

This article was downloaded by: [stefan thurner]

On: 06 May 2012, At: 07:48

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Quantitative Finance

Publication details, including instructions for authors and subscription information:
<http://www.tandfonline.com/loi/rquf20>

Leverage causes fat tails and clustered volatility

Stefan Thurner^{a b c}, J. Doyne Farmer^b & John Geanakoplos^{b d e}

^a Section of Science of Complex Systems, Medical University of Vienna, Spitalgasse 23, Vienna A-1090, Austria

^b Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

^c IIASA, Schlossplatz 1, A-2361 Laxenburg, Austria

^d James Tobin Professor of Economics, Yale University, New Haven, CT, USA

^e Ellington Capital Management, Old Greenwich, CT, USA

Available online: 04 May 2012

To cite this article: Stefan Thurner, J. Doyne Farmer & John Geanakoplos (2012): Leverage causes fat tails and clustered volatility, *Quantitative Finance*, 12:5, 695-707

To link to this article: <http://dx.doi.org/10.1080/14697688.2012.674301>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Leverage causes fat tails and clustered volatility

STEFAN THURNER*†‡§, J. DOYNE FARMER‡ and JOHN GEANAKOPOLOS‡¶||

†Section of Science of Complex Systems, Medical University of Vienna, Spitalgasse 23, Vienna A-1090, Austria

‡Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

§IIASA, Schlossplatz 1, A-2361 Laxenburg, Austria

¶James Tobin Professor of Economics, Yale University, New Haven, CT, USA

||Ellington Capital Management, Old Greenwich, CT, USA

(Received 31 January 2011; in final form 29 February 2012)

We build a simple model of leveraged asset purchases with margin calls. Investment funds use what is perhaps the most basic financial strategy, called ‘value investing’, i.e. systematically attempting to buy underpriced assets. When funds do not borrow, the price fluctuations of the asset are approximately normally distributed and uncorrelated across time. This changes when the funds are allowed to *leverage*, i.e. borrow from a bank, which allows them to purchase more assets than their wealth would otherwise permit. During good times, funds that use more leverage have higher profits, increasing their wealth and making them dominant in the market. However, if a downward price fluctuation occurs while one or more funds is fully leveraged, the resulting margin call causes them to sell into an already falling market, amplifying the downward price movement. If the funds hold large positions in the asset, this can cause substantial losses. This in turn leads to clustered volatility: before a crash, when the value funds are dominant, they damp volatility, and after the crash, when they suffer severe losses, volatility is high. This leads to power-law tails, which are both due to the leverage-induced crashes and due to the clustered volatility induced by the wealth dynamics. This is in contrast to previous explanations of fat tails and clustered volatility, which depended on ‘irrational behavior’, such as trend following. A standard (supposedly more sophisticated) risk control policy in which individual banks base leverage limits on volatility causes leverage to rise during periods of low volatility, and to contract more quickly when volatility becomes high, making these extreme fluctuations even worse.

Keywords: Systemic risk; Clustered volatility; Fat tails; Crash; Margin calls; Leverage

JEL Classification: E32, E37, G01, G12, G14

1. Introduction

Recent events in financial markets have underscored the dangerous consequences of the use of excessive credit. At the most basic level the problem is obvious: if a firm buys assets with borrowed money, then under extreme market conditions it may owe more money than it has and defaults. If this happens on a sufficiently wide scale, then it can severely stress creditors and cause them to fail as well.

We show here that a special but extremely widespread kind of credit called collateralized loans with margin calls has a more pervasive effect: when used excessively it can cause default and crashes, but it also leaves a signature

even when there is no default or crash. These kinds of loans have already been identified as a major culprit in the recent crisis, and in previous near crises as well.^T But we show here that they create a dynamic in asset price fluctuations that manifests itself at all time scales and to all degrees. The extraordinary crisis of the last couple of years is just one extreme (but not extremal) point on a continuum.

By taking out a collateralized loan a buyer of stocks or mortgage-backed securities can put together a portfolio that is worth a multiple of the cash he has available for their purchase. In 2006 this multiple or ‘leverage’ reached 60 to 1 for AAA-rated mortgage securities, and 16 to 1 for

*Corresponding author. Email: stefan.thurner@meduniwien.ac.at

^TFor previous equilibrium-based analyses of leverage that show that prices crash before default actually occurs, see Geanakoplos (1997, 2003), Fostel and Geanakoplos (2008), Brunnermeier and Pedersen (2009) and Geanakoplos (2010).

what are now called the toxic mortgage securities. The outstanding volume of these leveraged asset purchases reached many trillions of dollars. Leverage has fluctuated up and down in long cycles over the last 30 years.

Conventional credit is for a fixed amount and a fixed maturity, extending over the period the borrower needs the money. In a collateralized loan with margin calls, the debt is guaranteed not by the reputation (or punishment) of the borrower, but by an asset that is confiscated if the loan is not repaid. Typically, the loan maturity is very short, say a day, much shorter than the length of time the borrower anticipates needing the money. The contract usually specifies that, after the daily interest is paid, as long as the loan-to-asset-value ratio remains below a specified threshold, the debt is rolled over another day (up to some final maturity, when the threshold ratio might be changed). If, however, the collateral asset value falls, the lender makes a margin call and the borrower is expected to repay part of the debt and so roll over a smaller loan to maintain the old loan-to-value threshold. Quite often, the borrower will obtain the cash for this extra downpayment by selling some of the collateral. The nature of the collateralized loan contract thus sometimes turns buyers of the collateral into sellers, even when they might think it is the best time to buy.

Such an effect is well known to every hedge fund manager who uses leverage. It was discussed informally by Gennotte and Leland (1990) and Shleifer and Vishny (1997) where it was noted that leveraged value investors may cause mispricings to increase when they hit margin limits. Geanakoplos (2003) presents an explicit model of leverage and prices and solves for the equilibrium leverage and prices when every agent is fully rational. There the margin call effect is compounded by an increase in the endogenous equilibrium haircut. However, the model extends only for three periods, with just two possible shocks each period. We provide a quantitative dynamic model with arbitrarily many periods and continuous shocks of all sizes. This allows us to study how value investors decrease volatility under most circumstances, but occasionally dramatically increase volatility and generate crashes. It also allows us to examine the statistical signature of leveraged trading and to compare statistical measures of returns, such as kurtosis, in our model with measures obtained from actual data. In our model, fat tails and clustered volatility are statistically testable properties. By contrast, in the purely descriptive commentary of Shleifer and Vishny (1997), they can only suggest that price changes will be bigger than the shocks to fundamentals.

Needless to say, the higher the loan to value, or, equivalently, the higher the *leverage* ratio of asset value to cash downpayment, the more severe will be the feedback mechanism. A buyer who is at his threshold of λ times

leverage loses $\lambda\%$ of his investment for every 1% drop in the asset price, and on top of that will have to come up with $\$(\lambda - 1)/\lambda$ of new cash for every \$1 drop in the price of the asset. When there is no leverage, and $\lambda = 1$, there is no feedback, but as the leverage increases, so does the feedback.

The feedback from falling asset prices to margin calls to the transformation of buyers into sellers back to falling asset prices creates a nonlinear dynamic to the system. The nonlinearity rises as the leverage rises. This nonlinear feedback would be present in the most sophisticated rational expectations models or in the most simple-minded behavioral models: it is a mechanical effect that stems directly from the nonlinear dynamics caused by the use of leverage and margin calls. We therefore build the simplest model possible and then simulate it over tens of thousands of periods, measuring and quantifying the effect of leverage on asset price fluctuations.[†]

Our model provides a new explanation for the fat tails and clustered volatility that are commonly observed in price fluctuations (Mandelbrot 1963, Engle 1982). Clustered volatility and fat tails emerge in the model on a broad range of time scales, including very rapid ones and very slow ones. Mandelbrot and Engle found that actual price fluctuations did not display the independent and normally distributed properties assumed by the pioneers of classical finance (Bachelier 1964, Black and Scholes 1973). Although their work has been properly celebrated, no consensus has formed on the mechanism that creates fat tails and clustered volatility. The mechanism we develop here supports the hypothesis that they are caused by the endogenous dynamics of the market rather than the nature of information itself—in our model, information is normally distributed and i.i.d., but when leverage is used, the resulting prices are not.

Previous endogenous explanations assume the presence of a kind of trader who exacerbates fluctuations. Traders in these models are of at least two types: value investors, who make investments based on fundamentals, and trend followers, who make investments in the direction of recent price movements.[‡] Trend followers are inherently destabilizing, and many would dispute whether such behavior is rational. Value investors, in contrast, are essential to maintain a reasonably efficient market: they gather information about valuations, and incorporate it into prices. Thus, in this sense, value investing is rational. In typical models of this type, investors move their money back and forth between trend strategies and value strategies, depending on who has recently been more successful, and fat tails and clustered volatility are generated by temporary increases in destabilizing trend strategies.

The mechanism that we propose here for fat tails and clustered volatility only involves value investors, who are

[†]The nonlinear feedback that we describe here, which is driven by investors selling into a falling market, is in this sense similar to the model of hedging by Gennotte and Leland (1990); they also discuss how such feedbacks can cause crashes.

[‡]See Palmer *et al.* (1994), Arthur *et al.* (1997), Brock and Hommes (1997, 1998), Caldarelli *et al.* (1997), Lux (1999) Lux and Marchesi (1999) and Giardina and Bouchaud (2003). See also Friedman and Abraham (2009), who induce bubbles and crashes via myopic learning dynamics.

stabilizing in the absence of leverage. We do not claim that our mechanism for making fat tails is the only possible mechanism—indeed, it likely coexists with the myopic learning mechanism reviewed above, and may also coexist with other mechanisms, such as fat tails in exogenous information arrival.

An important aspect of our model is that even though the risk control policies used by the individual bank lenders are reasonable from a narrow, bank-centric point of view, when a group of banks inadvertently acts together, they can dramatically affect prices, inducing nonlinear behavior at a systemic level that gives rise to excessive volatility and even crashes. Attempts to regulate risk without taking into account systemic effects can backfire, accentuating risks or even creating new ones.[†]

The wealth dynamics in our model illustrate the interaction between evolutionary dynamics that occur on very long time scales, and short-term dynamics that occur on timescales of minutes. In our model, different agents use different levels of leverage. Agents who use more leverage produce higher returns and attract more investment capital, and as time goes on the most aggressive investors accumulate more wealth. Whereas funds normally damp price fluctuations by buying when the price falls, if they are fully leveraged, the margin call caused by a small downward price fluctuation can force them to sell into a falling market. In the early stages of a bubble, when the wealth of the funds is low, their positions are small, their impact on the market is low, and this is relatively harmless. However, in the later stages, when the combination of fund wealth and leverage is large, the impact is correspondingly large, and a relatively small downward price movement can trigger a crash.

The above scenario illustrates the evolutionary pressure driving funds toward higher and higher leverage. During stable periods in the market, funds that use large leverage grow at the expense of those who do not, and acquire more and more market power, while funds that do not employ sufficient leverage lose investment capital. Even if fund managers are aware of the danger of using leverage, the pressure of short-term competition may force them to do so. Regulating leverage is thus good for everyone, preventing behavior that all are driven to yet no one desires.

The leverage effect that we explore here is just one example of many types of nonlinear positive feedback that are often referred to as ‘pro-cyclical behavior’ in the economics literature. Other examples include stop-loss orders, exercise of put or call options, trend-following and dynamic hedging strategies. All of these have the common feature that they generate additional buying or selling in the direction the price is already moving, thereby amplifying a pre-existing trend. Furthermore, with the exception of trend following, these are all essentially

mechanical effects that, once contracts are in place, can lead to the amplification of price movements without any further decision making. Our work here is in the spirit of the pioneering paper of Kim and Markowitz (1989), who simulated dynamic hedging strategies believed to be involved in the crash of 1987 and demonstrated their effect on time series of prices. The destabilizing effects of derivatives have been studied by Brock *et al.* (2009) and Caccioli *et al.* (2009).

We wish to emphasize that we do not claim here that excessive borrowing by hedge funds caused the liquidity crisis of 2007 onwards. This work is instead designed to illustrate the general problems associated with leverage. The heavy-tailed price movements we demonstrate here, which are caused by selling into a falling market, should be observed in any situation where there are collateralized loans with margin calls, whether or not the borrowers are value investors.[‡]

2. The model

In our model, traders have a choice between owning a single asset, such as a stock or a commodity, or owning cash. There are two types of traders, noise traders and funds. The noise traders buy and sell nearly at random, with a slight bias that makes the price weakly mean-revert around a perceived fundamental value V . The funds use a strategy that exploits mispricings by taking a long position (holding a net positive quantity of the asset) when the price is below V , and otherwise staying out of the market. The funds can augment the size of their long position by borrowing from a bank at an interest rate that, for simplicity, we fix at zero, using the asset as collateral. This borrowing is called leverage. The bank will of course be careful to limit its lending so that the value of what is owed is less than the current price of the assets held as collateral. Default occurs if the asset price falls sufficiently far before the loan comes due in the next period.

In addition to the two types of traders there is a representative investor who either invests in a fund or holds cash. The amount she invests in a given fund depends on its recent historical performance relative to a benchmark return r^b . Thus successful funds attract additional capital above and beyond what they gain in the market and similarly unsuccessful funds lose additional capital.

2.1. Supply and demand

The total supply of the asset is N . At the beginning of each period $t \geq 1$, all agents observe the unit asset price $p(t)$. As is traditional, all the traders in our model are

[†]Another good example from the recent meltdown illustrating how individual risk regulation can create systemic risk is the use of derivatives.

[‡]The failure of Long Term Capital Management in 1998 was an example of a near-crisis caused by the precise mechanism discussed here. Some other types of investment strategies, such as trend-following or portfolio insurance, cause nonlinear feedback in prices, which is further amplified by leverage.

perfectly competitive; they take the price as given, imagining that they are so small that they cannot affect the price, no matter how much they demand.

2.1.1. Noise traders. The noise traders' demand is defined in terms of the cash value $\xi_{nt}(t)$ they spend on the asset, which follows an autoregressive random process of order 1 of the form

$$\log \xi_{nt}(t) = \rho \log \xi_{nt}(t-1) + \sigma \chi(t) + (1-\rho) \log(VN),$$

where χ is normally distributed with mean zero and standard deviation one. The noise traders' demand is

$$D_{nt}(t, p(t)) = \frac{\xi_{nt}(t)}{p(t)}.$$

When there are only noise traders the price is set such that $D_{nt}(t, p(t)) = N$. This choice of the noise trader process guarantees that with $\rho < 1$ the price is a mean-reverting random process with $E[\log p] = \log V$.

When there are only noise traders the log price follows an AR(1) process and so is normally distributed. However, the returns $r(t) = \log(p(t)/p(t-1))$ are somewhat more heavy tailed than normal. Furthermore, the mean reversion introduces a slight amount of clustered volatility. Both of these facts are shown in the appendix. For the purposes of this paper we fix $V=1$, $N=1000$, $\sigma=0.035$ and $\rho=0.99$. The choice of $\rho \approx 1$ ensures that the deviation from normality is minimal; with $\rho=0.99$ the typical fluctuation in volatility is about 1%.

2.1.2. Funds. We add a second class of demanders called funds. The funds in our model are *value investors* who base their demand $D_h(t)$ on a mispricing signal $m(t) = V - p(t)$. The perceived fundamental value V is held constant and is the same for the noise traders and for all funds. As shown in figure 1, each fund h computes its demand $D_h(t)$ based on the mispricing. As the mispricing increases, the dollar value $D_h(t)p(t)$ of the asset the fund wishes to hold increases linearly, but the position size is capped when the fund reaches the maximum leverage. Funds differ only according to an aggression parameter β_h that represents how sensitive their response is to the signal m .

Each fund begins with the same wealth $W_h(0) = 2$. After noting the price $p(t)$, at each date $t \geq 1$ hedge fund h computes its wealth $W_h(t) = W_h(t, p(t))$, as described in the next section. The fund must split its wealth $W_h(t)$ between cash $C_h(t)$ and the value of the asset $D_h(t)p(t)$,

$$W_h(t) = W_h(t, p(t)) = D_h(t)p(t) + C_h(t).$$

When the fund is borrowing money, $C_h(t)$ is negative and represents the loan amount. If $W_h(t, p(t)) \geq 0$, the fund's demand $D_h(t) = D_h(t, p(t))$ can be written as

$$m(t) < 0 : \quad D_h(t) = 0, \quad (1)$$

$$0 < m < m_h^{\text{crit}} : \quad D_h(t) = \beta_h m(t) W_h(t) / p(t), \quad (2)$$

$$m \geq m_h^{\text{crit}} : \quad D_h(t) = \lambda_{\text{MAX}} W_h(t) / p(t). \quad (3)$$

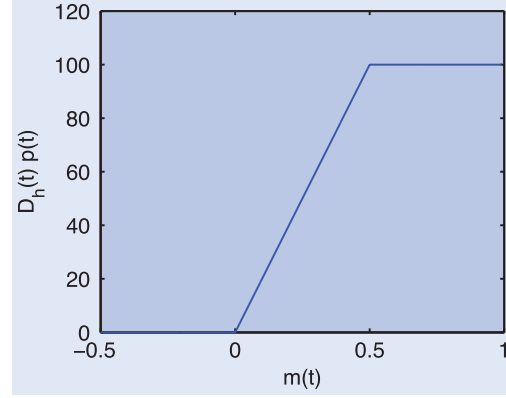


Figure 1. Demand function $D_h(t)p(t)$ of a fund (measured in dollars) vs. the mispricing signal $m(t) = V - p(t)$. The investor does nothing when the asset is overpriced, and she buys more and more as the asset becomes underpriced, until she hits her leverage limit at $m = m^{\text{crit}}$. After this her demand remains flat.

In (1) the asset is over-priced and the fund holds nothing. In (2) the asset is underpriced but the mispricing is not too large. The fund takes a position whose monetary value is proportional to the mispricing $m(t)$, the fund's wealth $W_h(t)$, and the aggression parameter β_h , which can vary from fund to fund. In (3) the asset is even more underpriced so that the fund has reached its maximum leverage $\lambda_h(t) = \lambda_{\text{MAX}}$. This occurs when $m(t) \geq m_h^{\text{crit}} = \lambda_{\text{MAX}} / \beta_h$.

The leverage λ_h is the ratio of the dollar value of the fund's asset holdings to its wealth, i.e.

$$\lambda_h(t) = \frac{D_h(t)p(t)}{W_h(t)} = \frac{D_h(t)p(t)}{(D_h(t)p(t) + C_h(t))}. \quad (4)$$

The fund is required by the bank it borrows from to maintain $\lambda_h(t) \leq \lambda_{\text{MAX}}$. If $\bar{\lambda}_h(t) = D_h(t-1)p(t)/W_h(t) > \lambda_{\text{MAX}}$, the fund will have to sell the asset in order to bring leverage $\lambda_h(t)$ under the maximum allowed. This is called meeting a margin call.

Note that a $k\%$ change in the asset price from $p(t-1)$ to $p(t)$ causes a $\bar{\lambda}_h(t)k\%$ change in wealth $W_h(t)$, hence the name 'leverage'. A fund that satisfied its leverage limit at time $t-1$ might face a margin call at time t either because $\bar{\lambda}_h(t) > 1$ and $p(t)$ falls below $p(t-1)$, causing $W_h(t)$ to fall by a larger percentage than the asset price, or because $W_h(t)$ falls below $W_h(t-1)$ due to redemptions, described in the next section.

If $W_h(t) < 0$, the fund defaults and goes out of business. The fund sells all its assets, demanding $D_h(t) = 0$, and returns all the revenue to pay off as much of its borrowed money as it can to its bank lender. The bank bears the loss of the default. For simplicity, we assume the bank has deep pockets and, despite the loss, continues to lend to other funds as before. After a period of time has passed, the defaulting fund reemerges again as a new fund, as we shall describe below.

Prices are set by equating the demand of the funds plus the noise traders to the fixed supply of the asset

$$D_{nt}(t, p(t)) + \sum_h D_h(t, p(t)) = N.$$

2.2. Fund wealth dynamics

The funds' wealth automatically grows or shrinks according to the success or failure of their trading. In addition, it changes due to additions or withdrawals of money by investors, as described below. If a fund's wealth goes below a critical threshold, here set to $W_h(0)/10$, the fund goes out of business,[†] and after a period of time, T_{reintro} , has passed it is replaced by a new fund with wealth $W_h(0)$ and the same parameters β_h and λ_{MAX} . In the simulations we use $T_{\text{reintro}} = 100$ time steps.

A pool of fund investors, who are treated as a single representative investor, contribute or withdraw money from each fund based on a moving average of its recent performance. This kind of behavior is well documented.[‡]

Let

$$r_h(t) = \frac{D_h(t-1)(p(t) - p(t-1))}{W_h(t-1)}$$

be the rate of return by fund h on investments at time t . The investors make their decisions about whether to invest in the fund based on $r_h^{\text{perf}}(t)$, an exponential moving average of these performances, defined as

$$r_h^{\text{perf}}(t) = (1-a)r_h^{\text{perf}}(t-1) + ar_h(t). \quad (5)$$

The flow of capital in or out of the fund, $F_h(t)$, is given by

$$\tilde{F}_h(t) = b[r_h^{\text{perf}}(t) - r^b][D_h(t-1)p(t) + C_h(t-1)], \quad (6)$$

$$F_h(t) = \max(\tilde{F}_h(t), -[D_h(t-1)p(t) + C_h(t-1)]), \quad (7)$$

where b is a parameter controlling how sensitive the percentage contributions or withdrawals are to returns and r^b is the benchmark return of the investors. The investors cannot take out more money than the fund has.

Funds are initially given wealth $W_0 = W_h(0)$. At the beginning of each new time step $t \geq 1$, the wealth of the fund changes according to

$$W_h(t) = W_h(t-1) + [p(t) - p(t-1)]D_h(t-1) + F_h(t). \quad (8)$$

In the simulations in this paper, unless otherwise stated we set $a = 0.1$, $b = 0.15$, $r^b = 0.005$, and $W_0 = 2$.

The benchmark return r^b plays the important role of determining the relative size of hedge funds vs. noise traders. If the benchmark return is set very low, then funds will become very wealthy and will buy a large quantity of the asset under even small mispricings, preventing the mispricing from ever growing large. This effectively induces a hard floor on prices. If the benchmark return is set very high, funds accumulate little wealth and play a negligible role in price formation. The interesting behavior is observed at intermediate values of r^b where the funds' demand is comparable to that of the noise traders.

2.3. A few remarks about the model

2.3.1. Lack of short selling and its consequences. We have intentionally avoided short selling because short positions are inherently riskier than long positions. With an unleveraged long-only position it is not possible to lose more than one owns. In contrast, with a short position it is possible to lose an arbitrarily large amount, even without leverage. Because we wanted to be able to switch off excess riskiness completely, we intentionally kept short selling out of this model.

The disadvantage of this approach is that it makes the model explicitly unrealistic in ways that need to be taken into account when interpreting the results. When the asset is overpriced, long-only funds are entirely out of the market, which can cause strong asymmetries in the properties of prices. Since the funds normally damp excursions from fundamentals, it can mean that the volatility is higher when the asset is overpriced than when it is underpriced. Mixing the two cases together would result in artificially induced heavy tails and give an artificial impression of clustered volatility. The predictions of the model are only relevant when the asset is underpriced and we therefore condition our analyses on the asset being underpriced.

2.3.2. Trend following in wealth dynamics. The wealth dynamics of the funds involves a representative investor who takes her money in or out of the fund based on its recent performance. We introduced this into our model because it guarantees a steady-state behavior, with well-defined long-term statistical averages. Without this the wealth of the funds grows without bound, since the funds consistently profit at the expense of the noise traders. This causes the price to eventually 'freeze' with the value V as a floor due to the fact that any underpricing is immediately corrected by the funds. Since the wealth dynamics we have chosen is a form of trend following, it unfortunately introduces some confusion about the source of the heavy tails that we observe here. As we explain later, based on various experiments we are confident that the wealth dynamics of the investors is not the source of the heavy tails.

3. Simulation results

3.1. Wealth dynamics

In figure 2 we illustrate the wealth dynamics for a simulation with 10 funds whose aggression parameters are $\beta_h = 5, 10, \dots, 50$. They all begin with the same low wealth $W_h(0) = 2$; at the outset they make good returns and their wealth grows quickly. This is particularly true for the most aggressive funds; with higher leverage they make

[†]Using a positive survival threshold for removing funds avoids the creation of 'zombie funds' that persist for long periods of time with almost no wealth.

[‡]Some of the references that document or discuss the flow of investors in and out of mutual funds include Chevalier and Ellison (1997), Remolona *et al.* (1997), Sirri and Tufano (1998, Busse (2001) and Del Guercio and Tka (2002).

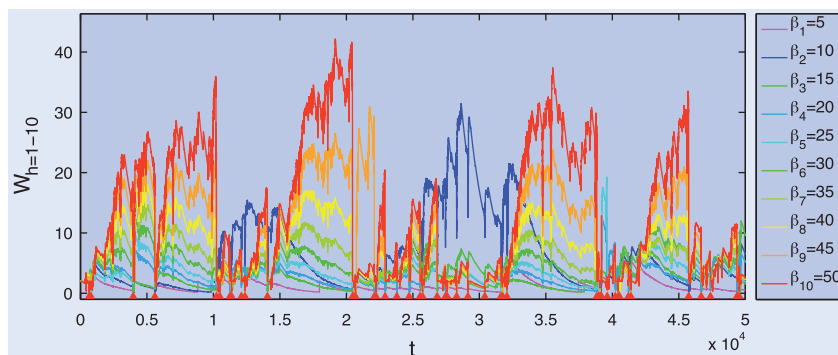


Figure 2. Wealth time series $W_h(t)$ for 10 funds with $\beta_h = 5, 10, \dots, 50$ and $\lambda_{\text{MAX}} = 20$ for all funds. Times at which (at least) one fund collapses are marked by triangles.

higher returns so long as the asset price is increasing. As their wealth grows the funds have more impact, i.e. they themselves affect prices, driving them up when they are buying and down when they are selling. This limits their profit-making opportunities and imposes a ceiling of wealth at about $W = 40$. There are a series of crashes that cause defaults, particularly for the most highly leveraged funds. Twice during the simulation, at around $t = 10\,000$ and $25\,000$, crashes wipe out all but the two least aggressive funds with $\beta_h = 5$ and 10 . While funds $\beta_3 - \beta_{10}$ wait to get reintroduced, fund β_2 manages to become dominant for extended periods of time.

3.2. Returns and correlations

The presence of the funds dramatically alters the statistical properties of price returns. This is illustrated in figure 3, where we compare the distribution of logarithmic price returns $r(t) = \log p(t) - \log p(t-1)$ for three cases: (1) noise traders only; (2) hedge funds with no leverage ($\lambda_{\text{MAX}} = 1$); and (3) substantial leverage, i.e. $\lambda_{\text{MAX}} = 10$. With only noise traders the log returns are (by construction) nearly normally distributed. When funds are added without leverage the volatility of prices drops slightly, but the log returns remain approximately normally distributed. When we increase leverage to $\lambda_{\text{MAX}} = 10$, however, the distribution becomes much more concentrated in the center and the negative returns develop fat tails. (Recall that since the funds are long-only, they are only active when the asset is undervalued, i.e. when the mispricing $m > 0$. This creates an asymmetry between positive and negative returns.) As shown in figure 3(b), for $\lambda_{\text{MAX}} = 10$ the cumulative distribution for the most negative returns roughly follows a straight line on a double logarithmic scale, suggesting that it is reasonable to approximate the tails of the distribution as a power law of the form $P(r < -R \mid m > 0) \sim R^{-\gamma}$.

The exponent γ may be regarded as a measure of the concentration of extreme risks, and a low value of γ implies fat tails. The transition from normality to fat tails

occurs more or less continuously as λ_{MAX} varies. This is in contrast to the conjecture of Plerou *et al.* (Plerou *et al.* 1999, Gabaix *et al.* 2003, 2006, Plerou and Stanley 2008) that γ has a universal value $\gamma \approx 3$. In figure 3(c) we measure γ as a function of λ_{MAX} . As λ_{MAX} increases, γ decreases, corresponding to heavier tails.[†] This trend continues until $\lambda_{\text{MAX}} \approx 10$, where γ reaches a floor at $\gamma \approx 2.5$. (The reason this floor exists depends on the particular choice of parameters here, and will be explained later.) A typical value measured for financial time series, such as American stocks (Plerou *et al.* 1999, Cont 2000), is $\gamma \approx 3$. In our model this corresponds to a maximum leverage $\lambda_{\text{MAX}} \approx 7.5$. It is perhaps a coincidence that 7.5 is the maximum leverage allowed for equity trading in the United States, but in any case this demonstrates that the numbers produced by this model are reasonable.

In figure 4 we show the log returns $r(t)$ as a function of time. The case $\lambda_{\text{MAX}} = 1$ is essentially indistinguishable from the pure noise trader case; there are no large fluctuations and little temporal structure. The case $\lambda_{\text{MAX}} = 10$, in contrast, shows large, temporally correlated fluctuations. The autocorrelation function shown in panel (c) is similar to that observed in real price series. This suggests that this model may also explain clustered volatility (Engle 1982).

4. How leverage induces nonlinear feedback and clustered volatility

4.1. When do the funds sell?

The fat tails of price movements in our model are explained by a combination of the nonlinear positive feedback caused by leveraging, which causes crashes, and the wealth dynamics of the value funds. When the funds are unleveraged, they will always buy into a falling market, i.e. when the price is dropping they are

[†]We measured γ using a Hill estimator (Hill 1975) based on the largest 10% of the returns. The value of γ when $\lambda = 1$ should be infinite, in contrast to the measured value. Large values of γ are difficult to measure correctly, whereas small values are measured much more accurately.

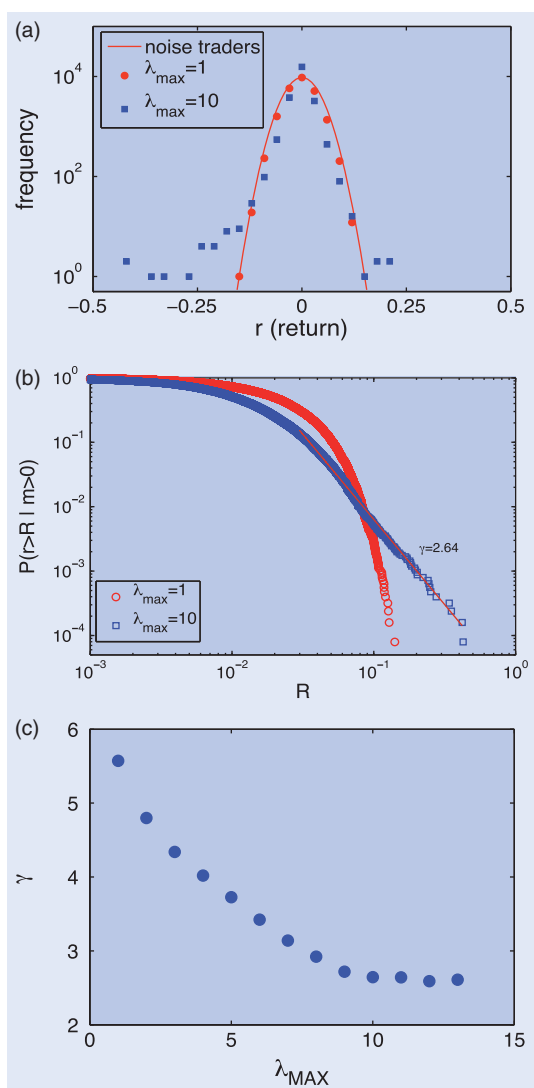


Figure 3. The distribution of log returns r . (a) The density of log returns $p(r | m > 0)$ on a semi-log scale. The results are conditioned on positive mispricing $m > 0$, i.e. only when the funds are active. The unleveraged case (red circles) closely matches the noise trader only case (red curve). When the maximum leverage is raised to 10 (blue squares) the body of the distribution becomes thinner but the tails become heavy on the negative side. This is seen from a different point of view in (b), which plots the cumulative distribution for negative returns, $P(r > R | m > 0)$, on a log-log scale. For $\lambda_{\max}=10$ we fit a power law to the data across the indicated region and show a line for comparison. In (c) we vary λ_{\max} and plot fitted values of γ , illustrating how the tails become heavier as the leverage increases. Same β values as in figure 2.

guaranteed to be buyers, thus damping price movements away from the fundamental value. When they are leveraged, however, this situation is sometimes reversed—if they are maximally leveraged they sell into a falling market, thus amplifying the deviation of price movements away from the fundamental value.

This is easily understood by differentiating the fund's demand function given in equations (2) and (3) with respect to the mispricing. Ignoring the slow-moving fund deposits and redemptions $F(t)$, write $W(t) = D(t-1)$

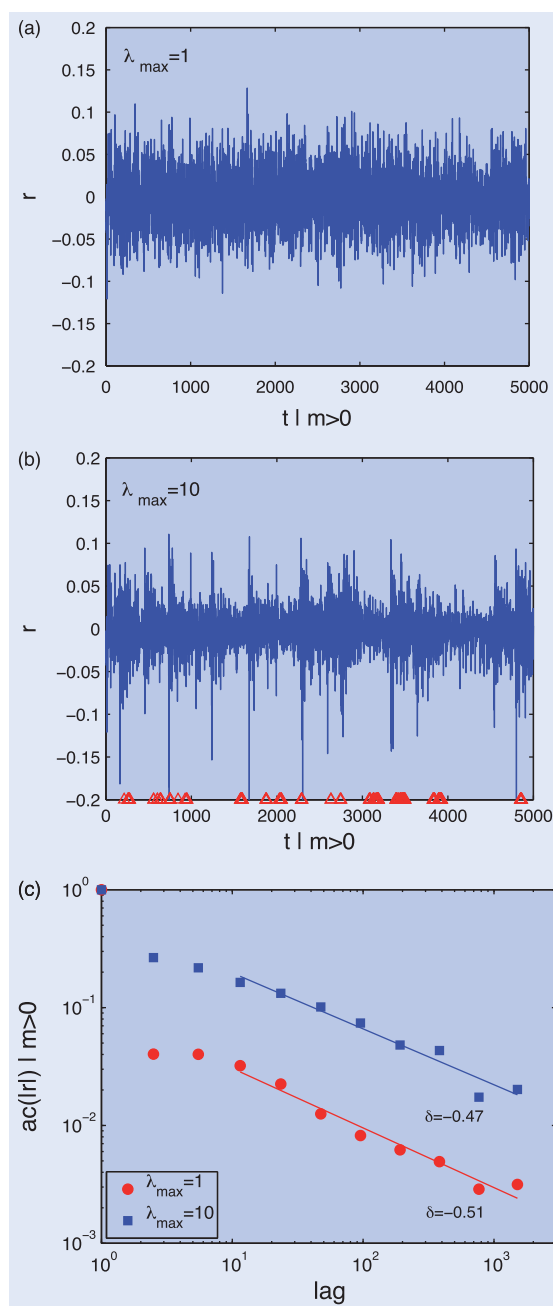


Figure 4. Log return time series (a) $\lambda_{\max}=1$ and (b) $\lambda_{\max}=10$. Triangles mark margin calls in the simulation, indicating a direct connection between large price moves and margin calls. (c) Autocorrelation function of the absolute values of log returns for (a) and (b) obtained from a single run with 100,000 time steps. This is plotted on a log-log scale in order to illustrate the power-law tails. (The autocorrelation function is computed only when the mispricing is positive.) Same β values as in figure 2.

$p(t) + C(t-1)$. Recalling that $m = V - p(t)$ and differentiating gives

$$\text{for } m < m^{\text{crit}} : \quad dD/dm = \beta \left[D(t-1) + \frac{C(t-1)V}{(V-m)^2} \right],$$

$$\text{for } m > m^{\text{crit}} : \quad dD/dm = \frac{\lambda_{\max} C(t-1)}{(V-m)^2}.$$

As long as the fund always remains unleveraged, the cash $C(t-1)$ is always positive and the derivative of the demand with the mispricing is always positive. This means the fund always buys as the price is falling. In contrast, when the fund is leveraged, then $C(t-1)$ is negative. This means that the fund is always selling as the price is falling when it is above its leverage limit, and depending on the circumstances, it may start selling even before then.

To visualize this more clearly, consider the derivative at the value of the mispricing at the last period, $m = V - p(t-1)$. At that point, ignoring redemptions, we can assume that $D(t-1)$ and $C(t-1)$ are chosen so that for $m < m^{\text{crit}}$ the fraction of the wealth held in the asset is $Dp/W = \beta m$ and the fraction held in cash is $C/W = 1 - \beta m$. Similarly, if the fund is over its leverage limit we can assume that the fraction of the wealth held in the asset is $Dp/W = \lambda^{\text{MAX}}$, and the fraction held in cash is $C/W = 1 - \lambda^{\text{MAX}}$. This implies that the rate of buying or selling under an infinitesimal change in the mispricing from the last period is

$$\begin{aligned} \text{for } m < m^{\text{crit}} : \quad \frac{dD/dm}{W} &= \frac{\beta(V - \beta m^2)}{(V - m)^2}, \\ \text{for } m > m^{\text{crit}} : \quad \frac{dD/dm}{W} &= \frac{\lambda_{\text{MAX}}(1 - \lambda_{\text{MAX}})}{(V - m)^2}. \end{aligned}$$

When the fund is leveraged, then $1 - \lambda_{\text{MAX}} < 0$, and the second term is negative, so when $m > m^{\text{crit}}$ the fund always sells as the mispricing increases. If $\beta m^2 > V$, then the fund may sell as the mispricing increases even when $m < m^{\text{crit}}$.

This is illustrated in figure 5, where we plot the derivative of the fund's demand function, dD/dm , as a function of the mispricing m . First consider the case where the maximum leverage is one ($\lambda_{\text{MAX}} = 1$). The fund buys as the mispricing increases as long as the mispricing is small enough that the leverage is under the leverage limit, i.e. for $m < m^{\text{crit}} = \lambda_{\text{MAX}}/\beta = 0.1$. When the mispricing becomes greater than this it simply holds its position. In contrast, with a maximum leverage of two the critical mispricing increases to $m^{\text{crit}} = 0.2$. The fund now buys as the mispricing increases over a wider range of mispricings, but switches over to selling when $m > m^{\text{crit}}$. When the leverage is further increased to three, this effect becomes even stronger, i.e. the fund sells even more aggressively while the price is falling.

Even when there is no cap on leverage, for a sufficiently large mispricing the fund eventually becomes a seller as the mispricing increases. This is a consequence of the fact that we chose the demand function to be proportional to wealth. When the mispricing becomes large enough the decrease in wealth overwhelms the increase in the mispricing, so the fund sells even without a margin call from the bank. This can be viewed as a kind of risk reduction strategy on the part of the fund.

By altering the margin call policy of the bank it is possible to eliminate the systemic risk effect entirely.

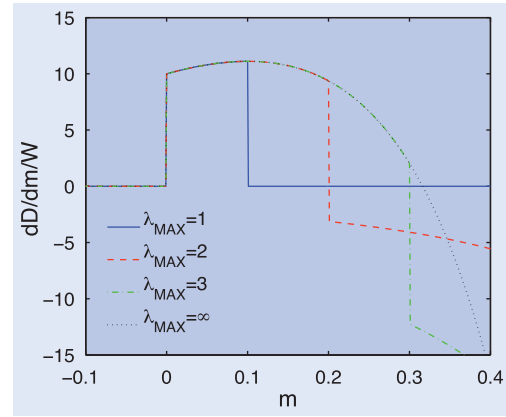


Figure 5. Change in demand dD/dm as a result of an infinitesimal increase in the mispricing from the mispricing $m = V - p$ prevailing the last period. $dD/dm > 0$ means the fund buys when the mispricing increases, and $dD/dm < 0$ means the fund sells when the mispricing increases. Here the fund has $\beta = 10$ and $W = 1$. When the maximum leverage $\lambda_{\text{MAX}} = 1$ and $m < \lambda_{\text{MAX}}/\beta = 0.1$ the fund responds to an increased mispricing by buying, and when $m > 0.1$ the fund neither buys nor sells (it holds its position). When $\lambda_{\text{MAX}} = 2$ the fund continues to buy until $\lambda_{\text{MAX}}/\beta = 0.2$, but sells for $m > 0.2$. When $\lambda_{\text{MAX}} = 3$ the fund sells heavily for $m > 0.3$. Finally, when $\lambda_{\text{MAX}} = \infty$, the fund gradually begins selling anyway.

Suppose, for example, that rather than demanding debt repayment, the bank simply takes ownership of the shares of the fund,[†] and that the demand function of the fund never causes them to sell into a falling market. In this case there is no nonlinear feedback and no systemic risk effect.

4.2. Nonlinear amplification of volatility

If the fund is leveraged, once the mispricing becomes great enough it transitions from being a buyer to being a seller. When the fund is below the leverage limit it damps volatility, for the simple reason that it buys when the price falls, opposing and therefore damping price movements. It is easy to show that with a reasonably low leverage limit λ_{MAX} , when $\lambda < \lambda_{\text{MAX}}$ the expected volatility $E[r_t^2]$ is damped by a factor approximately $1/(1 + (\beta/N)(C_h + D_h V)) < 1$ relative to the volatility for noise traders alone, where N is the total number of shares of the asset.

When funds reach their maximum leverage this reverses and funds instead amplify volatility. To remain below λ_{MAX} the fund is forced to sell when the price falls. The volatility in this case is amplified by a factor approximately $1/(1 - (\lambda_{\text{MAX}}/N) V) > 1$. This creates a positive feedback loop: dropping prices cause funds to sell, which causes a further drop in prices, which causes funds to sell. This is clearly seen in figure 4(b), where we have placed red triangles whenever at least one of the funds is at its maximum leverage. All the largest negative price changes occur when leverage is at its maximum. The amplification of volatility by leverage is illustrated in figure 6(b) where

[†]This actually happened when the Bear-Stearns hedge funds went out of business; the bank attempted to sell the underlying assets, but the liquidity was so low that they gave up and simply held them.

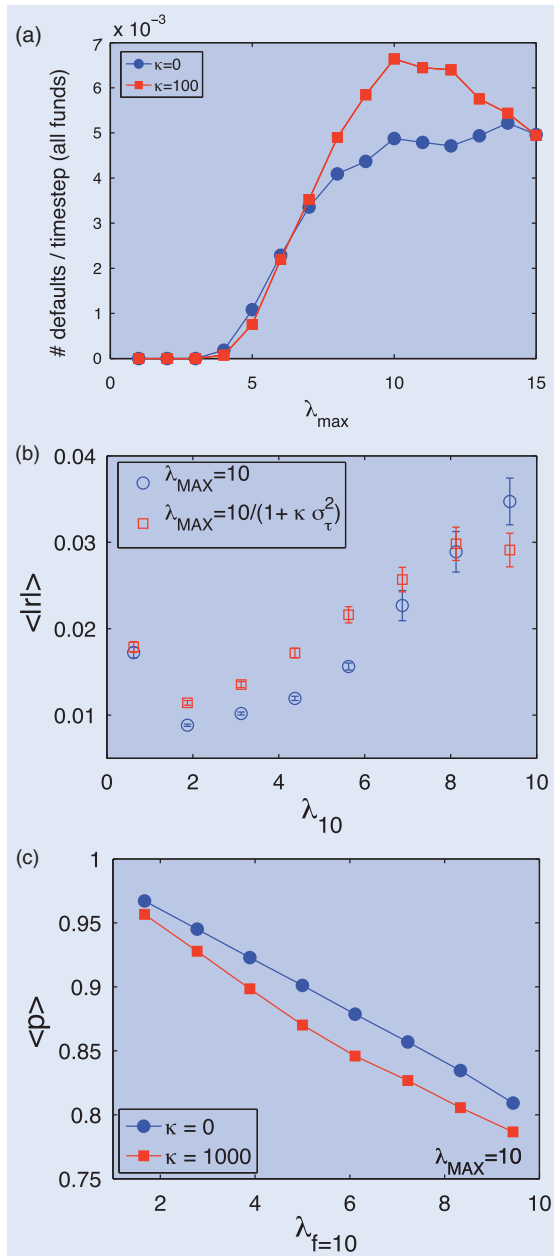


Figure 6. An illustration of how prices and volatility depend on leverage and leverage policy. We explore two different bank leverage policies. In the first policy the maximum leverage λ_{MAX} is held constant (blue circles) and in the second it is varied (red squares) so that maximum leverage decreases when historical volatility increases according to equation (9). There are 10 funds with the same β values as in figure 2. Panel (a) shows the default rates as a function of maximum leverage, and panels (b) and (c) show the average volatility and price as a function of the average leverage λ_{10} of the most aggressive fund with the maximum leverage fixed at $\lambda_{\text{MAX}}=10$. Volatility is computed as the average absolute value of the logarithmic price returns. Use of a volatility-dependent leverage can increase defaults, increase volatility, and drive prices further away from fundamentals, even though the maximum leverage is always less than or equal to its value under the fixed leverage policy.

we show that the average volatility is an increasing function of the average leverage used by the most aggressive fund.

Thus we see that, under normal circumstances where the banks impose leverage limits, the proximate cause of the extreme price movements is the margin call, which funds can meet only by selling and driving prices further down. Of course, we are not saying banks should not maintain leverage at a reasonable level; we are only saying that if they all maintain leverage at a similar level, many funds may receive margin calls at nearly the same time, inducing an instability in prices. As we have already pointed out, this can be averted by using alternative risk control policies.

4.3. How the leverage cycle drives volatility clustering

The underlying cause of volatility clustering in this model is the leverage cycle. To see how this occurs, assume that we begin at a point where the wealth of all funds is small (such as following a major market crash in which all funds default). In the early stages, all funds tend to accumulate wealth, with aggressive funds growing faster than cautious funds.[†] As shown in the previous section, the overall increase in the wealth of funds lowers volatility. In addition, the increase in the wealth of the most aggressive funds drives up the overall use of leverage.

Eventually, a substantial downward fluctuation in noise trader demand happens to occur at the same time that one or more wealthy, aggressive funds are fully leveraged. This triggers a large sell-off by the aggressive funds, which drives prices down, and generates a crash. After the crash the overall wealth of funds is substantially diminished, and as a result volatility goes back up.

This is illustrated in figure 7, which shows the time sequences of asset price returns before and after a crash. Before the crash the overall wealth in funds is large and as a consequence volatility is low; after the crash many or most of the funds are wiped out and volatility is once again high. The crash illustrated in figure 7 is just one of many, all of which follow a similar pattern: Averaging over the 500 time steps before and after a crash, and using standard deviation as the measure of volatility, the average volatility before a crash is 0.018 ± 0.003 and the average volatility after a crash is 0.032 ± 0.003 , i.e. on average it is nearly twice as much. This is the basic mechanism underlying the clustered volatility driven by the leverage cycle.

Note that, in this model, the deviation from normality of the noise traders, which is needed in order to drive prices toward their fundamental value, causes weak clustered volatility. One might suspect that leverage is merely amplifying this effect. This is not the case: as we

[†]There are two reasons why aggressive funds grow faster than passive funds. The superior returns achieved by using leverage both make the funds already under management grow faster and attract new investors. As the wealth of the funds grows sufficiently large, their market impact also grows, decreasing returns. This can drive the returns of the less aggressive funds below the benchmark return r^b and cause them to lose investment capital. This explains the pattern seen in figure 2, in which less aggressive funds grow in the period right after a crash but then eventually shrink.

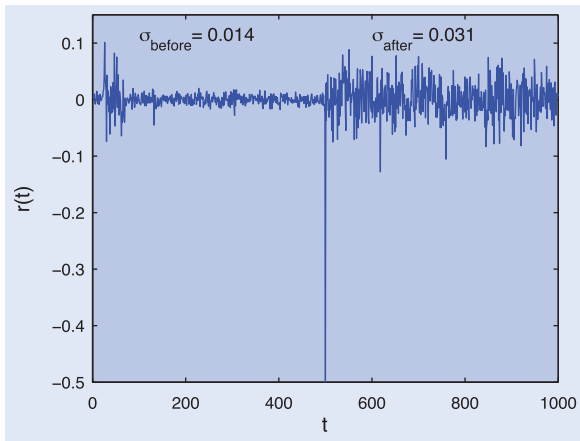


Figure 7. A typical example illustrating how the leverage cycle drives volatility clustering. Before the crash the most aggressive funds are wealthy, and under normal conditions their tendency to buy on downward price moves damps volatility. The crash occurs because the noise traders make a downward demand fluctuation while one or more of these funds is fully leveraged, causing a large sell off in an already falling market. After the crash many funds are wiped out, causing volatility to rise again (same parameters as in previous figures).

demonstrate above, the primary causes of clustered volatility are the leverage-induced crashes and the wealth dynamics. While crashes are indeed triggered by small fluctuations of the noise traders, this amplification is highly selective, and the wealth dynamics also plays an important role. Thus it is not accurate to say that the leverage merely amplifies the clustered volatility caused by the noise traders.

Another suspicion might be that the clustered volatility is driven by the trend following behavior of the investors, who pull their money in and out of the funds based on past performance. This is true in the sense that the investors' wealth dynamics affects the wealth of the funds, which in turn modulates the amplification of volatility. However, this effect would occur even if no money were moved in or out of the funds, due to the profits and losses of the funds themselves. To test this we have varied the parameter a , which sets the timescale over which investors average the returns of the funds (see equation (5)). When we vary a from 0.3 to 0.05 we see little change in the observed behavior, illustrating that the trend following of the investors is not essential.

4.4. Evolutionary pressure to increase leverage

There is also longer term evolutionary pressure driving leverage up which comes from the wealth accumulation process in this model, as illustrated in figure 8. More aggressive funds use higher leverage. During times when there are no crashes, more aggressive funds make better returns, attract more investment, and accumulate more wealth, and are thus selected over less aggressive funds. Thus, on average, during good times the average leverage used by the funds tends to increase, until there is a crash, which preferentially wipes out the most aggressive funds and resets the average leverage to a lower level.

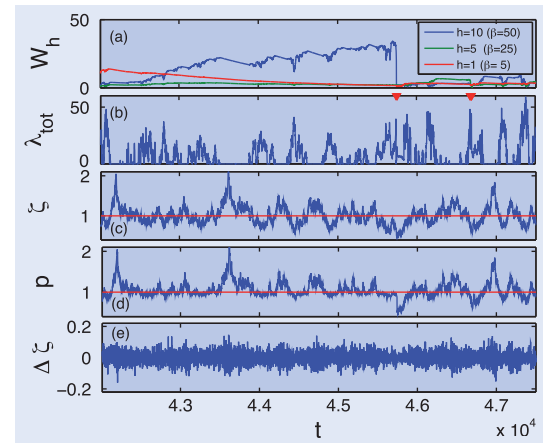


Figure 8. Anatomy of two crashes. Crashes are indicated by red triangles. From top to bottom we see (as a function of time): (a) The wealth W_h of three representative funds whose aggression levels β_h range from highest to lowest. The wealth of the most aggressive fund builds in the period leading up to the crash. (b) The average leverage λ_{tot} , calculated by summing the demand and wealth in equation (4) over all funds. (c) The noise trader demand ξ . (d) The price p . (e) The fluctuations in the noise trader demand, $\Delta\xi$.

The wealth dynamics are illustrated in the top panel of figure 8. The total leverage is shown in the panel below it. The leverage comes in bursts as mispricings develop, but the size of these bursts tends to become larger as the relative wealth of the more aggressive funds increases.

The next panel illustrates the noise trader demand, which is a weakly mean-reverting random process, and the panel below it illustrates the price. During positive excursions of the noise trader demand the asset is overpriced, the funds stay out of the market, and the price is equal to the noise trader demand (measured in dollars). When the noise trader demand becomes negative the asset is underpriced, and under normal circumstances the fund demand prevents the price from moving very far below V . This is true until the first crash occurs, where the floor of demand provided by the funds collapses due to margin calls, and the funds sell instead of buy as the price falls, temporarily driving the price down even faster than where it would be with noise traders alone.

Crashes are typically not caused by unusually large fluctuations in noise trader demand, as shown in the next two panels. When the crashes occur the value of change in the noise trader demand, $\Delta\xi(t) = \xi(t) - \xi(t-1)$ is nothing out of the ordinary, and indeed for the examples given here it is not even one of the larger values in the series. Nonetheless, the associated change in price, $\Delta p(t) = p(t) - p(t-1)$, is highly negative. Also, one can see that, while there is a very small amount of clustered volatility due to the mean reversion in the demand fluctuations of the noise traders, this is enormously amplified in the price fluctuations.

To illustrate the evolutionary pressure towards higher leverage explicitly, we performed simulations holding all but one fund at a constant leverage λ_{MAX} and sweeping the maximum leverage of the last fund, as illustrated

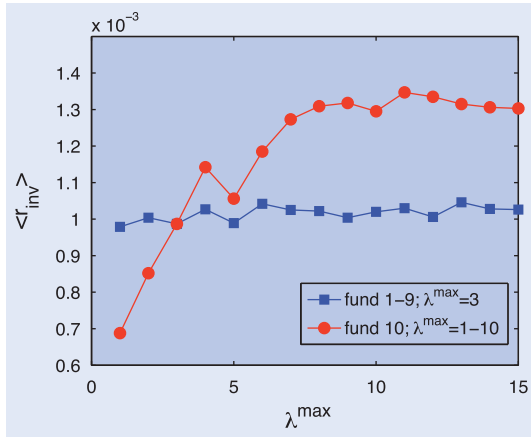


Figure 9. Demonstration of evolutionary pressure to increase leverage. The maximum leverage is held constant at $\lambda^{\text{MAX}}=3$ for nine funds while λ^{MAX} varies from 1 to 10 for the remaining fund. The vertical axis shows the returns to investors, with $\beta=20$ for all funds and $T_{\text{reintro}}=10$. The fund with a higher leverage limit obtains better returns, and so attracts more capital. Averages are taken over 50 independent simulations of 100 000 time steps each.

in figure 9. For example, if the nine funds have $\lambda_{\text{MAX}}=3$, a fund with $\lambda_{\text{MAX}}>3$ generates higher returns, as seen in the figure, and thus accumulates wealth and becomes the dominant fund. In a real-world situation this would of course put pressure on other fund managers to increase their leverage. There is thus evolutionary pressure driving leverage up.

4.5. Better individual risk control can backfire for the system as a whole

In an attempt to achieve better risk control, banks often vary the maximum leverage based on the recent historical volatility of the market, lowering maximum leverage when volatility has been high and raising it when it has been low. This is prudent practice when lending to a single fund. But this can be counterproductive when all the funds might be deleveraging at the same time.

In figure 6 we investigate an alternative leverage policy in which lenders tighten leverage restrictions whenever there is increased historical volatility. Maximum leverage is adjusted according to the relation

$$\lambda_{\text{adjust}}(t) = \max\left[1, \frac{\lambda_{\text{MAX}}}{1 + \kappa\sigma_\tau^2}\right], \quad (9)$$

where $\kappa=100$ is the bank's responsiveness to volatility, and σ_τ^2 is the asset price variance computed over an interval of $\tau=10$ time steps. For low values of maximum leverage the number of defaults is about the same, but for higher maximum leverage, in the range $7 < \lambda_{\text{MAX}} < 15$, the number of defaults is greater with the variable leverage policy. The reason for this is simple: lowering the maximum leverage across all funds can cause massive selling at just the wrong time, creating more defaults rather than less. Once again, an attempt to improve risk

control that is sensible if one bank does it for one fund can backfire and create more risk if every bank does it with every fund.

This kind of policy also has another important unintended consequence. During times of low volatility leverage goes up. This in turn drives volatility up, which forces leverage back down. Thus, in such a situation there are stochastic oscillations between leverage and volatility which, on average, drive volatility up and drive prices further away from fundamentals. This is illustrated in figures 6(b) and (c), where we plot the average volatility and the average price as a function of the leverage of the most leveraged fund. Note that the amplification of volatility occurs even though, under the variable maximum leverage risk control protocol, the maximum volatility is always less than or equal to its value under the fixed protocol.

5. Conclusion

The use of leverage in the economy is not just an esoteric matter relating to funds: it is unavoidable. It is the mechanism through which most people are able to own homes and corporations do business. Credit (and thus leverage) is built into the fabric of society. The current financial crisis perfectly illustrates the dangers of too much leverage followed by too little leverage. Like Goldilocks, we are seeking a level that is 'just right'. This raises the question of what that level is (Hackbarth *et al.* 2006, Peters 2011).

We are not the first to recognize the downward spiral of margin calls (Geanakoplos 1997, 2003, 2010, Shleifer and Vishny 1997, Gromb and Vayanos 2002, Fostel and Geanakoplos 2008, Brunnermeier and Pedersen 2009). After the Great Depression the Federal Reserve was empowered to regulate margins and leverage. However, the model we have developed here provides a quantifiable and testable framework to explore the consequences of leverage and its regulation. Recent empirical work has found a correlation between leverage and volatility (Adrian and Shin 2008), but our work suggests a more subtle relationship. We make the falsifiable prediction that high leverage limits, such as we had in reality until very recently, cause increased clustering of volatility and fat tails, and that these effects should go up and down as leverage goes up and down. This can work in parallel with other effects that generate heavy tails, such as myopic learning.

During good times leverage tends to creep up, creating a dangerous situation leading to a sudden crash in prices. We have shown that when individual lenders seek to control risk through adjusting leverage, they may collectively amplify risk. Our model can be used to search for a better collective solution, perhaps coordinated through government regulation.

Standard methods currently used by bankers and regulators to monitor risk by measuring recent volatility are thus doomed to failure, for it is precisely when the volatility is low that the risk of collapse is greatest.

Low volatility and high leverage go hand and hand in our model. This connection was seen by Geanakoplos (2003, 2010), but there the causality ran from low volatility to high leverage: lenders are willing to give higher leverage when they think asset prices will not go down much. In the present model the causality also runs in the reverse direction. The reason for that is our assumption that the funds have a more stable view of the fundamental value than the noise traders. When the funds' total purchasing power increases they will stabilize the price against small downward fluctuations in noise traders' demand for the asset. Higher leverage increases the funds' purchasing power and thus tends to reduce (downward) price volatility. When the funds are successful their purchasing power grows because of their earnings and because investors shift money into funds. Furthermore, the distribution of wealth across funds shifts towards the more highly leveraged funds, because they have relatively higher earnings and because their relative success attracts more investors. This increases their average leverage. Thus volatility is often very low and leverage very high after a sustained run of price increases.

At a broader level, the present work shows how attempts to regulate risk at a local level can actually generate risks at a systemic level. The key element that creates the risk is the nonlinear feedback on prices that is created due to repaying loans at a bad time. This mechanism is actually quite general, and also comes into play with other risk control mechanisms, such as stop-loss orders and many types of derivatives, whenever they generate buying or selling in the same direction as price movement. We suspect that this is a quite general phenomenon that occurs in many types of systems whenever optimization for risk reduction is done locally without fully taking collective phenomena into account.

Acknowledgements

We thank Barclays Bank, Bill Miller, National Science Foundation grant 0624351 and Austrian Science Fund grant P17621 for support. We would also like to thank Duncan Watts for stimulating the initiation of this project, and Alan Kirman and Chris Wood for useful discussions.

References

Adrian, T. and Shin, H.S., Liquidity and leverage. Technical Report 328, Federal Reserve Bank of New York, 2008.

Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R. and Taylor, P., Asset pricing under endogenous expectations in an artificial stock market. In *The Economy as an Evolving Complex System II*, edited by W.B. Arthur, S.N. Durlauf, and D.H. Lane, pp. 15–44, 1997 (Addison-Wesley: Redwood City).

Bachelier, L., In *The Random Character of Stock Market Prices*, edited by P.H. Cootner, 1964 (MIT Press: Cambridge).

Black, F. and Scholes, M., The pricing of options and corporate liabilities. *J. Polit. Econ.*, 1973, **81**(3), 637–659.

Brock, W.A. and Hommes, C.H., Models of complexity in economics and finance. In *System Dynamics in Economic and Financial Models*, edited by C. Hey, J. Schumacher, B. Hanzon, and C. Praagman, pp. 3–41, 1997 (Wiley: New York).

Brock, W.A. and Hommes, C.H., Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J. Econ. Dynam. Control*, 1998, **22**, 1235–1274.

Brock, W.A., Hommes, C.H. and Wagener, F., More hedging instruments may destabilize markets. *J. Econ. Dynam. Control*, 2009, **33**, 1912–1928.

Brunnermeier, M. and Pedersen, L., Market liquidity and funding liquidity. *Rev. Financ. Stud.*, 2009, **22**(6), 2201–2238.

Busse, J.A., Look at mutual fund tournaments. *J. Financ. Quantit. Anal.*, 2001, **36**, 53–73.

Caccioli, F., Marsili, M. and Vivo, P., Eroding market stability by proliferation of financial instruments. *Eur. Phys. J. B*, 2009, **71**, 467–479.

Caldarelli, G., Marsili, M. and Zhang, Y.-C., A prototype model of stock exchange. *Europhys. Lett.*, 1997, **40**, 479–484.

Chevalier, J. and Ellison, G., Risk taking by mutual funds as a response to incentives. *J. Polit. Econ.*, 1997, **105**, 1167–1200.

Cont, R., Empirical properties of asset returns: Stylized facts and statistical issues. *Quantit. Finance*, 2000, **1**, 223–236.

Del Guercio, D. and Tka, P.A., The determinants of the flow of funds of managed portfolios: Mutual funds vs. pension funds. *J. Financ. Quantit. Anal.*, 2002, **37**, 523–557.

Engle, R.F., Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica*, 1982, **50**, 987–1008.

Fostel, A. and Geanakoplos, J., Leverage cycle and the anxious economy. *Am. Econ. Rev.*, 2008, **98**(4), 1211–1244.

Friedman, D. and Abraham, R., Bubbles and crashes: Escape dynamics in financial markets. *J. Econ. Dynam. Control*, 2009, **33**(4), 922–937.

Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E., A theory of power-law distributions in financial market fluctuations. *Nature*, 2003, **423**, 267–270.

Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E., Institutional investors and stock market volatility. *Q. J. Econ.*, 2006, **121**, 461–504.

Geanakoplos, J., Promises, promises. In *The Economy as an Evolving Complex System, II*, edited by W.B. Arthur, S.N. Durlauf, and D. Lane, pp. 285–320, 1997 (Addison-Wesley: Redwood City).

Geanakoplos, J., Liquidity, default, and crashes. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, edited by M. Dewatripont, L.P. Hansen, and S.J. Turnovsky, pp. 170–205, 2003 (Cambridge University Press: Cambridge).

Geanakoplos, J., The leverage cycle. In *NBER Macroeconomics Annual 2009*, vol. 24, edited by D. Acemoglu, K. Rogo and M. Woodford, pp. 1–65, 2010 (University of Chicago Press: Chicago).

Gennotte, G. and Leland, H., Market liquidity, hedging and crashes. *Am. Econ. Rev.*, 1990, **80**(5), 999–1021.

Giardina, I. and Bouchaud, J.-P., Bubbles, crashes and intermittency in agent based market models. *Eur. Phys. J. B*, 2003, **31**, 421–437.

Gromb, D. and Vayanos, D., Equilibrium and welfare in markets with financially constrained arbitrageurs. *J. Financ. Econ.*, 2002, **66**, 36–407.

Hackbarth, D., Miao, J. and Morellec, E., Capital structure, credit risk and macroeconomic conditions. *J. Financ. Econ.*, 2006, **82**(3), 519–550.

Hill, B.M., A simple general approach to inference about the tail of distribution. *Ann. Statist.*, 1975, **3**, 1163–1174.

Kim, G. and Markowitz, H.M., Investment rules, margin and market volatility. *J. Portfol. Mgmt.*, 1989, **16**(1), 45–52.

Lux, T., The socio-economic dynamics of speculative markets: Interacting agents, chaos, and the fat tails of return distributions. *J. Econ. Behav. Organiz.*, 1999, **33**, 143–165.

- Lux, T. and Marchesi, M., Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 1999, **397**, 498–500.
- Mandelbrot, B.B., The variation of certain speculative prices. *J. Business*, 1963, **36**, 394–419.
- Palmer, R., Arthur, W.B., Holland, J.H., LeBaron, B. and Taylor, P., Artificial economic life: A simple model of a stock market. *Physica D*, 1994, **75**, 264–274.
- Peters, O., Optimal leverage from non-ergodicity. *Quantit. Finance*, 2011, **11**, 1593–1602.
- Plerou, V., Gopikrishnan, P., Nunes Amaral, L.A., Meyer, M. and Stanley, H.E., Scaling of the distribution of price fluctuations of individual companies. *Phys. Rev. E*, 1999, **60**, 6519–6529.
- Plerou, V. and Stanley, H.E., Stock return distributions: Tests of scaling and universality from three distinct stock markets. *Phys. Rev. E*, 2008, **77**, 037101.
- Remolona, E.M., Kleiman, P. and Gruenstein, D., Market returns and mutual fund flows. *FRBNY Econ. Policy Rev.*, 1997, 33–52.
- Shleifer, A. and Vishny, R.W., The limits of arbitrage. *J. Finance*, 1997, **52**, 35–55.
- Sirri, E.R. and Tufano, P., Costly search and mutual fund flows. *J. Finance*, 1998, **53**, 1589–1622.

Appendix A: Properties of the noise trader process

In this appendix we show that, for the parameters we use here, the noise trader process by itself has slightly non-normal returns and weak clustered volatility. Assume no funds, so that the dynamics are determined solely by the noise traders. For convenience make the change of

notation $y_t = \log \xi(t)$, and for convenience let $V = N = 1$ (which only shifts the mean). The price is given by $D_{nt} = \xi(t)/p(t) = N = 1$ and the log price is $\log p(t) = \log \xi(t) = y(t)$. The price dynamics become

$$y(t+1) = \rho y(t) + \sigma \chi(t+1). \quad (\text{A1})$$

The noise $\chi(t)$ is normally distributed with zero mean. The return $r(t)$ is

$$r(t+1) = y(t+1) - y(t) = (\rho - 1)y(t) + \sigma \chi(t+1). \quad (\text{A2})$$

Squaring this and averaging gives

$$E[r(t)^2 | y(t)] = (\rho - 1)^2 y(t)^2 + \sigma^2.$$

Thus the volatility varies conditionally on $y(t)$, implying that there is some clustered volatility even when only noise traders are present. An estimate of the typical size of $y(t)^2$ can be made by squaring equation (A1) and making use of stationarity, which yields $E[y(t)^2] = \sigma^2/(1 - \rho)^2$. Substituting in equation (A2) gives a typical relative variation in volatility of $(1 - \rho)^2/(1 - \rho^2)$, which for $\rho = 0.99$ is about 0.005. Thus the variation in volatility for the pure noise trader process is small for the parameters we use here, and vanishes in the limit $\rho \rightarrow 1$. The fact that the variance fluctuates means that the time series $r(t)$ is not identically distributed, and the marginal distribution $P(r)$ is a Gaussian mixture, which is slightly more heavy-tailed than a normal distribution.