



Volume 33, No. 4
April 2009

ISSN 0165-1889

JOURNAL OF Economic Dynamics & Control

Editors
JAMES BULLARD
CARL CHIARELLA
TIMOTHY COOLEY
CARIS H. HOMMES
PETER N. IRELAND
MICHEL JUILLARD

Advisory Editors
ROGER CRANE
WOUTER J. DEN HAAN
KENNETH L. JUDD
DAVID KENDRICK
DAVID LUENBERGER
ROBERT S. PINDEYCK
BERNARD RUSTEM
STEPHEN J. TURKOVSKY
VOLKER WIELAND

Associate Editors
KLAUS ADAM
ETAN ALTMAN
LANIER BENKARD
JOYDEEP BHATTACHARYA
MICHAEL BINDER
JEAN BOVIN
WILLIAM BRANCHI
FABRICE COLLARD
HERBERT DAVID
FRANK DE KONG
DAVID DE LA CROIX
JIN-CHUAN DUAN
JOHN DUFFY
GEORGE W. EVANS
ESTER FAIA
JESUS FERNANDEZ-VILLaverde
JEAN-MICHEL GRANDMONT
STEVEN R. GRENADER
RICHARD HARTIL
ALAIN HAURIE
VICKY HENDERSON
BURTON HOLIFIELD
SELAHATTIN IMKONKORÇU
LARRY KARP
MASAAKI KIJIMA
ALAN KIRMAN
PAUL KLEIN
PETER KORT
YUKI KIWOKI
HUW LLOYD-ELLIS
THOMAS LUBIK
MARIO MIRANDA
SALVADOR ORTIZ-UGIERA
CHRISTOPHER OTFROK
HASHEM PESARAN
BARBARA ROSSI
JUAN RUBIO-RAMÍREZ
THOMAS RUTHERFORD
SUDIPTO SARKAR
KARL SCHMEDDERS
GERHARD SORGER
STEPHEN SPEAR
LEIGH TESSARON
JAN TUINSTRRA
NICOLAAS J. VREIND
NOMI WILLIAMS
FERNANDO ZAPATERO

CONTENTS

G. Glomm, J. Jung and C. Tran, Macroeconomic implications of early retirement in the public sector: The case of Brazil	777
A. Gamba and M. Tesser, Structural estimation of real options models	798
D. Ladley and K.R. Schenk-Hoppé, Do stylised facts of order book markets need strategic behaviour?	817
A. Matsumoto, Note on Goodwin's 1951 nonlinear accelerator model with an investment delay	832
Y.S. Kim, Exchange rates and fundamentals under adaptive learning	843
A.C. Ghent, Comparing DSGE-VAR forecasting models: How big are the differences?	864
C.B. Durdú, Quantitative implications of indexed bonds in small open economies	883

CONTENTS CONTINUED ON BACK COVER

www.elsevier.com/locate/jedc

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Bubbles and crashes: Gradient dynamics in financial markets

Daniel Friedman^{a,*}, Ralph Abraham^b^a Economics Department, University of California, Santa Cruz, CA 95064, USA^b Mathematics Department, University of California, Santa Cruz, CA 95064, USA

ARTICLE INFO

Article history:

Received 5 June 2007

Accepted 22 October 2008

Available online 3 December 2008

JEL classification:

C63

C73

D53

Keywords:

Bubbles

Escape dynamics

Time varying risk premium

Constant-gain learning

Agent-based models

ABSTRACT

Fund managers respond to the payoff gradient by continuously adjusting leverage in our analytic and simulation models. The base model has a stable equilibrium with classic properties. However, bubbles and crashes occur in extended models incorporating an endogenous market risk premium based on investors' historical losses and constant-gain learning. When losses have been small for a long time, asset prices inflate as fund managers increase leverage. Then slight losses can trigger a crash, as a widening risk premium accelerates deleveraging and asset price declines.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Since their origin, financial markets have suffered from sporadic bubbles and crashes—episodes in which asset prices rise dramatically for no obvious reason, and later plummet (e.g., [Penso de la Vega, 1688/1996](#); [Mackay, 1841/1996](#)). Recent examples include Japan's stock and land price bubbles in the late 1980s, and the US dot.com and telecom bubbles in 2000. Such episodes are important as well as dramatic. As shorelines and river valleys are shaped largely by “100 year events,” so are financial markets, and the economy more generally. For example, the US Securities and Exchange Commission, the segregation of commercial banking from investment banking, and active monetary policy all arose in reaction to the 1929 US stock market crash and subsequent Great Depression (e.g., [Kindleberger, 1978/1989/2000](#)).

Despite their intrinsic interest, financial bubbles and crashes as yet have no widely accepted theoretical explanation. One reason is simply that they are so sporadic. They seldom recur in the same country or market sector within the same generation of participants, so the data are problematic. A second reason is that established theoretical models maintain the assumption of financial market equilibrium. That assumption is difficult to reconcile with the dramatic episodes.

The present paper introduces new models and techniques for studying bubbles and crashes. The focus is on professional fund managers whose payoffs are the risk-adjusted returns they earn on their portfolios. Payoff maximization is not well defined outside equilibrium, so we assume that the managers continuously adjust their risk exposure so as to move up the current payoff gradient. Another non-standard feature is constant-gain learning (e.g., [Cho et al., 2002](#)), also known as exponential average expectations. An endogenous cost of risk is obtained from applying such learning to realized losses.

* Corresponding author. Tel.: +1831 459 4981; fax: +1831 459 5700.

E-mail addresses: dan@ucsc.edu (D. Friedman), abraham@vismath.org (R. Abraham).

Where possible, our models adapt and streamline standard ingredients. For example, we assume a single source of systematic risk, ignore inflation and taxes, and blur the distinctions between cash flow, earnings and dividends.

Section 2 spotlights some relevant strands of literature. Our models draw inspiration from one of the older strands, due to Keynes, Minsky and Kindleberger (KMK), and also from newer strands, especially the emerging agent-based approach.

Section 3 presents the basic model, beginning with the static ingredients and gradient dynamics. We characterize analytically a unique symmetric equilibrium and its comparative statics in the simplest version of the model. An agent-based simulation model illustrates and extends these results. The simulations converge reliably and smoothly to the equilibrium over a very wide range of parameter configurations.

Bubbles and crashes first appear in Section 4. It presents an extension incorporating some KMK-inspired features such as mean-reverting, manager-specific luck (or perhaps skill); investors who are constant-gain learners; and an endogenous risk cost. Analytic approximations suggest typical behavior and lead to conjectures about bubbles and crashes. Simulations confirm such episodes for a wide range of parameter configurations. The intuition is that when losses have been small for a long time, asset prices inflate as fund managers adopt riskier portfolios. When losses occur, as they eventually must, the more exposed managers get hit harder and sell faster. This puts downward pressure on asset prices, and an increasing risk premium accelerates the decline in asset price. This vicious cycle can drive asset price below fundamental value.

Bubbles and crashes are more prevalent in simulations when manager-specific luck is more volatile and longer-lived, when investors have shorter memories, when the economy grows faster and the discount rate is lower, and when current asset prices are higher.

Section 5 summarizes the results and suggests avenues for future research. Appendix A collects mathematical proofs, Appendix B describes extensions of the model, and Appendix C presents the KMK perspective. Additional material, including source and executable code for the simulations, can be found at <http://www.vismath.org/research/landscapedyn/models/markets>.

2. Existing literature

Modern financial economists define the fundamental value V of an asset as the expected present value, given all available information, of the net cash flow the asset generates. The accepted definition of a bubble is a deviation of market price P from V . Crashes are episodes when $B = P - V$ rapidly decreases from a positive value to a zero (or negative) value.

Beyond these simple definitions, consensus is elusive. Most early accounts of bubbles and crashes, e.g., Penso de la Vega and Mackay, emphasize the accompanying bursts of optimism and pessimism, and often seem to assign a causal role to “market psychology.” Absent some insight into (or preferably predictions of) how the bursts of optimism and pessimism arise, this approach does not seem very fruitful.

Some economists deny that bubbles exist, and assert that financial markets are always in equilibrium in the sense that $B = P - V = 0$. They explain famous historical episodes, such as Tulipmania in 17th century Netherlands, as just unusual moves in the fundamental value (e.g., Garber, 1989). Since V is not directly observable, and because the episodes are so sporadic, it is hard to prove (or disprove) this view.

The “rational bubble” models of the 1980s proposed a rather different view (e.g., Blanchard and Watson, 1982; Tirole, 1982). The models allow no intertemporal arbitrage opportunity from one period to the next and traders have the same beliefs, but with an infinite horizon there might be a gap between P and V that grows at an exponential rate. A diverse collection of later papers ascribe bubbles to problems with information aggregation (e.g., Friedman and Aoki, 1992) or to interactions of rational traders with irrational traders (e.g., DeLong et al., 1990; Youssefmir et al., 1998; Brock and Hommes, 1998).

LeRoy (2004) concludes his integrative survey as follows.

We have considered four categories of accounts . . . [for recent apparent bubble and crash episodes]. As explanations, all four categories have problems. . . . Within the neoclassical paradigm there is no obvious way to derail the chain of reasoning that excludes bubbles. An alternative to the full neoclassical paradigm is to think about bubbles in a rational-agent setting—in particular to define fundamentals using the present-value relation—but to break off the analysis arbitrarily at some point rather than following the reasoning to implausible conclusions. The problems with this alternative are obvious: How does one write down formal models in such a setting? Where does one break off the analysis? Which conclusions from neoclassical analysis are to be accepted? We have no answers to these questions. . . . (p. 801)

The present paper resolves LeRoy's conundrum by modelling financial markets that are not always in equilibrium. The agents always seek profit, and most of the time the market is near a steady state, but investors' ongoing learning processes occasionally push the market far from equilibrium. Our modeling choices are guided partly by the empirical literature, and partly by the perspective of Minsky (1975) and Minsky (1982), who drew on ideas from Keynes (1936), later elaborated in Kindleberger (1978/1989/2000). Thus bubbles in our model are touched off by unusual runs of luck by some managers. As the bubble inflates, losses are rare and dazzled investors allow risk premiums to shrink, leading to still higher asset prices. When losses finally do occur, risk premiums expand and asset prices decline, producing wider losses, wider risk premiums, and further losses. Appendix C spells out the KMK perspective on bubbles and crashes.

Our model was also influenced by recent mathematical literature. Drawing on the work of Freidlin and Wentzell (1984), evolutionary game theorists such as Young (1993) and Kandori et al. (1993) showed that some particular transitions among multiple equilibria are much more likely than others in the presence of low-amplitude noise. Sargent (1999) used similar methods to show that even when there is a unique equilibrium, there can be some particular “dominant escape path” that temporarily takes the economy far away from equilibrium. See Williams (2004) for a general exposition.

The key ingredient is constant-gain learning: the weight assigned to the most recent observations remains constant over time. Such perpetual learning is optimal in an environment where unobserved parameters drift over time, but not in a stationary environment (where the optimal weight goes to zero as experience accumulates). Of course, the parameters of an economy will typically shift over time when participants are learning, so constant-gain learning tends to justify itself. It often approximates actual human learning (e.g., Cheung and Friedman, 1997). As we will see in Section 4.1, constant-gain learning is implemented by taking an exponential average of historical data, a common practice among financial analysts. Yahoo Finance, for example, routinely displays exponential average returns at various gains.

There is an analytic downside, however. As Williams (2004, p. 10) notes, “in most cases even the simplest specifications require numerical methods for solution.” For that reason, our models are primarily computational.

Agent-based computational finance has grown rapidly in recent years; LeBaron (2006) and Hommes (2006) each survey more than 100 papers. In this approach, financial markets are modelled as interacting groups of learning, boundedly rational agents, and behavior is described mainly by running computer simulations rather than by solving equations or proving theorems. Hommes begins his survey with the analytic model of Zeeman (1974), which used catastrophe theory to characterize periodic financial market crashes. The dynamics arise from the interaction of two trader types, called fundamentalists (who buy when $B = P - V < 0$ and sell when $B > 0$) and chartists (who buy when P increases and sell when it falls). More recent papers introduce more trader types and explicit learning or evolution, and often try to match quantitative empirical regularities. For example, the Santa Fe Artificial Stock Market (Arthur et al., 1997) uses the genetic algorithm (Holland, 1975) so that agents explore a large finite (discrete) space of technical strategies. With a sufficiently slow update rate the market price converges to fundamental value, while with faster update rates it does not converge and exhibits realistic features such as high trading volume, clustered volatility and leptokurtotic returns distributions.

Brock and Hommes (1998) is another prominent example of the genre. They model evolutionary competition among two to four simple linear forecasting rules, and obtain chaotic price fluctuations. The last part of Hommes's survey discusses what happens as the number of active forecasting strategies (or trading rules) gets large; again, there can be chaotic price fluctuations. See also De Fontnouvelle (2000) and Brock et al. (2005) for other examples of complex dynamics arising from information and learning processes.

3. A base model

This section first lays out and justifies the main static elements of the model. Next it presents the basic dynamic elements, and notes the steady state equilibria. Then it presents simulations illustrating the stability of equilibrium.

3.1. Portfolios and managers' objectives

To begin, assume that there is a single riskless (“safe”) asset with constant return R_0 and a single risky asset with variable return R_1 . Standard theoretical literature often refers to the risky asset as the market portfolio or the unit beta portfolio. The safe asset can be thought of as insured deposits or government securities.

The agents in the model are portfolio managers,¹ each of whom chooses a single ordered variable $x \in [0, \infty)$ that represents the leverage on the risky asset. Thus $x = 1$ means fully invested in the risky asset, while $x > 1$ means leveraged investment (borrowing the safe asset) and $x < 1$ means that the fraction $1 - x$ of the manager's funds are invested in the safe asset.² The manager's portfolio has size $z \geq 0$.

Let F denote the cumulative distribution of choices x , weighted by portfolio size. Then the mean choice among portfolio managers (i.e., normalized asset demand) is $\bar{x} = \int_0^\infty xF(dx) = \int_0^\infty xf(x)dx$, where the middle expression is a Stieltjes integral and the last expression is valid when F has a density f .

For a given realized yield R_1 on the risky asset, the manager obtains gross return $R_G(x) = (1 - x)R_0 + xR_1$. The manager's cost of funds is the risk free rate R_0 plus a risk cost $c(x)$. Two standard interpretations are that the risk cost reflects variance aversion, or the concavity of investors' utility functions. Two alternative interpretations are that an insurance agency (e.g.,

¹ The US financial sector currently comprises about \$43.5 trillion of financial assets, of which households (and non-profit institutions) hold about \$40.5 trillion (US Flow of Funds Accounts, December 7, 2006, Tables L.1 and L.100). Several large pieces, including many sorts of deposits and non-corporate equity, are not traded in financial markets. Most of the tradeable assets are professionally managed. These include the largest piece, about \$11.6 trillion in pension funds, as well as \$4.6 trillion in mutual funds and \$1.1 trillion in life insurance reserves. Only \$5.3 trillion is held directly in corporate equities, of which a large (but undocumented) part is also managed professionally. Fund managers also dominate two of the most rapidly growing segments, private equity and hedge funds.

² The constraint $x \geq 0$ says that fund managers cannot short the market portfolio as a whole. This constraint could play a role in the some of the extensions discussed in Section 5, but appears to be inconsequential in the basic model.

the FDIC) charges a premium which increases in the expected loss claims, or that investors self-insure. All four interpretations seem consistent with the specifications to follow.

The risk cost up to second order is $c(x) = \frac{1}{2}c_2x^2 + c_1x + c_0$. Under the maintained assumption of negligible trading costs it turns out that $c(0) = c_0 = 0$ and $c'(0) = c_1 = 0$ (see Section 4.2), while $c_2 = c''(0) \geq 0$ can be interpreted as the market price of risk (see Section 3.4). In the basic model c_2 is an exogenous constant, but it can vary in extensions of the model. Either way, the net return $R(x)$ enjoyed by a manager choosing leverage x is the gross return less the risk-adjusted cost of funds, so

$$R(x) = x(R_1 - R_0) - \frac{1}{2}c_2x^2. \tag{1}$$

The manager's objective or payoff function ϕ depends positively on the net return. Extensions of the model allow it also to depend on the portfolio size and on relative performance, but the basic model ignores such complications and simply sets

$$\phi(x, F) = R(x). \tag{2}$$

The current distribution F of choices by managers affects managers' payoff via the R_1 term in (1), as explained next.

3.2. Asset price and return

The fundamental value V of a share of the asset is the present value of the per share earnings, i.e., revenues less economic costs, including the reinvestment costs necessary to maintain growth but excluding the rental rate of owned capital. Thus earnings are the residual cash flow available to the owners of the underlying real assets. In this simple model, earnings are synonymous with dividends, profit, return to capital, and net cash flow.

In the basic model, earnings are a continuous stream that grows forever at a constant rate g_s . Future earnings are discounted at some rate $R_s > g_s$, discussed below. The number of shares is normalized so that per share earnings are 1.0 at time 0. The initial fundamental value is the integral of the discounted earnings stream, $V(0) = \int_0^\infty 1e^{g_s t} e^{-R_s t} dt = (R_s - g_s)^{-1}$. At time $t > 0$ the fundamental value is similar except that the earnings stream starts at $e^{g_s t}$, so

$$V(t) = V(0)e^{g_s t} = e^{g_s t} / (R_s - g_s). \tag{3}$$

Asset supply comes from fundamental-oriented market participants such as issuers of stocks and bonds, and perhaps individual investors. We do not model them in detail, but simply assume that the net asset supply function has constant elasticity $a > 0$ so, after suitable normalization, it is $S = (P/V)^a$. Normalized asset demand by fund managers is, as already noted, $D = \bar{x}$. Solving $S = D$, we can write the price of the risky asset as

$$P = V\bar{x}^\alpha, \tag{4}$$

where $\alpha = 1/a > 0$. Thus asset price is equal to, less than or greater than fundamental value whenever normalized demand \bar{x} for the risky asset is equal, less than or greater than 1.0. An interpretation is that the fund managers exert buying pressure whose intensity is parametrized by α .

It is now straightforward to calculate R_1 , the return on the risky asset. By definition, it is the dividend yield plus the capital gains rate. The dividend yield is simply earnings per dollar invested, $e^{g_s t} / P(t) = (R_s - g_s)\bar{x}^{-\alpha}$. Use the notation $\dot{y} = dy/dt$ and take the log-derivative of (4) to obtain the capital gains rate $\dot{P}/P = \dot{V}/V + \alpha\dot{\bar{x}}/\bar{x} = g_s + \alpha\dot{\bar{x}}/\bar{x}$. Hence the realized yield on the risky asset is

$$R_1 = (R_s - g_s)\bar{x}^{-\alpha} + g_s + \alpha\dot{\bar{x}}/\bar{x}. \tag{5}$$

The first term is the dividend yield, the second term captures capital gains due to economic growth, and the third term reflects capital gains due to financial market activity. Note that R_1 is higher when mean leverage \bar{x} is lower or is increasing more rapidly. It is equal to the discount rate R_s when $\bar{x} = 1$ and is steady.

What is the discount rate R_s ? One component is the riskless rate $R_0 \geq 0$, which reflects investors' marginal rate of time preference. We write

$$R_s = R_0 + d_R, \tag{6}$$

where the term $d_R \geq 0$ represents all other factors. These factors include g_s , since economic growth is known and economy-wide, and we impose the constraint $R_s - g_s > 0$ to ensure that fundamental value is well defined in (3).

3.3. Gradient dynamics

As noted in the Introduction, neoclassical financial models assume that asset prices are always in equilibrium. Agents in such models choose portfolios to maximize the expected present value of terminal wealth, or utility. Expectations are well defined given a known equilibrium price path (possibly stochastic), but it is hard to reconcile such knowledge with bubbles and crashes. Our approach is instead to specify how agents adapt to any price history, equilibrium or not, and to find conditions under which the process leads towards or away from equilibrium.

So how might a portfolio manager adjust leverage x , given payoff $R(x)$? For discrete unordered choices, the standard process is replicator dynamics (e.g., Fudenberg and Levine, 1998), but when the choice variable x is continuous and ordered as here, mean field or gradient dynamics are standard (e.g., Sonnenschein, 1982; Aoki, 2004; Friedman and Ostrov, 2008).

Gradient dynamics are especially natural for fund managers, since they adjust leverage x mainly by selling or buying the risky asset. To the extent that the risky asset is not perfectly liquid, the per-share trading cost increases with the net amount traded in a given short time interval. If the increase is linear, then the adjustment cost (net trade times per share trading cost) is quadratic. It turns out that such quadratic adjustment costs are the key condition to obtain exact gradient dynamics (Friedman and Yellin, 1997, Proposition 1), rather than approximate gradient or sign-preserving dynamics.

We shall assume gradient adjustment without explicitly modelling trading frictions, since the frictions presumably are small relative to realized returns in (1). Thus portfolio managers continuously adjust their leverage choice x , moving up the payoff gradient at a rate $\dot{x} = dx/dt$ proportional to the slope $\phi_x = \partial\phi/\partial x$. Normalizing the rate to 1.0 gives us the expression

$$\dot{x} = \phi_x(x, F). \tag{7}$$

To obtain population-level dynamics in simplest form, assume for the moment that the fund does not retain the gross return but instead passes it through to its clients, who never withdraw or invest additional funds. Then the fraction $F(x, t)$ of managed funds that have leverage x or less changes at time t only because managers are adjusting their leverage. The rate of change $F_t = \partial F/\partial t$ is simply the density $F_x = \partial F/\partial x$ of funds with leverage x times their (downward) adjustment speed $-\phi_x$. Thus we obtain the master equation

$$F_t(x, t) = -F_x(x, t)\phi_x(x, F). \tag{8}$$

Appendix B begins with a generalization of (8) that permits retained earnings, distributions with mass points, and other complications.

Plugging (1) and (5) into (2), we obtain the fund manager's payoff function

$$\phi(x, F) = x[(R_s - g_s)\bar{x}^{-\alpha} + g_s - R_0 + \alpha\dot{\bar{x}}/\bar{x}] - \frac{1}{2}c_2x^2, \tag{9}$$

with gradient

$$\phi_x = (R_s - g_s)\bar{x}^{-\alpha} + g_s - R_0 + \alpha\dot{\bar{x}}/\bar{x} - c_2x. \tag{10}$$

3.4. Equilibrium

We focus on symmetric steady states of (8), i.e., states that do not change over time and in which all fund managers choose the same leverage $x = \bar{x}$. To put it more formally, we seek degenerate stationary solutions to the master equation: for all $t \geq 0$, $F(x, t) = 0$ for $x < \bar{x}$ and $F(x, t) = 1$ for $x \geq \bar{x}$. Allowing retained earnings fortunately does not change the analysis.

In steady state we must have $\phi_x = 0$ and, of course, $\dot{\bar{x}} = 0$ at that point. Inspection of (10) shows that one possibility is that $c_2 = 0$ and $x = \bar{x} = 1$ for all fund managers, so (5) collapses to $R_1 = R_s = R_0$ and (1) collapses to $R(x) = 0 \forall x$. This trivial equilibrium makes sense when there really is no risk.

One obtains a more interesting symmetric steady state at $x = \bar{x}$ when the marginal risk cost c_2x equals the marginal steady state net return $R_1 - R_0 = (R_s - g_s)\bar{x}^{-\alpha} + g_s - R_0$. We now show that there is a unique such steady state x^* , and derive expressions for how it varies in the underlying parameters c_2 , α , g_s , R_0 and d_R .

Proposition 1. *Given fixed positive parameters c_2 , α , g_s , R_0 , and d_R such that $R_s = R_0 + d_R > g_s$, there is a unique $x^* > 0$ such that the degenerate distribution at x^* is a symmetric steady state solution to the master equation (8). Moreover, x^* decreases in c_2 and increases in d_R . It increases in α and g_s and decreases in R_0 iff $x^* < 1$.*

Appendix A contains a formal proof and specific formulas for the comparative statics; here we just give the intuition. The key condition is that the payoff gradient is zero at $x = x^*$. Set $x = \bar{x}$, $\dot{\bar{x}} = 0$ and $\phi_x = 0$ in (10), and rearrange slightly to obtain

$$(R_s - g_s)x^{-\alpha} + g_s - R_0 = c_2x. \tag{11}$$

As shown in Fig. 1, the right hand side of (11) is a ray from the origin with positive slope c_2 . The left hand side is a hyperbola with the y -axis as the vertical asymptote and the line $y = g_s - R_0$ as the horizontal asymptote. Clearly there is a unique point of intersection $x^* > 0$.

Using our assumption $d_R = R_s - R_0 \geq 0$, it is not difficult to show that $x^* = \bar{x} > 1$ and $P > V$ when $c_2 > 0$ is sufficiently small. Here the high steady state asset price P reflects fund managers' desire to leverage their portfolios given the low risk cost. Likewise, larger values of the c_2 parameter imply $x^* = \bar{x} < 1$ and $P < V$. There is some intermediate value of c_2 such that mean leverage is $\bar{x} = 1$ and $P = V$. One can see from (10) or (11) that this implies $R_s = R_1 = R_0 + c_2$. In this "long-run" equilibrium, c_2 looks like the standard risk premium, e.g., the market price of risk in the Capital Asset Pricing Model.

Implicitly differentiating (11) with respect to parameters such as c_2 , one obtains expressions such as $\partial x^*/\partial c_2 = -x/[c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] < 0$. The last expression shows that increasing c_2 always decreases x^* by some proportion.

The proportion is larger when c_2 is small and x^* is large. The next section shows that when c_2 is endogenous, this proportional effect leads to bubbles and crashes. But first we present a simulation model that confirms the symmetric equilibrium even when earnings are retained and time is discrete. The model also illustrates the stability of the symmetric equilibrium and provides a benchmark for assessing the impact of endogenous risk premiums.




Fig. 1. Equilibrium risk position x^* . Parameter values determining the hyperbola are $R_0 = 3\%$, $R_s = 6\%$, $g_s = 0$ and $\alpha = 2.0$. The two rays are for alternative c_2 values 1% and 10%. The equilibrium risk position is the horizontal component of the intersection of the ray and hyperbola.

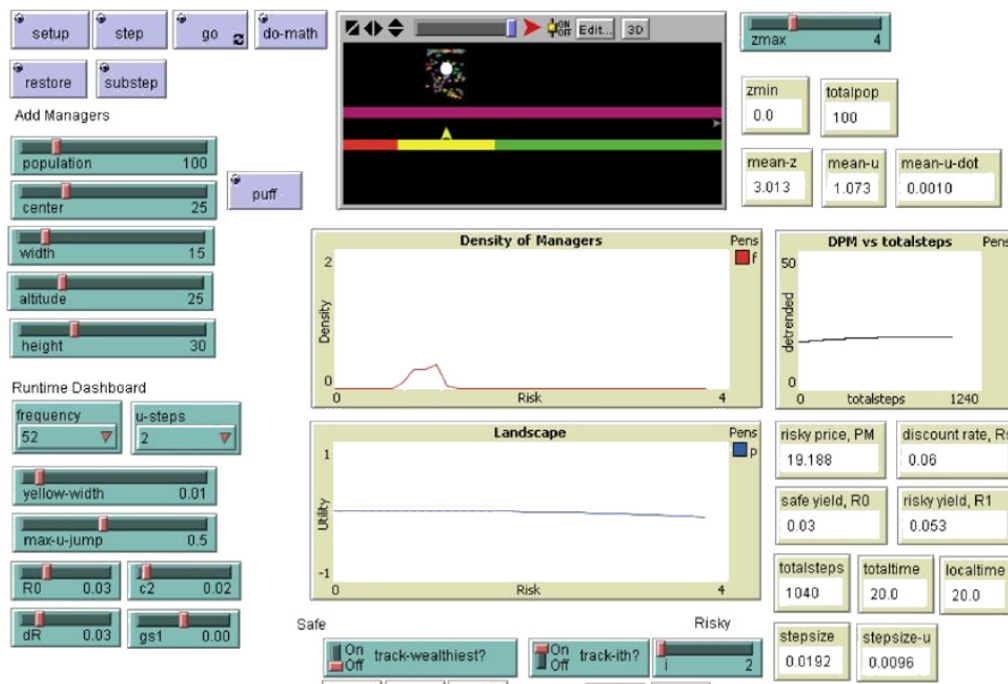


Fig. 2. User interface for base simulation.

3.5. Simulation results

Using Netlogo (<http://ccl.northwestern.edu/netlogo/>), we created a simulation model that closely parallels the basic analytical model just presented. Sliders allow the user to select parameter values and display options. Full documentation as well as executable code can be found at <http://www.vismath.org/research/landscapedyn/models/markets/>.

In brief, the simulation consists of fund managers $i = 1, \dots, M$ whose leverage x_i (horizontal coordinate) and portfolio size z_i (vertical coordinate) are floating point numbers that adjust in discrete time. The simulations drop the unrealistic pass-through assumption on earnings and instead assume that managers retain all earnings (and absorb any losses). The user chooses the frequency (daily, weekly, monthly, quarterly, or annual) and the number of time steps per period (up to 64). With N steps per year, the managers' annual returns R are computed as in Eqs. (1)–(5) and are adjusted to $r = (1 + R)^{1/N} - 1$ per time step.

Fig. 2 shows a sample simulation of the basic model, using parameter values $c_2 = 0.02$, $R_0 = 0.03$, $d_R = 0.03$, $g_s = 0$ and $\alpha = 2.0$. The simulation is at weekly frequency (Freq = 52) with just one time step per period (u-steps = 1). The initial population of managers is $M = 100$, uniformly distributed in the (x, z) rectangle $[0.2, 1.4] \times [0.4, 1.6]$, set via the sliders

Table 1
Simulation results for Model 0.

Parameters		Clumped steady state prediction			Simulation mean ± std. dev.	
Variable	Value	<i>V</i>	<i>P</i> *	<i>x</i> *	<i>P</i>	<i>x</i>
Baseline	See note	16.67	13.37	0.896	13.37 ± 0.01	0.896 ± 0.001
Population	15	16.67	13.37	0.896	13.37 ± 0.02	0.896 ± 0.001
	100	16.67	13.37	0.896	13.36 ± 0.02	0.895 ± 0.001
Ro	0.01	25.00	18.76	0.866	18.75 ± 0.02	0.896 ± 0.001
	0.05	12.50	10.45	0.914	10.45 ± 0.01	0.896 ± 0.001
dR	0.01	25.00	14.65	0.765	14.64 ± 0.01	0.765 ± 0.000
	0.05	12.50	12.50	1.000	12.49 ± 0.02	1.000 ± 0.001
gs	−0.04	10.00	8.59	0.927	8.59 ± 0.01	0.927 ± 0.000
	0.04	50.00	32.79	0.810	32.77 ± 0.06	0.810 ± 0.001
c2	0.01	16.67	23.83	1.196	23.47 ± 0.49	1.187 ± 0.013
	0.2	16.67	6.47	0.623	6.47 ± 0.00	0.623 ± 0.000

Note: Baseline values are Pop = 30, Ro = dR = 0.03, gs = 0 and c2 = 0.05. Means ± standard deviation are taken over 10 centuries of weekly observations, with the first two decades deleted to reduce the impact of the initial population distribution.

labelled “population”.³ The “step” button allows the user to see one time period (e.g., week) at a time, while the “go” button allows the simulation to proceed steadily. Other display options are controlled by green sliders, as explained in the on-line user’s manual. The tan graphs and monitors display simulation results.

In the sample simulation, the managers move smoothly towards $x^* \approx 1.08$ and maximal z , as asset price (displayed on the screen as DPM or detrended PM) converges to $1.08^2 / (0.06 - 0) \approx 19.4$. The figure shows that they are fairly close within two decades.

Proposition 1 suggests smooth convergence towards a symmetric choice at x^* and associated asset price P^* , and not just for the case shown in Fig. 2, but for a wide range of parameter values. To test this suggestion, we moved each of the five main parameters in turn, up to a rather high (but still admissible) level, and down to a rather low level, e.g., $c_2 = 0.01, 0.20$ compared to the baseline $c_2 = 0.05$. For each configuration we ran 10 weekly simulations each for 100 years. Table 1 shows the steady state predictions, and the observed results (dropping the first 20 years, which depend more on arbitrary initial conditions). In every case the observed standard deviation becomes quite small and the observed mean is quite close to the steady state prediction.

4. Endogenous risk

Dynamics become more interesting when the risk cost c_2 is endogenous. To that end, we incorporate into our basic model three new features suggested by the empirical evidence and by the KMK perspective.

4.1. Streaks, losses, and learning

The first new feature is that each manager has streaks in which she underperforms or outperforms the market. The idea is that different portfolios with the same leverage can have different realized returns; some managers temporarily are luckier (or more attuned to new opportunities) than others. Thus, instead of receiving the uniform return in (1), manager i receives net return

$$R_i(x) = x(R_1 - R_0 + \varepsilon_i) - \frac{1}{2}c_2x^2. \tag{12}$$

The manager’s idiosyncratic component ε_i is Ornstein–Uhlenbeck, i.e., mean reverting in continuous time.⁴ If the most recent known value is $\varepsilon_i(t - h)$, then the current value is the random variable

$$\varepsilon_i(t) = e^{-\tau h} \varepsilon_i(t - h) + \sqrt{\frac{1 - e^{-2\tau h}}{2\tau}} \sigma v, \tag{13}$$

³ To spell it out, the “center” slider sets the middle of the x (here called u) coordinate as a percent of screen width, here 4.0. Likewise “altitude” sets the middle of the initial z distribution, and “width” and “height” control the bounds on the rectangle. The “puff” button allows the user to build up arbitrary distributions from a sequence of uniform rectangular distributions.

⁴ A standard result in the empirical literature is that the idiosyncratic component of fund managers’ performance is mean reverting; see, e.g., Bollen and Busse (2005).

for some given volatility parameter $\sigma > 0$ and decay parameter $\tau > 0$, and an independent realization v from the unit normal distribution (Feller, 1971, p. 336). Baseline values are $\sigma = 0.20$, roughly the historical annualized volatility on the S&P500 stock index, and $\tau = 0.7$, implying a half-life of about 1 year for the idiosyncratic component.

The second new feature is loss aversion, where loss is defined as negative gross return.⁵ Since $R_{Gi} = (R_1 - R_0 + \varepsilon_i)x_i + R_0$ is the gross return that manager i currently earns on her portfolio, her loss is $L_i = \max\{0, -R_{Gi}\}$, the shortfall from zero.

Constant-gain learning is the third new feature. Investors seem to judge managers by the overall historical track record, with greater emphasis on more recent results.⁶ The natural formalization is an exponential average. In continuous time the exponential average loss for manager i is

$$\hat{L}_i(t) = \eta \int_{-\infty}^t e^{-\eta(t-s)} L_i(s) ds, \tag{14}$$

where the parameter η is the memory decay rate. Using the definition (14) and a little calculus, the reader can verify that over a time horizon h in which L is constant (or only observed once), the exponential average is updated from the previous exponential average loss $\hat{L}_i(t-h)$ as follows:

$$\hat{L}_i(t) = e^{-\eta h} \hat{L}_i(t-h) + (1 - e^{-\eta h}) L_i(t). \tag{15}$$

Update rule (15) defines “constant-gain learning” with gain $(1 - e^{-\eta h})$.

We now specify perceived risk c_2 as proportional to market-wide perceived losses,

$$c_2 = \beta \hat{L}_T(t), \tag{16}$$

where the parameter $\beta > 0$ reflects investors’ sensitivity to perceived loss, and $\hat{L}_T(t)$ is the perceived loss $\hat{L}_i(t)$ averaged across managers i weighted by portfolio size z_i . Baseline parameter values are $\beta = 2$ and $\eta = 0.7$.⁷

4.2. Equilibrium

Steady states in the current model are more intricate than in the base model. The new parameters for memory decay (η) and sensitivity to losses (β) help determine the risk cost c_2 , and it (along with the streak decay (τ) and volatility (σ) parameters) affects the perceived losses. As a first step towards characterizing steady states, the next proposition computes expected (hence, in steady state, perceived) loss for given c_2 .

Proposition 2. *In steady state with given c_2 , a manager with leverage x incurs expected loss $q(x|c_2) = (x\sigma/\sqrt{2\tau})\psi(z^0(x))$, where $z^0(x) = (-\sqrt{2\tau}/\sigma)[R_0(1/x - 1) + g_s + (R_s - g_s)(x^*)^{-\alpha}]$ and x^* is defined from c_2 in Proposition 1.*

The “wedge” function ψ is the definite integral of the cumulative unit normal distribution Φ ; see Appendix A for an explicit formula and a proof of the proposition.

Corollary. *The expected loss is zero and has derivative zero at $x = 0$. It is a convex increasing function for $x > 0$.*

The Corollary justifies the approximation first used in Eq. (1) that the risk cost $c(x)$ is quadratic with $c(0) = c'(0) = 0$. The intuition is that a tiny exposure to the risky asset requires at most a second order tiny amount of insurance, because portfolio returns will be in the vicinity of the safe return R_0 and outright losses are virtually impossible. The proof of the Corollary in Appendix A works even when an arbitrary distribution function replaces Φ in the formula for ψ , and the net return contains an arbitrary risk premium. The specific formulas of Proposition 2 come from the Ornstein–Uhlenbeck process, but the qualitative features are rather general.

Proposition 2 allows us to approximate a steady state as follows. Fix the exogenous parameters $\alpha, \beta, \eta, \sigma, \tau, g_s, R_0$ and d_R , and choose a reasonable initial estimate \hat{c}_2 of steady state c_2 . Suppose that all managers choose leverage $x = x^*(\hat{c}_2)$, where $x^*(\cdot)$ is defined implicitly in Eq. (11). Using the function $q(\cdot|c_2)$ defined in Proposition 2, note that $q(x^*(\hat{c}_2)|\hat{c}_2)$ approximates the steady state average of \hat{L}_T . Inserting that approximation into Eq. (16), one obtains a more refined estimate \check{c}_2 of steady state c_2 . On iterating (or using Newton’s Method), one obtains consistent values $c_2^{**} = \beta q(x^*(c_2^{**})|c_2^{**})$ and $x^{**} = x^*(c_2^{**})$.

For example, with baseline parameters we have $c_2^{**} = 0.058$ and $x^{**} = 0.866$. To check that these values are consistent, note that $(R_s - g_s)x^{-\alpha} + g_s - R_0 = (0.06)(0.866)^{-2} - 0.03 = 0.050 = 0.058 * 0.866 = c_2x$ so Eq. (11) holds, and that $\beta q(x^{**}|c_2^{**}) = \beta x^{**}(\sigma/\sqrt{2\tau})\psi(z^0(x^{**})) = 2 * 0.866 * 0.2 * 1.4^{-0.5}\psi(z^0(0.866)) = 0.293\psi(-0.501) = 0.293 * 0.198 = 0.058 = c_2^{**}$.

⁵ Loss aversion is a staple of the new behavioral finance literature; see for example Shefrin (2002) and Camerer et al. (2003). But finance theorists and practitioners have always known that investors respond more to downside risk (or losses) than to variance per se; see for example Levy and Markowitz (1979).

⁶ Investors in mutual funds chase returns, especially those funds that recently were top performers (Chevalier and Ellison, 1997; Sirri and Tufano, 1998; Karceski, 2002). Pension funds are less extreme in chasing top performance, but are harsher on funds that incur losses (Del Guercio and Tkac, 2002). Chevalier and Ellison show that fund managers, especially those of new funds, respond by increasing their risk stance when their returns trail their peers’. Sirri and Tufano note that the cross-sectional effects are supplemented by an industry effect: inflows/outflows from the equity mutual fund sector respond to bull and bear markets. Their estimates imply that inflow drops 70% in an average recession and increases 50% following an average bull market run.

⁷ Copeland et al. (2005, figure 6.10) show an average ex post premium of about 5%, but with considerable variability. The figure shows spells of several years in negative territory and several years in the double digit range.

We used $z^0(0.866) = -((2 * 0.7)^{0.5}/0.2) * (0.03 * (1/0.866 - 1) + 0 + (0.06) * (0.866)^{-2}) = -0.501$ and numerically integrated the cumulative unit normal distribution to obtain $\psi(-0.501) = 0.198$.

4.3. Dynamics

Of course, the distribution clumped at x^{**} cannot be an exact solution of the model with $\sigma > 0$, since the model is stochastic. The real question is whether the actual distribution remains nearby. To answer that question, we incorporate the new features into the simulation described in the previous section.

Fig. 3 shows typical asset price behavior in the extended model with idiosyncratic returns and an endogenous price of risk, using the baseline parameter values.

In the simulation, the managers circulate constantly, mostly in the range $0 < x < 3$, but their average fluctuates around $x^{**} \approx 0.866$. Asset price usually bounces around the predicted steady state value $P^{**} = Vx^{**\alpha} = (R_s - g_s)^{-1}(x^{**})^2 \approx (0.06)^{-1}0.866^2 \approx 12.5$, but occasionally it rises much higher or falls much lower.

The simulation software allows us to re-examine dramatic price movements after they occur, using the rewind button. The first simulation tried while writing this passage used baseline parameter values. The asset price P rose to $14.7 > P^{**} = 12.5$ by the end of year 12, and year 13 began with the managers somewhat extended at $\bar{x} = 0.94 > 0.866 = x^{**}$. Losses had begun to increase, and c_2 rose modestly from 6.2% at the beginning of the year to 9.0% by April. At that point, P began to decline at an accelerating pace. By the end of that unlucky year, the market was in free fall with c_2 approaching 100% and P half its former value. The market bottomed out in the first quarter of year 14 with P under 3 and \bar{x} below 0.4. A gradual recovery then brought P near $P^{**} = 12.5$ by year 21, where it hovered for the next several decades.

Dissecting dozens of similar examples suggests that the most likely way to leave the vicinity of the steady state begins with a streak of good luck for a few of the larger traders. In the absence of very bad luck for most other traders, asset price P exhibits a steady to rising trend. If the trend persists, the perceived loss $\hat{L}_T(t)$ declines and so does the risk premium c_2 . As shown in Fig. 1, this increases x^* and, as managers increase their average risk stance \bar{x} , it also increases P . As \bar{x} gets large relative to x^{**} , the effect attenuates, and then the market is vulnerable to a streak of moderately bad luck for some of the larger investors. Once P starts to decline, the process goes into reverse and accelerates. Returns turn negative and losses mount, so c_2 rises and x^* declines, dragging down P and leading to more negative returns and losses. The vicious cycle slows when x^* gets so low that it becomes relatively unresponsive to further increases in c_2 (recall Fig. 1). As \bar{x} stabilizes at a low level, returns turn positive, perceived losses decline, and P gradually heads back towards its steady state value x^{**} .

4.4. Statistical results

To investigate these impressions, we examined 18 variants of the baseline parameter configuration. As in Table 1, we ran 10 centuries of weekly simulations for each variant. To maintain comparability over time and across simulations with different growth rates, we work with detrended asset price, replacing P by $Pe^{-g_s t}$. Somewhat arbitrarily, we defined a crash as a decline in detrended P of at least 50% from its highest point within the last half year. (By comparison, the maximum drawdown of the Nikkei index was a bit less than 50% from December 1989 to September 1990, so it did not quite crash by our definition. Likewise, the initial decline in Nasdaq from its 5048 peak in March 2000 was less than 50%, but the index fell 59% from September 2000 to March 2001, which does meet our definition of a crash.)

Table 2 reports the results. Under baseline parameters, average asset price and risk position are modestly higher than predicted by the symmetric steady state, and crashes occur on average only once every other century. Changing the population size has no effect on the steady state predictions, but it does change the simulation averages slightly, and has a



Fig. 3. Detrended asset price. A 100 year simulation of Model 1 using baseline parameters $M = 30$, $R_0 = d_R = 0.03$, $g_s = 0$, $\alpha = 2$, $\sigma = 0.20$, $\beta = 2.0$ and $\tau = \eta = 0.7$.

Table 2
Simulation results for Model 1.

Parameters		Clumped steady state prediction			Simulation mean ± std. dev.		Crashes/century
Variable	Value	V	P^*	x^*	P	x	
Baseline	*	16.67	12.49	0.8658	14.8 ± 3.4	0.94 ± 0.11	0.5
Population	15	16.67	12.49	0.8658	13.8 ± 4.1	0.90 ± 0.14	1.6
	100	16.67	12.49	0.8658	15.3 ± 1.6	0.96 ± 0.05	0
R_0	0.01	25.00	16.15	0.8037	14.2 ± 3.3	0.92 ± 0.11	0.9
	0.05	12.50	10.31	0.9083	14.9 ± 2.8	0.94 ± 0.09	0.2
d_R	0.01	25.00	14.30	0.7563	17.4 ± 3.9	0.83 ± 0.10	0.7
	0.05	12.50	11.37	0.9536	13.2 ± 2.5	1.02 ± 0.10	0.3
g_s	-0.04	10.00	8.11	0.9005	9.2 ± 1.5	0.96 ± 0.08	0
	0.04	50.00	30.76	0.7843	38.8 ± 14.1	0.86 ± 0.17	1.3
sigma	0.05	16.67	23.66	1.1914	23.1 ± 1.2	1.18 ± 0.03	0
	0.4	16.67	8.69	0.7223	14.7 ± 3.6	0.93 ± 0.12	0.6
tau	0.1	16.67	7.52	0.6717	57.1 ± 6.9	1.78 ± 0.51	1.7
	3	16.67	17.86	1.0352	16.4 ± 1.2	0.99 ± 0.04	0
eta	0.1	16.67	12.47	0.8651	14.9 ± 2.3	0.94 ± 0.07	0
	3	16.67	12.48	0.8654	13.5 ± 4.5	0.88 ± 0.18	4.5
beta	1	16.67	15.20	0.9549	18.2 ± 3.8	1.04 ± 0.11	0.2
	5	16.67	9.65	0.7608	11.0 ± 2.6	0.81 ± 0.11	1.4
alpha	1	16.67	13.42	0.8051	15.8 ± 1.9	0.95 ± 0.11	0
	4	16.67	11.81	0.9175	10.4 ± 5.7	0.86 ± 0.14	6.1

Note: Baseline values are as in Fig. 3. Mean ± standard deviation entries are taken over 10 centuries of weekly observations.

substantial effect on crash frequency. Raising population size to 100 virtually eliminated crashes, while lowering it to 15 tripled their frequency.⁸

Recall that the (detrended) fundamental value is $V = (R_s - g_s)^{-1}$, so the direct effect of a shift in either the R_0 or the d_R component of the discount factor R_s is a shift V in the opposite direction. The table shows that these direct effects on predicted (detrended) asset price P^{**} are partially offset via x^{**} . Actual average P barely responds to R_0 but tracks the d_R predictions fairly well, and crash frequencies shift modestly in the expected direction. The underlying growth rate parameter g_s operates similarly, except that the impact via x^* reinforces the impact via V , and the average price when $g_s = 0.4$ is considerably higher than the already high forecast and has a very high standard deviation.

Recall that the unconditional variance of a manager's luck is $\sigma^2/(2\tau)$. Lowering the instantaneous volatility σ to 0.05 reduces the risk cost and substantially raises steady state risk stance x^{**} and price P^{**} , and the actual averages stay rather close to these predictions. Raising the decay rate τ to 3 produces roughly similar predictions and actual results. Raising σ to 0.40 moves everything in the opposite direction, and increases crash frequency and variability slightly above the baseline. The most interesting exercise here is lowering the decay rate to 0.1. The symmetric steady state prediction reflects the greater risk arising from the 7 year half-life of luck. However, in the simulations we see large, long-lived bubbles punctuated by occasional crashes (17 over the observed millennium), resulting in extremely high and variable risk stance and asset price.

Results for the remaining three parameters are also dramatic. When investors have a long memory (low gain), crashes disappear and prices are steadier (and a bit higher) than in the baseline. With short memories (high gain, half-life about 3 months), serious crashes hit every couple of decades. Changes in β , investors' sensitivity to realized losses, affect predicted and actual performance in the direction one might expect but not as strongly; the direct effect on risk cost is partially offset by endogenous response via x^{**} . Finally, the predicted impact of the elasticity or price pressure parameter α is as one might expect, while the actual impact is somewhat stronger. In particular, crashes are frequent and prolonged when $\alpha = 4$.

⁸ This is reminiscent of Egenter et al. (1999)'s findings on finite size effects.

4.5. Summary and interpretation

The simulation results suggest that Model 1 has three distinct modes of behavior. The first mode, stable convergence, is similar to behavior in Model 0. From reasonable initial states, the market converges to a neighborhood of the symmetric steady state and remains there. Lucky (or unlucky) streaks are not large enough to break away. Parameters conducive to this mode include a large population M , low volatility σ , high decay rate τ for luck, and a low η . In terms of the KMK perspective, the last parameter suggests that investors are not easily dazzled and the other three suggest that unusual opportunities (modelled as positive individual luck) are small and fleeting relative to the size of the market.

Convergence in this first mode often is to prices and risk positions somewhat higher than in the symmetric steady state prediction. Apparently the prediction is biased to the extent that managers are heterogeneous rather than symmetric. In the simulation, the luckier managers tend to have larger portfolios z and choose larger x than in the symmetric steady state, pushing up weighted average risk position x^* and asset price P .

A second behavioral mode, occasional bubbles and crashes, can be seen even in the baseline configuration. (The symmetric steady state prediction ignores crashes, which tends to offset the biases due to ignoring heterogeneity and bubbles.) As described in Section 4.3, this mode is reminiscent of escape dynamics as well as the KMK perspective. It prevails for moderate parameter values, covering most of the range a priori considered plausible.

A third behavioral mode, violent thrashing, prevails for some extreme parameter configurations. For example, with very large α or η , the observed asset price is usually far above the symmetric steady state value or far below, with rapid and erratic transitions. Normalcy is rare in this mode.

5. Discussion

We use simulations, supplemented by some analytics and statistics, to explore a perspective on financial market bubbles and crashes originating in the writings of Keynes, Minsky and Kindleberger (KMK). We begin with Model 0, a base model featuring a gradient process by which fund managers adjust the riskiness of their portfolios.

Model 0 is very stable. Analytic results suggest, and simulations confirm for a wide range of exogenous parameters, that the asset price converges quickly and reliably to a level proportional to the fundamental value. The proportion is a decreasing function of the risk cost parameter c_2 . The model never bubbles or crashes.

Model 1, an extension of Model 0, features an endogenous risk premium driven by constant-gain learning. It also has a unique steady state, but its dynamics are quite different. Although asset price usually is near its steady state value, there are recurrent episodes in which it rises substantially (typically 20–50% above normal levels) and then crashes (often to a third or less of normal levels within a few months). Such episodes occur over a wide range of “realistic” parameter values.

The episodes become rarer when parameter configurations give investors longer memories or give fund managers smaller (or more fleeting) streaks of luck. In opposite configurations the episodes become more common and, with extreme parameter values, normalcy becomes rare.

The occasional bubble and crash episodes seem related to escape dynamics, which identifies the “particular most likely way” in which the economy temporarily leaves the vicinity of a steady state (Williams, 2004, p. 7). However, Model 1 lacks the linear structure and explicit belief formation of the macrotheory literature, and its continuous action space puts it outside the evolutionary game theory literature cited in Section 2.2. It seems that new efforts, well beyond the scope of the present paper, are needed for analytic characterization of Model 1 dynamics.

Model 1 seems to vindicate the KMK perspective. It seems clear that bubbles and crashes would only be intensified by incorporating other realistic features emphasized by KMK, such as the heterogeneous expectations, rank-sensitive managers, and exogenous shifts in economy-wide growth opportunities.

Much work remains. One can extend Model 1 to incorporate such realistic features and, as noted in Appendix B, we have already included many of them in a Model 2. Preliminary results for Model 2 are encouraging but followup work is required. Also, other realistic features await exploration.

The most important task, however, is cross-validation. Simulation models are most valuable when they work in tandem with analytic results, empirical studies and/or experiments with human subjects. We hope that our work inspires new analytical, empirical and experimental work that deepens understanding of bubbles and crashes.

Acknowledgments

We are grateful to the National Science Foundation for support under Grant SES-0436509, and to Economics Department seminar audiences at UC Davis and Simon Fraser University. Buz Brock, Tom Copeland, Steve LeRoy and Dan Ostrov offered helpful advice and pointers to the literature. For invaluable and tireless research assistance, we thank Paul Viotti, Andy Sun, Todd Feldman, Matt Draper and Don Carlisle. The final version owes much to three anonymous referees of this journal. We retain sole responsibility for remaining idiosyncracies and errors.

Appendix A. Mathematical proofs

Proposition 1. *Given fixed positive parameters $c_2, \alpha, g_s, R_0,$ and d_R such that $R_s = R_0 + d_R > g_s,$ there is a unique $x^* > 0$ such that the degenerate distribution at x^* is a symmetric steady state solution to the master equation (8). Moreover, x^* decreases in c_2 and increases in $d_R.$ It increases in α and g_s and decreases in R_0 iff $x^* < 1.$*

Proof. We first show that Eq. (11) has a unique solution $x^*.$ Rewrite the equation as

$$U(x) = (R_s - g_s)x^{-\alpha} + g_s - R_0 - c_2x = 0. \tag{17}$$

Note that U is a continuous real valued function which is positive (indeed unbounded) as $x \searrow 0$ and is negative as $x \rightarrow \infty.$ Hence by the Intermediate Value Theorem, $U(x) = 0$ at some intermediate value $x^*.$ Since $U'(x) = -\alpha(R_s - g_s)x^{-\alpha-1} - c_2 < 0$ for all $x \in (0, \infty),$ it follows that $U(x) = 0$ has at most one solution, i.e., x^* is unique.

The next step is to show that (17) is necessary and sufficient for a symmetric (degenerate) solution to the master equation. That step proceeds exactly as in Friedman and Ostrov (2008, Proposition 2). It is omitted here because it requires several technical details tangential to the concerns of the present paper.

The master equation ignores the impact of retained earnings, which affect the growth rates of the weights $z = f(x)$ at different values of $x;$ see Eq. (20). Here we consider only distributions clumped at a single point so retained earnings have no impact and the master equation holds without modification.

To complete the proof, write $R_s = R_0 + d_R,$ differentiate (17) with respect to the given parameter and solve to obtain

$$\begin{aligned} \partial x^* / \partial R_0 &= (x^{-\alpha} - 1) / [c_2 + \alpha(R_s - g_s)x^{-\alpha-1}], \\ \partial x^* / \partial d_R &= x^{-\alpha} / [c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] > 0, \\ \partial x^* / \partial g_s &= (1 - x^{-\alpha}) / [c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] = -dx^* / dR_0, \\ \partial x^* / \partial c_2 &= -x / [c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] < 0, \end{aligned}$$

and

$$\partial x^* / \partial \alpha = -(R_s - g_s) \ln x / [c_2 x^\alpha + \alpha(R_s - g_s) / x].$$

Inspection shows that x^* is increasing in α and in g_s and decreasing in R_0 iff $x^* < 1.$ □

The “smoothed wedge” function ψ used in the next proposition is the definite integral of the cumulative unit normal distribution function $\Phi.$ Expressed in terms of the normal density $\Phi'(y) = (1/\sqrt{2\pi})e^{-y^2/2},$ it is

$$\psi(x) = \int_{-\infty}^x (x - y)\Phi'(y) dy = \int_{-\infty}^x \Phi(y) dy. \tag{18}$$

The last expression in (18) is obtained via integration by parts. The graph of ψ lies slightly above the graph of the simple wedge function $w(x) = \max\{0, x\}.$

Proposition 2. *In steady state with given $c_2,$ a manager with leverage x incurs expected loss $q(x|c_2) = (x\sigma/\sqrt{2\tau})\psi(z^0(x)),$ where $z^0(x) = (-\sqrt{2\tau}/\sigma)[R_0(1/x - 1) + g_s + (R_s - g_s)(x^*)^{-\alpha}]$ and x^* is defined from c_2 in Proposition 1.*

Proof. Recall that a loss is defined as the shortfall from 0 of gross returns, $R_{Gt} = (R_1 - R_0 + \varepsilon_t)x + R_0.$ The unconditional distribution of ε_t (obtained as the $h \rightarrow \infty$ limit in (13)) is normal with mean 0 and standard deviation $\sigma/\sqrt{2\tau}.$ Drop the non-steady state term in (5) to obtain $R_1 = (R_s - g_s)(x^*)^{-\alpha} + g_s.$ Plug this into the expression above for R_{Gt} to conclude that its unconditional distribution F is normal with mean $\mu \equiv [(R_s - g_s)(x^*)^{-\alpha} + g_s - R_0]x + R_0$ and standard deviation $s \equiv x\sigma/\sqrt{2\tau}.$ That is, gross returns can be expressed as $r = \mu + sz$ where z is a unit normal random variate. Since gross returns are negative for realizations of z that fall below $z^0 \equiv -\mu/s = (-\sqrt{2\tau}/\sigma)[R_0(1/x - 1) + g_s + (R_s - g_s)(x^*)^{-\alpha}],$ the expected loss is

$$\int_{-\infty}^0 [0 - r]F'(r) dr = s \int_{-\infty}^{z^0} [z^0 - z]\Phi'(z) dz = s\psi(z^0) = (x\sigma/\sqrt{2\tau})\psi(z^0). \quad \square \tag{19}$$

Corollary. *The expected loss is zero and has derivative zero at $x = 0.$ It is a convex increasing function for $x > 0.$*

Proof. Eq. (19) gives the expected loss as $Q(x) = ax\psi(z^0(x)),$ where $a = \sigma/\sqrt{2\tau} > 0$ and $z^0(x) = -b/x + k$ for $b = R_0/a > 0.$ It is immediate from (18) that $\psi'(y) = \Phi(y) \geq 0,$ $\psi''(y) = \Phi'(y) \geq 0,$ and $\psi(y) \rightarrow 0$ as $y \rightarrow -\infty.$ It now follows that $Q(0) = 0$ and $Q'(-\infty) = 0.$ Straightforward computations show that $Q'(x) = a[\psi(z^0) + (b/x)\psi'(z^0)] \geq 0$ and $Q''(x) = (ab^2/x^3)\psi''(z^0) \geq 0.$ Hence q is increasing and convex. By L'Hospital's rule, $Q'(0) = ab\Phi'(-\infty) = 0.$ □

Remark. The Corollary is true in considerable generality. The normality assumption plays essentially no role, nor do specification details for $z_0(x).$ The proof works with minor modifications as long as the tail probability $F(z)$ is $o(1/z)$ as $z \rightarrow -\infty,$ which holds for all distribution functions F commonly used in the literature.

Appendix B. Model extensions and robustness

B.1. Modified master equation

Let $Z(x, t)$ denote the total value at time t of all managed portfolios with leverage $\leq x$, and let $Z(t) = Z(\infty, t)$ denote the overall total value. The distribution described in the master equation is just its normalization $F(x, t) = Z(x, t)/Z(t)$. When the overall total $Z(t)$ is not constant over time, the master equation becomes

$$F_t(x, t) = -F_x(x, t)\phi_x(x, F) + \left[\int_0^x Z_{xt}(y, t) dy - F(x, t)Z_t(t) \right] / Z(t). \tag{20}$$

Here subscripts denote partial derivatives, and $Z(x, t)$ includes birth and death rates as well as the fickle investor effects in (22). To verify, differentiate the identity $F(x, t) = \int_0^y Z_x(y, t) dy / Z(t)$ in the case that $Z(x, t)$ has a density. For other cases (where there are mass points), take limits of cases with a density.

B.2. Logit regression

We drew 248 random parameter vectors from the uniform distribution on the truncated rectangle $1 \leq \alpha \leq 4, 1 \leq \beta \leq 5, 0.1 \leq \eta, \tau \leq 3, 0.05 \leq \sigma \leq 0.4, -0.04 \leq g_s \leq 0.04, 0.01 \leq R_0, d_R \leq 0.05, R_s = R_0 + d_R > g_s$, with population kept constant at $M = 30$. For each parameter vector we simulated a century of weekly data. We dropped centuries with more than 20 crashes, since the parameter vector in that case seems to lie outside the relevant region. A supplementary regression (available on request) suggests that the 20-crash rule is roughly equivalent to imposing the parameter constraint $-15.5d_R + 6g_s + 9\sigma - 1.5\tau + 0.6\eta + 0.2\beta + 1.4\alpha \leq 4.6$.

The dependent variable in Table 3 is crash-lagged, an indicator (binary valued). When a crash is detected at week t , i.e., the detrended asset price $DPM(t)$ is less than half its maximum over weeks $\{t - 1, \dots, t - 26\}$, then $\text{Crash-lagged}(t - 12) = 1$. That is, the dependent variable approximates the start of the crash by saying it occurred 12 weeks before the crash is confirmed. The explanatory variables in the logit regression consist of the parameter vector, quadratic terms in the vector components, and RDPM-M, a real-valued variable. In week t , RDPM-M is the ratio of the current detrended asset price, DPM, to its mean value over years 30–100. (The first 29 years are affected by the initial conditions.)

To interpret the RDMP coefficient, suppose that the initial probability of a crash is 1%, so initial log odds are $\ln(0.01/0.99) \approx -4.6$. The coefficient estimate of 5.11 implies that, other things equal, were asset price P to rise 20%, then the log odds would increase by about $0.2 * 5.11 \approx 1.0$ to -3.6 , implying a crash probability of $p = e^{-3.6}/(1 + e^{-3.6}) \approx 2.7\%$, i.e., the probability would increase by about 170%.

We checked robustness by running regressions with asset price normalized by fundamental value V instead of the mean value, with more or less stringent definitions of crashes, and with the dependent variable lagged more or fewer periods than 12. The results are all roughly similar; in some specifications the R_0 and β coefficients gain significance, often at the expense of the price coefficient or the α or τ coefficients. The coefficients on asset price tend to be a bit smaller when using V , but they typically remain quite significant economically and statistically. We initially considered all quadratic as well as all linear parameter terms, but Table 3 retains only the most significant quadratic terms.

The coefficient estimates generally confirm earlier conclusions: crashes are more likely with lower d_R, g_s and τ , and with higher σ, η and α . The quadratic coefficients indicate that the initially strong impact of σ and the negative impact of τ both taper off when σ is towards the upper end of its range. They also indicate that β 's impact becomes positive when η is larger,

Table 3
Logistic regression for crash in next 12 steps.

Variable	Coefficient	(Std. err.)
Intercept	-17.356**	(0.6295)
R0	-0.014	(4.3763)
dR	-15.893**	(4.1536)
gs	-9.887**	(4.6520)
sigma	17.136**	(2.6679)
tau	-2.952**	(0.2423)
eta	0.782**	(0.1597)
beta	-0.084	(0.0804)
alpha	2.003**	(0.1175)
sigma_sq	-21.840**	(6.3198)
sigma_tau	3.992**	(0.8009)
eta_beta	0.122**	(0.0500)
tau_gs	6.161**	(2.4888)
RDPM_M	5.113**	(0.2406)

Note: Dep. Var.: = 1 if crashed within next 12 steps, = 0 otherwise. RDPM calculation based on average of DPM. *significant at 5%. **significant at 1%.

and the negative impacts of g_s and τ tend to offset each other. The asset price coefficient indicates that crashes are considerably more likely when asset prices are high. As explained in appendix, the estimate of about 5.1 implies that a crash probability of 1% would more than double if, other things equal, the asset price increased by 20%.

B.3. Extension: fickle investors

Managed funds routinely reinvest positive returns and seldom ask investors to cover negative returns. Hence, other things equal, the fund grows at the fund's gross rate of return, viz., $\dot{z}_i^R/z_i = R_{Gi} = (R_1 - R_0 + \varepsilon_i)x_i + R_0$.

More importantly, as noted in the empirical facts section, investors chase returns. Managers with large perceived losses $\hat{L}_i(t)$ and small perceived returns should lose investors, and those with small losses and large returns should gain. This can be formalized in many ways. For simplicity and consistency with available evidence, we say that the defection rate is proportional to the perceived loss. Thus the outflow rate is $\dot{z}_i^O/z_i = -\delta\hat{L}_i$, where the defection parameter $\delta \geq 0$ reflects how strongly investors respond by withdrawing part or all of their funds. The outflow of funds initially goes to a pool z_0 not allocated to any fund manager.

Recruitment of new investors depends on relative perceived returns. Recycle the exponential average technique from (14) and apply it to net returns in (12) to get perceived net returns,

$$\hat{R}_i(t) = \eta \int_{-\infty}^t e^{-\eta(t-s)} R_i(s) ds. \tag{21}$$

The empirical papers cited at the end of Section 2 suggest that a disproportionate share of new investment goes to managers with the very best perceived returns. To capture this effect, we use a logit specification, with fund inflows proportional to $e^{\lambda\hat{R}_i}$ rather than to \hat{R}_i itself. When the logit parameter $\lambda = 0$, inflows are unrelated to perceived returns, and larger values of the parameter indicate greater sensitivity to relative performance.

The inflow rate is also proportional to z_0 , the pool of funds available for investment. Thus we obtain the following expression for fund inflows: $\dot{z}_i^I/z_i = \rho z_0 e^{\lambda\hat{R}_i}$. In the current version of the simulation, the parameter $\rho > 0$ is inversely proportional to the sum (or integral) over managers of $e^{\lambda\hat{R}_i}$, so the overall outflow rate is constant from the unallocated pool z_0 . To capture the actual tendency of investors to allocate more funds in bull markets, the parameter ρ could be made less responsive to the sum of the $e^{\lambda\hat{R}_i}$.

Putting the three terms together, the overall rate of increase in fund i in Model 2 is

$$\dot{z}_i = [\dot{z}_i^R + \dot{z}_i^O + \dot{z}_i^I]z_i = [R_0 + (R_1 - R_0 + \varepsilon_i)x_i - \delta\hat{L}_i + \rho z_0 e^{\lambda\hat{R}_i}]z_i. \tag{22}$$

B.4. Scale effects, exit and entry, and more

As noted at the end this Appendix, evidence suggests that large size tends to depress a manager's returns. We can accommodate that feature by including a quadratic penalty for risk exposure xz_i . Thus net returns in (12) become

$$R_i(x) = x(R_1 - R_0 + \varepsilon_i) - \frac{1}{2}c_2x^2 - \kappa(xz_i)^2, \tag{23}$$

where the default value of parameter κ is 0.01. That is, a standard size fund fully invested in risky assets incurs costs of 1% per annum.

Funds do not last forever, especially hedge funds. They tend to liquidate when investors leave. One could make the hazard rate a smoothly decreasing function of size or growth rate, but to keep things simple we simply decree a minimum size z_{LIM} , and say that fund i disappears (with its funds going to the unallocated pool z_0) whenever $z_i < z_{LIM}$. Otherwise the firm survives.

New funds appear occasionally, especially when the pool of unallocated funds is large. To keep things simple, we count the number of funds liquidated over the last year, and at the beginning of the new year we create the same number of new funds. The size z and initial risk stance x of each new fund is independently drawn from the uniform distribution on $[x_{SS} - W/2, x_{SS} + W/2] \times [z_L, z_U]$. Default choices are $z_L = 2z_{LIM}$, $z_U = 2z_L$, and width $W = x_{SS}/2$, where x_{SS} is the steady state leverage in Proposition 1 given the current price of risk c_2 .

Model 3 incorporates effects for size ("gravity") and entry and exit. Since the total mass of investor funds is no longer constant, the master equation must be modified slightly. This is spelled out in Appendix B. Models 2 and 3 exhibit bubbles and crashes similar to those of model 1.

Further extensions of the model have been contemplated but not implemented so far. One could include a Markov process that occasionally shifts the underlying state of the economy s . For example, there could be three states (say poor, good and great) and phase 1 of a KMK bubble could be touched off by a transition to the great state, with maximal g_s . Another extension would include other realistic components of the fund manager's objective function (2) that reflect absolute size or relative vs. absolute returns and losses. It might then be possible to explore the impact of contrarian strategies and short-selling.

Fund managers care about relative as well as absolute performance. The Wall Street Journal publishes rankings four times a year, and agencies such as Lipper Analytics and Morningstar do so more frequently. Higher rank brings managers larger bonuses and more competing job offers, and also increases their compensation by attracting more investment inflows.

On the other hand, large size tends to depress a manager's returns. Chen et al. (2004) find that a 2 standard deviation increase in fund size implies that annual returns decline by almost a percentage point. The reasons include illiquidity (it is more costly to redeploy a large portfolio than a small one) and some sorts of organizational costs.

Appendix C. The KMK perspective

Minsky (1975) and Minsky (1982), drawing on themes of Keynes (1936), developed a distinctive view of bubbles and crashes, later elaborated in Kindleberger (1978/1989/2000). Although never fully formalized, this KMK perspective helps identify features of financial markets that can make them vulnerable to bubbles and crashes.

The KMK perspective can be summarized informally as a sequence of phases. Phase 0 is normalcy. Financial market participants share a broad consensus on the earnings prospects for tradeable assets. Asset prices closely track fundamental values, and investors earn normal returns, commensurate with perceived risk.

Phase 1 begins when an unusual opportunity arises, financial or real. Two famous early examples: some investors saw tremendous profit opportunities for selling strikingly colored varieties of tulips to rising middle class families in early 17th century Netherlands. The South Sea Company seemed poised in the early 1720 to purchase the British national debt, opening unprecedented financial opportunities. More recently, in the late 1980s innovative Japanese car and consumer electronic manufacturers gained world leadership in efficiency and quality; and in the late 1990s the rapid rise of the Internet created a variety of new business opportunities.

Normally, shared experience leads to rough consensus on the value of available opportunities. However, opportunities sufficiently different from earlier events—the unusual opportunities—can easily lead to a divergence of opinion. Optimists may think the unusual opportunity will lead to once-in-a-lifetime profits for those who seize it, while pessimists may believe that it will produce normal profits at best. Well-known Internet optimists included Mary Meeker and Henry Blodgett, who predicted that dozens of startup companies would each be worth hundreds of billions of dollars. Pessimists (including most economists) argued that, although the Internet might attract a substantial share of commerce, it would tend to lower profit margins and that few of the startup companies would ever generate much shareholder value.

Phase 2 begins if and when the optimists reap impressive profits. For example, the market value of Netscape shares increased sixfold in 5 months from the initial offering in August 1995. Such returns attract trend-following investors, who in turn attract financial innovators. Venture capital firms mushroomed in the late 1990s, inundated by new investors, and day trading became popular.

Optimists get the new investment inflows. The flip side, often overlooked, is that pessimists either play along or else get left out. One of the authors observed top managers at major US bank during the energy boom of the late 1970s as they decided whether to expand energy lending, despite warnings that the sector was overextended. The clinching argument was that the bank had to make the loans to remain a major player. Perhaps the classic example is Sir Isaac Newton his role as Master of the Mint. The immortal physicist sold the Mint's South Sea shares at a decent profit in April 1720 but then came under increasing to match other investors' returns. In midsummer, he bought a large block of shares just as the bubble reached its maximum. The point of Keynes' famous beauty contest metaphor and comment on "levels of play" is that sophisticated pessimists should sometimes mimic optimists.

Crucially, asset quality deteriorates as the bubble inflates. Recent experience encourages some investors to pay high prices for promises that can only be fulfilled in good times, and financial market innovators offer a ready supply of such promises. "Sub-prime" home loans are a recent example: the borrower has little equity, and will be able to make promised payments only if home prices continue to rise briskly and refinancing remains easy to obtain. The financial innovations and lending standards induced by a bubble tend to make the financial sector increasingly vulnerable to unfavorable developments.

Phase 3 begins when the supply of dazzled new investors and financial innovation is exhausted, as must happen eventually in our finite world. A minor event then can touch off a cascade, as implicit (or explicit) defaults trigger further defaults and losses. It is hard to remember what event in March 2000 ended the runup of the NASDAQ index to over 5000, or what stopped Japan's Nikkei index just short of 40,000 in January 1990. But once asset prices started to decline, many leveraged investors had to sell, and the decline accelerated. Such declines corrode collateral, and borrower defaults can cause lender defaults, so a financial crash can be contagious. Phase 3 generally runs faster than phase 2.

A national or international recession may result. In modern jargon, the KMK story is that Phase 3 financial distress increases uncertainty, which increases the value of deferral options. Real investment therefore declines as financial distress spreads, and multiplier effects produce a recession. Countercyclical monetary and fiscal policies are intended to prevent such recessions, or reduce their severity, by shielding basically sound organizations from contagion and reducing uncertainty.

Phase 4 begins when asset prices are so low that savvy investors purchase again and the "bear market bottoms out." The NASDAQ was a good buy at 1200 in Summer 2002. With effective bankruptcy laws, the losses accrued in phase 3 are quickly

parcelled out, productive assets are redeployed, and recovery begins promptly, e.g., as in the US following the Savings and Loan debacle of the 1980s. Consensus beliefs return, and financial assets are again grounded in reality. Phase 0, normalcy, begins anew and often lasts for decades.

By contrast, a protracted political struggle ensues when it is unclear who must bear the losses accrued in a crash,⁹ as in Japan recently, or in Latin America in the 1970s and 1980s. Phase 4 then can be quite long and painful. (The Great Depression of the 1930s arguably involved inept countercyclical policy as well as inadequate bankruptcy laws.)

In our interpretation, the KMK perspective rests on a learning process distorted by financial market imperfections. Abraham Lincoln was right that you cannot fool all the people all the time, but you can fool lots of them occasionally. Once a bubble starts inflating, financial markets give investors little economic or psychological incentive to slow it down. The eventual crash completes the learning process, and inoculates investors. Hence bubbles tend not to repeat themselves: it takes a rather different novel opportunity, probably much later or in some distant location, to touch off the next episode.

References

- Aoki, M., 2004. *New Approaches to Macroeconomic Modeling: Evolutionary Stochastic Dynamics, Multiple Equilibria, and Externalities as Field Effects*. Cambridge University Press, New York.
- Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R., Tayler, P., 1997. Asset pricing under endogenous expectations in an artificial stock market. In: Arthur, W.B. (Ed.), *The Economy as an Evolving Complex System*. Addison-Wesley, Reading, MA.
- Blanchard, O.J., Watson, M.W., 1982. Bubbles, rational expectations and financial markets. NBER Working Paper, No. 945.
- Bollen, N.P.B., Busse, J.A., 2005. Short-term persistence in mutual fund performance. *Review of Financial Studies* 18 (2), 569–597.
- Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22 (8–9), 1235–1274.
- Brock, W.A., Hommes, C.H., Wagener, F., 2005. Evolutionary dynamics in markets with many trader types. *Journal of Mathematical Economics* 41 (1–2), 7–42.
- Camerer, C.F., Loewenstein, G., Rabin, R. (Eds.), 2003. *Advances in Behavioral Economics*, Princeton University Press, Princeton, NJ.
- Chen, J., Hong, H., Huang, M., Kubik, J.D., 2004. Does fund size erode mutual fund performance? The role of liquidity and organization. *The American Economic Review* 94 (5), 1276–1302.
- Cheung, Y.-W., Friedman, D., 1997. Individual learning in normal form games: some laboratory results. *Games and Economic Behavior* 19 (1), 46–76.
- Chevalier, J., Ellison, G., 1997. Risk taking by mutual funds as a response to incentives. *Journal of Political Economy* 105 (6), 1167–1200.
- Cho, I.K., Williams, N., Sargent, T.J., 2002. Escaping nash inflation. *Review of Economic Studies* 69, 1–40.
- Copeland, T.E., Weston, J.F., Shastri, K., 2005. *Financial Theory and Corporate Policy*, fourth ed. Pearson/Addison Wesley, New York.
- De Fontnouvelle, P., 2000. Information dynamics in financial markets. *Macroeconomic Dynamics* 4, 139–169.
- DeLong, J.B., Schleifer, A., Summers, L.H., Waldmann, R.J., 1990. Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance* 45, 379–395.
- Del Guercio, D., Tkac, P.A., 2002. The determinants of the flow of funds of managed portfolios: mutual funds versus pension funds. *Journal of Financial and Quantitative Analysis* 37, 523–557.
- Egenter, E., Lux, T., Stauffer, D., 1999. Finite-size effects in Monte Carlo simulations of two stock market models. *Physica A* 268 (1), 250–256.
- Feller, W., 1971. *An Introduction to Probability Theory and its Applications*. Wiley Series in Probability and Mathematical Statistics, vol. 2, third ed. Wiley, New York.
- Freidlin, M.I., Wentzell, A.D., 1984. *Random Perturbations of Dynamic System*. Springer, New York.
- Friedman, D., Aoki, M., 1992. Inefficient information aggregations as a source of asset price bubbles. *Bulletin of Economics Research* 44, 251–279.
- Friedman, D., Ostrov, D., 2008. Conspicuous consumption dynamics. *Games and Economic Behavior* 64 (1), 121–145.
- Friedman, D., Yellin, J., 1997. Evolving landscapes for population games, University of California, Santa Cruz draft manuscript, available at: (<http://www.vismath.org/research/landscapedyn/models/articles/>).
- Fudenberg, D., Levine, D.K., 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Garber, P., 1989. Tulipmania. *Journal of Political Economy* 97 (3), 535–560.
- Holland, J.H., 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor.
- Hommes, C.H., 2006. Heterogeneous agent models in economics and finance. In: Tesfatsion, L., Judd, K. (Eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. Elsevier, New York.
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 61 (1), 29–56.
- Karceski, J., 2002. Returns-chasing behavior, mutual funds, and beta's death. *Journal of Financial and Quantitative Analysis* 37, 559–594.
- Keynes, J.M., 1936. *The General Theory of Employment, Interest and Money*. Macmillan, London.
- Kindleberger, C.P., 1978/1989/2000. *Manias, Panics and Crashes: A History of Financial Crises*. Basic Books, New York.
- LeBaron, B., 2006. Agent-based computational finance. In: Tesfatsion, L., Judd, K. (Eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. Elsevier, New York.
- LeRoy, S.F., 2004. Rational exuberance. *Journal of Economic Literature* 42 (3), 783–804.
- Levy, H., Markowitz, H.M., 1979. Approximating expected utility by a function of mean and variance. *American Economic Review* 69 (3), 308–317.
- Mackay, C., 1841. *Extraordinary popular delusions and the madness of crowds*, (reprinted in Wiley's 1996 investment classics volume).
- Minsky, H.P., 1975. *John Maynard Keynes*. Columbia University Press, New York.
- Minsky, H.P., 1982. *Can 'IT' Happen Again? Essays on Instability and Finance*. M.E. Sharpe Inc., NY.
- Penso de la Vega, J., 1688. *Confusion e confusions* (reprinted in Wiley's 1996 investment classics volume).
- Sargent, T., 1999. *The Conquest of American Inflation*. Princeton University Press, New York.
- Shefrin, H., 2002. *Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing*. Oxford University Press, Oxford.
- Sirri, E.R., Tufano, P., 1998. Costly search and mutual fund flows. *Journal of Finance* 53, 575–603.
- Sonnenschein, H., 1982. Price dynamics based on the adjustment of firms. *American Economic Review* 72 (5), 1088–1096.
- Tirole, J., 1982. On the possibility of speculation under rational expectations. *Econometrica* 50 (5), 1163–1181.
- Williams, N., 2004. *Escape dynamics in learning models*. Manuscript, Princeton University, Economics Department.
- Young, P., 1993. The evolution of conventions. *Econometrica* 61 (1), 57–84.
- Youssefmir, M., Huberman, B.A., Hogg, T., 1998. Bubbles and market crashes. *Computational Economics* 12 (2), 97–114.
- Zeeman, E.C., 1974. On the unstable behavior of the stock exchanges. *Journal of Mathematical Economics* 1 (1), 39–44.

⁹ We are indebted to Axel Leijonhufvud for this observation.