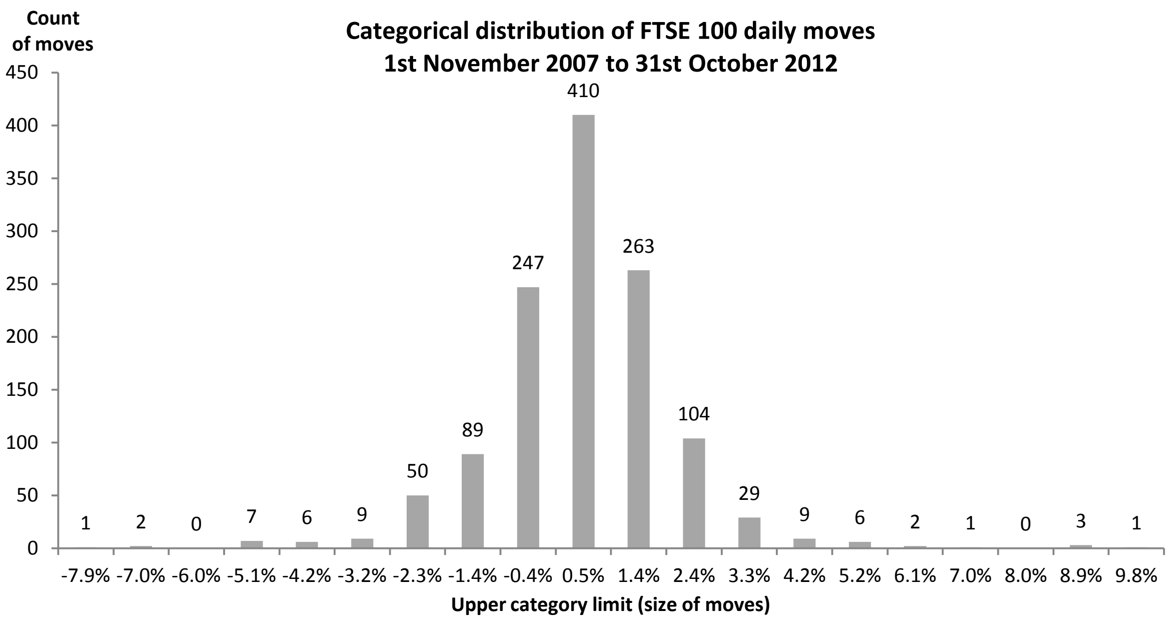

Figure 8

A visual inspection of the data in the figure above reveals that it is for all practical purposes normally distributed in this period. So rather than using simple binary switches we could simply distribute our index moves directly using the data above.

A detailed examination of figure 8 shows there is a slight bias to the downside. The actual respective starting and ending index values are 6586.10 and 5782.7 in this period so the index did indeed drift lower. Bearing in mind that the stock market, for many investors, is a long term investment vehicle we might look to a slightly longer period for a more definitive dataset. Figure 9 below shows the distribution of moves over the period 1st May 1984 to 28th September 2012.

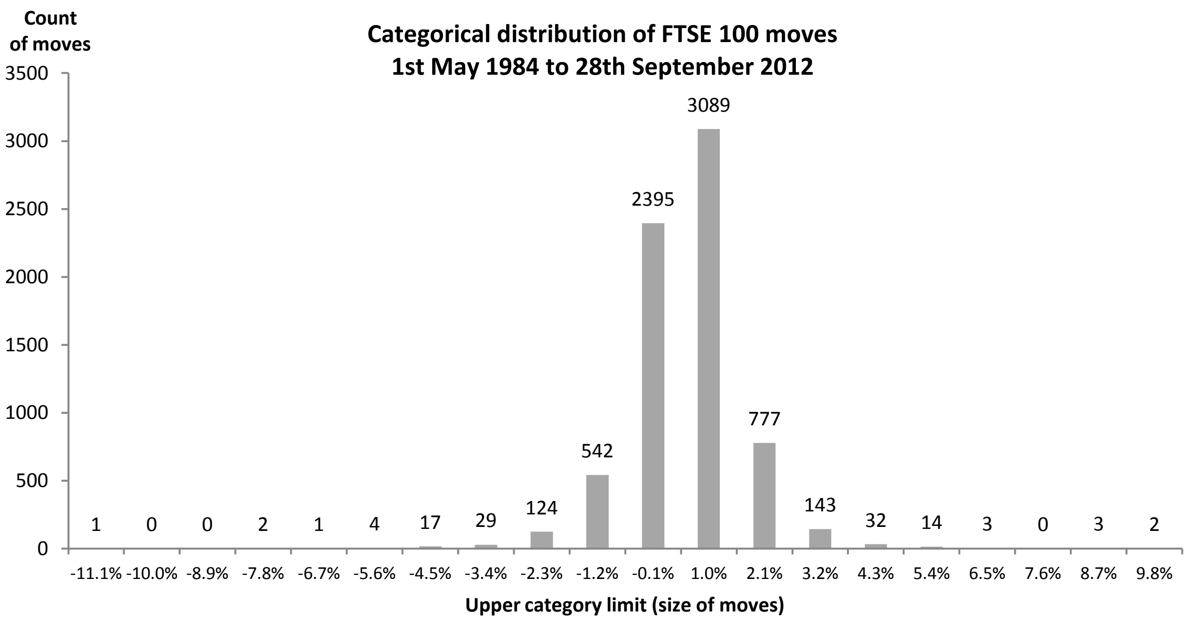

Figure 9

From figure 9 we can clearly see the expected slight positive drift in the data that was absent from figure 8.

Analysing the data over these two time periods gives us a set of figures that we can use to calibrate the simulation – which really just amounts to replacing the actual index data with simulation generated data based on whatever set of observations we use to calibrate it.

The figures in the table[8] below show two sets of calibration data.

Description	Calibration period	
	1st Nov 2007 to 31st Oct 2012	1st May 1984 to 28th Sep 2012
Size of move upwards	0.010657205	0.007927
Size of move downwards	-0.01061137	-0.00806
Size of sub Monday move	-0.00702	-0.00515
Number of moves upwards	629	3741
Number of moves downwards	630	3421
Number of sub-Monday moves	104	559
Percent moves upwards	0.498415	0.521176
Percent moves downwards	0.499208	0.476595
Percent sub-Monday moves	0.082409	0.077877

Table 1

We can now calibrate our artificial index with something more subtle than the basic 1 for up and 0 for down with drift process then run that simulation for a number of iterations to give us an idea of the accuracy of the model. The results of 10 iterations of the model, without adjustment for the sub-Monday variant, are described in the figure below.

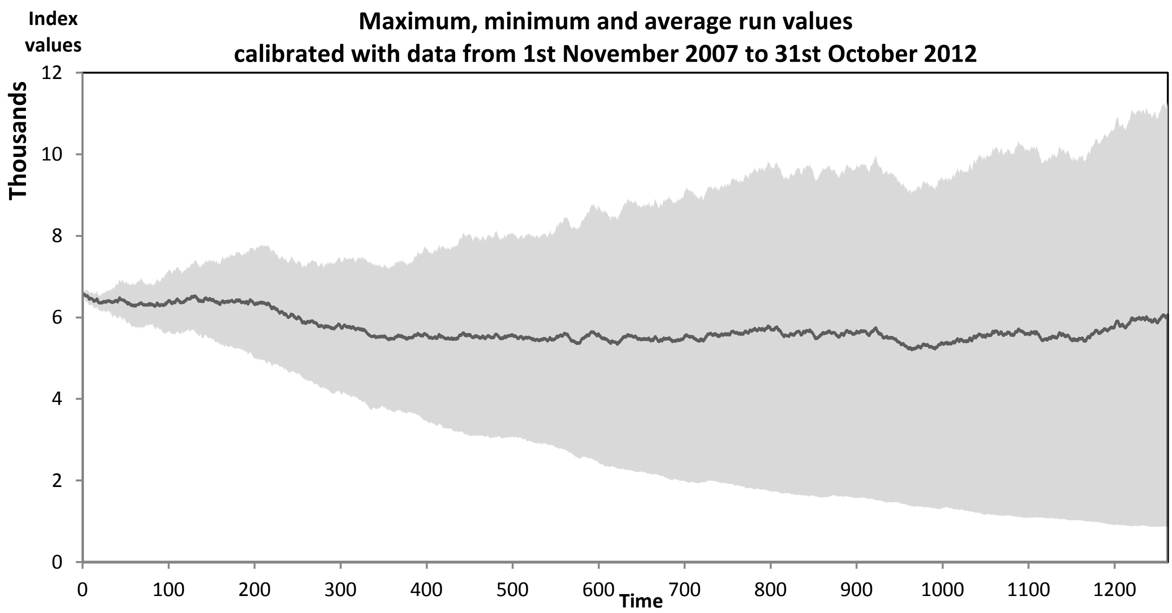

Figure 10

[8] There are also a small number of occasions when the closing price of the index showed no change from one day to the next these are 3 (0.24%) and 16 (0.22%) for the shorter and longer periods respectively. These no-change days have been left out of the model for simplicity.

Figure 10 shows the maximum values of all ten iterations of the model, the top of the shaded area, the lowest values from all ten iterations of the model, the lower limit of the shaded area, and the average of these two values, the grey line along the centre of the shaded area.

As the model contains unseeded pseudo random values that determine when the moves up or down occur no additional sensitivity analysis has been carried out; the model will naturally vary as the pseudo random values change with each iteration. The outcomes illustrated in figure 10 are not intended to be definitive and have been produced with the sub-Monday variant switched off. After extensive testing with the sub-Monday variant behaviour included the model tends to underestimate the actual behaviour of the index to a greater degree than that shown in figure 8.

We can now take this analysis a step forward and compare the average value from figure 10 against the actual data from the period used to calibrate the model. This is illustrated in figure 11 below.

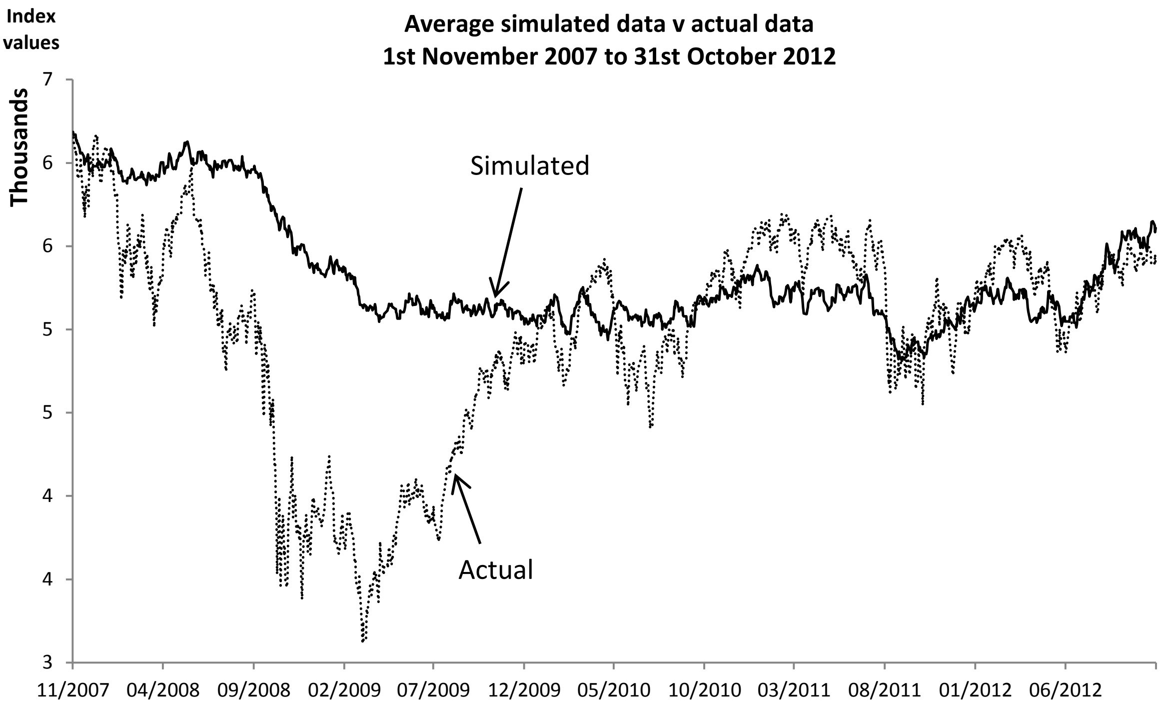

Figure 11

How accurate is the simulation in figure 11? Visually the relationship is close for much of the time though as the simulated data is an average it would not be expected to follow the peaks and troughs of the actual series and it obliges in this regard. Statistically the relationship is remote; the basic correlation is 0.31 (31%) and the coefficient of determination is 0.1 (10%) between the two series.

This index design can present an outperformance by the sub-Monday rule, when that behaviour is enabled, given that a bias has been built into this version of the index.

To project the relationship forward in time or switch it to another index the initial value of the index in the simulation needs to be set to the value of the known data series e.g. a five year forward projection for figure 11 needs the initial index value to be set at 5782.7 which is the closing value of the FTSE 100 on 31st October 2012. The model can be calibrated with either data set shown in table 1 or any other that the user chooses.

Review and discussion

The model reveals the advantages of buying according to the sub-Monday variation compared to buying according to the benchmark. The efficient markets theory precludes these findings and the findings are therefore anomalous with the efficient markets hypothesis.

The strength of the outperformance revealed in figures 5 and 6 is total (noting the exception for the initial observation given above). Given the complete dominance of the buying sub-Mondays strategy it is not possible to dismiss these observations as statistically significant but practically insignificant as French (1980) does with similar variants. Over the period illustrated the sub-Monday strategy outperforms the buy all rule, in terms of returns, by 20.44% over the five years shown. The practical significance of this result cannot be overestimated. It would leave the investor who achieved it in a position to outperform both the index and almost the entire body of professional investors.

The model design using actual data, without structural modification, can be varied to model any five year period in the history of the market shown. This involves replacing the data set (closing index prices), the buy sub-Monday signal and possibly the run length.

Either model design can simulate any five year period from any other market without further modification simply by replacing the data sets and varying the run length.

With some straightforward structural modification the model can be varied to model varying time periods and variants from the EMH. A selection of variations that outperform the FTSE 100 index and the S&P 500 index over different time periods is shown in Rafferty (2012).

It is a feature of financial analyses that they tend to be data rich. This is both a blessing and a curse when considering a system dynamics approach to modelling stock markets. This is a positive feature in that any model can be calibrated exactly using the available data. It is a negative feature in that the availability of the data exposes any variations, however slight, between the simulated behaviour and the actual behaviour. The latter can be due either to integration errors associated with the method of incorporating flows into stocks (Euler's method has been used here) or simply software limitations. Figure 7 highlights these differences and prompts the recommendation that such simulations are used to model periods or indices where real data is unavailable.

Differences between simulated results are an issue when exact data values are known or expected. They are not as substantive an issue when the model is being used to create forecasts or indications of general behaviour where exact, point by point, results are not the prime requirement.

It is not entirely clear why the system dynamics paradigm has largely ignored financial systems behaviour. However, the specific issues introduced by integration errors together with the weight of intellectual consensus in favour of the EMH, which would make system dynamics largely redundant if valid, may partly explain the lack of system dynamics work in this vast and lucrative field of research.

Philosophically the research presented here demonstrates, by reference to empirical observation, the power of inductive research over well regarded theory in the form of the EMH; the sub-Monday anomaly was discovered accidentally when searching the data for other patterns.

Conclusions

The concept that was being tested by this research, that there are significant cyclical patterns in market behaviour, is found to be valid. Therefore, useful simulations of stock market behaviour can be developed in standard system dynamics software and two such designs have been demonstrated here.

The exact nature, causal relations, of the feedbacks that are in operation is yet to be uncovered.

The simulation reveals that the efficient markets hypothesis is flawed in regard to the sub-Monday buying behaviour shown. This flaw goes beyond any simple statistical imperfection.

Care needs to be exercised when simulating behaviour with varying time steps if exact results are required.

References

Askar M, Farghaly A and Resk M, 2007, Modelling the dynamics of the Egyptian stock market, proceedings of the 2007 International System Dynamics Society conference

Black F and Scholes M, 1973, The Pricing of Options and Corporate Liabilities, The Journal of Political Economy, Vol. 81, No. 3, pp. 637-654

Cao M and Wei J, 2005, Stock market returns: A note on temperature anomaly, Journal of Banking & Finance, Vol 29, pp. 1559–1573

Chiarella C and Gao S, 2002, Modelling the value of the S&P 500 – A system dynamics perspective: Working paper No. 115, University of Technology Sydney

Connolly R A, 1991, A posterior odds analysis or the weekend effect, Journal of econometrics, 49

Constantinides G M, Harris M and Stulz R, 2003, Handbook of the economics of finance, Elsevier science

Dawkins R, 2006, The selfish gene, 3rd edition, Oxford University Press

Draper P and Krishna P, 2002, Explaining the Monday returns, Journal of Financial research, Vol. 25, No. 4, pp. 507-520

Dubow B and Monteiro N, 2006, Measuring market cleanliness, Financial Services Authority, Occasional Paper Series, 23, FSA

Fama E F, 1970, Efficient capital markets: A review of theory and practical work, The journal of Finance, Vol. 25, No. 2

French K R, 1980, Stock returns and the weekend effect, Journal of financial economics, 8, pp. 55-69

Garrett I, Kamstra M, Kramer L, 2004, Winter Blues and Time Variation in the Price of Risk, Working paper 2004-8, Federal Reserve bank of Atlanta

Getmansky M, 2003, Limits to arbitrage: Understanding how hedge funds fail, Proceedings of the 2003 International System Dynamics Society conference

Greenspan A, 1996, Remarks by Chairman Alan Greenspan at the annual dinner and Francis Boyer Lecture of The American Enterprise Institute for Public Policy Research, United states Federal Reserve Bank

Malkiel, B. G., 2003, The efficient market hypothesis and its critics, Journal of economic perspectives, Vol. 17, No. 1

McStravick A, 2000a, Stock Exchange Notice: Consultation on ex-dividend regime, N33/00, London Stock exchange

McStravick A, 2000b, Stock Exchange Notice: Confirmation of ex-dividend regime, N41/00, London Stock exchange

Monteiro N, Zaman Q and Leitterstorf S, 2007, Updated measurement of market cleanliness, Financial Services Authority, Occasional Paper Series, 25, FSA

Provenzano D, 2002, An artificial stock market in a system dynamics approach, Proceeding of the 2002 International System Dynamics Society conference

Rafferty M T, 2012, Ten simple rules that beat the Stockmarket: Outperforming the FTSE 100 May 1984 to October 2012, CreateSpace publishing

Soros G, 1994, The alchemy of finance: Reading the mind of the market, John Wiley and sons

Sterman J D, 2000, Business Dynamics: Systems thinking and modelling for a complex world, Mc-Graw Hill