

IMPLAUSIBLE EQUILIBRIUM SOLUTIONS IN ECONOMICS AND FINANCE

HARBIR LAMBA¹

Abstract

When modelling the aggregate behavior of a population over long periods of time the standard approach is to consider the system as always being in equilibrium — using averaging procedures based upon assumptions of rationality, utility-maximization and a high degree of independence amongst the agents. These analytically tractable models are fully determined by a small number of variables and the future evolution of the system is decoupled from its past. Previous studies [Akerlof and Yellen 1985, Scharfstein and Stein 1990] suggested that allowing a subset of the agents to be non-maximizing might shift the equilibrium solution but a shortcoming in those analyses prevented the possibility of non-equilibrium solutions.

We show how equilibrium assumptions and solutions can be ‘stress-tested’ by embedding the models within frameworks capable of far-from-equilibrium dynamics. This procedure is applied to Brownian motion (Random walk) asset pricing which serves as our prototypical equilibrium model. The introduction of a simple, yet plausible, form of non-maximizing coupling between the agents results in endogenous fluctuations that destabilize the equilibrium solution. These fluctuations typically consist of long periods of increasingly severe mispricing followed by rapid, unannounced, corrections. Importantly, this occurs at realistic parameter values and is consistent with the observed price dynamics of real markets. Furthermore the mispricing phase may be mistaken for a trend in the fundamentals of the equilibrium model. We argue that the ‘boom-bust’ nature of the destabilizing endogenous dynamics is likely to be prevalent in economic systems at all scales.

¹Department of Mathematical Sciences, George Mason University, MS 3F2, 4400 University Drive, Fairfax, VA 22030 USA. E-mail: hlamba@gmu.edu .

JEL Classification: C62; D01; D03; D50; D53; E32; G01

Keywords: Far-from-equilibrium; Instability; Hysteresis; Behavioral economics; Asset pricing; Financial Crises; Animal spirits; Business Cycle

I Introduction

“We can model the euphoria and the fear stage of the business cycle. Their parameters are quite different. We have never successfully modelled the transition from euphoria to fear.”

— Alan Greenspan, *Financial Times*, March 27th 2009.

All the mainstream schools of economic thought and, in turn, modern finance have been strongly motivated by the mathematical elegance and predictive success of Newtonian mechanics and statistical physics. This has resulted in mathematical models constructed so that the solution is a stable, usually unique, equilibrium described by a small number of variables and whose future evolution is independent of its past states. Indeed the notion of equilibrium has become a unifying concept in modern economics although its precise meaning varies with context. Here we shall consider those equilibrium models that attempt to describe systems involving large numbers of interacting agents over significant periods of time. Some examples, discussed briefly below, include the supply-demand curves of microeconomics, the Brownian motion description of asset prices, and the DSGE models used in macroeconomics and central banking.

In the physical sciences, equilibria are most often studied as the limiting values of dynamical systems where the short-term transient behavior may not even be of interest. However, economic systems have no such end-state [Robinson 1974] and so the only remaining (legitimate) justification for equilibrium models requires the assumption that equilibrating processes operate (and dominate) over timescales shorter than or comparable to any exogenous changes. This implies the system is

at all times close to being in equilibrium and is referred to as a *quasi-equilibrium process* in statistical physics.

In the context of large systems of interacting agents we shall define an *equilibrium model* as one where there is an assumed absence (or near-perfect cancellation) of endogenous dynamics so that the state of the system at any moment is determined solely by the current values of a small number of aggregate or exogenous variables. In particular it precludes any dependence upon prior states of the system. This assumption is valid for the case of gas molecules in a closed box that are reacting to slow changes in temperature — the pressure is determined by the Ideal Gas Law and the *current* temperature.² This happy state of affairs is guaranteed by the Laws of Thermodynamics that provide both a unique statistical equilibrium and a physical mechanism for reaching it.

It is useful to extend the analogy with thermal equilibrium in two ways. Firstly, if the exogenous changes are too large or rapid, compared to the strength of the equilibrating processes, then the box will move a long way from equilibrium with its surroundings. However, there is a second potential source of disequilibrium in economic systems that can be compared to a box of special molecules [Buchanan 2007] that do not obey the First and Second Laws of Thermodynamics (Conservation of Energy and Increasing Entropy respectively). These particles can spontaneously speed up or slow down, reveal a temporary preference for one side of the box over the other, remember where they have been, or choose to follow other particles. For a boxful of such molecules there also exist *disequilibrating* forces that, when strong enough, can destabilize the equilibrium by generating significant endogenous dynamics. We shall use a simple model, with a small number of key observable parameters, to try and quantify the battle between equilibrating/disequilibrating/exogenous forces in a typical financial market.

Equilibrium models have the virtue of being analytically tractable but this very tractability places severe constraints on the nature of the solutions. This

²At a fixed volume, the pressure is proportional to the temperature.

means that real-world phenomena that do not comfortably fit into the equilibrium paradigm require ad hoc modelling adjustments that are often the cause of much heated debate between different economic schools of thought.

As our introductory example we consider the supply-demand curves used for price determination. The price value at the intersection of such curves (representing differentiable functions) is taken to be the equilibrium price that balances the opposing forces of supply and demand. In the simplest model, if a shock occurs to either curve and is then removed the price should very quickly settle back to its pre-shock value (the critiquing of this model has a long and distinguished history but we shall not reproduce it here). In other words an instantaneously-acting equilibrating force (or ‘invisible hand’) restores the previous equilibrium and no memory of the shock remains — this picture fits perfectly within our definition of an equilibrium model.

In Section V we shall argue that barriers to instantaneous equilibration, including various forms of hysteresis, invalidate the overly simplified notion of ‘balance of forces’ underlying the model. The entirety of producing and consuming agents can be regarded as forming an endogenous structure with multiple feasible states that can absorb some of the forces and also cause endogenous dynamics. In this scenario, after a shock is removed the system will not necessarily return to its previous state and the system retains memory. Note that this cannot be corrected by replacing the differentiable functions defining the supply-demand curves by ones that are merely continuous or even discontinuous. Instead one has to recognise that there are no longer single-valued functions or curves — instead there are supply and demand ‘smudges’ with the current price determined by the historical record but lying within some achievable intersecting set.

Our second example concerns financial markets and the wild, occasionally devastating, fluctuations that can occur. The Brownian motion (Random Walk) description of asset pricing is an equilibrium model asserting that exogenous Brownian information streams are correctly translated into Brownian price changes.

Brownian processes have no memory and their increments are independent and Gaussian-distributed — in particular the probability of extreme increments (ie. in the tails of the increment distribution) decays exponentially fast as a function of their magnitude. However these properties are flatly contradicted by the two most significant ‘stylized facts’ [Mantegna and Stanley 2000] of financial markets. These are *volatility clustering* where periods of excess volatility exist and *fat tails* where the observed frequency of large price moves can be an order of magnitude of orders of magnitude larger than predicted by the equilibrium model. The consequences of this latter discrepancy cannot be overstated. At the time of writing the global banking system is still being bailed out from bad loans and investments (and levels of leverage) that were justified using equilibrium-based risk-assessment models such as VaR (Value-at-Risk).

A third example concerns the rational expectations DSGE (Dynamic Stochastic General Equilibrium) models of macroeconomics.³ They fail to explain anomalies in important variables such as the output gap that display both fat tails [Fagiolo, Napoletano, Piazza and Roventini 2009] and the periodicities in ‘animal spirits’ referred to as the Business Cycle. This is often explained away by positing correlations and non-Gaussian statistics within the exogenous shocks but there is no compelling reason for believing this is the case. There is a growing literature on agent-based macroeconomics (see [Tesfatsion 2006, LeBaron and Tesfatsion 2008]) and an interesting recent study [De Grauwe 2012] demonstrates that both effects can be generated endogenously if one replaces the usual rationality assumptions on the agents with bounded rationality plus limited learning capabilities. This is very much in the spirit of what will be presented here although we shall go further and use a fully agent-based (bottom-up) model to test the stability of an equilibrium

³These models use averaged variables and have unique solutions but often incorporate delays into the system variables since macroeconomic policy changes cannot be expected to take effect immediately. However such delay equations can always be rewritten without the delay by introducing extra variables.

solution.

Economic systems are aggregations of many heterogeneous agents and the a priori requirement of history-independent equilibrium solutions has profound consequences (see [Kirman 1992]). Most of these models can be derived by assuming that agents display enough homogeneity to be ‘averaged’ or scaled-up and replaced by a representative agent who is memoryless, perfectly rational (usually in the sense of maximizing some hypothetical utility function) and correct about future expectations. This averaging procedure, when applied to expectations, is known as the Rational Expectations Hypothesis and implies that while individual agents’ expectations may be wrong they all use all available information and, on average, agree with the expectations assumption being used.⁴

In this paper we use a heterogeneous agent model of a financial market to demonstrate as starkly as possible that there are significant, systemic, real-world effects that must not be averaged away. As the magnitude of these effects is increased the system remains close to the equilibrium solution but only until a certain (very low) level is reached. At this point the positive feedback mechanisms are strong enough to counteract any equilibrating processes and far-from-equilibrium endogenous dynamics replace the now completely unobservable equilibrium solution.

As we shall see, for long periods of time *the far-from-equilibrium solution may be mistaken for an equilibrium solution whose parameters are changing*. There may be partial corrections along the way but the endogenous state eventually (and unpredictably) shifts via a rapid cascade that can even overshoot and cause a mispricing in the opposite direction. This whole process cannot be represented by *any* equilibrium model⁵ and to make matters even worse, the boom-and-bust aspect of

⁴Such averaging assumptions are the economic analogue to the Central Limit Theorem in Statistics which states that the mean of independent random variables (from a distribution with finite first and second moments) is normally distributed. The assumption of independence is crucial!

⁵See the quote by Alan Greenspan at the start of this section.

the endogenous dynamics indicates a high degree of nonlinearity that cannot be approximated by simply linearizing about an equilibrium solution. In other words a genuinely nonlinear mathematical description is needed.

The cause of positive feedbacks/instabilities (and non-averagability) that we shall focus upon is *herding behavior* whereby some agents have some additional motivation (ranging from completely ‘irrational’ to hyper-rational) to prefer the position taken by the majority of agents. While there are many possible causes of herding behavior, the effects should usually be similar and can be introduced very simply using the framework developed in Section II.

Boundedly rational behavior [Simon 1982] and contagion effects are often qualitatively blamed for market bubbles and other extended mispricings. Two widely-cited quantitative studies ([Akerlof and Yellen 1985, Scharfstein and Stein 1990]) showed that introducing non-maximizing changes into a subset of the agents does indeed shift the value of the equilibrium solutions in their respective models (the second study is especially relevant here in that it considered herding effects among investment managers concerned about their relative performance to be the primary source of the deviation). However these analyses were performed entirely within an equilibrium framework that precluded the possibility of non-equilibrium dynamic solutions⁶.

Although it might appear that these theoretical results provide some hope for the development of equilibrium models that adequately incorporate bounded rationality and/or herding there are two fundamental problems. Firstly, in an equilibrium model the only timescales that exist are those imposed exogenously. Thus in order to shed light on how long a mispricing might last and how quickly it will be corrected there must, at the very least, be an additional non-equilibrium exogenous model attached to the equilibrium one. However, if the causes of the mispricing are endogenous then the modelling process should respect that fact.

⁶Similarly [Banerjee 1992] and [Bikhchandani, Hirshleifer and Welch 1992] introduced models of herd behavior via sequential decision-making but again within the equilibrium paradigm.

The second problem is that in both analyses the logical steps were performed in the wrong order. As usually occurs in equilibrium modelling, perfect utility-maximization was assumed and then an averaging procedure (not explicitly stated) generated the equilibrium solution, and only then was utility-maximization weakened. However the averaging procedure requires utility-maximization and independence between the agents in order to be valid. In other words an assumption was made, *then it was used*, and then it was weakened. In order to avoid this logical error we shall not use averaging or assume the existence of an underlying stable equilibrium. Instead we will perturb a (bottom-up) agent-based model that reproduces the equilibrium solution at special parameter values corresponding to the perfect maximizing and independence assumptions.

The non-equilibrium modeling framework in the next Section was introduced in [Lamba and Seaman 2008a] and [Lamba 2010] and is based upon earlier related models ([Cross, Grinfeld, Lamba and Seaman 2005, Cross, Grinfeld, Lamba and Seaman 2007]). The primary motivation for this earlier work was to incorporate various systemic defects such as perverse incentives and investor psychology into an otherwise efficient market and thus establish a causal link between endogenous agent behavior at the micro level and the non-Gaussian price statistics observed in financial markets. One clear conclusion was that herding can indeed induce ‘fat-tails’ consistent with observed power-law decays of real asset price changes. Thus herding is a possible, indeed likely, contributor to an extremely important phenomenon that is inconsistent with standard financial models.

The motivation here is different, although the modeling details are similar. The simulations in Section IV should be considered a ‘stress-test’ of an extant equilibrium pricing model that is achieved by weakening a particular subset of the assumptions from the outset. Note that this allows us to carry over those assumptions from the standard pricing model that are not being weakened.

Of course there are many other agent-based models that can also approximate the stylized facts (see [Hommes 2006] for a survey of the earlier models)

and many of these also incorporate some form of herding or imitation. However our model has certain features that make it especially suitable for the task at hand. Most importantly, at the ‘default’ parameters the averaging procedures required to generate the equilibrium solution are mathematically justified — and so in this default state the model fits easily within the orthodox modeling paradigm. Secondly, we introduce a *separation of timescales*. This is a standard device for systems that have different components operating at very different rates and so we shall have ‘fast’ agents who react near instantaneously to new information and the far more interesting ‘slow’ agents who have both a limited form of memory and evolving strategies over long periods. This allows us to introduce herding (as well as many other behavioral effects if desired) in a simple way that more accurately captures their incremental nature.

The paper is organized as follows. The modeling framework is described fully in Section II together with an explanation of how various non-standard motivations, such as perverse incentives and the findings of behavioral economics, can be approximated. Then in Section III the results of numerical simulations are provided for realistic estimates of the model parameters. These demonstrate that the qualitative and quantitative changes introduced by such ‘imperfections’ are consistent with observations of real markets. Section IV contains the most significant numerical results. Here the herding strength is used as a bifurcation parameter to establish how much herding is required to destabilize the equilibrium market solution. The answer would appear to be an order of magnitude lower than a plausible estimate of herding in actual markets. In Section V links are established between our modeling approach and various concepts from economics and finance. Some implications of the findings and suggestions for further research are given in Section VI.

II The Model

Our starting point is the *One-Dimensional Random Walk*. This is a stochastic process $S(n)$ defined at discrete times $n = 0, 1, 2, \dots$ where $S(0) = 0$ and $S(n) = S(n-1) + X(n)$. The random variables $X(n)$ for $n = 1, 2, 3, \dots$ are all independent and identically distributed (i.i.d.) and take the values ± 1 with probability $\frac{1}{2}$ (so that their mean is 0 and variance is 1). To use a physical analogy, $S(n)$ corresponds to the position at time n of a particle jumping perfectly randomly either one unit to the left or the right.

In spite of its simplicity the one-dimensional random walk has some highly counter-intuitive properties⁷. These stem from the fact that a random walk does not ‘revert-to-the-mean’ in the way that most naturally-occurring processes do. The process $S(n)$ has no memory of where it has been and experiences no restoring force back towards 0 however far away it may be — this follows immediately from the *independence* of the $X(n)$ (instead $S(n)$ has the Martingale Property which states that the expected value of $S(n)$ at any future time n equals its last known value).

The basic random walk $S(n)$ can be generalized by replacing the $X(n)$ with different i.i.d. random variables that take more than just two values and can also have a non-zero mean allowing for a possible drift or bias in one direction or the other. Such random walks became the standard description of asset prices [Fama 1965, Malkiel 1973]. The justification was that the arrival of independent information, corresponding to the $X(n)$, is translated by the influence it has on market participants into a random walk process for the asset price $S(n)$.

Discrete-time random walks were soon superseded by continuous-time stochastic processes. The first change we need to make is to replace the evolving sum

⁷For example the expected value of the random variable that is the time taken to first hit position -1 is easily shown to be infinite. By symmetry the expected time taken to first hit $+1$ is also infinite — even though one of these two events must happen at $n = 1$!

of the independent random variables $X(n)$ by a continuous-time *Brownian motion* $B(t)$ that represents arriving exogenous information. Pricing models are then derived as *solutions* to stochastic differential equations (SDEs) as follows. The simplest, and most common, pricing model assumes that the price $p(t)$ at time t solves the SDE

$$(1) \quad dp = ap \, dt + bp \, dB$$

where dp is the change in price and dB is the change in a standard (drift-free, variance 1) Brownian motion *input* $B(t)$ over an infinitesimal time interval dt . The constant a is the average rate of return per unit time and b multiplies the effect of the completely random exogenous information/shocks $B(t)$ by a volatility factor. The reason for the factors p in the right-hand terms is that both the drift and random shocks will have a multiplicative rather than an additive effect on the price.⁸

Note that there are two critical assumptions being made. Firstly that the exogenous information stream $B(t)$ is indeed a Brownian motion with independent, uncorrelated, Gaussian-distributed, increments. This is unlikely to be perfectly true in practice but this assumption will not be relaxed here. The second assumption is that $B(t)$ is translated perfectly and instantly into independent price increments by the participants. Or, putting it another way, Brownian information passes through the ‘filter’ consisting of all the market participants and comes out unchanged as independent, proportional, instantaneous price changes. This assumption will be weakened by adding an additional term to (1) corresponding to price changes caused by endogenous effects.

The solution to (1), found using the Itô Calculus, is the *geometric Brownian*

⁸If the last term in (1) is missing then we have the ODE $dp = ap \, dt$, or equivalently $\frac{dp}{dt} = ap$, whose (deterministic) solution is the growing or decaying exponential $p(t) = p(0)e^{at}$ depending upon the sign of a .

motion (gBm)

$$(2) \quad p(t) = p(0) \exp \left(\left(a - \frac{1}{2} b^2 \right) t + b B(t) \right).$$

The SDE (1) is very special in that it has an explicit solution. Even more unusually the solution (2) depends only upon the *current* value of $B(t)$ and not the entire history of the Brownian process $\{B(s)\}_{s=0}^t$ up to that point. Thus the variable $p(t)$ can be considered an archetype for other variables in economic models that are assumed to be the outcome of instantaneously equilibrating processes that are path-independent.

This very close correspondence between the standard assumptions of mainstream economics and the most basic pricing model in finance is of course no accident. Separating the history of a mathematical variable from its future evolution makes it far easier to generate models with analytic closed-form solutions (such as the celebrated Black-Scholes formula [Black and Scholes 1973] for pricing financial derivatives that uses the result (2)). Equilibrium-based thinking is a direct consequence of this sacrificing of accuracy for mathematical expediency.

It will be more convenient to work with the log-price variable $r(t) = \ln p(t)$ which obeys an equivalent, but simpler, SDE using the same Brownian input $B(t)$. For constant drift a and volatility b the log-price is given by the solution

$$(3) \quad r(t) = at + bB(t)$$

to the SDE

$$(4) \quad dr = a dt + b dB.$$

For any given time interval h the changes in log-price $r(t+h) - r(t)$ have a Gaussian distribution $\mathcal{N}(ah, b^2h)$. However this is in poor agreement with reality. The ‘stylized facts’ of financial markets (see [Mantegna and Stanley 2000, Cont 2001]) are a set of statistical observations that appear to hold across all asset classes, independent of geography, history and trading rules. Two are especially

important. The first is the presence of ‘fat-tailed’ price returns whereby the occurrence of the largest price changes (as measured over intervals of hours up to months or years) follows an approximate power-law decay contradicting the exponential decay of Gaussian distributions. Thus the probability of the largest price moves is underestimated by many orders of magnitude. The second phenomenon is volatility clustering, whereby large price moves (in either direction) are more likely to occur shortly after other large price moves.

Without loss of generality, we may choose $a = 0$ so that $r(t)$ becomes the log-price relative to the risk-free interest rate plus any risk premium. We may also, by rescaling time, choose $b = 1$. In order to do computations we now discretize time so that the solution at the end of the n^{th} time interval of length h is given by

$$(5) \quad r(n) = r(n-1) + \sqrt{h}\eta(n)$$

where $\eta(n) \sim \mathcal{N}(0,1)$ is the discretized Brownian process, i.e. $\sum_{k=1}^n \eta(k) = B(nh)$. Note that at the times $t = 0, h, 2h, \dots$ this numerical solution coincides with (3) and so they are effectively equivalent.

Almost by definition, the price changes caused by the (exogenous) Brownian information stream $\eta(n)$ are effected by agents in the marketplace who a) act very fast and b) are motivated by the arrival of new information. This perfect information-to-price translation is justified by supposing that it is *as if* [Friedman 1953] all agents are continuously, instantaneously and correctly (on average) maximizing their respective utility functions. In reality the presence of transaction costs and the immense computational effort required will mean that for a subset of M agents trading or shifting of opinions occurs over much longer timescales. We shall call them ‘slow agents’. We do not assume that slow agents are of uniform size (in terms of their trading positions or influence) and thus weight the i^{th} slow agent by her size $w_i > 0$ and define $W = \sum_{i=1}^M w_i$.

Only the slow agents will be simulated directly because of the long timescales that are our concern here. In the numerics to follow, h will be chosen to corre-

spond to approximately 1/10 of a trading day. Fast agents include institutions (or individuals) that regularly trade the asset over a timescale of less than a day, and/or are motivated primarily by new information. Slow agents on the other hand will typically shift investment positions or opinions over weeks, months or longer.

In standard pricing models the precise details of the market-clearing mechanism are left unstated and it is merely assumed to be perfectly liquid and efficient. Since this assumption will not be weakened here we do not describe it either — only the slow agents are simulated and we assume that the fast agents provide sufficient liquidity so that trades do not need to be matched in order to be executed⁹.

We now make some assumptions that, it must be emphasized, are not fundamental to the modeling philosophy. They merely keep the model simple and are sufficient for the purpose at hand. Firstly, we assume that over the n^{th} time interval the i^{th} slow agents can only be in one of two states, the state $s_i(n) = +1$ meaning that she owns w_i units of the asset, and the state $s_i(n) = -1$ meaning that she owns none of the asset¹⁰. We can thus define the quantity $\sigma(n) = \frac{1}{W} \sum_{i=1}^M s_i(n) w_i$ which is a linear measure of the aggregate demand of the slow agents. We shall refer to σ as the *sentiment*. Note that σ varies between ± 1 with $\sigma(n) = -1$ (very bearish) when none of the slow agents own the asset, $+1$ when they all do (very bullish), and 0 when there are equal strengths in each state. Movements in σ are assumed to affect the price $p(t)$ proportionally via the change in demand. This is equivalent to an arithmetic shift in the log-price and the discretized pricing formula (5)

⁹The reader is directed to [Lamba 2010] for a discussion of how the model can be modified for those extreme market conditions where liquidity cannot be assumed. Briefly, under these circumstances the model reduces to a simulation of an order book and must be run using very small intervals h and all the agents, both slow and fast, must be explicitly included.

¹⁰In reality a slow agent may choose to gradually increase or decrease their holdings, short the market, or buy derivatives, but this complicates the dynamics of the slow agents without providing further insights. Also, some of the slow agents may not even represent active traders. They may be analysts who affect other investors with their weight w_i being their influence upon the market.

becomes

$$(6) \quad r(n) = r(n-1) + \sqrt{h}\eta(n) + \kappa\Delta\sigma(n)$$

where $\kappa > 0$ and $\Delta\sigma(n) = \sigma(n) - \sigma(n-1)$. The parameter κ is a measure of the total market depth of the slow agents. Thus one can interpret (6) as stating that price changes have an exogenous component $\sqrt{h}\eta(n)$ due to new information, and an endogenous one, $\kappa\Delta\sigma(n)$, caused by changes in the market sentiment. The value of κ will be larger in a market where endogenous effects are stronger relative to exogenous ones.

The difference between the equilibrium pricing model (5) and (6) is of course the extra endogenous term $\kappa\Delta\sigma(n)$. Whereas the solution to (5) is necessarily a Brownian log-price the solution to (6) will depend upon the dynamical properties of σ . The extra term reflects the fact that before information can be turned into price changes it must first pass through a ‘filter’ of imperfect agents.

The equilibrium pricing model (5) can be recovered from (6) in two different ways. Firstly, we can simply set $\kappa = 0$ so that there are no slow agents (or that they have no effect on the price) and the extra term in (6) simply vanishes. Secondly (and this is fundamental to our argument) we can suppose that the slow agents are always in near-equilibrium making $\sigma(n) \approx 0 \forall n$ so that the last term in (6) becomes negligible. This is equivalent to averaging away the effects of imperfect agents. However, the pricing formula (6) allows for the *possibility* that endogenous non-equilibrium dynamics amongst the slow agents will affect the price. In other words we have embedded the standard equilibrium model (5) within a less restrictive one — one that is capable of dynamical solutions.

This embedding is crucial if we are to avoid the logical inconsistency present in the analyses of [Akerlof and Yellen 1985, Scharfstein and Stein 1990]. We shall arrange things so that at a certain parameter value there is indeed a valid averaging argument that proves $\sigma(n) \approx 0 \forall n$. This parameter value corresponds precisely to the orthodox modelling assumptions. However this averaging will not necessarily

be valid at nearby parameter values *and we will not assume that it is*. Instead we perform numerical simulations at varying parameter values and observe the nature of the solutions.

Before we carry out this program we shall introduce one further generalization of (5) by weakening the assumption that the fast (information driven) agents correctly transform new information into price changes. This is achieved by adding a (for now unspecified) function $f(\bullet)$ that modifies the effect on prices of new information entering the market changing (6) to

$$(7) \quad r(n) = r(n-1) + f(\bullet)\sqrt{h}\eta(n) + \kappa\Delta\sigma(n).$$

with the fast agents acting perfectly if $f(\bullet) \equiv 1$.¹¹

Equation (7) does not constitute a closed system because no rules governing the switching of the slow agents between the -1 and $+1$ states have been specified yet. There are many types of rule or trading strategy that could be used, involving any desired combination of pure utility function maximization, bonus/commission maximization, inductive learning, imitation among a network of slow agents, technical trading, bounded rationality, ‘gut instinct’ or animal spirits, imperfect time discounting, profits, losses, relative performance of other investment options, market volatility, fear, greed, margin calls, the weather, and so on. Prior studies such as the Santa Fe model (see [LeBaron, Arthur and Palmer 1999]) have indeed used complicated ecosystems of trading strategies and there is much to commend this approach. However for our current purposes, simplicity is more appropriate than complexity.

We shall use an approach based upon moving price thresholds developed in [Lamba and Seaman 2008a] and [Lamba 2010] which is deceptively simple but

¹¹As will be demonstrated numerically, this extra change has very little effect upon the stability investigation that is our main concern. But weakening this assumption about the fast agents induces volatility clustering in a natural way [Lamba 2010] and so we include it just for completeness.

is capable of introducing various real world influences, market ‘imperfections’, bounded rationality and psychological biases (see Section A below). At the start of the n^{th} time interval, the i^{th} slow agent is represented by its state, -1 or $+1$, and an evolving open price interval $I_i(n) = (L_i(n), U_i(n))$ where $L_i(n) < r(t) < U_i(n)$ (Figure I). The endpoints L_i and U_i are referred to as the lower and upper thresholds respectively. If, at the end of that time interval, the price has crossed *either* threshold agent i is deemed to be no longer comfortable with her current investment position, switches states, and the interval I_i is updated so that the price is again an interior point.

For example, an agent who is $+1$ may be switching to either take profits or cut losses depending on which threshold is breached and the price at which they bought. An agent who is -1 and then switches may be motivated by, say, a sell-off which makes them think the stock is now cheap, or an urge to follow the momentum of a rising stock or attempt to catch up with a benchmark. Or consider a market analyst who has a buy rating on the asset but may switch when either the price becomes too high or when the price falls far enough that they (belatedly) change their opinion.

Finally, from (7) the action of switching causes a small jump in the sentiment σ of $\pm 2w_i/W$ and so in the log-price of $\pm 2\kappa w_i/W$. Each of the M slow agents has their own interval straddling the current price with the price and all of the thresholds evolving at each timestep.

Since the intervals I_i are also evolving, it just remains to define the dynamics of the L_i and U_i . These thresholds for each agent will change (usually slowly) between switchings and correspond to that agent’s evolving *strategy*¹². Or equivalently the intervals can be thought of as agents’ comfort zones within which that agent is still satisfied with their current investment position — or stated opinion in the case of a market analyst. Note that in the case of an algorithmic

¹²The strategy is ‘two-sided’ and takes into account the possibility of prices both rising and falling as any complete strategy should.




Figure I: A representation of the model showing two agents in opposite states. Agent i is in the $+1$ state and is represented by the (interval between) the two circles, and agent j is in the -1 state and is represented by the two crosses.

‘black box’ or a hyper-rational utility maximizing investor these thresholds will be consciously/explicitly specified and known by that agent (but no other). A less-rational slow agent may not even be consciously aware of the threshold values but will sense when one of them is violated and act to switch.

Our use of thresholds is reminiscent of many models of market microstructure but there are important differences. Firstly, the microstructure models are concerned with the inner workings of exchanges and the operation of market mechanisms over very small timescales. Here, we just assume that the trades can be executed at close to the current price since we are concerned with mispricings and disequilibrium over timescales orders of magnitude larger than the time taken to execute the trades. Secondly, the orders in the microstructure models are typically generated according to some stochastic distribution whereas here we shall focus upon the processes leading to the creation of the orders. Or, rephrasing, we are concerned with *why* trades are placed rather than *how* they are executed.

The state of the equilibrium model (5) at any given moment is completely specified by *just one number*, the current price (or equivalently the current value of the Brownian driving process $B(t)$). The situation for the full threshold model is very different. To specify the current state of the system completely requires the additional knowledge of all of the $2M$ threshold values *and* the rules specifying

their dynamics. These endogenous ‘hidden variables’ add a great deal of potential complexity to the model but we shall make some simplifying assumptions that are sufficient for addressing the stability issues concerning us here.

Just for the moment, in order to apply a *valid* averaging argument and prove that the (equilibrium) solution to (5) can be reproduced at a special parameter value, we suppose that there are myriad influences on slow agents’ strategies that can be adequately represented by different, uncorrelated and independent, dynamics applied directly to every agent’s thresholds. These will be a mixture of purely rational independent analysis, non-rational thought processes, boundedly rational heuristics, correct and incorrect assumptions, and many other things in differing proportions for each agent. All these influences will move each upper and lower threshold either inwards or outwards (towards or away from the current price) thus making the agent more or less likely to switch states respectively. *If we assume that these myriad influences are statistically independent* then we can treat every threshold as following its own independent random walk, giving

$$(8) \quad L_i(n+1) = L_i(n) + \delta_i \mathcal{N}(0, h), \quad U_i(n+1) = U_i(n) + \delta_i \mathcal{N}(0, h).$$

where the quantity δ_i is the volatility of the threshold motion per unit time for the i^{th} agent. Finally, when an agent does in fact switch at the end of a time interval, their threshold values should also reset in a way that is independent of the other agents. Then, on average, equal numbers of slow agents will be switching at each timestep and, provided that $\sigma(0) = 0$ with identical threshold distributions for the agents in each state at time 0, σ remains close to 0 for all time. The result of these assumptions on the threshold dynamics is only very small, random, deviations from the Brownian equilibrium pricing model (5).

To summarize, the model so far sits quite comfortably within the consensus of mainstream economic schools of thought. The slow agents may not act instantaneously upon new information but they have differing expectations about the future that are independent and, on average, correct and the model simply repro-

duces the equilibrium price.

We have now, finally, reached the point where we can weaken some modelling assumptions. These take the form of changes to the slow agents that destroy their statistical independence and invalidate the above averaging procedure.¹³ We posit that agents who are in the minority state (i.e. those who are in state +1 if $\sigma < 0$ and state -1 if $\sigma > 0$) will have, to a varying degree, some pressure/motivation to switch and join the majority that builds up gradually over time¹⁴. There are several real-world effects that make this change to the dynamics a natural one to consider.

Firstly, any substantial change in σ will lead to a drift in price that may be (mis)interpreted as a fundamental trend causing the minority agents to react accordingly to the price signal. Secondly, there may be ‘rational herding’ by investment managers, say, who find themselves chasing a benchmark average so as not to lose their jobs, bonuses or investment capital¹⁵(see [Keynes 1936, Scharfstein and Stein 1990]). A third cause is the actions of momentum investors who are trying to detect and profit from a nascent trend (or bubble) as part of their investment strategy. Fourthly, there is the propensity of people, when faced with uncertainty, to believe that other people are better informed than they are. Next, people may find it preferable to risk failing conventionally than succeeding unconventionally ([Keynes 1936]). Finally, there may be purely psychological effects at work caused by the discomfort of being in a minority, especially within social or professional networks. We shall refer to all the above effects as causes of *herding*.

For simplicity we assume that each agent reacts to the current value of σ

¹³Erroneous assumptions about the independence of the probability of foreclosures in different housing markets were a major contributor to the financial crisis that started in 2008.

¹⁴This is not to say that truly contrarian investors do not exist but, almost by definition, their numbers are smaller and could be incorporated by making similar changes to some of the majority agents. Furthermore, contrarian-like behavior by some of the slow agents will emerge quite naturally within the model as a response to extreme mispricing and price corrections

¹⁵These effects will be amplified by the short time-horizons of such evaluation periods.

although in reality agents will have different perceived values of this quantity (or may be reacting, in part, subconsciously). Herding is introduced into the model by supposing that for agents *in the minority position only* the purely random threshold motions (8) are replaced by

$$(9) \quad \begin{aligned} L_i(n+1) &= L_i(n) + C_i h |\sigma(n)| + \delta_i \mathcal{N}(0, h) \\ U_i(n+1) &= U_i(n) - C_i h |\sigma(n)| + \delta_i \mathcal{N}(0, h) \end{aligned}$$

while those in the majority are still governed by (8). Note that this change simply adds an inward drift to the threshold dynamics. Drifting the thresholds inwards (towards the current price of course) reduces the time to the next switching and can be interpreted as an agent’s ‘comfort zone’ being squeezed by the majority opinion. Indeed this ability of inward-drifting thresholds to capture the incremental nature of herding pressures was one of the primary motivations for their inclusion in the model (Section A outlines several other advantages of using thresholds). The rate of drift is taken to be proportional to the length of the timestep, h , the magnitude of the imbalance $|\sigma(n)|$ and a constant $C_i \geq 0$ quantifying the herding effect for that agent. Note that herding is a positive-feedback effect — as σ moves away from 0 it, at least initially, provides a mechanism for the imbalance to increase. It also causes (global) coupling between agents’ strategies that violates the very strong conditions necessary for the rigorous application of averaging-type simplifications. This completes the description of the model.

A Justifications for moving thresholds

As mentioned above, once the equilibrium model (5) has been embedded within a more general non-equilibrium one, (6) or (7), there are various ways to specify the switching rules of the slow agents. Treating each slow agent as a dynamic interval (L_i, U_i) on the price line (that must at all times contain the current log-price) is certainly unusual and at first may seem rather unnatural. However using such pairs of price thresholds offers some compelling advantages.

Firstly there is the observation that many agents react to price changes over long periods or to previously specified price targets rather than to incoming information. Indeed investment advice and analysis is usually offered in the form of price triggers, as are the outputs of computerized trading algorithms. Sometimes, such as in a margin call, the agent has no choice over the pricing point.

A second important issue is that of transaction (or sunk) costs. These are often neglected via simplifying assumptions¹⁶ but they profoundly change the nature of agents' behavior. Even if one believes that agents are continuously maximizing their utility functions, this process must somehow be translated into an acceptably infrequent transaction rate since a continuous process of incremental adjustments would be ruinous. Provided that threshold switchings result in new threshold values that are a non-zero distance away from the current price, then moving price thresholds automatically convert continuous strategy updating into discrete trades/actions. The existence of sunk costs is closely linked to issues of hysteresis, memory-dependence and irreversibility that will be discussed further in Section V.

We now turn to behavioral economics. It has already been shown how the propensity for herding can be included by moving thresholds inwards for agents in the minority position. However other effects such as anchoring, loss-aversion [Benartzi and Thaler 1995] and the decision-making processes outlined in Prospect Theory [Kahneman and Tversky 1974] can also be replicated. Anchoring is almost automatic as thresholds are reset around the last trading price while loss aversion requires slightly more complicated rules that involve keeping track of an agent's profit or losses. An extreme, but unfortunately quite common, example helps demonstrate the idea.

Imagine an individual who buys a dot-com stock at the height of a tech bubble.

¹⁶It is amusing to note that many people who rely on such models are actually paid from transaction costs. And all-too-often the possibility of significant transaction fees can skew the information and research entering the market.

Immediately the price goes down but, due to loss aversion which is the emotional difficulty of acting to realize a loss, the lower threshold moves down even faster. It is likely that the upper threshold is moving downwards too but there is never enough of a temporary bounce in the stock price to cause a switch which would cut their losses. Eventually the stock price hits zero without the agent ever selling.

A significant body of work (see [Kahneman and Tversky 1979, Rubinstein 1998, Gigerenzer 2002, Barberis and Thaler 2003, Earl, Peng and Potts 2007]) has formulated plausible heuristics that agents may be using in practice. It should be possible to recreate many of these rules using moving thresholds¹⁷ and then observe the effects upon aggregate statistics — just as is done in the following Section for the herding phenomenon.

III Preliminary Numerics

In this Section we provide enough details for replication of the numerical results, including parameter values and the reasoning behind the choices made, and sufficient simulations to reveal the nature of the non-equilibrium solution. The stability of the equilibrium solution will then be investigated in Section IV. It must be emphasized that no fine-tuning of parameters is required. Indeed the most significant parameters can be simply estimated and we do so conservatively and as simply as possible.

Firstly, we choose a timestep h defined in units such that the variance of the external information stream $B(t)$ is unity for $h = 1$. This timestep should be small enough by at least an order of magnitude to catch all the switchings of the slow agents. An observed standard deviation in price returns of 0.6–0.7% suggests that $h = 0.000004$ should correspond to approximately 1/10 of a trading day. The price changes of 10 consecutive timesteps are then summed to give the daily price

¹⁷Furthermore, these rules can be superposed so that an agent can be a mix of rational, behavioral, and even technical trading tendencies.

returns.

The most important single parameter is κ since this quantifies the effect of slow agents upon the price in (7). We set $\kappa = 0.2$ since this implies that the difference in price between neutral ($\sigma = 0$) and polarized markets $\sigma = \pm 1$ is, from (7), $\exp(0.2) \approx 22\%$. This value may seem to be on the low side since it caps the size of any bubble or anti-bubble at just 22% from the fundamental price and many (anti-)bubbles are far larger. However there is an important feature of real-world bubbles that we are neglecting — new agents may enter the market and fuel the bubble, especially if they are speculating or momentum-trading. This effect will temporarily increase κ until the price correction occurs. We choose to ignore this important effect for two reasons. Firstly, it only comes into play once the equilibrium solution has already lost stability. Secondly, it acts as a further source of positive-feedback *away from* the equilibrium and so its omission does not undermine the central findings of this paper.

We next assume that all slow agents have equal weight $w_i = 1$. This could be replaced by a more realistic distribution, such as a Pareto distribution, but it does not significantly affect our numerical results and so we shall not do so here. We can also justify the choice of equal weights as follows. It is not necessary to assume that each slow agent corresponds to just one individual or institution — they could each refer to a subset of agents with very similar strategies or propensities in which case they may indeed be of approximately equal influence in the market.

This brings us to the important question of how many slow agents there should be. There exists a continuum limit as $M \rightarrow \infty$ and numerically we observe that even $M = 1000$ agents is sufficiently close to the continuum limit for our purposes. All the simulations below will use $M = 100000$.

Next we fix how the slow agents' thresholds are reset after a switching. If agent i switches at a log-price r^* then immediately afterwards the interval is reset to

$$(L_i, U_i) = (r^* - Z_L, r^* + Z_U)$$

where Z_L, Z_U are chosen from the uniform distribution on the interval $[0.05, 0.25]$. This corresponds to an initial strategy that requires price changes in the range $\approx 5\text{--}25\%$ before another switching (remembering of course that the threshold dynamics are altering the strategy as the system evolves).

For simplicity we shall suppose that all the slow agents have the same herding propensity, i.e. $C_i = C, \forall i = 1, \dots, M$ although the numerical results are once again essentially unchanged if one chooses a distribution of values centred around C . The above distribution of the thresholds after resetting allows us to estimate C by supposing that in a moderately polarized market with $|\sigma| = 0.5$ a typical minority agent (outnumbered 3–1) would switch due to herding pressure after approximately 80 trading days (or 3 months, a typical reporting period for investment performance). From (9) the calculation $80C|\sigma| = |\ln(1 - \frac{0.05+0.25}{2})|/0.00004$ gives $C \approx 100$. This single number C represents the herding level in the market. Recall that if C is set to 0 then we have no coupling between the agents and we will obtain the equilibrium solution.

The parameters δ_i represent the ‘volatility’ of each agent’s independent strategy (or expectations) and we once again for simplicity assume they are all equal. Note that if $\delta_i = 1$ then the volatility of the thresholds is the same as the volatility in price due to the information $\sqrt{h}\eta$ entering the system in (7). This is almost certainly too large since the slow agents are not the ones motivated primarily by new information and should alter their expectations more slowly. Thus we set $\delta_i = 0.2 \forall i$. Even if we chose all the $\delta_i = 0$ the simulations would not be much affected and so these parameters are not especially relevant here — the coupling between agents caused by the inward drift of the thresholds is far more significant than their random (averagable) jitter.

We now consider the fast agents. The function $f(\bullet)$ was introduced in (7) to weaken the assumption that the fast agents always accurately translate new information into price changes¹⁸. We shall suppose that at times of non-neutral market

¹⁸This is a simple but plausible mechanism for the generation of volatility clustering.

sentiment, and especially when σ is far from 0, there is excess activity by fast agents. This may be due to new agents entering the market or by too much significance being attached to new information by traders expecting a market correction. There is some evidence for this (see [Brown 1999]) and it also helps correct for the fact that in our simple-as-possible model no new agents are allowed and the slow agents are not allowed to own multiple units of stock. As in previous work on this model the simplistic but plausible choice $f(\bullet) = 1 + \alpha|\sigma|$ is made with $\alpha > 0$. We choose $\alpha = 1$ so that at the most extreme mispricings, new information moves the market twice as much as it would if the fast agents were correctly incorporating it. As shown in the next Section, the presence of the function $f(\bullet)$ has a negligible effect upon our stability results which suggests that ignoring new agents and/or allowing agents to buy multiple units of stock is indeed a valid approximation for stability purposes.




Figure II: Asset price of a simulation over 40 years with 100000 agents. Note that the price $p(t)$ is plotted, not the log-price $r(t)$.

Figure V shows a simulation over a (hypothetical) 40 year period with the above parameters. The initial states of the agents and their thresholds are random-




Figure III: A plot of σ against time for the simulation in Figure V.

ized. The thicker, more volatile, curve is the pricing output of (7) while the lighter curve is the output of the geometric Brownian pricing model (that, recall, can be recovered by either setting $\kappa = 0$ or $C = 0$). A typical characteristic is that mispricings develop slowly and then suddenly reverse and often end up overshooting. This can also be seen in Figure III which plots σ against time.

The incremental mispricing is caused by the endogenous dynamics of the slow agents under the action of herding effects. The sudden reversals occur when enough majority agents switch position to start a cascade process (or avalanche) — as agents switch they cause a change in price (due to the $\kappa\Delta\sigma$ term in (7)) that trips other agents' thresholds and so on. These cascades are highly unpredictable, both in their timing and magnitude, since they are very sensitive to the precise distribution of threshold values and the exogenous information stream. In particular there can be many smaller cascades (false tops or bottoms) before the one that definitively corrects the mispricing. This simple version of the model is symmetric with respect to rising and falling sentiment/prices. In practice there are some important asymmetries [Lamba and Seaman 2008b] but the model contains




Figure IV: Daily price returns for Figure .

enough flexibility in the vast range of possible threshold dynamics to incorporate these differences.

Figure IV shows the daily percentage price returns and for comparison the price changes from the equilibrium model are shown in Figure V. There are clearly a significant number of large price changes that cannot be explained by the Gaussian information stream — the equilibrium pricing model only has a handful of days over the 40 year period where the daily price change exceeds $\pm 2\%$. These outlier price changes of course correspond to the endogenous cascades.¹⁹

Finally a snapshot of the internal structure of the market is shown in Figure VI. The density (histograms) of the lower and upper thresholds of each type of agent (-1 and $+1$) are plotted relative to the current price (so that for each agent one threshold is negative and the other positive) at a moment when $\sigma \approx 0$ (the mismatch between the densities is even more severe when $|\sigma| \approx 1$). The two density

¹⁹The frequency and magnitude of these cascades is an important issue but it does not concern us here since they only occur after the equilibrium solution has already been lost. Interestingly there is a very close correspondence between the dynamics of these cascades and Queueing Theory that is described in detail in [Lamba 2010].




Figure V: Daily price returns of the equilibrium price from Figure .

plots are not identical (as they must be in the equilibrium solution) meaning that as the system evolves σ will once again move away from 0 because different numbers of agents will be switching in either direction.

Figure VI helps illustrate a fundamental fact about the full model — once the equilibrium solution is lost it will almost certainly never be recovered [Robinson 1962]. The sentiment σ will occasionally pass through 0 but this is correct pricing only in the sense that a stopped clock is correct twice a day.

In the Appendix we show how the full complexity of the model dynamics can be visualized in a different way. Instead of each agent being described by a state (+1 or -1) and two threshold values, we can represent them by a dot or a cross in two dimensions. Thus each agent is a signed particle moving in a two dimensional region and switches sign and jumps into the interior whenever it hits a boundary. Full details are given in the Appendix and the interested reader can explore the behavior of the model and vary the parameters at <http://math.gmu.edu/~harbir/market.html> .




Figure VI: The density of the thresholds along the price axis for agents in each of the two states. The difference between the two distributions is a ‘memory effect’ of the prior behavior of the system and affects the future evolution.

IV Stability simulations

Recall that if all the C_i are set to $C = 0$ (ie. the agents are uncoupled) then the equilibrium solution is recovered. We now address the central question of this paper which is to try and quantify how the level of disequilibrium depends upon the strength of the coupling between agents.

We quantify the level of disequilibrium in the system by recording the maximum value of $|\sigma|$ and averaging over 20 runs each for fixed values of C in the range $0 \leq C \leq 40$. All the other parameters and the initial conditions are kept unchanged from those used to generate Figures V–VI.

The results are shown in Figure VII. As can be seen, the equilibrium solution can be considered a reasonable approximation only for $C < 5$ which is more than an order of magnitude below our real-world estimate $C = 100$ from Section III. The loss of stability, measured in terms of the maximum deviation of σ from 0, is gradual and reaches the maximum the system will allow at around $C = 40$.




Figure VII: A plot showing the degree of disequilibrium $|\sigma|_{\max}$ averaged over 20 runs for varying levels of the herding parameter C .

Figure VII can be explained as follows. The drift in the threshold dynamics of the minority agents is a destabilizing (dis-equilibrating) influence, while the diffusion of the thresholds/strategies and the fact that the majority agents do of course eventually switch out of their position are stabilizing influences. Simple first-exit-time calculations (not presented here) to compute average switching rates in either direction confirm that beyond $C \approx 40$ the herding pressure is capable of sustaining the system far from equilibrium for long periods of time.²⁰ A mathematical treatment of these stability issues will be presented elsewhere.

The presence of $\alpha \neq 0$ in $f(\bullet)$ which modifies the influence of the fast agents might be responsible for the loss of equilibrium stability. This is ruled out by Figure VIII which shows that the results with $\alpha = 0$ are very similar. Thus it is the lack of independence between agents' behaviors caused by the herding term that is responsible for the loss of stability of the equilibrium Brownian solution.

Finally, we consider the parameters δ_i that govern the diffusion of slow agent thresholds. In Section III it was argued that this would likely be much lower

²⁰The existence of such competing forces is a common source of interesting dynamics in complex nonlinear systems.

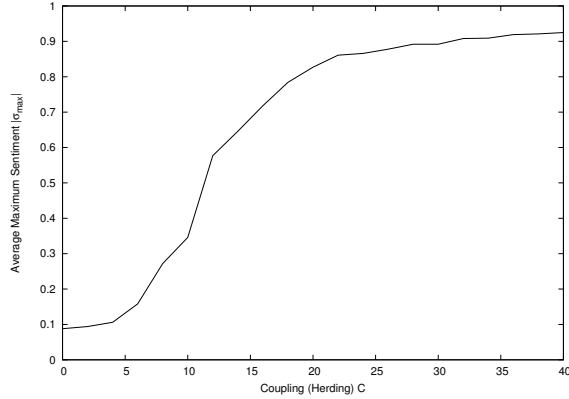


Figure VIII: The same plot as Figure VII but with $\alpha = 0$ so that the fast agents are perfectly translating new information into price changes.

than the price volatility and so was set to $\delta_i = 0.2 \forall i$. Figure IX sets $\delta_i = 1 \forall i$ so that their magnitudes are now comparable. A larger variance in the agents' (independent) changes in strategy should help mitigate the effects of herding and this is confirmed, although at $C = 100$ the system is still capable of reaching full disequilibrium $\sigma \approx \pm 1$.

V Connections to economic concepts and other non-equilibrium models

A Equilibria, memory and history-dependence

The notion of equilibrium described in the introduction is a strong one, precluding the possibility of multiple internal configurations. Nonetheless the absence of (non-trivial) endogenous dynamics is a prerequisite for models that are *reversible* and history-independent with temporary exogenous shocks having no permanent effect.

A key issue is one of *memory* at the micro-level. If individual agents have no




Figure IX: The threshold dynamics are made more volatile by increasing δ to 1.

memory then it becomes much more reasonable to assume that they will *near-instantaneously* reverse their actions and expectations when the external influence is removed. This type of assumption is crucial in justifying the quasi-equilibria of statistical physics. However economic actors are subject to many effects that cannot reasonably be modelled in this way.

Perhaps the most significant of the rational factors is the presence of transaction (sunk) costs (see [Dixit 1992, Piscitelli, Grinfeld, Lamba and Cross 1999, Göcke 2002]). These are expenses or other penalties incurred that cannot be recouped on reversing the action. As an example suppose that at the current widget price it is not profitable for a manufacturer to have a factory produce the widget. However when the price increases to β (perhaps due to a demand shock) the firm switches a factory over to widget production from something else, incurring costs such as re-tooling and factory down-time. If the price then falls back below β the factory will not immediately switch out of production but rather waits until the price falls below some value α . Thus if one only looks at the current price p and $\alpha < p < \beta$ it is not possible to know what the factory is producing — one also needs to know which of the threshold values α and β was last crossed. In other words a memory dependence/persistence effect has been introduced.

This type of memory dependence is referred to as *hysteresis*²¹ and the reader is directed to [Cross, Grinfeld and Lamba 2009] for a fuller description of the role of hysteresis in economics. The presence of many such factories, all with differing threshold values, results in many possible alternative internal configurations for the same price level. Each of these possibilities results in a different future evolution of the system which now displays both irreversibility and history-dependence²².

At an abstract level, the thresholds used to describe the slow agents in Section II are a mechanism for introducing memory/history into the modeling process (the fast agents, by assumption, only react to new information and have no such mechanism). To see this, an agent who has been in the minority state for a long time and has been influenced by a herding pressure, will on average have thresholds that are much closer together and so be more likely to switch in the near future. This information is propagated from one timestep to the next along with the agent's current state.²³

It is now time to revisit the concept of equilibrium, allowing for the possibility of multiple endogenous configurations and endogenous effects that involve long timescales. Some people might say that the market model simulations from Section III are still in an equilibrium until just before a sharp reversal. But this is like saying that a geological fault-line is in equilibrium until just before an earthquake. This is true, in that there is a balance of both external and internal forces, but the statement that the fault-line is in equilibrium just after the earthquake is

²¹In economics the term 'hysteresis' also refers to the persistence of deviations due to shocks at the macro-level.

²²When such effects are observed in macroeconomics a common equilibrium-based explanation is the presence of a *unit-root* since, if a system is only marginally stable it will take a long time to return to its former state after a disturbance. However marginal stability implies that a system is close to instability and this should be far more worrying to economists than irreversibility!

²³As noted in the Introduction the incremental nature of such deviations from rationality are especially problematic for equilibrium models that have no endogenous timescales.

then equally true! The ‘before’ and the ‘after’ can be mimicked deceptively well by unique equilibrium models, but not the transition!²⁴ Again, the issue can be reinterpreted in terms of timescales. Any earthquake zone looks like it is in a stable equilibrium on all those days when there are no earthquakes. Indeed any dynamical process will look like an equilibrium if it is studied over timescales that are inappropriately short.

It is a widespread misconception that balancing forces, such as aggregate supply and demand, must result in a unique equilibrium. As noted in the Introduction, much of mainstream economics is based upon analogies with Newtonian physics and classical mechanics but it is only in the simplest physical systems that a ‘balance of forces’ necessarily produces a unique equilibrium. If the system has an internal structure then multiple internal configurations may be feasible for the same set of external conditions. The internal state that is chosen will be history-dependent and the system may shift very rapidly from one internal state to a different one (with lower energy in the case of an earthquake). It should be noted that multiple equilibrium models do exist in mainstream economics with the initial conditions determining which equilibrium is achieved. However the situation now is far worse — the set of feasible configurations cannot be enumerated in advance and depends on the path taken by the process.

B Rational expectations and efficient markets

The consequences of the mainstream acceptance of the hypothesis of memory-free, efficient, markets cannot be overstated. Although the concepts were introduced by Bachelier in his 1900 Ph.D. thesis, this work was largely forgotten until the 1960s when they became known collectively as the Efficient Market Hypothesis (EMH) (see [Fama 1965, Samuelson 1965, Fama 1970]).

The three versions of the EMH (weak, semi-strong and strong) all rely upon

²⁴The reader is directed to the quote at the start of the Introduction a second time.

two qualitatively different classes of assumption. Firstly, there are assumptions about the infrastructure of the market itself and the processing/dissemination of information. This information consists of, amongst other things, economic statistics, performance reports, geopolitical events, insider information and analysts projections.

The second class of assumptions relates to the market participants themselves, who are deemed to be perfectly rational, correctly incentivized, and capable of instantaneously incorporating new data into their differing strategies and expectations. However, heterogeneity of agents (or their expectations) is necessary to ensure that trading occurs in the absence of arbitrage opportunities.²⁵ Thus the final ingredient in the EMH description is the Rational Expectations Hypothesis (REH) (see [Muth 1961]) stating that the differing expectations driving trades, when used as predictions, are on average correct²⁶ and do not result in market mispricing. Additional assumptions, such as the absence of transaction costs, yield the standard formulae used for risk management and derivative pricing which form the bedrock of modern financial engineering ([Black and Scholes 1973]). It is this second class of assumptions that have been the focus of this paper.

As described in Section II, if one sets $f(\bullet) \equiv 1$ in (7), $C_i = 0 \forall i$, and further assumes that all threshold dynamics for the slow agents are the result of perfectly rational, independent, utility maximizing behavior (satisfying the REH) then one recovers an equilibrium market following geometric Brownian motion and satisfying the EMH. Once one weakens these assumptions to allow for more general dynamics (and motivations), the moving threshold model can be thought of as a ‘perturbation space’ within which one can explore the robustness/stability of the default EMH model.

²⁵It is the *differences* between agents’ needs, wants, surpluses and deficits that drive almost all economic activity. However the first necessary step in most equilibrium models is to average away these very differences.

²⁶As opposed to being consistent with the modeling assumptions which is the form of the REH most often used in macroeconomic modeling.

Even if agents are not perfectly rational and other factors influence the threshold dynamics, the pricing should remain correct provided that the REH still holds. However the presence of just one REH-violating perturbation invalidates its use. As was shown in the numerical simulations of Sections III and IV, real-world effects that induce herding do exactly that, giving rise to price dynamics that differ greatly over long time periods from the equilibrium model by introducing a form of coupling between agents' actions.

Even markets in which bounded rationality exists can be 'efficient' in the sense that investors cannot earn above-average returns without taking on above-average risk. Indeed, this is a minimal requirement for any predictive market model (see [Malkiel 2003]). In [Cross et al. 2007] it was shown that there is no statistically significant difference between the investment performance of agents with differing herding propensities C_i (and hence nothing to be gained by adaptively changing their reaction to herding). Furthermore, when transaction costs are taken into account the traders with the highest values of C_i , who tended to trade more frequently, performed the worst, in agreement with previous studies ([Odean 1999]).

One potential criticism of the model is that the fast agents are assumed to be reacting to *new* information and converting it into price *changes*. However there may also exist true fundamental fast agents who are aware of the mispricing and view this as a trading opportunity. This would act as an additional equilibrating (negative-feedback) mechanism. This brings us to the very important issue of the limits of arbitrage, both in the model and in real markets.

Firstly, as regards the model, σ is assumed to be precisely known by all the slow agents. This is purely for simplicity and agents may have widely-differing perceived values that are only approximately (or on average) correct. Also, it is important to note that no agents are assumed to know the 'correct' geometric Brownian motion price. It is calculated and plotted in Section III but this is only a visual aid — none of the agents are aware of it. As it stands the model is a

caricature, albeit one that can be made arbitrarily more complicated and realistic. As this complexity grows, any potential model-specific opportunities for arbitrage that might exist will reduce and so now we discuss the limits to arbitrage in real markets.

Arbitrageurs and/or fundamentalist investors have the potential to counteract herding effects. However there are severe limitations in practice. Firstly there is the noise-trader problem (see [Schleifer 2000] and [Shleifer and Vishny 1997]) — arbitrageurs typically have very short time-horizons but mispricings can last a very long time. Secondly there are speculative traders and short term momentum-traders who may actually make the mispricing worse. Thirdly, it is extremely difficult in practice to be sure what the fundamental price actually is. There is no visible, unambiguous, information stream and all trends may be misinterpreted as rational — especially by those who subscribe to the EMH. Some evidence for this may be found in the wide variations over time of even the most basic measures of value such as the P/E ratio of a stock or price moves versus changes in dividends [Shiller 1981].

The possibility that positive feedback effects such as herding can be counteracted by fundamental valuations should not be dismissed. Indeed many *potential* bubbles may get deflated before positive feedback effects take hold. However the results of Section IV strongly suggest that the battle between equilibrating and dis-equilibrating forces is far less one-sided than is commonly assumed.

VI Implications and Concluding Remarks

The numerical results of Section IV demonstrate that a hypothetical, but recognizably orthodox, equilibrium pricing model becomes unstable and is replaced by more complex and far richer non-equilibrium dynamics when non-maximizing behavior is introduced into a subset of the agents. This appears to contradict the findings of [Akerlof and Yellen 1985, Scharfstein and Stein 1990] but as was

pointed out in the Introduction there is a logical inconsistency in those analyses that can be corrected by using a fully bottom-up agent-based model.

An important aspect of the scientific method is that a model/theory can be overturned by contradictory experiments alone — no explanation or replacement is required. The model in this paper is not being offered as an alternative for asset pricing (although it could indeed be used that way). Rather, it should be viewed as an experiment that tests the degree to which the stability of the equilibrium solution can survive changes to the modelling assumptions.

This does not by itself prove that such far-from-equilibrium dynamics occurs in any actual market, but there are three key observations that suggest this is indeed the case. Firstly, the price statistics are much closer to the stylized facts of financial markets than those of the equilibrium solution. Of particular importance is the ‘fat-tailed’ frequency of extreme price changes that are caused by sudden shifts in the internal state of the system. Secondly, the model has a very small number of easily-estimated parameters and the equilibrium solution loses stability over a very large region of the parameter space that includes simple but conservative estimates of the parameters. Thirdly, while the model itself is quantitative and new it shares important features with more qualitative, but long-standing, critiques of equilibrium assumptions ([Robinson 1974, Minsky 1992] as well as numerous others before and after).

The profound differences between equilibrium and far-from-equilibrium solutions have equally profound consequences for their analysis that we now summarize:

- **Timescales** Any dynamical process will look like it can reasonably be approximated by an equilibrium model if it is studied over timescales that are too short. For systems that really are always close to equilibrium there is no internal timescale and no memory and so a ‘snapshot’ approach is justified — the last few sets of data points may well be enough to accurately describe the state of the system. However this is not true for far-from-equilibrium

systems. In addition to the timescale introduced by barriers to equilibrium there will also be longer timescales present that are an emergent property determined by the competition between equilibrating and dis-equilibrating forces. For example, in the market model presented here the herding influence is assumed to act over a timescale of weeks or months for any given agent yet the model produces incremental mispricings that can persist for several years before suddenly reversing. A similar effect may well be occurring in the multi-year ‘Business Cycles’ and fluctuations in ‘Animal Spirits’.

- **Unpredictability** All economic systems have a vast number of intrinsically unknowable variables in terms of the knowledge, propensities, bounded rationality, expectations and constraints of the agents within them. When these do not average away they necessarily add a great deal of potential complexity and uncertainty to the future evolution of the system that is generated by the system itself.²⁷ This is especially relevant to the problem of estimating the probability of extreme events in a highly-leveraged financial system or debt-fueled economy.
- **Model Recalibration** Using new data to update or recalibrate a model is an obvious thing to do in any quantitative field and often works extremely well. However if the model being used is not an accurate enough representation of reality then it may cause enormous difficulties. As noted above, the far-from-equilibrium solutions involve long timescales. Over this time a recalibrated equilibrium model may be changed significantly while appearing to do an acceptably good job over much shorter timescales and thus giving the user a completely unjustified sense of confidence in its predictions.²⁸ In

²⁷Furthermore, even very simple dynamical systems, evolving deterministically without any external influences, have been proved to be inherently unpredictable (see for example [Hirsch, Smale and Devaney 2004]). Yet, even in the absence of explicit solutions it is still often possible to quantify both the probabilities of events and the theoretical time limit on *meaningful* predictions.

²⁸The equilibrium models will look like they are working until suddenly they don't. The quote

macroeconomics the problem is magnified because the model output may be used to direct monetary policy and so feed back into the system itself.

The mathematical description of the model lies within a new class of stochastic partial differential equations (see the Appendix) for which closed-form analytic solutions almost certainly do not exist and this may be troubling to many. However, closed-form solutions are the exception rather than the rule in mathematics and the value of approximate numerical solutions in every other quantitative discipline is beyond dispute.²⁹

Equilibrium models in economics have been so resistant to their critics because they are mathematically expedient and most of the time tomorrow really does turn out to be a lot like yesterday! All the day-to-day seemingly minor deviations of actual economic activity from economic theory, unforeseen but enormous economic and financial disasters, and the findings of behavioral economics have not yet coalesced into a single, coherent, organized alternative.

For the field of behavioral economics there is an additional obstacle. It has largely failed to impact orthodox modeling because in an averaged model there is nowhere left to incorporate the rules and heuristics that have been uncovered. But in a non-averaged model these findings are exactly what is needed to define plausible rules for the endogenous dynamics. The moving threshold model presented here can be used as a ‘laboratory’ in which the interactions between various forms of bounded rationality and utility-maximizing can be systematically explored. In fact, because the threshold approach is so flexible it should be possible to generate an entire ecosystem of agents with as much complexity as desired — although complexity often comes at the expense of clearly delineated assumptions or conclusions.³⁰

at the start of the paper is referred to one last time.

²⁹For example, the circumference of an ellipse has no closed-form expression but satellites and planets are always going to move in ellipses.

³⁰The use of thresholds to specify agent dynamics is a very natural choice but is not required. Other types of rules and their corresponding perturbations may yield additional insights.

The framework has also been used to investigate the effects of technical traders and chartists within a financial market since the most widely-used strategies are very easy to recreate using thresholds (especially those involving resistance/support levels). Intuitively, if significant numbers of agents are all reacting in a similar way to certain patterns in the price then the effects may be comparable to those induced by herding in this study. In [Krejčí, Melnik, Lamba and Rachinskii 2013] this was indeed shown to be the case. The model was exactly the same as used here except that the herding effect was removed and instead the slow agents were given realistic draw-up/draw-down technical trading strategies. Once again the equilibrium solution was replaced by far-from-equilibrium boom-bust dynamics and in this particular case it was also possible to rigorously derive the fat-tailed distribution of price returns and other relevant statistics.³¹ Furthermore it was proved that the system has a *critical point* (see [Sornette 2009] for example) which in this context is a critical proportion of agents using such strategies. Once this critical value is exceeded then a system-wide crash becomes inevitable.

The results of both the herding model and the technical trading version help illustrate a more general point. Averaged models in economics as well as their SDE counterparts in finance assume, via equilibration assumptions, that arriving exogenous information will pass through the endogenous ‘filter’ essentially unchanged. However this filter possesses various forms of hysteresis, stickiness, memory effects, contagion processes, bounded rationality, perverse incentives, limits to arbitrage, and lack of independence between agents that can cause very long-lasting distortions. We have demonstrated that boom-bust dynamics is one possibility that arises naturally and robustly without requiring any exogenous cause. However,

³¹There is no clear consensus as to whether technical trading actually produces excess profits [Brock, Lakonishok and LeBaron 1992, Chan, Jegadeesh and Lakonishok 1996, Malkiel 2003] but if some strategies do indeed work then another interesting question is why? It may be caused by the presence of one or more systemic deviations from full rationality or it may in fact be a self-fulfilling prophesy caused by sufficient numbers of technical traders using that rule themselves. Questions such as this can also be investigated.

other kinds of endogenous dynamics may occur in different economic situations. A systematic exploration of all the possibilities would be invaluable in helping to determine which deviations from equilibrium theory actually might require an exogenous explanation.

To this end, it is worth restating more abstractly the process leading from (5) to the non-equilibrium model. This started with the embedding of an equilibrium model that implicitly relies upon averaging and/or a representative agent into a larger framework that explicitly allows for the possibility of endogenous dynamics. Crucially, if the rules governing the endogenous agent dynamics were uncorrelated and independent then the aggregate behavior of the system was unchanged. Then less restrictive and more realistic micro-founded rules were introduced to describe the effect under investigation and these were treated as perturbations to the equilibrium model.

It should be possible to carry out such a program on many other equilibrium models. Preliminary results from a (very simple) rational expectations DSGE model suggested endogenous multi-year fluctuations (analogous to those seen in Figure III) that now correspond to changes in ‘Animal Spirits’ and Business Cycles.

**Department of Mathematical Sciences, George Mason University, MS 3F2,
4400 University Drive, Fairfax, VA 22030 USA. E-mail: hlamba@gmu.edu**

VII Appendix: The model as a stochastic particle system with switching and re-injection

In the original model setup the state of each slow agent is defined by a pair of numbers (an interval) on the real line. An equivalent formulation is to treat each agent’s state as being a single point in the real plane where the position of the i^{th} agent is $(x_i, y_i) = \left(\frac{U_i + L_i}{2} - r_t, \frac{U_i - L_i}{2} \right)$ — the x -coordinate is the distance of $r(t)$

from the center of that agent's interval while the y -coordinate is the semi-width of the interval.

The slow agents/particles move within the set $D \subset \mathbb{R}^2$, $D = \{(x, y) : -y < x < y, y > 0\}$ (see Figure X). The M signed particles (with states $+1$ or -1) move within D subject to three different motions. Firstly there is a stochastic bulk forcing the x -direction that acts upon every particle and consists of the exogenous Brownian forcing B_t and the change in $\kappa\Delta\sigma$. Secondly, each particle has its own independent two-dimensional diffusion process. Thirdly, for agents *in the minority state only*, there is a downward (negative y -direction) drift that is proportional to the imbalance. Note that the particles do not interact locally or collide with one another.

When a particle hits the boundary ∂D it is re-injected into D with the opposite sign according to some predefined probability measure. Finally, when a particle does switch the position of the other particles is kicked in the x -direction by a (small) amount $\pm \frac{2\kappa w_i}{W}$ where the kick is positive if the switching particle goes from the -1 state to $+1$ and negative if the switch is in the opposite direction.

In other words, the bulk stochastic motion is due to exogenous noise changing the price; the individual diffusions are caused by strategy-shifting of the slow agents; the downward drift of minority agents is due to herding effects; the re-injection and switching are the agents changing investment position; and the kicks that occur at switches are due to the change in sentiment affecting the asset price via the linear supply/demand-price assumption.

This particle system provides an interesting physical analog for the model and perhaps for economic systems in general. The boom-and-bust nature of the non-equilibrium dynamics described in this paper can be observed dynamically via a browser-based interactive simulation at <http://math.gmu.edu/~harbir/market.html>.

One does not need to assume that all the slow agents are of equal size, have equal strategy-diffusion, and equal herding propensities. But if one does set $w_i =$

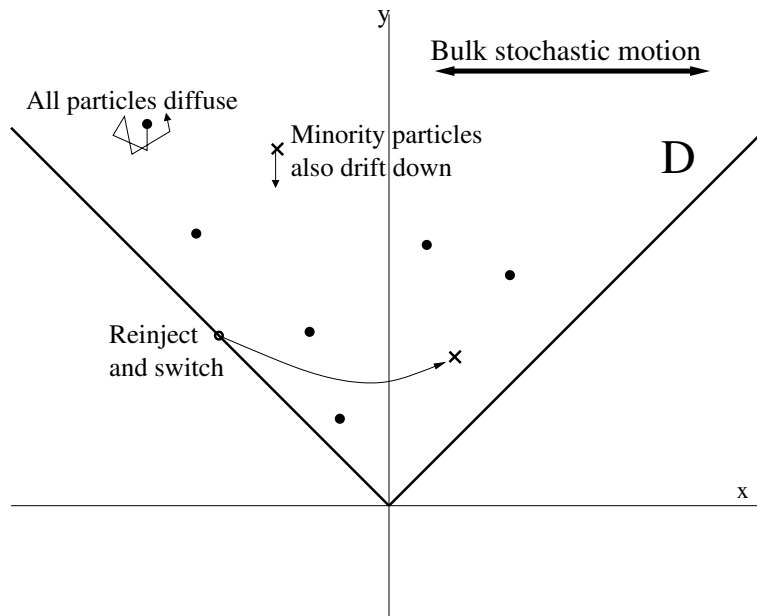


Figure X: All the M signed particles are subject to a horizontal stochastic forcing and they also diffuse independently. Minority particles also drift downwards at a rate proportional to the imbalance. When a particle hits the boundary it is re-injected with the opposite sign and a kick is added to the bulk forcing that can trigger a cascade (see text).

1, $\delta_i = \delta$ and $C_i = C \forall i$ then, in the continuum limit $M \rightarrow \infty$ the system can be defined as a pair of evolving stochastically-forced density functions with coupled source terms and boundary conditions.

References

Akerlof, G. and J. Yellen, "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?," *The American Economic Review*, LXXV (1985), 708–720.

- Banerjee, A.V., “A simple model of herd behavior,” *The Quarterly Journal of Economics*, CVII (1992), 797–817.
- Barberis, Nicholas and Richard Thaler, “A survey of behavioral finance,” *Handbook of the Economics of Finance*, I (2003), 1053–1128.
- Benartzi, Shlomo and Richard H Thaler, “Myopic loss aversion and the equity premium puzzle,” *The quarterly journal of Economics*, CX (1995).
- Bikhchandani, S., D. Hirshleifer, and I. Welch, “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of political Economy*, C (1992).
- Black, F. and M. Scholes, “The pricing of options and corporate liabilities,” *J. Polit. Economy*, LXXXI (1973), 637–659.
- Brock, W., J. Lakonishok, and B. LeBaron, “Simple technical trading rules and the stochastic properties of stock returns,” *Journal of Finance*, XLVII (1992), 1731–1764.
- Brown, G.W., “Volatility, sentiment and noise traders,” *Financial Analysts Journal*, LV (1999), 82–90.
- Buchanan, M., *The social atom* (New York, NY, USA: Bloomsbury, 2007).
- Chan, L.K.C., N. Jegadeesh, and J. Lakonishok, “Momentum strategies,” *Journal of Finance*, LI (1996), 1681–1713.
- Cont, R., “Empirical properties of asset returns: stylized facts and statistical issues,” *Quantitative Finance*, I (2001), 223–236.
- Cross, R., M. Grinfeld, and H. Lamba, “Hysteresis and economics,” *IEEE Control Systems Magazine*, XXIX (2009), 30–43.

- ____, _____, _____, and T. Seaman, “A threshold model of investor psychology,” *Phys. A*, CCCLIV (2005), 463–478.
- ____, _____, _____, and _____, “Stylized facts from a threshold-based heterogeneous agent model,” *Eur. J. Phys. B*, LVII (2007), 213–218.
- Dixit, A.K., “Investment and hysteresis,” *J. Econ. Persp.*, VI (1992), 107–132.
- Earl, P.E., T.-C. Peng, and J. Potts, “Decision-rule cascades and the dynamics of speculative bubbles,” *Journal of Economic Psychology*, XXVIII (2007), 351 – 364.
- Fagiolo, Giorgio, Mauro Napoletano, Marco Piazza, and Andrea Roventini, “Detrending and the distributional properties of US output time series,” Manuscript, LEM Working Paper Series, 2009.
- Fama, E.F., “The Behavior of Stock Market Prices,” *Journal of Business*, XXXVIII (1965), 34–105.
- ____, “Efficient capital markets: A review of theory and empirical work,” *J. Finance*, XXV (1970), 383–417.
- Friedman, Milton, “The methodology of positive economics,” *Essays in positive economics*, III (1953).
- Gigerenzer, G., *Adaptive thinking: Rationality in the real world* (Oxford University Press, USA, 2002).
- Göcke, M., “Various concepts of hysteresis applied in economics,” *J. Econ. Surveys*, XVI (2002), 167–188.
- Grauwe, Paul De, “Booms and busts in economic activity: A behavioral explanation,” *Journal of economic behavior & organization*, LXXXIII (2012), 484–501.

Hirsch, Morris W, Stephen Smale, and Robert L Devaney, *Differential equations, dynamical systems, and an introduction to chaos*, Vol. 60 (Academic press, 2004).

Hommel, Cars H, “Heterogeneous agent models in economics and finance,” *Handbook of computational economics*, II (2006), 1109–1186.

Kahneman, D. and A. Tversky, “Judgement under Uncertainty: Heuristics and Biases,” *Science*, CLXXXV (1974), 1124–1131.

Kahneman, Daniel and Amos Tversky, “Prospect theory: An analysis of decision under risk,” *Econometrica: Journal of the Econometric Society*, XLVII (1979), 263–291.

Keynes, J.M., *The general theory*.

Kirman, A., “Who or what does the representative agent represent?,” *Journal of Economic Perspectives*, VI (1992), 117–136.

Krejčí, P., S. Melnik, H. Lamba, and D. Rachinskii, Analytical Solution for a class of network dynamics with mechanical and financial applications“ (2013).” <http://arxiv.org/abs/1309.4050>.

Lamba, H., “A queueing theory description of fat-tailed price returns in imperfect financial markets,” *Eur. Phys. J. B*, LXXVII (2010), 297–304.

— and T. Seaman, “Rational expectations, psychology and learning via moving thresholds,” *Phys. A*, CCCLXXXVII (2008), 3904–3909.

Lamba, Harbir and Tim Seaman, “Market statistics of a psychology-based heterogeneous agent model,” *International Journal of Theoretical and Applied Finance*, XI (2008), 717–737.

- LeBaron, B., W.B. Arthur, and R. Palmer, “Time series properties of an artificial stock market,” *Journal of Economic Dynamics and Control*, XXIII (1999), 1487–1516.
- LeBaron, Blake and Leigh Tesfatsion, “Modeling macroeconomies as open-ended dynamic systems of interacting agents,” *The American Economic Review*, XCVIII (2008), 246–250.
- Malkiel, B.G., *A Random Walk Down Wall Street: The Time-Tested Strategy for Successful Investing* (1973).
- Malkiel, B.G., “The efficient markets hypothesis and its critics,” *Journal of Economic Perspectives*, XVII (2003), 59–82.
- Mantegna, R. and H. Stanley, *An Introduction to Econophysics* (CUP, 2000).
- Minsky, H.P., “The financial instability hypothesis,” *The Jerome Levy Institute Working Paper*, LXXIV (1992).
- Muth, J.F., “Rational expectations and the theory of price movements,” *Econometrica*, VI (1961).
- Odean, Terrance, “Do Investors Trade Too Much?,” *American Economic Review*, LXXXIX (1999), 1279–1298.
- Piscitelli, L., M. Grinfeld, H. Lamba, and R. Cross, “Exit-entry decisions in response to aggregate shocks,” *Appl. Econ. Lett.*, VI (1999), 569–572.
- Robinson, Joan, *Essays in the theory of economic growth* (Macmillan London, 1962).
- _____, *History versus equilibrium* (Thames Polytechnic, 1974).
- Rubinstein, A., *Modeling Bounded Rationality* (Cambridge MA: MIT Press, 1998).

- Samuelson, P., “Proof That Properly Anticipated Prices Fluctuate Randomly,” *Industrial Management Review*, VI (1965), 41–49.
- Scharfstein, D. and J. Stein, “Herd Behavior and Investment,” *The American Economic Review*, LXXX (1990), 465–479.
- Schleifer, A., *Inefficient Markets* Clarendon Lectures in Economics (OUP, 2000).
- Shiller, Robert, “Do stock prices move too much to be justified by subsequent changes in dividends?,” *American Economic Review*, LXXI (1981).
- Shleifer, A. and R.W. Vishny, “The limits of arbitrage,” *Journal of Finance*, LII (1997), 35–55.
- Simon, Herbert Alexander, *Models of bounded rationality: Empirically grounded economic reason*, Vol. 3 (MIT press, 1982).
- Sornette, Didier, *Why stock markets crash: critical events in complex financial systems* (Princeton University Press, 2009).
- Tesfatsion, Leigh, “Agent-based computational economics: A constructive approach to economic theory,” *Handbook of computational economics*, II (2006), 831–880.