Draft1/v2 January 2007

Bifurcations and bubbly outcomes: The local instability of financial markets

Roger J Bowden[1]

Abstract

Concerns about 'short termism' associated with hedge funds, and other manifestations of diminishing risk aversion, have revived interest in speculative bubbles and related episodes of financial instability. Earlier models of rational bubbles fell short in explaining why bubbles appear and collapse; nor can their dynamics generate local dependences in volatility or the mean of returns. It is shown that social reflexivity of investor opinions in response to information flows is capable of creating equilibrium bifurcations in the mapping from signals to states, and hence bubbly outcomes. Even serially uncorrelated information results in episodic locally-trapped states, that exhibit serially dependent returns and path dependence. Structural modelling of social influence connects local stickiness propensities to parameters and states of investor unanimity and risk aversion. As to hedge funds, it is suggested that the real risk lies in convergent behaviour.

Key Words: Attracting equilibrium, hedge funds, market noise, path dependence, rational bubbles, short termism, social inference, volatility clustering.

JEL reference numbers: C22; D82-4; G14

[1] School of Economics and Finance, Victoria University of Wellington, P.O. Box 600 Wellington; roger.bowden@vuw.ac.nz; roger.bowden@kiwicap.co.nz.

I Introduction

Speculative bubbles have long fascinated market commentators, academics, and the public at large. But are they really bubbles, and what is a bubble anyway? Some writers have considered that 'bubbles ain't bubbles', but simply manifestations of properly revised market opinions about sustainable investment fundamentals. Thus Garber (2000) mounted a case that the famous Dutch tulip episode of 1634-37 could be justified in terms of the underlying economic value of the variegated bulbs. Even without expectations of further price rises, a holder could sustain a loss on the purchase of a single bulb simply because the bulb could be propagated. Certainly, usage has cast its net wide enough to encompass situation where it is perfectly reasonable to explain price rises as the operations of fundamentals. In recent times, central bankers[2] have periodically warned about 'housing bubbles', while the Economist magazine put it even more strongly[3]: "The worldwide rise in house prices is the biggest bubble in history. Prepare for the economic pain when it pops" . But a good case can be mounted that it is perfectly possible to explain house prices in terms of the fundamental three I's of the housing market: incomes, immigration and interest rates. Moreover, houses are not very liquid assets with heavy transactions costs, and households typically take a much longer view of things. If one is to look for bubbles, it might be better to think in terms of shorter run holding periods and to look at more organised trading exchanges for homogeneous financial assets.

Concerns about 'short termism' in stock markets do indeed appear to share some semantic connotations with bubbles. Short termism itself refers to a preoccupation with short holding period returns, and a feeling that this may cause investors to devote an undue amount of attention to end of period price. In turn this predisposes the market to infection by extraneous influences that may wash out over longer holding periods. Models of rational bubbles, in terms of rational expectations equilibria (REE), concentrated on this aspect. Taylor (1977) noted that expectional equations had a predisposition to generate extraneous solutions. Other early work along the same lines can be found in McCallum (1977), Flood and Garber (1982) and Blanchard and Watson (1982), working in contexts such as hyperinflation or foreign exchange markets. Gourieroux *et al* (1982) formalised the solution in terms of a decomposition principle according to which extraneous information creates a non

[2] Examples are US FRB 's Alan Greenspan (*New York Times* 2005/12/25) and Alan Bollard of New Zealand in a series of pronouncements between 2004-6, for which see www.rbnz.govt.nz.
[3] Issue June 16, 2005

fundamental solution; the observed path consists of a sum of the fundamental and non fundamental price levels. Bowden (1988) surveys much of this work.

Rational expectations models of this kind were originally labelled as bubble models, but a better case could be made that they are simply models of market noise. A difficulty is that the rational noise so generated has very little structure, so that excess returns cannot be earned on the average, even locally. On the other hand, the popular view of speculative bubbles is that they are more structured in nature. They take root, grow and collapse. The bubble regimes are sticky while they last, and may well generate episodes where asset returns are themselves serially correlated. However, the serial correlation is purely local. It may be difficult for the market to learn its precise form in time to exploit the inefficiency. This suggests that bubbles are regimes or phases, and bubbles form or collapse because of a phase transition.

In the rational bubble literature, phases can exist, but their transitions are not modelled – they just occur randomly. Likewise, the econometric literature has attempted to model phase transitions in terms of Markov switching or mixed distributions (e.g. Peria (1999), Pericoli and Sbracia (2001), Lim and Martin (1998), Lim et al (1998)). However as with the original regime switching model of Hamilton (1991), these remain black box models as to mechanism; one might hope to identify sensitive values for key variables that might govern transition probabilities, but that is all.

The models of Lux (1995) and Kojima (2000) attempt to go a stage further and supply some of the necessary structure. They have a physical reference to analogous models of excited gases in thermodynamics, but they share a common ground in what Topol (1991) characterised as models of incomplete information and limited rationality. In such models, agents attempt to assess or guess the attitudes of others as way of reinforcing their own information. This introduces a social reference to the prediction of market outcomes. Phelps (1983) had used the social reference to model inflationary expectations. Earlier still, the famous beauty context paradigm of Keynes (1936) referred to much the same effect. Keynes likened players in financial markets to panel members who had to guess which girl the judging panel as a whole would vote for. Models of limited rationality should not be confused with models of irrational bubbles, most commonly contagion, in which bubbles arise because investors tend to act in concert as a result of the pyschology of crowds or herd dynamics (e.g Kindelberger (1996)).

In the Kojima model, agents settle down to a distribution of opinions driven by their own assessment of the available information, on the one hand, and their assessment of what the market is likely to think, on the other. Kojima asssumes that agents tend to align with each other to minimise a 'disgreement function' for a given amount of opinion dispersion (measured as entropy) between them. The physical reference aside, the alignment is what Phelps had in mind, essentially a combination of sources for the agent's final opinion, one stemming from the available external information, the other from what the agent thinks the market as a whole will make of it. However, the Kojima model stops short of explaining just why agents settle down into such a mutual equilibrium; the 'disagreement function' that is being minimised remains rather metaphysical in its origins. But as suggested by Bowden (1988) in the general context of social inference, the resultant effect has the character of an equilibrium to an inferential game. Indeed, the structural view of equilibrium developed in the present paper works in terms of randomised reaction correspondences recognisable in terms of cooperative game theory,

The present contribution combines a number of the above influences and insights into a more structured model of phases and phase transitions. Individual opinions depend not only upon the external information, but also how the investor thinks other investors will interpret it. The social reference is constructed in such a way as to be consistent with either the imperfect information model or with a herd effect. A reaction correspondence is established between individual opinions and market opinion. The precise nature of this depends upon factors like opinion dispersion, social reliance, and also investor inertia, interpreted here as a disposition to 'sit on their hands', perhaps driven by fear of regret. The correspondence in turn implies an equilibrium mapping from the external information to the market outcomes, in the form of excess returns or the price level.

The resulting mapping exhibits phase bifurcations, depending on key parameter values that can be related back to investor behaviour. This tells us something about why bubbles might form and why they will eventually collapse. The bifurcations are a topological fixed point property of executable market opinions. Provided only that stronger positive opinions mean stronger positive market outcomes, the same properties should be true of observed price or return outcomes. The model is capable of generating the same outcomes as those of Kojima, but yields more insight as to why they might happen and the stochastic structures that result. It is also more closely related to models of market valuation. In particular, the conventional rational bubbles model is nested within it as one of the bifurcations - noisy but not bubbly. Bubbles, when they occur, are locally non REE, and are sticky – once there, the jump-out is not immediate. In other words, bubbles are locally trapped states.

This leads to the interesting result that even if the underlying external information flow is pure innovation (white noise), the outcomes are not. Thus a serially uncorrelated input can create a serially correlated output, even without a lagged transfer function. The effect arise because the mapping from external information to market outcome is not one to one and continuous, but bifurcates as point to set mapping, depending on parameters and also according to the current state. As a result, the dynamics is path dependent: recent history matters. Bubble phases have temporal structure. The ground states that they revert to have no temporal structure; they are noisy REE, but not bubbly.

The scheme of the paper is as follows. Section II sets the stage by reviewing the standard REE model that generates rational noise, going on to look at the distinction between noise and bubbles. Section III establishes the structured bubble, showing how it emerges from the Phelps and related models, and exploring variants. Section IV analyses the equilbrium bifurcations that characterise the solution structure of the model, showing how it nests the REE model of market noise, and examining influences on growth and collpase. Section V offers some concluding remarks. It reviews empirical implications for models of volatility clustering, and how the conceptual framework can help with issues of market or systemic stability, such as whether hedge funds help or hinder in this respect.

II Rational expectations and market noise

This section adapts some of the earlier literature on bubbles to the context of a generalised financial market, where investor have to price a given asset in terms of a valuation model. The same framework could potentially handle a variety of financial markets: the stock market, foreign exchange market, commodity markets, and perhaps even real estate markets. The underlying pricing model is simple and stylised: time is discrete, and investors price the asset as the discounted value of expected rewards over the holding period at a given cost of capital (r). Expectations are rational, and formed as of information available at the start of the period. The underlying model so generated is a standard geometric Brownian motion in asset prices, though we work in discrete time. The objective is to show that market noise can arise, and to examine its stochastic character. A suitable starting point is what happens in the absence of extraneous noise.

2.1 The fundamental solution

Let (m_t) denote an ongoing reward process: dividends, coupons, rentals as appropriate. If r is an investor cost of capital, then the fundamental solution p_{mt} is defined by

(1) $$p_{mt} = \frac{1}{1+r} E[m_{t+1} + p_{m,t+1} \mid \Im_t^m],$$

where $\Im_t^m$ is the filtration generated by the *(m_t)* process. It is assumed that the ex post fundamental solution for any time *t* can be written as the discrete time geometric Brownian motion process define by:

(2) $$p_{mt} = (1+r)p_{m,t-1} - E[m_t \mid \Im_{t-1}^m] + p_{m,t-1}\sigma_m \varepsilon_{mt},$$

In equation (2), ε_{mt} is the normalised innovation in the *(m_t)* process such that $m_t = E[m_t \mid \Im_{t-1}^m] + \sigma_m \varepsilon_{mt}$. Note that in this context $\Im_t^m = \Im_t^{m,p_m}$,

so that the price process itself *(p_mt)* contains no further information than what is contained in the fundamentals, and vice versa: one can back out all the information about the fundamental process from past observed prices.

Along the path defined by (2), the *ex post* return is given by

$$r_t = \frac{m_t + p_{mt} - p_{m,t-1}}{p_{m,t-1}} = r + \frac{1 + p_{m,t-1}}{p_{m,t-1}} \sigma_m \varepsilon_{m,t}.$$

Along this path, $E[r_t \mid \Im_{t-1}^m] = r$; so that on the average, no excess returns can be derived, given the available information.

2.2 Introducing extraneous information

Suppose now that the information filtration gets infected by an extraneous source of information with innovations i_t, which is statistically independent of the fundamentals innovations process (ε_{mt}). The sequence of 'news' i_t contains no true information about asset fundamentals, but is accepted by the market as though it does. It might come from the pronouncements of market commentators who have acquired fortuitous credibility. Or such a process might implicitly exist as backed out from the insufficiency of fundamental news to explain past fluctuations in prices.

In this case, an augmented investor information set results, which can be denoted by $\Im_t^{m,i}$. The resulting share price sequence (P_t) is defined by

(3) $$P_t = \frac{1}{1+r} E[m_{t+1} + P_{t+1} \mid \Im_t^{m,i}].$$

An equivalent formulation in this case is that $\Im_t^{m,i} = \Im_t^{m,P}$. The infection sequence *(i_t)* can be backed out from observations on the market price (P_t), augmented with the fundamentals sequence (m_t).

Taylor (1977) and Gourieroux *et al* (1982) showed in such a case that the price process (P_t) decomposes into two components, a separation principle. The first is the fundamentals solution, defined as in (2) above, which is driven only by the fundamentals (ε_{mt}). The second is the extraneous solution (x_t) say, which is driven only by the extraneous information i_t, and can be taken to be generated by the Markov process

(4) $$x_t = (1+r)x_{t-1} + i_t \ .$$

Some writers refer to the process x_t as defined by (4) as a 'bubble', but we shall comment further on this usage below.

Achieving strict separation between fundamentals and extraneous components can be achieved but a more natural approach for the geometric Brownian motion process is to assume conditional separation, which applies to the present time period but not to the entire past. Thus consider an asset price evolution defined by

(5) $$P_t = (1+r)P_{t-1} - E[M_t \mid \Im_{t-1}^m] + P_{t-1}(x_t - (1+r)x_{t-1}) + P_{t-1}\sigma\varepsilon_{mt} .$$

The system (4), (5) is recursive. The extraneous solution x_t remains independently determined via (4) together with the extraneous innovation series, and P_t then follows via (5). Both past extraneous elements and fundamentals will jointly affect the scale factor P_{t-1} , so that complete separation does not hold. On the other hand, the current innovations i_t and ε_{mt} will enter separately in their effect on the current observed stock price P_t [1] .

Rational expectations and market efficiency will continue to hold conditional on information available at time t:

(6) $$r = \frac{E[P_{t+1} + m_{t+1} \mid \Im_t^{m,i}]}{P_t} .$$

An implication is that ignoring the extraneous component can be costly, even though one might suspect that in reality the flow of information i_t is vacuous. If the market as whole believes that it has informational merit, then it will end up driving the market. As noted in Evans (1983) and Bowden (1988), rational expectations models of this kind can be interpreted as collective non-cooperative Nash equilibria.

From (5) and (6), the actual holding period return over period *t* is given by

(7a) $$r_t = r + g_t + \sigma\varepsilon_{mt}\frac{1+P_{t-1}}{P_{t-1}},$$

where

(7b) $$g_t = x_t - (1+r)x_{t-1}.$$

Recalling that r is an investor required rate of return, the magnitude g_t can be interpreted as an excess return earned from the extraneous part of the solution, x_t . In the noisy REE model, $g_t = i_t$ which is serially uncorrelated and unpredictable.

Notice also that the emphasis in the above has partly shifted away from the future flow of fundamentals, e.g. as in the PV of future dividends, towards a focus on the end of period price P_t . It is the infection of the latter by the current extraneous information (i_t) , via the conditioning operator in (6), that creates 'short termism'. The latter could be defined as a myopic focus on calling the market just one period ahead rather than the entire future horizon.

2.3 Bubbles versus noise

In the noisy REE model as so far exposited, infection by extraneous information does not interfere with the market efficiency proposition that on the average, no excess returns can be earned, so that $E[g_t \mid \Im_{t-1}^{m,i}] = 0.$ Once the extraneous solution is incorporated, the stock price P_t can certainly wander well away from the fundamental path p_{mt} , just by the operation of a drifting random walk like (4). But excess returns along such a path have no serial structure; they remain unpredictable.

However, departures of this kind do not exhibit the typical features of what many people might think of as a bubble, namely a rapid expansion phase that appears to have a more definite dynamic structure or generating mechanism, typically followed by an equally rapid deflation. Such paths may well exhibit local serial correlation or stickiness, so that the excess returns g_t in equation (7a) are potentially predictable even if the predictability lasts only so long as the life of the bubble. Moreover, the expansion phase commonly appears to be characterised by a form of social influence or contagion, so that some attention has to be paid to a generating role for interactions among investors, in which each is trying to guess the future actions of others.

The distinction therefore proposed is one between 'rational market noise' and a 'bubble', both subsumed under the more general label of market noise. Rational market noise has no serial temporal structure; bubbles do, even if it is only a local and temporary. The ensuing discussion builds on this distinction by establishing a class of structures that may from time to time give rise to local dependence or memory, lapsing to rational market noise as a ground state.

III Structured bubble models

3.1 Review

Models in which agents have to pay attention to the presumed stances of others have appeared in the general rational expectations literature, notably in connection with monetary and inflation equilibria. Thus in the Phelps (1983) inflation model, firms pay attention to a monetary shock as the basic informational input and form their final asking prices by combining their assessment of the monetary shock with their expectations of price movements in society at large. The latter are to be formed in a consistent way according to the same thought processes of other firms. A rather similar model was the basis for the well known critique of monetary policy by Lucas (1972). Inputs in the Lucas price setting model are the firm's own private demand signals together with their views about inflation at large, i.e. what other firms are doing collectively. Monetary shocks can stimulate output if there is a divergence between the two, i.e. a temporary expectational disequilibrium.

Similar ideas can be utilised to inform models of investor attitudes. Market players might mutually guess their way to an equilibrium in which each investor is trying to allow for the presumed attitudes of other investors, in addition to their own private assessment of the available information. A celebrated analogy of this kind is Maynard Keynes' beauty contest (1936 p.158), wherein judges have to assess not so much their own preferences, but also the average preferences of the judging panel. Lux (1995) and Kojima (2000) have recast interactive models of this kind into a framework provided by the statistical physics of excited gases.

In the Kojima formulation, investors form discrete attitudes - either 'buy' or 'sell'. A 'state', indexed by k, is a particular configuration of buy/sell among each of N investors. Thus $\boldsymbol{h}^k$ might stand for a vector of length N, each of whose elements is either +1 (buy) or (-1) (sell). Each investor's attitude is informed by (i) a private response to an informational signal (which is public information); and (ii) a feedback term from the presumed attitudes of the other investors. Thus all investor attitudes are formed simultaneously. Of course, it is unreasonable to expect all investors to know at all times the precise intentions of the others. Quite apart from practical informational impossibility, they are all to be jointly determined. So what Kojima does is to assume that the given states, or attitude configurations, settle down to maximise investor agreement[2]. The settling down is captured by finding the probability distribution of state vectors that maximise agreement between investors, but subject to a given entropy in the $\boldsymbol{h}^k$ as the underlying measure of uncertainty. The distribution of states that does this is the Boltzmann distribution. Kojima utilises the average $\bar{h}^k = \frac{1}{N}\sum_{i=1}^{n} h_i^k$ as the excess demand pressure in each state, and goes on to find an expression for the expected value of $\bar{h}^k$ over all possible states *k*. Price change is determined by this expected value. The latter obeys a symmetric logistic expansion or contraction, provided by the hyperbolic tangent function *tanh*, which is defined between ± 1. The dynamics of price change are conditioned by the given entropy in the distribution of states. If this is low enough, the rate of price change will start increasing up to a stationary point, so an expansion in price first accelerates then remains at a constant velocity thereafter. If for some reason the entropy then rises past the critical point, the system will move into a contraction phase, and the bubble will collapse just as quickly as it rose.

The physical analogy is interesting and suggestive, but there are some inherent problems:

(a) Models sourced more or less directly from statistical physics do not mesh too well with asset valuation or rational expectations. The analysis needs to take place within a nesting framework relative to asset price rational expectations equilibria or disequilibria. In other words one should be able to relate back to system (1) - (7) of section 2.

(b) Entropy is an inherently black box concept. But investors are hardly inanimate gas molecules; they have different opinions, different risk aversions, and cleave to different degrees with the herd. Moreover, such differences may themselves be endogenous, with a role to play in the establishment or disestablishment of bubbles. The correspondence of entropy with the microeconomics of market price determination, agent behaviour and interaction needs to be made clear, or alternatively a notion of mutual uncertainty among investors developed that is more identifiable in such terms.

The resolution of such difficulties is the departure point for subsequent development. However the underlying framework owes much to the discussions quoted to this point.

3.2 Equilibrium and investor opinion

The working space in what follows concerns investor opinions, or more precisely executable investor attitudes or judgements. Investors are presumed to ‘put their money where their mouth is’, so that opinions give rise to orders to buy, sell, or in some variants to sit tight and do nothing, as a third alternative. Working in opinion space is a convenience that avoids the necessity to choose a specific market clearing mechanism or microstructure. One need only assume that a monotonic relationship exists under which higher values of aggregate executable opinions generate higher values of their object (see below).

The object of opinions in the coming period t could be taken to be

(a) The excess returns (g_t);

(b) The level of prices warranted by the extraneous information (x_t);

(c) Some intermediate fractional difference $x_t^d = (1-(1+r)L)^d x_t ; \quad 0 < d < 1$.

The bulk of the development that follows is concerned with opinions about excess returns (*growth bubbles*). Active investors such as hedge funds and other speculators are more likely to be drawn into trading by the prospect of excess returns. However, extensions are indicated to case (b) (*levels bubbles*).

It is convenient to define an executable opinion or judgement space consisting of the extended unit interval H =[-1,1] such that an opinion $h \in [-1,1]$. In the most basic version (see below), positive values induce buy orders and negative values sell orders. Depending on the context, the opinion can be either private to individual i as h_i; or public, denoted simply as h.

It is supposed that investors have respective weights w_i in aggregate executable opinion, reflecting disposable wealth or some other indicator of order size or pricing influence. If $h_i ; i = 1,2,...N$ represent individual opinions among a set of N investors then aggregate opinion is taken to be $h = \sum_{i=1}^{N} w_i h_i$. For example, suppose that aggregate market orders, are proportional to aggregate opinion h, as hC. Then investor i contributes $w_i h_i C$ to the aggregate order book.

The desired insights concern properties that predispose a system to develop sticky states that arise from multiple fixed points. Properties of this kind are really topological in character rather than dependent upon specific metrics that might arise in connection with particular market clearing processes. To focus discussion on such aspects, it is assumed that a topological conjugacy[4] φ exists as between investor opinions and their outcome objects, so that one maps into the other in a smooth fashion, regardless of the precise underlying market

[4] Devaney (1989) is a well focussed source of some of the dynamic concepts used in the present paper, including fixed points, bifurcations and other path properties.

clearing mechanism. The function φ can be taken to be monotonic positive, so that stronger opinions generate higher excess return outcomes, if the opinions are about excess returns.

The point of topological conjugacy is that fixed points will correspond as between executable opinions and outcomes. For instance, suppose the objects of opinions are excess returns, g. If H is the space of executable opinions and G is the space of excess returns, and given functions f: H→H and q: G→G, then there exists a homeomorphism φ : H→G such that $\varphi \circ f = q \circ \varphi$. This means that a fixed point h^* in opinion space such that $f(h^*)=h^*$ will generate a corresponding fixed point $g^* = q(g^*)$ in returns space. Working in opinion space makes it easier to examine conditions under which consequential fixed points will arise in outcome space, without the necessity to construct a particular model of market clearing.

3.3 Equilibrium opinions

Investor opinions originate in the flow of new extraneous information and its presumed effects. Let $\eta \in H$ denote an informational signal that influences investor opinions. It is assumed that this signal is public information – the analysis can be extended to a set of purely private informations (as η_i), but the emphasis in what follows will be on differential reactions to common information.

A Phelps opinion equilibrium can be then characterised in game theoretic terms as a reaction correspondence of the general form:

$$(8a) \qquad h_i \leftrightarrow \alpha\eta + (1-\alpha)h; \;\; i = 1,2,...N$$

$$(8b) \qquad h = \sum_i w_i h_i; \quad w_i \geq 0, \sum_{i=1}^{N} w_i = 1$$

where N is the number of agents, assumed large. Expression (8a) denotes an equilibrium correspondence between investor *i*'s opinion, and a compound input in forming that opinion on the right hand side. The latter can be represented as a linear weighted combination of the informational signal η and the common opinion *h,* and it is assumed that $0 \leq \alpha \leq 1$.

If the above correspondence has an equality in (8a), then the only solution is trivially $h_i = h = \eta$, i.e. complete commonality in opinions and primary information signal. This would be a model of investor homogeneity. In general, however, people can have different opinions:

(a) There might be individual differences in how to interpret the given public informational signal. Thus the effective information acted upon by agent i might be $\eta_i = \eta + \delta_i$ with $\sum_i w_i \delta_i = 0.$

(b) Similarly, there could be differences in the degree to which individuals allow for common opinion, or adhere to the herd: thus $\alpha_i = \alpha + \upsilon_i$ with $\sum_i w_i \upsilon_i = 0.$

(c) Investors might differ in their assessment of market opinion, so that investor i will assess this as $\hat{h}_i$ with the average at h.

(d) Other executable differences can exist in the extent to which investors translate the available common information into investment decisions, e.g. variations in risk aversion could mean that investors are less responsive to any given signal.

Thus in general equality will not hold in (8a), so the task becomes to specify that nature of the equilibrium correspondence in the presence of investor heterogeneity. It is also desirable to show that such equilibria are causal, e.g. by exhibiting an iterative structure in which investors form their estimates of the common opinion recursively, even if not in real time. These aspects will be considered in turn.

3.4 Models of heterogeneous opinion equilibria

A first approach might be to suppose that in equilibrium, an investor can be either a buyer or a seller, with probability governed by his or her assessment of the common opinion. One could view this as analogous to a randomised response in game theory.

To fix ideas, suppose that the odds that investor i, taken at random, is a buyer versus a seller are log linear in the right hand side of (8a). Thus if $\lambda = prob(h_i = +1) / prob(h_i = -1)$ then the log odds are linear in the compound signal:

(9) $$\log \lambda = \mu[\alpha\eta + (1-\alpha)h] .$$

In turn, the above log odds model is formally equivalent to a discrete choice structure of the form:

(10) $$\begin{aligned} h_i &= +1 \Leftrightarrow \alpha\eta + (1-\alpha)h \geq u_i \\ h_i &= -1 \Leftrightarrow \alpha\eta + (1-\alpha)h < u_i . \end{aligned}$$

Here the u_i are random drawings from a logistic distribution function with zero mean:

$F(u) = \frac{1}{1+e^{-2\mu u}}$, with standard deviation $\sigma_u = \frac{1}{\mu}\frac{\pi}{2\sqrt{3}}$. The classification disturbances u_i can be taken to proxy individual variation in the willingness of investors to become buyers or sellers. The parameter μ measures opinion concentration. Higher values of μ mean less dispersion in investor attitudes and a sharper response to the compound signal $\alpha\eta + (1-\alpha)h$. The entropy associated with $F(u)$ is given by $2 - ln(2\mu)$, so the odds parameter μ can be taken as an inverse entropy measure.

From expression (8b) we have

$$h = \sum_i w_i h_i = \sum_i w_i \, sgn(\alpha\eta + (1-\alpha)h - u_i).$$

Now

$$E_u[sgn(\alpha\eta + (1-\alpha)h - u)] = F(\alpha\eta + (1-\alpha)h) - (1 - F(\alpha\eta + (1-\alpha)h)).$$

One can therefore write:

(11) $\quad sgn(\alpha\eta + (1-\alpha)h - u_i) = 2F(\alpha\eta + (1-\alpha)h) - 1 + \xi_i,$

where the ξ_i comprise a set of independent but non-identically distributed residuals with bounded variation and zero mean.

In addition, suppose that the market influence weights w_i and the number of individuals are large enough such that $\sum_{i=1}^{N} w_i^2 \approx 0$, interpreted as an approximation that becomes zero as $N \to \infty$, uniformly in N. This can be taken to mean that no single investor has too large a weight in the formation of market prices, so that the strong law of large numbers applies as the number of investors becomes large. It follows from expression (11) that $\sum_{i=1}^{N} w_i \xi_i \approx 0$, almost everywhere, as N becomes large. Thus in the limit as $N \to \infty$,

(12) $\quad h = 2F(\alpha\eta + (1-\alpha)h) - 1.$

With the logistic choice for F, we end up with an implicit equilibrium condition for aggregate executable opinion:

(13) $\quad h = tanh(\mu[\alpha\eta + (1-\alpha)h]).$

Kojima (2000) ends up with the same expression indirectly, as the expectation of the Boltzmann distribution that arises when the individual state vectors $\boldsymbol{h}^k$ settle into an alignment that maximises investor agreement for a given entropy. The present formulation entails fewer arbitrary assumptions and is more readily interpreted in terms of investor behaviour (see below).

Figure 1a,b illustrates the nature of the fixed point in (13), centred for the particular case η=0. Non-zero values of η displace the graph horizontally. Two cases are illustrated, case (a) $\mu(1-\alpha)>1$, and case (b) $\mu(1-\alpha)<1$. For case (a) it will be apparent that three equilibria exist. Those marked h_U , h_L are attracting, while the equilibrium at 0 is repelling (unstable). For case (b) only the one equilibrium exists at the attracting point $h = 0$.

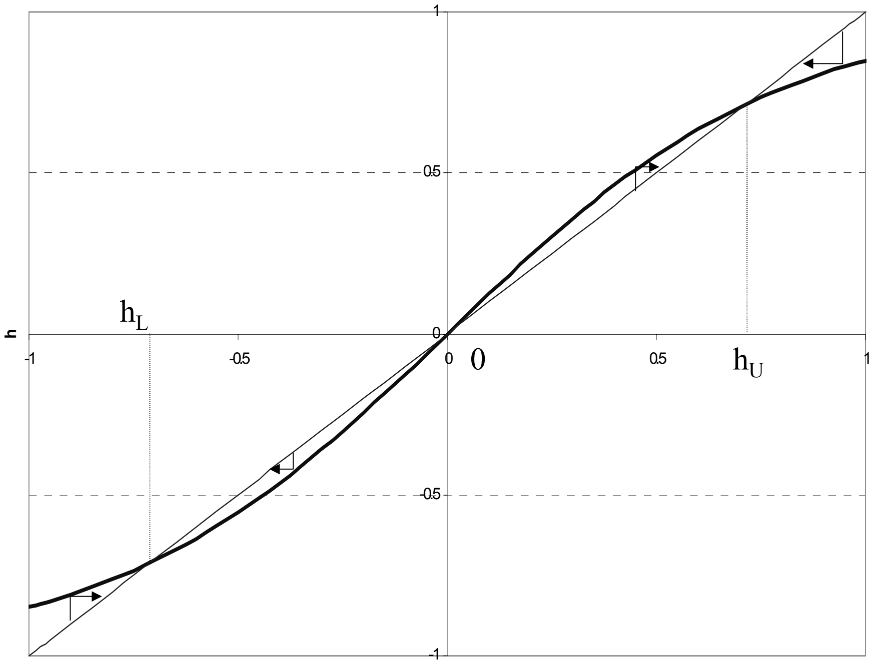

Figure 1: Equilibria - case (a): $\mu(1-\alpha)>1$

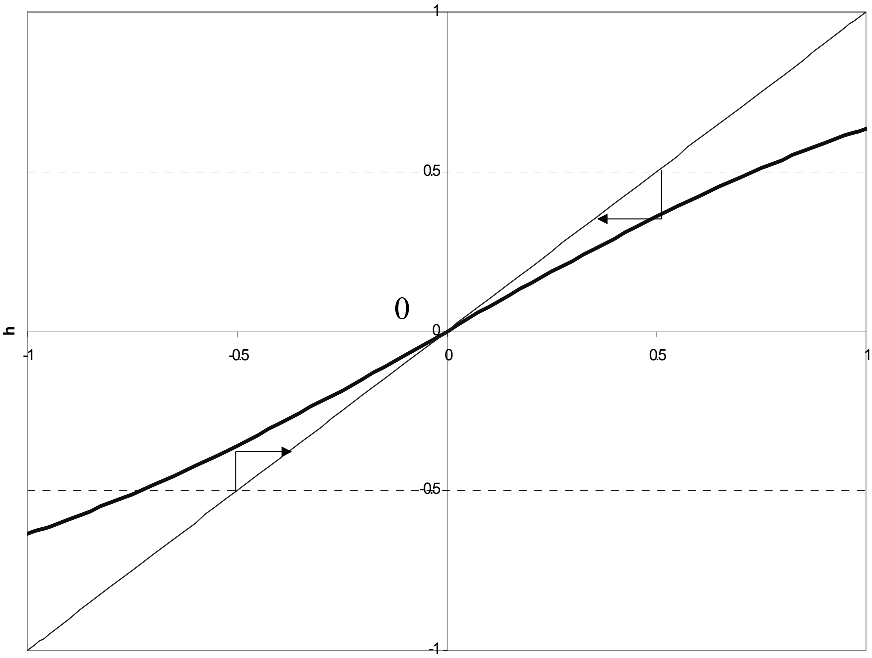

Figure 1b: Equilibrium case (b): $\mu(1-\alpha)<1$

The key parameters are evidently α and μ. The difference 1-α can be taken as an average bandwagon or commonality effect. The parameter μ has concentration (inverse dispersion) references[3]. Stronger herd effect in themselves create a propensity for bubbles, but this can be offset can be offset by higher opinion dispersion or disagreements among investors.

It will be evident from the foregoing that the logistic assumption for *F(u)* is convenient, but not necessary. Instead of assuming that the seller-buyer odds have a particular form, one could start with the stochastic classification (10) and specify a suitable distribution function for the random thresholds u_i . For instance, one could specify a normal distribution. The end result in terms of equilibrium mappings is similar in general form. Note also that as specified above, there is a potential anomaly between the domain of opinions, as the extended unit interval H, and the unbounded support of the logistic or normal distribution functions. Problems of this kind can be resolved by squashing the logistic down to the interval H in a way that preserves as closely as possible the original shape of the density [4].

3.5 Recursive behaviour

If preferred, one could think of the equilibria as the outcomes of a mental process in which given the informational signal η, investors project themselves into the collective mind to obtain a preliminary estimate of the collective estimate *h* to improve their own signal. Knowing that other investors will be thinking in the same way, the aggregate can then be fed into a new iterative round, according to a scheme like

$$h^r = tanh(\mu[\alpha\eta + (1-\alpha)h^{r-1}]) \; ; \; r = 1,2,3,\ldots$$

This amounts to the iterative solution to a game where the collective is the dominant player. It will yield one or other of the attracting points, e.g. h_U or h_L in case (a) above and the attracting point $h = 0$ in case (b).

It is of interest to point out an analogy with Delphic forecasting, a committee based forecasting system in which each member improves his or her own forecast based on knowing what the committee as a whole thinks. Agent *i* revises his private estimate η based on the difference between the committee's view and his own. The results are then aggregated to form the next round committee estimate. On this view, financial markets could be interpreted as Delphic forecasting committees in the large.

3.6 Equilibrium extensions

The structural choice model (10) can be modified to allow for three regimes, including a middle option of staying put, with $h_i = 0$. Thus in place of expression (10) above, one could specify the trinomial tick process as:

$$(14)\qquad \begin{aligned} &h_i = +1 \Leftrightarrow \alpha\eta + (1-\alpha)h > u_i + c \\ &h_i = -1 \Leftrightarrow \alpha\eta + (1-\alpha)h < u_i - c \\ &h_i = 0 \text{ otherwise} . \end{aligned}$$

For any value of u_i there is a dead zone of $\pm c$ such that the individual will take no action. We will refer to c as the 'inertia constant'. Investor regret is a possible reason, the fear of being wrong whatever one does (Bell (1982), (1983), Loomes and Sugden (1982)). Regret can be interpreted as a two sided risk premium relative to doing nothing. Barberis *et al* (2001) note that following short term adverse price movements, investors tend to wait rather than sell stocks at a loss, effectively sitting on the fence until things hopefully become clearer. In place of expression (12), the value of h can be solved from the implicit equation:

$$(15)\qquad h = F(\alpha\eta + (1-\alpha)h - c) - (1 - F(\alpha\eta + (1-\alpha)h + c) .$$

If F is the logistic distribution, then

$$(16)\qquad h = \frac{1}{2}[tanh(\mu(\alpha\eta + (1-\alpha)h - c)) + tanh(\mu(\alpha\eta + (1-\alpha)h + c))] .$$

Figure 2 illustrates for three values of the threshold constant c: low (L), middle (M) and high (H). The case depicted is for η=0; non-zero values simply translate everything horizontally. A range of equilibria is now possible, starting with three as above for $c = 0$, splitting to five for slightly higher values of c (bold curve), and becoming just the one for higher values of c. For the latter, point A is an attracting point from beneath but a repelling point from above, while B is an attracting point from either direction.

Higher values of the inertia constant c increase the stability of the system. To see this note that at $\eta = 0$ and $h = 0$, the derivative with respect to h of the right hand side of equation (16) is given by

$$(17)\qquad \mu(1-\alpha)(1 - tanh^2(\mu c)) ,$$

compared with just $\mu(1-\alpha)$ with $c = 0$. The lower slope means that the system is less prone to developing non-zero fixed points, or in this context, bubbles. Investor doubt has a moderating effect on bubble propensity.

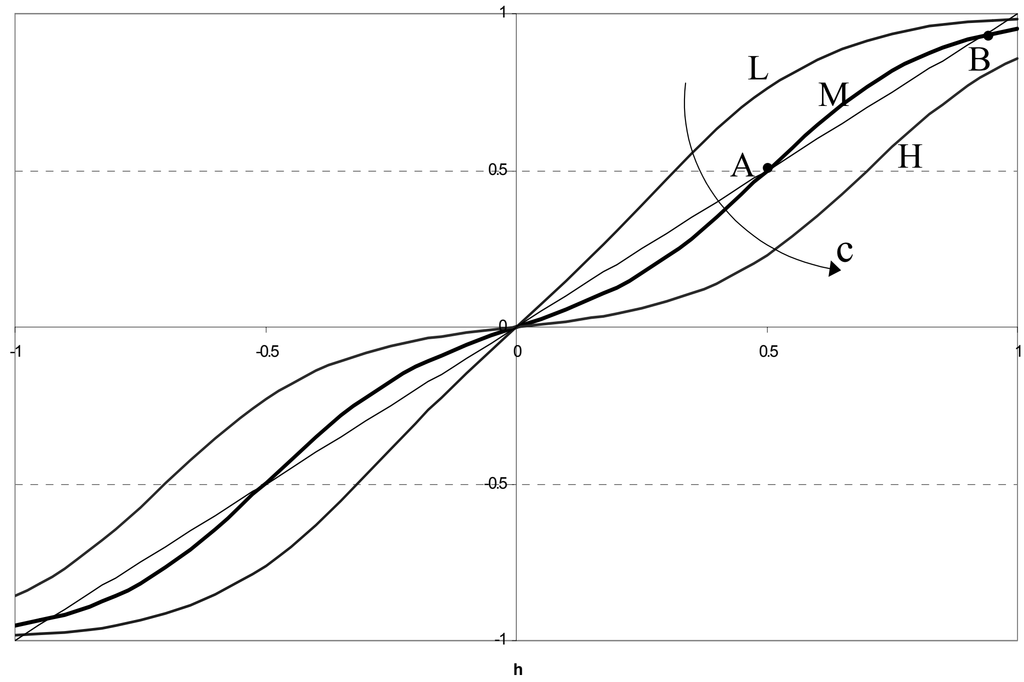

Figure 2: Multiple equilibria in the presence of an intermediate no-action zone

3.7 Levels bubbles

In the outcome space, as in section 2, news about excess returns are automatically also news about bubble levels. In other words suppose we define $i_t^g = g_t - E_{t-1}[g_t]$ as the innovations in the excess returns process and $i_t^x = x_t - E_{t-1}[x_t]$ as the innovations in the levels process. Because $x_t = (1+r)x_{t-1} + g_t$, it follows that $i_t^x = i_t^g$. Thus we could imagine agents thinking in terms of any given information signal in opinion space as applying equally to levels. Equilibrium could be defined in terms of current bubble levels, so that equilibrium conditions like (13) or (16) of section III might apply to opinions about levels, rather than rates of growth. Pursuing the analogy further, one could end up with bifurcations in terms of bubble size (x) instead of bubble growth (g). The market would get locked into stationary states, where it is overvalued in terms of fundamentals, but things do not get worse.

It is certainly possible to proceed in this way, but there is a difficulty. The economic rationale for the earlier equilibrium conditions was derived in terms of buy-sell-hold dichotomies as in (10) and (14). Decisions of this kind are more naturally made in the context of expected excess returns rather than expected levels. The latter have no real economic meaning. Expected rises in x trigger a buy decision only to the extent that they are large enough to exceed the cost of capital benchmark *r*. But in that case agents are really thinking in terms of excess returns *g* and one is back to the growth framework of section III.

Thus at this stage we note only the possibility of applying a similar analysis to end up with scenarios in which the bubbles might grow to a stationary point and then simply remain there, apart from minor fluctuations. The need would be to rationalise such equilibria in terms of the economics of investment decisionmaking.

3.8 The outcome in price space

Once the equilibrium opinion h is determined, the observed excess return for period *t* can be obtained via the homeomorphism $g_t = \varphi(h_t)$ assumed to connect executable opinion with price or price-based outcomes, in the case an excess return. The price that generated this return is

(18) $P_t = (1+r)P_{t-1} - m_t + P_{t-1}g_t$.

Using $m_t = E[m_t \mid \Im^m_{t-1}] + \sigma_m \varepsilon_{mt}$, it will be apparent that the dynamics expressed by (18) is homologous with the noisy REE model (5) of section 2. The difference lies in the more complex behaviour of the bubble-like component g_t. This is no longer necessarily white noise, but acquires a temporal structure, as the next section makes clear.

IV Equilibrium behaviour

This section examines the way that the collective opinion equilibrium depends on the informational input η. For this purpose the two regime model (13) will be used. Thus the objective is to examine the solutions for *h* as η varies to the implicit equation:

(19) $h = tanh(\mu[\alpha\eta + (1-\alpha)h])$.

4.1 Bifurcations

The required dependence *h(η)* turns out to be a point to set mapping rather than a proper function, and it is best to examine its behaviour in terms of the inverse $\pi : h \to \eta$, which is a proper function over the domain H = [-1,1] . Inverting equation (19), we obtain

(20) $\eta = -(\frac{1-\alpha}{\alpha})h - \frac{1}{2\mu\alpha}log(\frac{1+h}{1-h})$,

with

(21) $$\frac{d\eta}{dh} = \frac{1}{\mu\alpha}[\frac{1}{1-h^2} - \mu(1-\alpha)] .$$

$$\frac{d^2\eta}{dh^2} = \frac{1}{\mu\alpha}\frac{2h}{(1-h^2)^2} .$$

As in section 3.3, there are two cases to consider:

Case (a): $\mu(1-\alpha) > 1$. From (21), the slope is evidently locally negative around $h = 0$ and the function changes from concave to convex with the sign of *h*. Thus if $\mu(1-\alpha) > 1$ there is evidently backward bending behaviour.

Case (b): $0 < \mu(1-\alpha) \le 1$. In this case there is no backward bending and the function remain monotonic with positive slope, even though the convexity changes through the origin.

Figure 3 depicts the map of the function (20) with the axes interchanged, so that *h* is on the vertical axis. For case (b) there is a unique value of *h* for any value of η. But for case (a) there is a zone around η = 0 in which for any value of η there are three possible values of h.

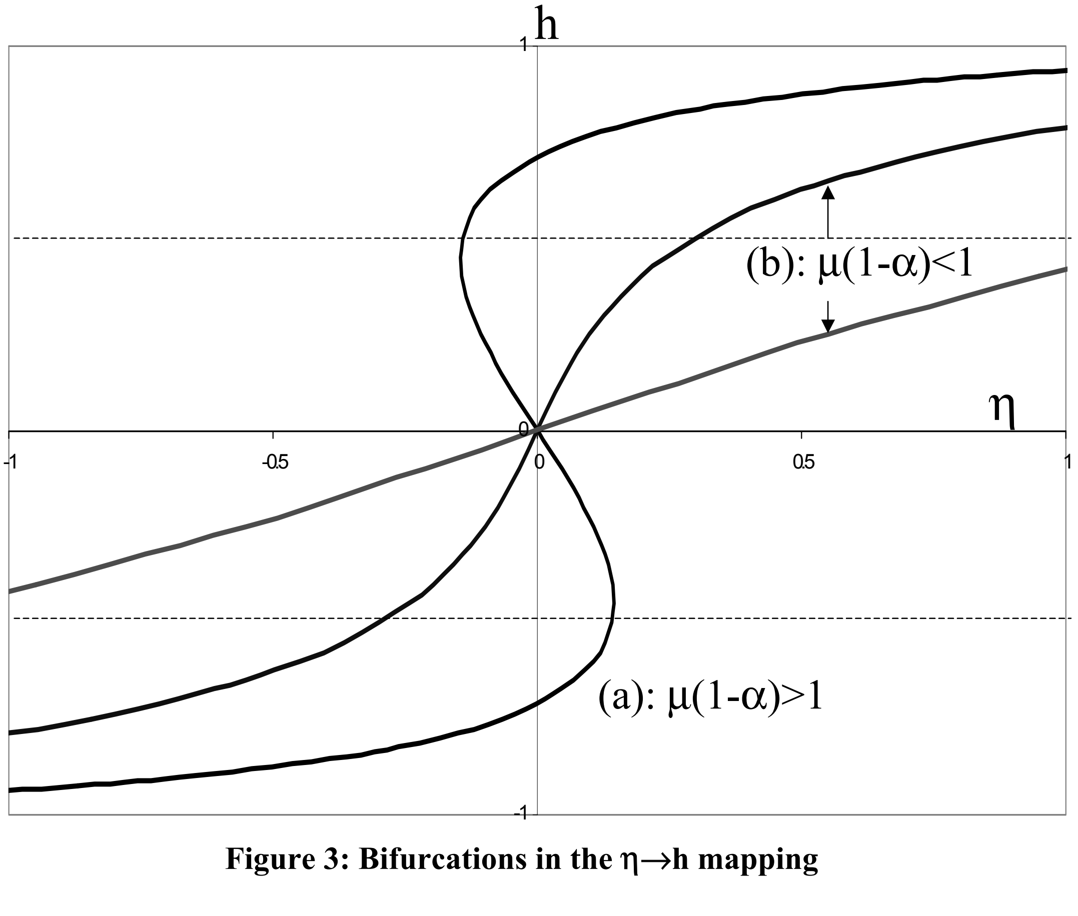

Figure 3: Bifurcations in the η→h mapping

Figure 4 takes a closer look at case (a), with $\mu(1-\alpha) > 1$ There are three zones to consider. Zones R_A and R_B consist of attracting equilibria. Smaller changes in η will see the system stay in one or other of these zones, except where $\eta < \eta_A$ or $\eta > \eta_B$, in which case the solution will jump to the complementary attracting set, upper to lower or vice versa. The zone marked R_C indicates the set of repelling equilibria. If the system finds itself in this zone, any change in η will cause the equilibrium to change radically to one of the attracting sets, as indicate by the direction of the arrows.

If inertia zones exist as in section 3.6 above, then the bifurcation map becomes more complex. For intermediate levels of the inertia constant *c* (corresponding to curve M in figure 3), the curve intersects the vertical axis at five points in the region near η=0, rather than just three. The intermediate points (c.f. point A in figure 2) are attracting from below but repelling from above.

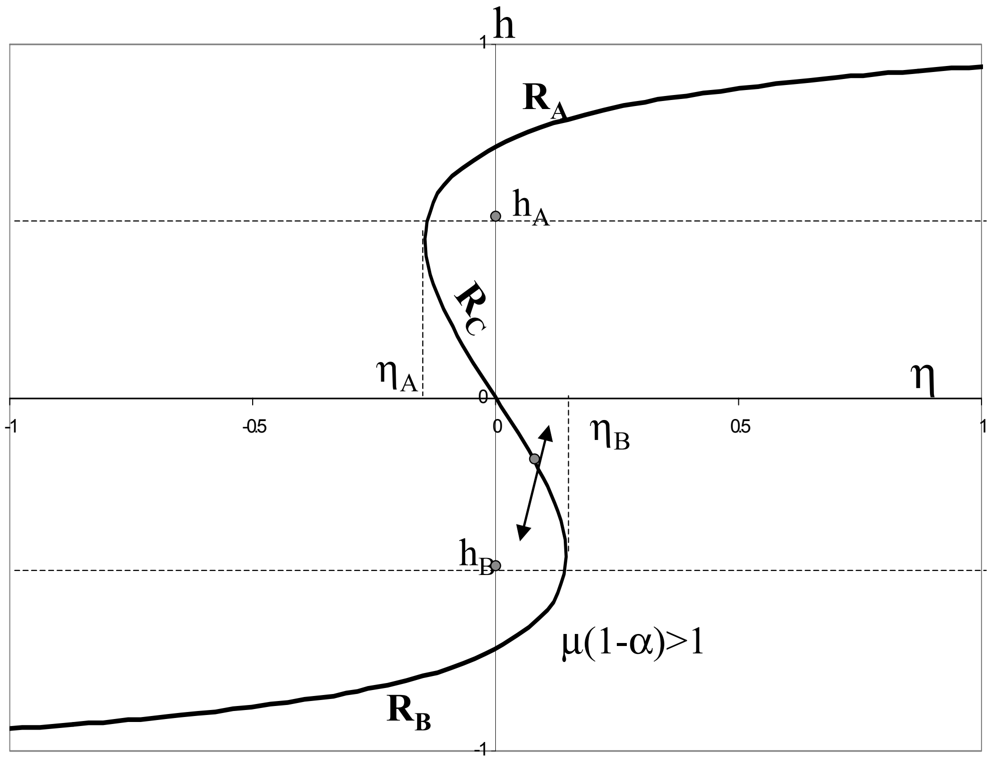

Figure 4: Attracting and repelling zones

4.2 Path dependence and serial correlation

Preceding development suggests that opinion bubbles have a certain amount of inertia. If the system has been depressed in a negative bubble state R_B, it will tend to remain there even if the informational signal starts looking more promising i.e. η becomes positive. However after a certain critical point is reached, marked in figure 4 as η_B, aggregate opinion will suddenly swing around to the upper regime R_A and a positive bubble establishes itself. Likewise, if the system is initially in a positive bubble state R_A it will tend to remain there and it will take a larger negative movement to knock it away from this zone to a negative bubble state R_B.

A feature of interest is that even if the informational signal η is serially uncorrelated, the output state h is now locally serially correlated. This is at odds with the usual econometrics of serial correlation, which leads us to expect that any continuous function of a serially uncorrelated variable should itself be serially uncorrelated. The general proposition fails in the present context because the mapping that relates η to h is not a function, let alone a continuous one. It is instead a point to set mapping, and it is this that creates the path dependence. The following result makes the point more formally.

Theorem1

Suppose that

(i) *The informational signal process (η_t) is serially independent. For each time t, a density $f_t(\eta)$ exists, is continuous, and is symmetric about $\eta = 0$.*

(ii) *The parameters $\mu > 0$ and $\alpha > 0$ are constant over time, and $\mu(1-\alpha) > 1$.*

Then:

(a) *The distribution of h_t is path dependent and exhibits local serial dependence.*

(b) *The values of h_t fall with probability 1 into one or the other of two disjoint regimes: upper (R_A) with positive values, and lower (R_B) with negative values.*

(c) $E[h_t \mid h_{t-1} \in R_A] > 0$ *and* $E[h_t \mid h_{t-1} \in R_B] = -E[h_t \mid h_{t-1} \in R_A] < 0$.

Proof: See the Appendix.

4.3 Noise or bubbles?

The preceding development allows us to make a mathematical distinction between noise and bubbles. Starting with the two regime model:

(a) If $\mu(1-\alpha) \le 1$ then we can write $h_t = \pi(\eta_t)$ as a proper C^1 function, from which the equilibrium opinion values form a serially uncorrelated process. In turn, the homeomorphism between opinions and excess returns space will ensure that the same time series property holds for excess returns $g_t = \varphi(h_t)$. This again gives us the rational noise model of section 2.

(b) If on the other hand, $\mu(1-\alpha)>1$ then from theorem 1, episodes will exist wherein opinions and excess returns are not centred at zero and exhibit local serial correlation. If parameters remain unchanged while the informational signal η fluctuates, the only exit from such 'trapped' states is to the contrary state (deflation bubble to inflation bubble and vice versa). Even if $\eta=0$ in any period, so that no fresh informational stimulus appears, then the excess returns can remain positive or negative depending upon the state in the previous period. In this sense therefore bubbles do not necessarily depend upon the flow of information, once they are established.

4.4 Phase transitions: growth and collapse

Foregoing model development suggest that the ability of bubbles to develop or to sustain themselves depends on some key parameters:

(a) The dispersion of opinion (σ_u^2); or its inverse, the opinion concentration (μ).

(b) The average social reflexivity (loosely, herding propensity), represented in the model by (1-α).

(c) The existence of zones of indecision, or investor doubt, represented by the dead zone interval *c* in the three-regime model.

A general rule is that in the extended system, covering both two and three regime variants (10) and (14), bubbles can develop in response to information flows if

(19) $\mu(1-\alpha)(1-\tanh^2(\mu c))>1$.

A useful way to think of phase transitions is in terms of local changes in one or other of the above parameters. Expansion phases tend to develop when opinions have settled into concordance (μ higher); everybody knows it (α lower); and investors are more willing to jump off the fence (*c* smaller) and take a position. Collectively, this corresponds to herding behaviour.

Or, suppose the system is currently in an expansionary bubble state. As prices soar, some investors may come to feel uncertain about the directional sustainability, even if they feel there is no immediate downside threat. This could be manifested in a decision to sit tight - a watching brief - rather than to buy further units. Doubt means that the inertia parameter rises, resulting in a shift downwards from curve H to curve M in figure 2.

A first manifestation of this is a relatively small decline in the excess growth rate or price level to point B. Observing that the market has paused will give investors at large occasion to consider their own opinions, and in particular their assessment as to how other investors will interpret the pause. In turn, uncertainty as to the latter may induce investors to rely more on their own instincts so that the average α rises, i.e. the herd instinct (1- α) falls. Or the uncertainty as to the market opinion may differ between investors, which means that the opinion concentration parameter μ falls (see point (c) of section 3.3). Expression (20) is violated, and there is a phase transition (e.g. figure 1a to figure 1b). The market pause becomes a full scale collapse.

V Concluding remarks

The main concern of the paper has been to offer a structured model of asset market bubbles, as matter of insight rather than application. However theory should never take place in a vacuum, and it is may be useful to briefly explore some implications. The first concerns the modelling of returns with bubble episodes, an aspect for interest for option pricing and related position taking or risk management instruments. The second looks at a broader more policy oriented picture in terms of market or systemic stability.

6.1 Empirical modelling

The models of noise and bubbles point to a regime structure for asset returns or risk adjusted excess returns. Empirical modelling would require specification for the transition probabilities between regimes and some attention to mean dependences inside each of them. Findings of relevance can be summarised as follows.

(a) The system falls into distinct phases or states, some of them excited bubble regimes and others not bubbly (the ground states). For the basic dichotomous choice model (no investor inertia), there are two excited states and one ground state.

(b) Transition probabilities between states can depend upon (i) external information signals, (ii) current and past values of return outcomes, and (iii) changes in parameter values. The precise specification will depend upon the state pairs involved. Thus a transition from an upper positive bubble to a lower one might occur if the information signal reaches a critical adverse point (η_A in figure 4). A transition from an upper positive bubble to the ground state might take place if underlying parameter values change. Empirically, such a transition could be proxied in terms of observed asset returns or short memory cumulated past returns. Thus the more the bubble has grown, the more individuals are liable to change behaviour, and hence the higher the probability of collapse and return to the ground state.

(c) All states exhibit some market noise; background static, as it were. For much of the time this stays within more or less normal bounds. Volatility clustering is likely to occur in and around points of transition between phases.

(d) There is mean dependence of excess returns inside some states but not others. The theorem of section IV suggests that given the system is in the upper bubble state, the conditional expectation of g_t will depend upon g_{t-1}; similarly for the lower (negative) state.

6.2 Market stability and investor heterogeneity

Market instability is commonly rooted home to 'short termism'. This has earlier been characterised as an excessive preoccupation with end of period price and a short holding horizon, rather than the entire future stream of dividends that the end of period price theoretically represents. Translated to the current context and framework, a sensitivity of this kind would imply that the market is more acutely tuned to news, and by that token more liable to infection by news that is in reality extraneous.

A further concern is that an important class of agents that might exhibit such behaviour, namely the hedge funds, now have a considerable capacity to collectively move markets. It is hard to estimate the number and size of hedge funds worldwide; see Fung and Hsieh (2006), the source of comments that follow. However, the stylised facts seem reasonably clear. There are probably about 10,000 of them, collectively controlling assets of about USD10 trillion, and the number is growing, despite an attrition rate of 20-30% a year. They occupy a number of distinct niches as to preferred investment types and market niches, but a common denominator is that they show little preoccupation with risk. Indeed, managerial rewards are structured to encompass out-performance, rather than compliance or benchmarking with respect to with any model of risk-adjusted returns or risk management in general. However, the very number of them mitigates to some extent their ability to independently sway markets, at least relative to the days of Long Term Capital Management (see e.g. Dunbar (2000), Lowenstein (2001)). Currently only about 19% of hedge funds are estimated to have assets greater than $200m.

In terms of the modelling of the present paper, one would expect hedge fund managers not only to have low buy-sell inertia, but also to be acutely tuned into the market as a whole. A key issue would then be whether they are herd followers, or on the other hand more motivated to form and act on their own opinions - heterogeneity in the α_i is good for market stability. A possible answer arises out of managerial incentive structures. These are heavily oriented towards very generous absolute returns, often 10-25% of the growth in fund value above high water mark, with minimal or only administrative compensation otherwise. This might generate cyclical herding behaviour on the part of informed hedge funds. One might expect that in the earlier stages of a potential bubble, hedge funds would pick up the market mood and jump on board. However as the bubble starts to expand, they might be readier to abandon ship, backing their own independent opinions that this is simply a bubble and liable to collapse. The disagreement might grow (μ decline), ultimately past the critical point that precipitates a market collapse. This suggests that if there are many of them, but all motivated to develop independent opinions, then hedge funds are not in themselves necessarily destabilising.

Perhaps more of a concern arises from recent structural changes in the hedge fund industry, partly as a result of prudential developments in the finance industry. Because of the difficulty of accessing debt funds as private partnerships, more hedge funds are now structuring themselves as formal investment funds, with trust deeds and more public reporting responsibilities. And mainline institutions – insurance companies, mainline unit trusts, superannuation funds, and even banks - are forming associated or subsidiary hedge funds to accommodate a demand for portfolio enhancement. Managerial incentive structures for such operations are likely to be less geared towards independence and disagreement. There may also be a question as to whether sufficient investment and judgemental expertise can be spread so thinly around so many institutions. One might expect accentuated herding behaviour in such circumstances. When everybody piles on board the boat, and on the same side, a phase transition is likely to occur.

Appendix

Proof of the theorem in section IV

Suppose that

(iii) *The informational signal process (η_t) is serially independent. For each time t, a density $f_t(\eta)$ exists, is continuous, and is symmetric about $\eta = 0$.*

(iv) *The parameters $\mu > 0$ and $\alpha > 0$ are constant over time, and $\mu(1-\alpha) > 1$.*

Then:

(d) *The distribution of h_t is path dependent and exhibits local serial dependence;*

(e) *The values of h_t fall with probability 1 into one or the other of two disjoint regimes, upper (R_A) with positive values, and lower (R_B) with negative values.*

(f) $E[h_t \mid h_{t-1} \in R_A] > 0$ *and* $E[h_t \mid h_{t-1} \in R_B] = -E[h_t \mid h_{t-1} \in R_A] < 0$.

Proof

Figure 4 of the text reproduces the equilibrium mapping $\pi : \eta \to h$ with some additional notation delimiting the following zones. With reference to the figure, define subspaces as follows:

$R_A = \{ h \geq h_A, \eta \geq \eta_A \};$

$R_B = \{ h \leq h_B, \eta \leq \eta_B \};$ and

$R_C = \Omega - R_A - R_B$

where $\Omega = H \times H$.

The graph of upper attracting equilibria is a subset of R_A , that of lower contracting equilibria is contained in R_B, and the set of repelling equilibria is in R_C. Restricted to the domains to R_A and R_B , the mappings $\pi : \eta \to h$ become proper functions defined respectively by:

(A1) $h = \pi_A(\eta); \quad h \geq h_A, \eta \geq \eta_A$

(A2) $h = \pi_B(\eta); \quad h \leq h_B, \eta \leq \eta_B$.

Suppose that at time t-1, the system is in the lower attracting regime B. (Loosely we will sometimes say that $h_{t-1} \in R_B$, etc.). When at time *t* a new value of η is drawn, then with probability 1, the system can either stay in R_B (if $\eta_t \leq \eta_B$) or exits to the other attracting regime R_A (if $\eta_t > \eta_B$). It will exit to the repelling regime R_C with probability zero.

Thus conditional on $h_{t-1} \in R_B$, we must have

$h_t = \pi_B(\eta_t); \quad \eta_t \leq \eta_B$

$h_t = \pi_A(\eta_t); \quad \eta_t > \eta_B$.

One can therefore write:

$$E[h_t \mid h_{t-1} \in R_B) = \int_{-\infty}^{\eta_B} \pi_B(\eta) f(\eta) d\eta + \int_{\eta_B}^{\infty} \pi_A(\eta) f(\eta) d\eta .$$

However, $\pi_A(\eta) = -\pi_B(-\eta)$, $\eta_A = -\eta_B$, and $f(\eta) = f(-\eta)$. It follows that

(A3) $$E[h_t \mid h_{t-1} \in R_B) = \int_{-\infty}^{\eta_B} \pi_B(\eta) f(\eta) d\eta - \int_{-\infty}^{\eta_A} \pi_B(\eta) f(\eta) d\eta .$$

Now define versions of π_A, π_B with censored domains by:

$$\hat{\pi}_B(\eta) = \pi_B(\eta) \Leftrightarrow \eta_A \le \eta \le \eta_B$$
$$= 0, \ \ otherwise$$

$$\hat{\pi}_A(\eta) = \pi_A(\eta) \Leftrightarrow \eta_A \le \eta \le \eta_B$$
$$= 0, \ \ otherwise$$

Then from (A3),

$$E[h_t \mid h_{t-1} \in R_B) = \int_{\eta_A}^{\eta_B} \hat{\pi}_B(\eta) f(\eta) d\eta < 0 \ .$$

Similarly, suppose the system is initially in regime A. Using $\hat{\pi}_A(\eta) = -\hat{\pi}_B(-\eta)$, we obtain:

$$\begin{aligned} E[h_t \mid h_{t-1} \in R_A) &= \int_{\eta_A}^{\eta_B} \hat{\pi}_A(\eta) f(\eta) d\eta \\ &= -\int_{\eta_A}^{\eta_B} \hat{\pi}_B(\eta) f(\eta) d\eta \\ &> 0. \end{aligned}$$

References

Barberis, N, M. Huang and T. Santos (2001) 'Prospect theory and asset prices', *Quarterly Journal of Economics*, 116, 1-33.

Blanchard, O.J. and M. Watson (1982) 'Bubbles, rational expectations and financial markets', in P. Wachtel (ed.) *Crisis in the economic and financial structure*, Lexington, Mass.: Lexington.

Bell, David E. (1982). Regret in decision making under uncertainty, *Operations Research*, 30(5), 961-981.

Bell, D. E. (1983). Risk premiums for decision regret. *Management Science*, 29(10), 1156-1166.

Bowden, R.J. (1988) *Statistical games and human affairs*, New York: Cambridge University Press.

Devaney R.L. (1989) *An introduction to chaotic dynamical systems*, 2nd ed., Redwood City CA.: Addison-Wesley.

Dunbar, N. (2000) *Investing money: The story of Long Term Capital Management and the legends behind it*, Chichester: Wiley.

Evans, G. (1983) 'The stability of rational expectations in macroeconomic models', in R. Frydman and E.S. Phelps (eds.) *Individual forecasting and aggregate outcomes*: *"Rational expectations" re-examined'*, New York: Cambridge University Press

Fung, W.K.H. and D.A. Hsieh (2006) 'Hedge funds: An industry in its adolescence', *Economic Review,* Federal Reserve Bank of Atlanta, 91, fourth quarter, forthcoming.

Garber, P.M. (2000) *Famous first bubbles: The fundamentals of early manias*, Cambridge. Mass: MIT Press.

Kaizoji, T. (2000) 'Speculative bubbles and crashes in stock markets: An interacting-agent model of speculative activity', *Physica A*, 287, 493-506.

Kindleberger, C.P. (1996) *Manias, panics and crashes: A history of financial crises*, 3rd ed., New York: Wiley.

Loomes, G. & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92, 805-824.

Lowenstein, R. (2001) *When genius failed; The rise and fall of Long Term Capital Management*, London: Fourth Estate.

Lucas, R.E. (1972) Expectations and the neutrality of money', *Journal of Economic Theory*, 4, 103-124.

Lux, T. (1995) 'Herd behaviour, bubbles and crashes', *Economic Journal*, 105, 881-896.

Taylor. J.B. (1977) 'Conditions for unique solutions in macroeconomic models with rational expectations, *Econometrica*, 45, 1377-85

Flood, R.P. and P.M. Garber (1980) 'Market fundamentals versus price level bubbles: the first tests', *Journal of Political Economy*, 88, 745-770.

Gourieroux, C, J.J. Laffont, and A. Monfort (1982) 'Rational expectations in dynamic linear models: Analysis of the solutions', *Econometrica*, 50, 409-425.

Hamilton, J.D. (1989) 'A new approach to the economic analysis of non-stationary time-series and the business cycle', *Econometrica*, 57, 357-384.

Keynes, J.M. (1936) *The general theory of employment, interest and money*, London: Macmillan.

Lim. G.C. and V.L. Martin (1998) 'Forecasting jumps in exchange rates', working paper, Dept. of Economics, University of Melbourne.

Lim, G.C., Martin, V.L., and Teo, L.E. (1998) 'Endogenous jumping and asset price dynamics', *Macroeconomic Dynamics*, 2, 1-25.

McCallum, B. J. (1977) 'The role of speculation in the Canadian forward exchange market: Some estimates assuming rational expectations', *Review of Economics and Statistics*, 9,145-151

Peria, M.S.M. (1999) 'A regime-switching approach to studying speculative attacks', Policy Research working paper #2132, Washington DC: World Bank.

Pericoli, M. and M. Massimo (2001) 'A primer on financial contagion', Banca D'Italia, #407, June.

Phelps, E.S. (1983) 'The trouble with rational expectations', in R. Frydman and E.S. Phelps (eds.) *Individual forecasting and aggregate outcomes: "Rational expectations" re-examined'*, New York: Cambridge University Press.

Taylor, J.B. (1977) 'Conditions for unique solutions in stochastic macroeconomic models with rational expectations, *Econometrica*, 45, 1377-85.

Topol, R. (1991) 'Bubbles and volatility of stock prices: Effects of mimetic contagion', *Economic Journal*, 101, 786-800.

Endnotes

.

---

[1] An additive decomposition is $P_t = p_{mt} + z_t$, where p_{mt} is the fundamental process defined as in equation (1) or (2) and $z_t = az_{t-1} + P_{t-1}(x_t - ax_{t-1}) = az_{t-1} + P_{t-1}i_t$.

[2]Kojima calls the metric a ' disagreement', with a negative sign, but it has the format of an angle between two vectors and seems therefore to be a measure of similarity rather than dissimilarity. The intention is to end up with an energy measure and thence to apply an energy -minimising statistical distribution as in the derivation of the Boltzmann distribution.

[3] Note that the concentration parameter might itself be endogenous, depending upon *h*. Thus suppose opinion heterogeneity exists for α, η as in case (a) and (b) of section 3.3 above. In such a case

$$u_i = (\alpha_i - \alpha)(h - \eta) - \alpha(\eta_i - \eta) - (\alpha_i - \alpha)(\eta_i - \eta)$$

so that if α_i, η_i are statistically independent across individuals *i*, then

$$Var(u_i) = (h - \eta)^2 \sigma_\alpha^2 + \alpha^2 \sigma_{\eta_i}^2 + \sigma_\alpha^2 \sigma_{\eta_i}^2 ,$$

where $\sigma_\alpha^2, \sigma_{\eta_i}^2$ refer to the variances across individuals at any given point in time. This suggests that dispersion parameters such as μ or σ_u may themselves be of the character of state variables. States in which assessed public opinion *h* differs significantly from the given private signal η are potentially high dispersion states.

[4] If desired the range of u can be made to conform with the extended interval [-1,1]. For instance if *u* is zero mean logistic with dispersion parameter μ, then the transformation $y = tanh(\theta u)$ with $\theta < \mu$ will create a density of very similar form to the logistic, but constrained to [-1,1] with density formula given by

$$f(y) = \frac{1}{1 + \left(\frac{1-y}{1-y}\right)^{\mu/\theta}} \times \left[1 - \frac{1}{1 + \left(\frac{1-y}{1-y}\right)^{\mu/\theta}}\right] \times \frac{1}{\theta(1-y^2)}$$