Bubbles and Crashes: Escape Dynamics in Financial Markets *

Daniel Friedman[†] Ralph Abraham[‡]

January 10, 2006

Abstract

We develop a financial market model focused on fund managers who continuously adjust their exposure to risk in response to the payoff gradient. 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 adopt riskier portfolios. Then slight losses can trigger a crash, as a widening risk premium accelerates the decline in asset price.

Keywords: Bubbles, escape dynamics, time varying risk premium, constant gain learning, agent based models.

JEL codes: C63, C73, D53

*We are grateful to the National Science Foundation for support under grant SES-0436509. For research assistance, we thank Paul Viotti, Andy Sun, Matt Draper and Don Carlisle.

[†]Economics Department, University of California, Santa Cruz, CA, 95064. *dan@ucsc.edu*

[‡]Mathematics Department, University of California, Santa Cruz, CA 95064. *abraham@vismath.org*

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, 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 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. Where possible, it adapts and streamlines standard ingredients. For example, it assumes a single source of systematic risk, ignores inflation and taxes, and blurs the distinctions between cash flow, earnings and dividends. Unlike most previous financial market models, the focus is on professional fund managers who follow the profit gradient. Another non-standard feature is constant gain learning (e.g., Cho, Williams and Sargent, 2002), also known as exponential average expectations. That feature leads to an endogenous risk premium, and to investors who chase returns and flee from losses.

Section 2 touches on several strands of literature. Our model draws inspiration from one of the older strands, due to Keynes, Minksy and Kindleberger (KMK), as well as from the emerging agent-based approach. The section concludes by listing a set of empirical facts that influenced modeling choices.

Section 3 presents the basic model, beginning with the static ingredients and gradient dynamics. It characterizes analytically a unique equilibrium and its comparative statics. These are illustrated in an agent-based model that indeed appears quite stable dynamically over a very wide range of parameter configurations.

Section 4 lays out an extension that incorporates some KMK-inspired features such as mean-reverting, manager-specific skill (or luck), and investors who are constant-gain learners. Analytic approximations provide predictions of typical behavior and conjectures about failure modes, or "escape dynamics." Simulations of the extended agent-based model confirm sporadic bubbles and crashes for a wide range of parameter configurations. Section 5 presents further extensions of the model.

After a concluding discussion, an appendices collects technical details. Additional material, including source code 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 its expected present value, given all available information, of the net cash flows 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 doesn't 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.

At the other extreme, a "rational bubble" literature flowered in the 1980s (e.g., Blanchard and Watson, 1983, and 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., De Long et al, 1990; Huberman et al., 1998; Brock and Hommes, 1998).

LeRoy (2004) concludes his extensive 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.

2.1 The Keynes-Minsky-Kindleberger Perspective

Hyman Minsky (1975, 1982), drawing on themes of John Maynard Keynes (1936), developed a distinctive view of bubbles and crashes, later elaborated in Charles Kindleberger (1978/1989/2000). Although never fully formalized, this Keynes-Minsky-Kindleberger 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 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 five 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 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's 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,[1] 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 hinges on a learning process distorted by financial market imperfections. Abraham Lincoln was right that you can't 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 innoculates 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.

2.2 Escape Dynamics

Drawing on mathematical results such as 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

[1] I am indebted to Axel Leijonhufvud for this observation.

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."

2.3 Agent-based Simulation Models

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 modeled as interacting groups of learning, boundedly rational agents, and behavior is described mainly by computer simulations rather than by explicit solutions or 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 $B > 0$ 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 doesn't 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. It models evolutionary competition among two to four simple linear forecasting rules, and exhibits 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.

The model we present in this paper can be regarded as containing either two agent types (a priori identical fund managers interact with a second type of agent not explicitly modeled, as noted below in footnote 4) or a continuum of manager types ordered by ex-post choice of risk stance. See Friedman and Yellin (1997) for a general discussion regarding evolution of an ordered continuum of strategies.

2.4 Some Suggestive Empirical Facts

The US financial sector currently comprises about \$43.5 trillion of financial assets, of which households (and nonprofit 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 dominate two of the most rapidly growing segments, private equity and hedge funds.

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 returns decline by almost a percentage point. The reasons include liquidity (it is more costly to turn over a large portfolio than a small one) and some sorts of organizational costs.

Investors in mutual funds chase returns, especially those funds that recently were top performers (Sirri and Trufano, 1998; Karceski, 2002). Pension funds are less extreme in chasing top performance, but are harsher on funds that incur losses (Del Guerco and Tkac, 2002). Sirri and Trufano note that this cross-sectional effect is 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.

Underlying such behavior are variations in the risk premium. Copeland et al (2005, figure 6.10) shows 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.

3 A base model

This section first lays out 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_o 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, 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. 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 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 $R_G(x) = (1 - x)R_o + xR_1$. The manager's cost of funds is the risk free rate R_o plus a risk cost $c(x)$.[2] The risk cost up to second order is $c(x) = \frac{1}{2}c_2x^2 + c_1x + c_0$. We will see later that, under the maintained assumption of negligible trading costs, $c(0) = c_0 = 0$ and $c'(0) = c_1 = 0$, while $c_2 = c''(0) \geq 0$ can be interpreted as the market price of risk. In the basic model c_2 is an exogenous constant, but it can vary in extensions of the model. Either way, the risk-adjusted

[2]The standard interpretation is that the risk cost reflects the concavity of investors' utility functions. An alternative interpretation is that an insurance agency (e.g., the FDIC) charges a premium which increases in the expected loss claims, or that investors self-insure. Either interpretation seems consistent with the specifications to follow.

net return enjoyed by a manager chosing leverage x is

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

The manager's objective or payoff function ϕ depends positively on the net return. It may also depend on the portfolio size and on relative performance, but for now we ignore such complications and write

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

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.[3] 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, $V(0) = \int_0^\infty 1e^{g_st}e^{-R_st}dt = (R_s - g_s)^{-1}$. At time $t > 0$ the fundamental value is similar except that the earnings stream starts at e^{g_st}, so

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

Asset price moves with fundamental value but also responds to buying pressure. That is, other things equal, the asset price is driven higher when in aggregate the fund managers demand more of it.[4] Assume for tractability that the elasticity is a constant, $\alpha \geq 0$. With appropriate normalization, the price of the risky asset then is

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

[3] Earnings means 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.

[4] Some other traders must accommodate the fund managers; implicitly we assume that they are "fundamentalists" who are betting that price will return to V. These other traders could be issuers of stock, individual investors, foreigners, etc.

Thus asset price is equal, less than or greater than fundamental value whenever mean leverage $\bar{x}$ is equal, less than or greater than 1.0.

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 in this model 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? It starts with the riskless rate $R_o \geq 0$, which reflects investors' marginal rate of time preference. It includes g_s since economic growth is known and economy-wide. Including a third term e_s to account for other possible factors yields the specification

$$R_s = R_o + g_s + e_s. \tag{6}$$

To ensure that fundamental value is well defined in (3), we assume that $e_s > -R_o$. It will be convenient in comparative statics exercises to write $d_R = g_s + e_s$ and to analyze the separate impacts of R_o, d_R and g_s.

3.3 Gradient dynamics

Each manager adjusts 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 (Proposition 1 of Friedman and Yellin, 1997), rather than approximate gradient or sign-preserving dynamics. Hence gradient dynamics is a natural way to describe fund managers' behavior and financial market impact.[5]

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 proportional to the slope. Neglecting for now factors affecting the size of the manager's portfolio,[6] we obtain the master equation

$$F_t(x,t) = -F_x(x,t)\phi_x(x,F), \tag{7}$$

where here the subscripts denote partial derivatives.

The master equation can be interpreted as conservation of mass since the size of each fund is conserved as the manager reallocates between safe and risky assets. The equation explains changes over time in the fraction $F(x)$ of managed funds that have leverage x or less. That fraction increases at rate $F_t(x,t)$ as some managers decrease leverage from above x to below x. The right hand side of the equation is the net flux, the density $F_x(x,t)$ of funds with leverage x times the (leftward, hence the minus sign) velocity given by the payoff gradient ϕ_x at that point.

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

$$\phi(x,F) = x[(R_s - g_s)\bar{x}^{-\alpha} + g_s - R_o + \alpha\dot{\bar{x}}/\bar{x}] - \frac{1}{2}c_2x^2, \tag{8}$$

with gradient

$$\phi_x = (R_s - g_s)\bar{x}^{-\alpha} + g_s - R_o + \alpha\dot{\bar{x}}/\bar{x} - c_2x. \tag{9}$$

3.4 Equilibrium

We focus on steady states of (7), especially those in which all fund managers choose the same leverage $x = \bar{x}$.[7] In steady state we must have $\phi_x = 0$ and, of course, $\dot{\bar{x}} = 0$ at that

[5] The standard literature suggests no alternative to gradient dynamics; it focuses on equilibrium and neglects the adjustment process.

[6] That is, assume for now that the fund does not retain the gross (or net) return but remits it to its clients, who never withdraw or invest additional funds. Soon we will relax these unrealistic assumptions.

[7] That is, we seek solutions to the master equation such that for all $t \geq 0$, $F(x,t) = 0$ for $x < \bar{x}$ and $F(x,t) = 1$ for $x \geq \bar{x}$. We refer to such a distribution as clumped at $\bar{x}$.

point. Inspection of (9) shows that one trivial possibility is that $c_2 = 0$ and $x = \bar{x} = 1$ for all fund managers so (5) collapses to $R_1 = R_s = R_o$ 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 steady state with $x = \bar{x}$ when the marginal risk cost $c_2 x$ equals the marginal steady state net return $R_1 - R_o = (R_s - g_s)\bar{x}^{-\alpha} + g_s - R_o$. 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_o and R_s.

Proposition 1. Given fixed parameters $c_2 > 0, \alpha > 0$, $R_o \geq 0$, and $R_s > g_s$, there is a unique point $x^* > 0$ such that the distribution clumped at x^* is a steady state solution to the master equation (7). Moreover, x^* decreases in c_2 and R_o and increases in d_R. It increases in α and in g_s iff $x^* < 1$.

————figure 1 about here————

See Appendix A for a formal proof and for the comparative static expressions; 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 (9), and rearrange slightly to obtain

$$(R_s - g_s)x^{-\alpha} + g_s - R_o = c_2 x. \tag{10}$$

As shown in Figure 1, the right hand side of (10) 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 $y = g_s - R_o$ as the horizontal asymptote. Clearly there is a unique point of intersection $x^* > 0$.

Note that if c_2 is sufficiently low (and if $d_R = R_s - R_o \geq 0$) then $x^* = \bar{x} > 1$ and $P > V$. The high steady state asset price reflects fund managers' desire to leverage their portfolios given the low risk cost. Likewise, higher values of the c_2 parameter imply $x^* = \bar{x} < 1$ and $P < V$. There is some intermediate value of c_2 (it can be thought of as a long-run steady state) such that mean leverage is $\bar{x} = 1$ and $P = V$. One can see from (9) or (10) that this implies $R_s = R_1 = R_o + c_2$. In this equilibrium, c_2 looks like the standard risk premium, e.g., the market price of risk in the Capital Asset Pricing Model.

Implicitly differentiating (10) 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 (especially when α and $R_s - g_s$ are small). We will see in the next section that this leads to behavior reminiscent of KMK phases 2 and 3 when c_2 is endogenous.

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 $i = 1, ..., M$ whose risk stance x_i (horizontal coordinate) and portfolio size z_i (vertical coordinate)[8] are floating point numbers that adjust in discrete time. 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 specified above and adjusted to $r = (1 + R)^{1/N} - 1$ per time step.

———-figure 2 about here———-

Figure 2 shows a sample simulation of the basic model, using parameter values $c_2 = 0.02, R_o = 0.03, d_R = 0.03, g_s = 0$ 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 labelled "population", etc.[9] The managers move smoothly towards $x^* \approx 1.08$ and maximal z, as asset

[8] The simulations assume that managers retain all earnings (and absorb any losses). Appendix A notes the changes in the master equation required to accommodate this and other extensions described in later sections.

[9] 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 rectangluar distributions. 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.

price (displayed on the screen as DPM or detrended PM) converges to $1.08^2/(0.06-0) \approx 19.4$. Figure 2 shows that they are fairly close within two decades.

Proposition 1 suggests smooth convergence towards a clump at x^* and associated asset price P^* 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). Indeed, in every case the observed standard deviation becomes quite small and the observed mean is quite close to the steady state prediction.

–Table 1 about here——

4 Endogenous Risk

Dynamics become more interesting when we endogenize the risk cost c_2. To do so, we introduce three new features that play a prominent role in the KMK approach.

4.1 Streaks, losses, and learning

The first is that each manager can have streaks in which she underperforms or outperforms the market. The idea is that not every portfolio with given leverage has the same realized return; some managers are luckier (or more attuned to new opportunities) than others. Thus, instead of receiving the uniform return in (1), manager i receives net return

$$R_i(x) = x(R_1 - R_o + \epsilon_i) - \frac{1}{2}c_2 x^2. \tag{11}$$

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

$$\epsilon_i(t) = e^{-\tau h}\epsilon_i(t-h) + \sqrt{\frac{1-e^{-2\tau h}}{2\tau}}\sigma\nu, \tag{12}$$

for some given volatility parameter $\sigma > 0$ and decay parameter $\tau > 0$, and an independent realization ν 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, defined as negative gross return. Now $R_{Gi} = (R_1 - R_o + \epsilon_i)x_i + R_o$ is the gross return that manager i currently earns on his portfolio, so 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. The natural formalization is an exponential average. In continuous time the exponential average loss for manager i is

$$\hat{L}_i(t) = \gamma \int_{-\infty}^{t} e^{-\gamma(t-s)} L_i(s) ds, \tag{13}$$

where the parameter γ is the memory decay rate. Using the definition (13) 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:

$$\hat{L}_i(t) = e^{-\gamma h} \hat{L}_i(t-h) + (1 - e^{-\gamma h}) L_i(t). \tag{14}$$

Update rule (14) defines ”constant gain learning” with gain $(1 - e^{-\gamma h})$.

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

$$c_2 = \beta \hat{L}_T(t), \tag{15}$$

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 $\gamma = 0.7$.

4.2 Equilibrium

Steady states in the current model are more intricate than in the base model. The new parameters for memory decay (γ) 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 equilibrium, 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 $q(x|c_2) = (x\sigma/\sqrt{2\tau})\psi(z^o(x))$, where $z^o(x) = (-\sqrt{2\tau}/\sigma)[R_o(1/x-1) + g_s + (R_s - g_s)(x^*)^{-\alpha}]$ 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 equation (1) that the risk cost $c(x)$ is quadratic with $c(0) = c'(0) = 0$. The proof of the corollary in Appendix A works even when an arbitrary distribution function replaces Φ in the formula for ψ. Thus 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 $\alpha, \beta, \gamma, \sigma, \tau, g_s, R_o$ and R_s, and choose a reasonable initial estimate $\hat{c}_2$ of steady state c_2. Suppose that all managers clump at $x = x^*(\hat{c}_2)$, where $x^*(.)$ is defined implicitly in equation (10). Using the function $q(.|.)$ 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 equation (15), 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 $(R_s - g_s)x^{-\alpha} + g_s - R_o = (.06)(0.866)^{-2} - 0.03 = 0.050 = 0.058 * 0.866 = c_2 x$ so equation (10) holds, and that $\beta q(x^{**}|c_2^{**}) = \beta x^{**}(\sigma/\sqrt{2\tau})\psi(z^o(x^{**})) = 2 * 0.866 * 0.2 * 1.4^{-0.5}\psi(z^o(0.866)) = 0.293\psi(-0.501) = 0.293 * 0.198 = 0.058 = c_2^{**}$. We used $z^o(0.866) = -((2 * 0.7)^{0.5}/0.2) * (0.03 * (1/0.866 - 1) + 0 + (0.06) * (0.866)^{-2}) = -0.501$ 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^{**} can't 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.

—figure 3 about here———-

Figure 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. Using baseline parameter values in the first example tried while writing this passage, asset price P rose to $14.7 > P^{**} = 12.5$ by the end of year 12. 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.

The apparent mechanism is that a run of good luck for a few traders, and the absence of very bad luck for most others, puts P on 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 Figure 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, as in late phase 2 of KMK. At that point, the market is vulnerable to a string of moderately bad luck striking 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. As in late phase 3 of KMK, the vicious cycle slows when x^* gets so low that it becomes relatively unresponsive to further increases in c_2 (recall Figure 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 analysis

To investigate these impressions, and to check their robustness, we examined 18 variants of the baseline parameter configuration. As in Table 1, we ran 10 centuries of weekly simulations for each variant. Somewhat arbitrarily, we defined a crash as a decline in the asset price (detrended to remove the capital gains from g_s) by at least 50% from its highest point within the last half year.[10]

–Tables 2 and 3 about here––

Table 2 reports the results. Under baseline parameters, average asset price and risk position are modestly higher than predicted by the clumped steady state, and (severe) 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 substantial effect on crash frequency. Raising population size to 100 eliminated crashes, while lowering it to 15 tripled their frequency.

Recall that the (detrended) fundamental value $V = (R_s - g_s)^{-1}$, so shifts in the components R_o and d_R of the discount factor R_s shift V in the opposite direction. The table shows the impacts on predicted asset price P^{**} are partially offset via x^{**}. Actual average P barely responds to R_o 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

[10] The maximum drawdown of the Nikkei index was a bit less than 50% from December 1989 to September 1990, so it didn't 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.

decay rate to 0.1. The clumped 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 paramater α is as one might expect, while the actual impact is somewhat stronger. In particular, crashes are frequent and prolonged when $\alpha = 4$.

The logit regressions reported in Table 3 provide a complementary perspective. To interpret the coefficient estimates, suppose that the initial probability of a crash is 1%, so initial log odds are $\ln(.01/.99) \approx -4.6$. The coefficient estimate of 1.94 on RDPM-M then implies that, other things equal, were asset price P to rise 20%, then the log odds would increase by about $0.2 * 1.94 \approx 0.4$ to -4.2, implying a crash probability of $p = e^{-4.2}/(1 + e^{-4.2}) \approx 1.5\%$, i.e., the probability would increase by about 50%. The coefficient estimates on the parameters confirm that the effects noted earlier are statistically significant, most at the 1% (or even 0.01%) level. The exceptions are that σ's effects are significant at the 4% level and those of d_R are insignificant. We have also run regressions with asset price normalized by fundamental value V and by predicted price P^{**}, and with more or less stringent definitions of crashes. The results are roughly similar; in some specifications the d_R and σ coefficients gain significance, often at the expense of the price coefficient and the R_o and τ coefficients. We conclude that crashes are indeed more likely when asset price is high and that sampling error plays no important role in the crash frequencies reported in Table 2.

4.5 Summary and Interpretation

The simulation results suggest that the Model 1 has three distinct modes of behavior. The first mode is stable convergence, similar to behavior in Model 0. From reasonable initial states, the market converges to a neighborhood of the clumped 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, low volatility σ, high decay rate τ for luck, and a low γ. In terms of the KMK perspective, the last parameter suggests that investors are not easily dazzled and the other three suggest that unusual opportunities (modeled as positive individual luck) are small and fleeting relative to the size of the market.

In the baseline and several other parameter configurations we observe something close to stable convergence, but to prices and risk positions somewhat higher than in the clumped steady state prediction. That prediction seems downward biased to the extent that managers are heterogeneous rather than clumped. In the simulation, the luckier managers tend to have larger portfolios z and choose larger x than in the clumped steady state (CSS), 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 confituration. (The CSS prediction ignores crashes, so for this mode it has an upward bias that can offset the downward biases due to ignoring heterogeneity and bubbles.) As described in subsection 4.3, this mode has phases reminiscent of KMK and an underlying mechanism reminiscent of escape dynamics. It prevails for moderate parameter values, covering most of the range a priori considered plausible.

A third behavioral mode can be detected for parameter configurations at the opposite extreme to those conducive to the first mode. In particular, for very large α or γ, the observed asset price is usually high above its CSS value or far below, in alternating irregular cycles. KMK's normalcy phase 0 is relatively rare.

5 Model Extensions

The model still lacks some realistic features that might affect performance. Investors chasing the highest returns plays an important role in the KMK narrative, and can be seen in recent field data. Likewise for scale effects, and rank-based incentives for managers. This section sketches how such features are incorporated into the analytic and simulation models.

5.1 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., $\dot{z}_i^R/z_i = R_{Gi} = (R_1 - R_o + \epsilon_i)x_i + R_o$.

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_o not allocated to any fund manager.

Recruitment of new investors depends on relative perceived returns. Recycle the exponential average technique from (13) and apply it to net returns in (11) to get perceived net returns,

$$\hat{R}_i(t) = \gamma \int_{-\infty}^{t} e^{-\gamma(t-s)} R_i(s)ds. \tag{16}$$

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_o, the pool of funds available for investment. Thus we obtain the following expression for fund inflows: $\dot{z}_i^I/z_i = \rho z_o e^{\lambda\hat{R}_i}$. 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_o. 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 is

$$\dot{z}_i = [\dot{z}_i^R + \dot{z}_i^O + \dot{z}_i^I]z_i = [R_o + (R_1 - R_o + \epsilon_i)x_i - \delta\hat{L}_i + \rho z_o e^{\lambda\hat{R}_i}]z_i. \tag{17}$$

Remark. A much simpler specification, less consistent with the evidence, is replicator dynamics e.g., Weibull, 1995). Here $z_o = 0$ and

$$\dot{z}_i = k(\hat{R}_i(t) - \hat{R}_T(t))z_i, \tag{18}$$

where $\hat{R}_T$ is the z-weighted average of $\hat{R}_i$ across managers and k is a rate constant.

5.2 Scale effects, exit and entry, and more

As noted at the end of section 2, evidence suggests that large size tends to depress a manager's returns. We can accommodate that feature by including a quadratic penalty for the size xz_i of the risky portfolio. Thus net returns in (11) become

$$R_i(x) = x(R_1 - R_o + \epsilon_i) - \frac{1}{2}c_2x^2 - \kappa(xz_i)^2, \tag{19}$$

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 don't 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_o) whenever $z_i < z_{\text{LIM}}$. 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 $z_L = 2z_{\text{LIM}}$, $z_U = 2z_L$, 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 83.05 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 the Appendix. Models 82 and 83 exhibit bubbles and crashes similar to those of model 81.

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.

6 Discussion

We explore a perspective on financial market bubbles and crashes originating in the informal writings of Keynes, Minsky and Kindleberger (KMK). After writing out a basic formal model and a gradient process for adjustment, we extend the model to capture some financial market features that loom large in the KMK perspective.

The basic model is very stable. Analytic results, confirmed by simulations, show that for a wide range of exogenous parameters, 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. There is no scope for bubbles and crashes.

Model 1, the first extension of the basic model, features an endogenous risk cost driven by constant gain learning, our formalization of a feature central to the KMK perspective. This extension also has a unique steady state, but has much different dynamics. Although the model typically spends most of the time in the vicinity of the steady state, there are recurrent episodes in which overall market valuations rise substantially (typically 20-50% above normal levels) and then crash (often to a third or less of normal levels within a few months). Such bubble and crash episodes occur over a wide range of "realistic" parameter values.

The episodes become rarer when parameter configurations give investors longer memories or smaller (or more fleeting) unusual opportunities. In opposite configurations the episodes become more common, and with extreme parameter values normalcy becomes rare and the market alternates between bubbles and crashes.

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 macro theory literature, and its continuous action space puts it outside the evolutionary game theory literature cited in section 2.2. Evidently 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. Clearly bubbles and crashes would only be intensified by incorporating other KMK features, such as the heterogeneous expectations, rank-sensitive managers, and exogenous shifts in economy-wide growth opportunities.

Much work remains. We have not yet fully explored the extensions already programmed, intended to capture fickle investors and other realistic KMK features of financial markets. Perhaps other realistic features should be written into the simulations and explored.

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 laboratory experiments with human subjects. We hope that our work inspires new analytical, empirical and laboratory studies that deepen understanding of bubbles and crashes.

7 Appendix: Mathematical Details

Proposition 1. Given fixed parameters $c_2 > 0, \alpha > 0$, $R_o \geq 0$, and $R_s > g_s$, there is a unique point $x^* > 0$ such that the distribution clumped at x^* is a steady state solution to the master equation (7). Moreover, x^* decreases in c_2 and R_o and increases in d_R. It increases in α and in g_s iff $x^* < 1$.

Proof We first show that the equation (10) has a unique solution x^*. Rewrite the equation as

$$U(x) = (R_s - g_s)x^{-\alpha} + g_s - R_o - c_2 x = 0. \tag{20}$$

Note that U is a continuous real valued function which is positive (indeed unbounded) as $x \searrow 0$ and is negative as $x \to \infty$. 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 (20) is necessary and sufficient for a clumped solution to the master equation. That step proceeds exactly as in Friedman (2005, Proposition 2). It is omitted here because it requires several technical details tangential to the concerns of the present paper.

To complete the proof, write $R_s = R_o + d_R$, differentiate (20) with respect to the given parameter and solve to obtain
$\partial x^*/\partial R_o = -1/[c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] < 0$,
$\partial x^*/\partial d_R = x^{-\alpha}/[c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] = -x^{-\alpha}dx^*/dR_o > 0$,
$\partial x^*/\partial g_s = (1 - x^{-\alpha})/[c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] = (x^{-\alpha} - 1)dx^*/dR_o$,
$\partial x^*/\partial c_2 = -x/[c_2 + \alpha(R_s - g_s)x^{-\alpha-1}] < 0$, 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 iff $x^* < 1$. •

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

$$\psi(x) = \int_{-\infty}^{x} (x - y)\Phi'(y)dy = \int_{-\infty}^{x} \Phi(y)dy. \tag{21}$$

The second expression in (21) is obtained from the first 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 $q(x|c_2) = (x\sigma/\sqrt{2\tau})\psi(z^o(x))$, where $z^o(x) = (-\sqrt{2\tau}/\sigma)[R_o(1/x - 1) + g_s + (R_s - g_s)(x^*)^{-\alpha}]$ 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, $R_{Gi} = (R_1 - R_o + \epsilon_i)x + R_o$. The unconditional distribution of ϵ_i (obtained as the $h \to \infty$ limit in (12)) 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 for R_{Gi} to conclude that its unconditional distribution F is normal with mean $\mu \equiv [(R_s - g_s)(x^*)^{-\alpha} + g_s - R_o]x + R_o$ 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 $z^o \equiv -\mu/s = (-\sqrt{2\tau}/\sigma)[R_o(1/x - 1) + g_s + (R_s - g_s)(x^*)^{-\alpha}]$, the expected loss is

$$\int_{-\infty}^{0} [0 - r]F'(r)dr = s\int_{-\infty}^{z^o} [z^o - z]\Phi'(z)dz = s\psi(z^o) = (x\sigma/\sqrt{2\tau})\psi(z^o). \tag{22}$$

•

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

Proof. Equation (22) gives the expected loss as $Q(x) = ax\psi(z^o(x))$, where $a = \sigma/\sqrt{2\tau} > 0$ and $z^o(x) = -b/x + k$ for $b = R_o/a > 0$. It is immediate from (21) that $\psi'(y) = \Phi(y) \geq 0$, $\psi''(y) = \Phi'(y) \geq 0$, and $\psi(y) \to 0$ as $y \to -\infty$. It now follows that $q(0) = 0a\psi(-\infty) = 0$. Straightforward computations show that $q'(x) = a[\psi(z^o) + (b/x)\psi'(z^o)] \geq 0$ and $q''(x) = (ab^2/x^3)\psi''(z^o) \geq 0$. Hence q is increasing and convex. By L'Hospital's rule, $q'(0) = ab\Phi'(-\infty) = 0$. •

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

$$F_t(x,t) = -F_x(x,t)\phi_x(x,F) + [\int_0^x Z_{xt}(y,t)dy - F(x,t)Z_t(t)]/Z(t). \tag{23}$$

Here subscripts denote partial derivatives, and $Z(x,t)$ includes birth and death rates as well as the fickle investor effects in (17). 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.

8 Bibliography

Arifovic, Jasmina. (2000), ”Evolutionary Algorithms in Macroeconomic Models”, Macroeconomic Dynamics, Vol. 4(3), pp 373-414

Arthur, WB. JH Holland, B LeBaron, R Palmer, and P Tayler (1996) ” Asset Pricing Under Endogenous Expectations in an Artificial Stock Market”,in WB Arthur, et al, The Economy as an Evolving Complex System, Reading, MA: Addison-Wesley.

Blanchard, OJ and MW Watson (1983), ”Bubbles, Rational Expectations and Financial Markets”, NBER working paper, No. 945

Brock, WA and CH Hommes (1998) ”Heterogeneous beliefs and routes to chaos in a simple asset pricing model” Journal of Economic Dynamics and Control, Vol. 22(8-9), pp 1235-1274

Camerer, Colin. (1989). ”Bubbles and Fads in Asset Prices: A Review of Theory and Evidence”. Journal of Economic Surveys, Vol. 3(1), pp 3-41.

Campbell, John Y., and John H. Cochrane. (1999) ”By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior”, Journal of Political Economy, Vol. 107, pp 205-251.

Chen, Joseph, Harrison Hong, Ming Huang and Jeffrey D. Kubik. (2004) ”Does Fund Size Erode Mutual Fund Performance? The Role of Liquidity and Organization”, The American Economic Review Vol. 94 (5), Dec. 2004. pp.1276-1302

Cheung, Yin-Wong and Daniel Friedman. (1997), ”Individual Learning in Normal Form Games: Some Laboratory Results”. Games and Economic Behavior, Vol. 19(1), pp. 46-76

Cho, IK. N. Williams and TJ. Sargent (2002). ”Escaping Nash Inflation”, Review of Economic Studies Vol. 69. pp. 1-40.

Copeland, Thomas E., J.F. Weston, and K. Shastri.(2005) Financial Theory and Corporate Policy (4th edition). New York: Pearson/Addison Wesley, 2005.

DeLong, JB, A. Schleifer, LH Summers and RJ Waldmann(1990) ”Positive Feedback Investment Strategies and Destabilizing Rational Speculation” Journal of Finance 45, 397-395.

Del Guerico, Diane, and Paula A. Tkac. (2002). ”The determinants of the flow of funds of managed portfolios: mutual funds versus pension funds: Mutual Funds vs. Pension Funds." Journal of Financial and Quantitative Analysis, Vol. 37, pp. 523-557.

Farmer, J. Doyne, Patelli, Paolo, and Ilija I. Zovko. (2005). "The Predictive Power of Zero Intelligence in Financial Markets." The National Academy of Sciences.

Feller, William. (1971). "An introduction to probability theory and its applications", Vol. 2. Wiley Series in Probability and Mathematical Statistics, New York: Wiley, 1971, 3rd ed

Freidlin, M.I. and A.D. Wentzell, (1984) Random Perturbations of Dynamic System, New York, Springer-Verlag, 1984

Friedman, Daniel. (2005). "Conspicuous Consumption Dynamics". UCSC working paper No. 609 http://www.vismath.org/research/landscapedyn/models/articles/

Friedman, Daniel and Masanao Aoki. (1992) "Inefficient Information Aggregations as a Source of Asset Price Bubbles". Bulletin of Economics Research, Vol. 44, pp. 251-279.

Friedman, Daniel and Joel Yellin, (1997). "Evolving Landscapes for Population Games", University of California, Santa Cruz draft manuscript, available at:
http://www.vismath.org/research/landscapedyn/models/articles/

Garber, Peter. (1989) "Tulipmania", Journal of Political Economy, Vol. 97(3), pp. 535-60

Hommes, Cars H. (2006). "Heterogeneous Agent Models in Economics and Finance", in L. Tesfatsion and K. Judd, eds, Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics. NY: Elsevier.

Youssefmir , Michael; Bernardo A Huberman, and Tad Hogg, (1998) "Bubbles and Market Crashes" Computational Economics, Vol. 12(2), pp.97-114

Karceski, Jason. (2002) "Returns-Chasing Behavior, Mutual Funds, and Beta's Death", Journal of Financial and Quantitative Analysis, Vol. 37, pp. 559-594.

Keynes, John Maynard. (1936). The General Theory of Employment, Interest and Money. London: Macmillan

Kindleberger, Charles P. (1978/1989/2000). Manias, Panics and Crashes: A History of Financial Crises. NY: Basic Books.

LeBaron, Blake. (2006). Agent-based Computational Finance. in L. Tesfatsion and K. Judd, eds, Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics. NY: Elsevier.

LeRoy, Stephen F. (2004). "Rational Exuberance". Journal of Economic Literature, Vol. 42(3), pp. 783-804.

Mackay, Charles. (1841) Extraordinary Popular Delusions and the Madness of Crowds, 1841. (reprinted in Wiley's 1996 investment classics volume.)

Minsky, Hyman P. (1975). John Maynard Keynes. NY: Columbia University Press.

Minksy, Hyman P. (1982). "Can 'IT' Happen Again?" Essays on Instability and Finance. NY: M.E. Sharpe Inc.

Penso de la Vega, Josef, Confusion e confusions, (1688). (reprinted in Wiley's 1996 investment classics volume.)

Sargent, Thomas. (1999). "The Conquest of American Inflation" NY: Princeton University Press

Sirri, Erik R. and Peter Tufano. (1998), "Costly Search and Mutual Fund Flows", Journal of Finance, Vol. 53, pp. 575-603.

Tirole, Jean. (1982). "On the Possibility of Speculation under Rational Expectations", Econometrica, Vol. 50(5), pp.1163-81

Williams, Noah. (2004). Escape Dynamics in Learning Models. Manuscript, Princeton University Economics Department.

Young, Peyton. (1993) "The Evolution of Conventions". Econometrica, Vol. 61(1), pp.57-84

Zeeman, E. Christopher. 1974. "On the Unstable Behavior of the Stock Exchanges." Journal of Mathematical Economics. Vol. 1(1), pp. 39-44.

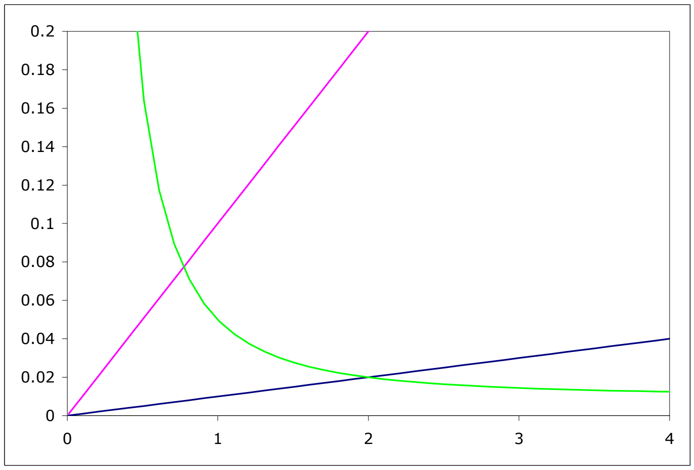

Figure 1: Equilibrium risk position x^*. Parameter values determining the hyperbola are $R_o = 3\%, R_s = 6\%, g_2 = 2\%$ and $\alpha = 2.0$. The two rays are for alternative c_2 values 1% and 10%. The equilibrium risk position is the horizontal component of the intersection of the ray and hyperbola.

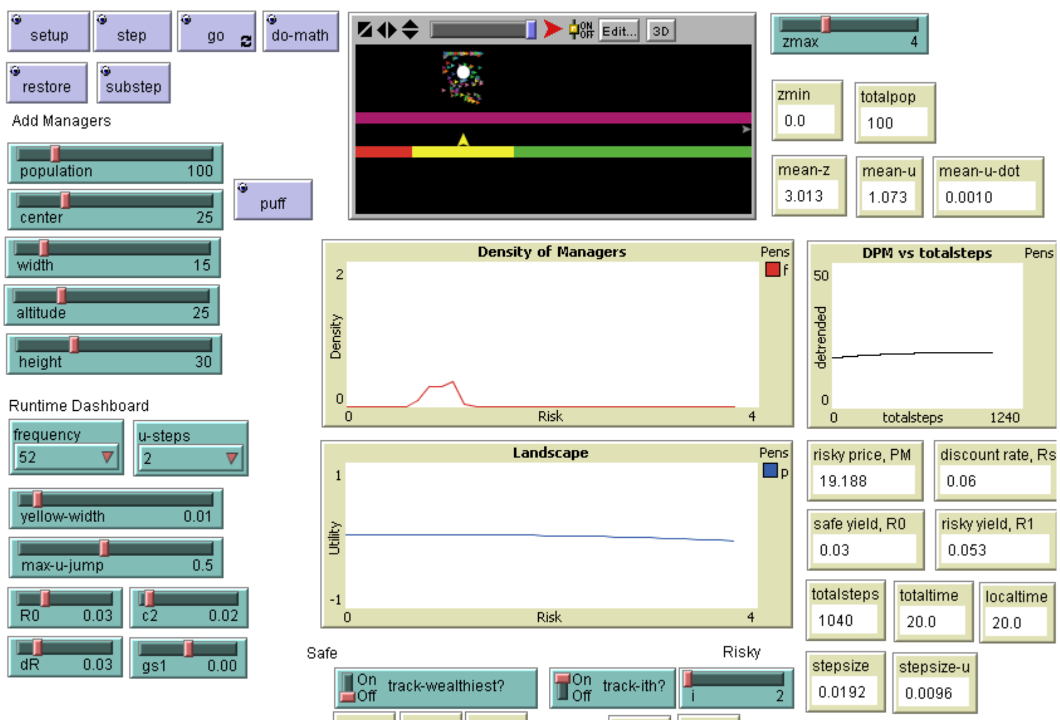

Figure 2: User interface for base simulation

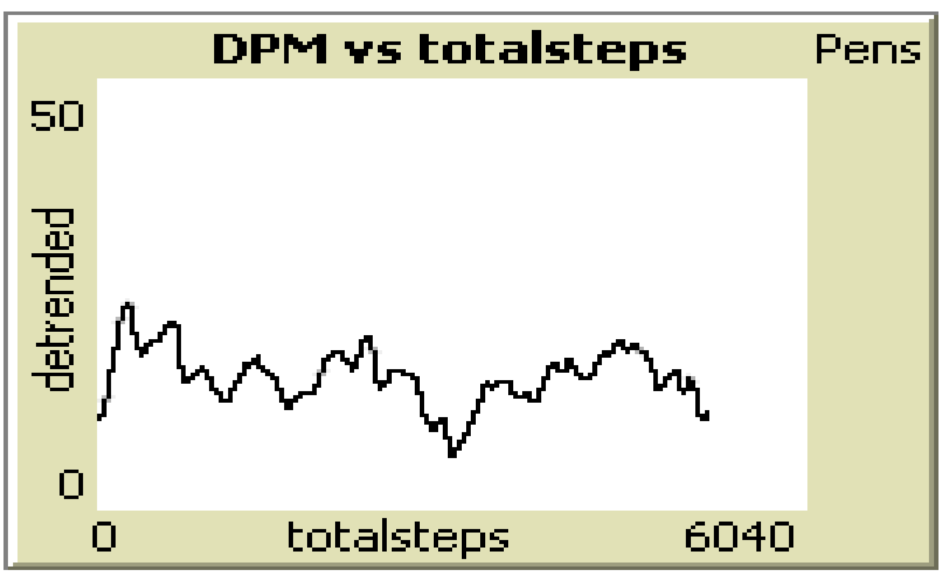

Figure 3: Detrended Asset Price. A 100 Year Simulation of Model 1 using baseline parameters $M = 30, R_o = d_R = 0.03, g_s = 0, \alpha = 2, \sigma = 0.20, \beta = 2.0$ and $\tau = \gamma = 0.7$.

Table 1. Simulation Results for Model 0

Parameters		Clumped Steady State Prediction			Simulation Mean ±Std. Dev	
Variable	**Value**	**V**	**P***	**x***	**P**	**x**
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.
Each mean ± standard deviation is taken over 41600 observations: 10 centuries of weekly observations, with the first two decades deleted to reduce the impact of the initial population distribution.

Table 2. Simulation Results for Model 1

Parameters		Clumped, Steady State Prediction			Simulation Mean		
variable	**value**	**V**	**P***	**x***	**P**	**x**	**Crashes / Century**
Baseline		16.7	12.5	0.87	14.8 ±3.4	0.94 ±0.11	0.5
Population	15.00	16.7	12.5	0.87	13.8 ±4.1	0.90 ±0.14	1.6
	100.00	16.7	12.5	0.87	15.3 ±1.6	0.96 ±0.05	0
Ro	0.01	16.7	16.2	0.80	14.2 ±3.3	0.92 ±0.11	0.9
	0.05	16.7	10.3	0.91	14.9 ±2.8	0.94 ±0.09	0.2
dR	0.01	16.7	14.3	0.76	17.4 ±3.9	0.83 ±0.10	0.7
	0.05	16.7	11.4	0.95	13.2 ±2.5	1.02 ±0.10	0.3
gs1	-0.04	16.7	9.2	0.74	9.2 ± 1.5	0.96 ±0.08	0
	0.04	16.7	18.0	1.04	38.8 ±14.1	0.86 ±0.17	1.3
sigma	0.05	16.7	23.7	1.19	23.1 ±1.2	1.18 ±0.03	0
	0.40	16.7	8.7	0.72	14.7 ±3.6	0.93 ±0.12	0.6
tau	0.10	16.7	7.5	0.67	57.1 ±26.9	1.78 ±0.51	1.7
	3.00	16.7	17.9	1.04	16.4 ±1.2	0.99 ±0.04	0
eta	0.10	16.7	12.5	0.87	14.9 ±2.3	0.94 ±0.07	0
	3.00	16.7	12.5	0.87	13.5 ±4.5	0.88 ±0.18	4.5
beta	1.00	16.7	15.2	0.95	18.2 ±3.8	1.04 ±0.11	0.2
	5.00	16.7	9.7	0.76	11.0 ±2.6	0.81 ±0.11	1.4
alpha	1.00	16.7	13.4	0.81	15.8 ±1.9	0.95 ±0.11	0
	4.00	16.7	11.8	0.92	10.4 ±5.7	0.86 ±0.14	6.1

Note: Baseline values are as in Figure 3.
Means ± standard deviation are over 10 centuries of weekly observations.

Table 3. Logistic regression for crashes.

Parameter	DF	Estimate	Standard Error	Wald Chi-Square	Pr > ChiSq
Intercept	1	-9.5614	1.0802	78.3497	<.0001
population	1	-0.1088	0.0212	26.3342	<.0001
R0	1	-48.8239	18.0990	7.2770	0.0070
dR	1	-29.2605	19.3792	2.2798	0.1311
gs1	1	34.2313	7.9701	18.4467	<.0001
sigma	1	4.0874	1.9500	4.3935	0.0361
tau	1	-2.6528	0.5230	25.7308	<.0001
eta	1	1.1475	0.1066	115.8002	<.0001
beta	1	0.4975	0.1084	21.0684	<.0001
alpha	1	1.2995	0.1236	110.5284	<.0001
RDPM_M	1	1.9366	0.1584	149.4600	<.0001

Note: The dependent variable Crash=1 in a given week if the price falls to less than 50% of current value within the next 26 weeks, and otherwise Crash=0. The data are all observations from Table 2, except a string of consecutive Crash=1 observations is counted only once. The explanatory variables are model parameters defined in the text, except for RDPM_M. It is the detrended weekly asset price (DPM) divided by its mean conditioned on the parameter values. (The means are from Table 2 with the first two decades deleted from each simulation, to reduce the impact of the initial population distribution.)