Adaptive Expectations and Stock Market Crashes

David M. Frankel

June 21, 2001

[1] Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 69978, Israel (email: dfrankel@post.tau.ac.il). I thank Simon Benninga and Ady Pauzner for their extensive comments on this paper. I also thank Gady Barlevy, Amir Kirsh, Assaf Razin, and seminar participants at Tel Aviv University for helpful comments.

Adaptive Expectations and Stock Market Crashes

Abstract

A theory is developed that explains how the stock market can crash in the absence of news about fundamentals, and why crashes are more common than frenzies. The theory is based on the interaction of rational traders and naive traders, who believe in a simple model in which stock prices follow a random walk with serially correlated volatility. From time to time, the rational traders dump their shares. They do so because the apparent rise in volatility leads the naive traders also to sell, driving prices down further. In contrast, frenzies cannot occur: if the rational traders were all to buy, driving prices up, the sudden increase in volatility would lead the naive traders to sell, pushing prices back down; hence, the actual incentive for a rational trader would be to *sell*.

Keywords: Stock market crashes, adaptive expectations.

1 Introduction

On October 19, 1987, the S&P 500 index fell by 20.5%. Evidence from option prices indicates that investors expect more crashes to occur.[1] What causes such jumps in prices? The explanation should reflect the fact that many traders were responding to prior price declines themselves, rather than to news about the economy or firm profitability.[2] It should also explain why crashes happen more often than comparable-sized frenzies, in which prices rise sharply.[3]

We present a new explanation for crashes that includes these elements. A crash results from the interaction of two groups of risk-averse investors. One group consists of rational traders who know the correct model and which equilibrium is being played. The second consists of naive traders who believe in a simple statistical model in which stock prices follow a random walk with serially correlated volatility. ("Volatility" refers to the variance of the stock return.) Their volatility predictions are formed adaptively, as a weighted sum of past realized and predicted volatilities.

The naive investors in our model are not fully rational. However, their beliefs are closer to the efficient-markets paradigm than the views of naive traders in some other models. For example, the "feedback traders" of De Long, Shleifer, Summers,

[1] Ait-Sahalia, Yared, and Wang [1] find that the high prices of out-of-the-money put options on the S&P 500 index are inconsistent with a model in which stock prices change continuously according to a Markov process. However, they are consistent with occasional downward jumps in stock prices. Thus, a complete explanation for observed option prices must include a theory of asymmetric price jumps.

[2] According to Shiller's postcrash survey [49, p. 386], declining prices on October 14-16 and the morning of October 19 were the news items that most influenced investors' views of the stock market on October 19, 1987. Cutler, Poterba, and Summers [16] show that the large one-day movements in the S&P 500 index in the postwar period (and in particular the crash of 1987) were not accompanied by important news events. After the 6.1% minicrash of September 3, 1946, the most-often cited reason for selling (43% of sellers) was "declining prices on September 3" (see the Securities and Exchange Commission [51], cited by Shiller [49, p. 373-4]).

[3] Since 1945, the S&P Composite index fell by 5% or more on 13 separate days; the average of these declines was 7.5%. The index rose by over 5% on only 5 days; the average was only 5.9%.

and Waldman [17] naively chase trends. They believe that price increases tend to be followed by more increases. In contrast, naive traders in our model believe that stock prices follow a random walk: that stock returns are serially uncorrelated. This view dates from Bachelier [2] and has been widely promulgated. In his best-selling textbook, Sharpe [48, p. 315] writes:

> Stock returns exhibit almost no serial correlation: the particular value of return in the last period provides little if any help in predicting the likelihood of various possible returns in the next period.

Malkiel makes the same point forcefully in his well-known book *A Random Walk Down Wall Street* [38].

On the other hand, our naive traders believe that volatility *is* serially correlated. This belief has been the consensus in academic circles since Mandelbrot [39, pp. 418-9] and Fama [24, pp. 85-7] and is taught in popular textbooks such as Brealey and Myers [11, p. 510] and Sharpe [48]. Sharpe also discusses how one can predict future volatility using historic volatility, and why it is worthwhile to put more weight on recent returns [48, p. 441]. The naive traders in our model do exactly this.

From day to day, the serial correlation of volatility is much stronger than the serial correlation of returns. For the S&P Composite Index from 1929 to 1999, the serial correlation of daily volatility was 0.23; in comparison, the serial correlation of daily returns was only 0.055.[4] In addition, the serial correlation of daily returns appears to be a statistical artifact caused by nontrading of some stocks in the index.[5] (This may

[4]For the Dow Jones Industrial Average over the same period, the analogous figures were 0.22 and 0.052, respectively.

[5]Evidence for this comes from MacKinlay and Ramaswamy [36], who compute daily autocorrelations in log returns for the S&P 500 index and for futures contracts on this index during the 1983-1987 period. They find an average autocorrelation of 6.04% for daily index returns versus -0.24% for daily futures returns [36, p. 148, panel E of table 2]. Since futures returns are essentially uncorrelated, the serial correlation of index returns would seem to be due to nontrading of some of the smaller stocks in the index and thus should not affect traders' behavior.

not be true for returns over longer horizons; see Barberis, Shleifer, and Vishny [4] for a survey.)

Adaptive expectations have also become the dominant approach among financial econometricians for modelling the dynamics of asset price volatility. In 1982, Engle [20] first proposed the ARCH (Autoregressive Conditional Heteroskedasticity) model, in which next period's volatility is a weighted sum of past realized volatilities. In 1986, Bollerslev [10] generalized this to GARCH (Generalized ARCH) by letting next period's volatility depend also on past *predicted* volatilities. In the past two decades, over 200 journal articles have used ARCH or GARCH to model the changing volatility of asset returns.[6] Our naive traders also use a GARCH model to predict future volatility.

A crash occurs in the following way. Each day the rational traders observe a common signal that acts as a coordinating device. For certain values of this signal, they dump their shares. The naive traders are risk averse and have short horizons. Thus, the apparent increase in volatility prompts them also to sell, which pushes prices down further - justifying the rational agents' prior decision to sell. Importantly, this mechanism does not give rise to frenzies. If rational traders were all to buy shares, the sharp price increase would also raise the naive traders' future volatility estimate, prompting them to sell and pushing prices down in the following period. Anticipating this, each rational trader would have an incentive to sell when the others were buying.

This model contains the main stylized facts surrounding crashes. Prices can jump discontinuously. Downward jumps are more common than upward jumps.[7] Some traders - the naive ones - sell in response to prior price changes. Finally, crashes are unexpected: until the crash signal is observed, noone knows a crash is about to happen. This mirrors findings of Bates [7] that option prices indicated no crash fears in the 2 months leading up to the 1987 crash.

The rational traders dump their shares in response to a common signal that serves

[6] Author's tabulation from Econlit.

[7] In this stylized model they are infinitely more common, though this would probably not hold if naive traders extrapolated the sign of stock returns (à la De Long, Shleifer, Summers, and Waldman [17]) as well as their volatility. We have not explored this avenue.

as a coordinating device. While the naive traders may also see this signal, they are not aware of its coordinating role. In principle, any information will do: statements of public figures, trade deficit numbers, changes in exchange rates, inflation reports, and so on. The information may also be completely unrelated to fundamentals. One example is technical analysis. In the 1987 crash, a third of investors reported being influenced by the price dropping through a long-term trend line (Shiller [49, p. 394]). Similarly, in the 6.1% minicrash of September 3, 1946, 13% of sellers reported being motivated by Dow theory, a system of technical analysis.[8]

There is also anecdotal evidence that technical analysis led a few large money managers to anticipate the 1987 crash. In an interview with Jack Schwager, Paul Tudor Jones (a well-known trader and then manager of Tudor Futures Fund) states:

> When we came in on Monday, October 19, we knew that the market was going to crash on that day. [Q: What made you so sure?] Because the previous Friday was a record volume day on the downside. The exact same thing happened in 1929, two days before the crash. [46, p. 133]

Jones's analysis influenced others as well. Stanley Druckenmiller (at the time manager of George Soros's Quantum Fund) describes his experience:

> That Friday [October 16, 1987] after the close, I happened to speak to Soros. He said that he had a study done by Paul Tudor Jones that he wanted to show me. [...] The study demonstrated the historical tendency for the stock market to accelerate on the downside whenever an upward-sloping parabolic curve had been broken - as had recently occurred. The analysis also illustrated the extremely close correlation in price action between the 1987 stock market and the 1929 stock market, with the implicit conclusion that we were now at the brink of a collapse. I was sick to my stomach when I went home that evening. I realized that I had blown it and that the market was about to crash. [47, pp. 198-9]

[8]See Securities and Exchange Commission [51], cited by Shiller [49, p. 373-4]

Indeed, technical indicators may be useful primarily because they allow coordination among a small subset of traders who are "versed in the art." This interpretation of technical analysis is also taken by Froot, Scharfstein, and Stein [27].

The crash occurs because selling by rational traders leads naive traders to sell in the following period. Without this postcrash selling, there would be no crash: the rational traders would expect the price to rise right after the crash, which would eliminate their incentive to sell. The naive traders sell because their reliance on GARCH leads them to overestimate postcrash volatility. This prediction of the model is borne out in the data from the 1987 crash. Figure 1 compares realized volatility to the volatility predicted based on the GARCH model.[9] GARCH vastly overestimates postcrash volatility.[10]

Figure 1: 1987 Crash, Realized Vs. Predicted (GARCH) Volatility

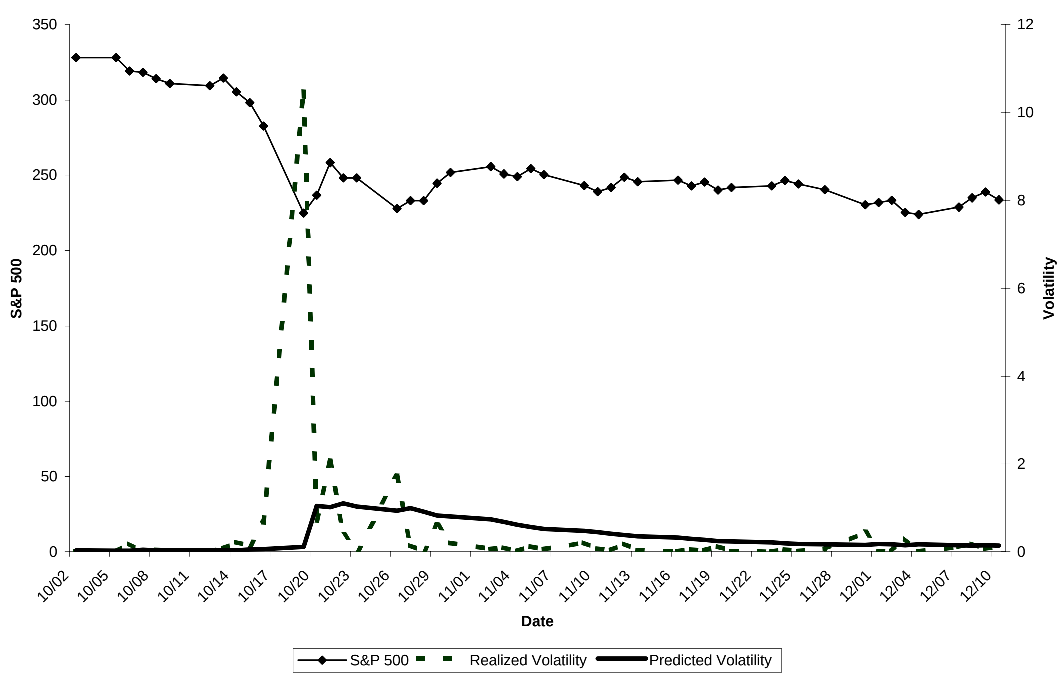

---

[9] All volatilities are annualized by multiplying by 250, the number of trading days in a typical year. Predicted volatility is based on GARCH estimation on the S&P Compositive Index for the period 1/3/28 through 10/16/87 (the last trading day before the crash on 10/19/87). The S&P Index is the solid line marked with diamonds. The plain solid line shows predicted volatility while the dashed line gives actual volatility.

[10] This phenonenon, known as the "threshold effect," was first documented in 1993 by Engle and Ng [23].

Since naive traders mispredict volatility in the aftermath of the crash, their trading strategies perform poorly in this period. Does this mean they will eventually be driven from the market? One reason for doubt is that crashes are rare. Of the survivors of the 1929 crash, few if any were alive to trade actively in the 1987 crash. Since new naive traders are continually entering the market, there is no reason to expect them to disappear in the long run.

2 Relevant Literature

Stock market crashes exemplify four properties of stock returns that have been studied in the literature:

1. There is *excess volatility*: changes in stock prices appear to be too large relative to subsequent changes in dividends (e.g., Shiller [49]).

2. Volatility is *stochastic*: some periods are relatively quiet while in others, prices swing wildly.

3. Returns are *skewed*: price declines tend to be larger than price increases.

4. Prices can *jump discontinuously* without jumps in fundamentals.[11]

Hong and Stein [31] study a model that explains all four properties. They assume risk-neutral, uninformed arbitrageurs who interact with investors who observe private signals of the value of the market portfolio. The investors are overconfident: each pays attention only to her own signal. They are also short-sale constrained: pessimistic investors are inactive (own no stock), so their signals are not revealed to the arbitrageurs. When optimistic investors get bad news, information is revealed about the signals of the inactive investors according to when they offer buying support. As prices fall, the lack of buying interest from inactive investors is interpreted by the arbitrageurs as a bad signal,

[11] This is a special case of property 1. We separate it because many models yield excess volatility but do not generate jumps.

which drives prices further down. In this sense, price declines can feed on themselves and may be large even if the initial bad news is minor.

Barlevy and Veronesi [5] and Gennotte and Leland [28] study models in which uninformed traders infer fundamentals from prices. A low price implies bad fundamentals, which lowers uninformed trader demand, implying that the price should be even lower, and so on. These models are static. The demand curve for stocks has upwards sloping segments, so there can be multiple equilibrium prices. Small changes in parameters can cause discontinuous changes in the equilibrium set: by changing a parameter slightly, an equilibrium price can simply disappear. A "crash" occurs if a high-price equilibrium disappears and only lower priced equilibria remain. The models don't explain skewed returns since a low-price equilibrium might disappear, causing a "frenzy." Another issue is whether there would still be jumps in suitable dynamic versions of these models.

There are several other models that, while not specifically aimed at explaining crashes, also generate some of the above properties. Grossman and Zhou [29] assume two types of risk averse investors who each maximize expected consumption utility. One type, the "portfolio insurers," have an additional constraint that their wealth must not fall before a certain level. As fundamentals worsen, the portfolio insurers sell stock at an accelerating rate, leading to an increase in volatility. This model does not explain news-free jumps, since a large change in prices cannot result from an arbitrarily small shock.

As in our model, De Long, Shleifer, Summers, and Waldman [17] assume two groups of investors. One group is fully rational. The other is composed of "feedback traders" who chase trends, buying after prices rise and selling after they fall. This leads rational traders to react to news in an exaggerated way: if the news is good, they drive prices exaggeratedly high since they know that the feedback traders will pay even more the next period. This model does not generate skewed returns.

Barberis, Huang, and Santos [3] link excess volatility to psychological evidence that prior gains make investors less risk averse since they are "playing with house money." An increase in stock prices leads investors to invest more in stocks, pushing prices yet higher; a price decline makes investors more risk averse, leading them to sell their shares.

This model does not explain skewed returns or news-free price jumps.

Our theory is also related to the "volatility feedback" effect first studied by French, Schwert, and Stambaugh [26], Malkiel [37] and Pindyck [43]. They point out that greater stock market volatility can lead to a higher risk premium and thus to lower stock prices.[12] Campbell and Hentschel [13] show that this effect can also give rise to negative skew: price declines are larger, on average, than price advances. They assume a fully rational agent who sees dividends that follow a GARCH process: a random walk with serially correlated volatility.[13] Large dividend shocks lead to lower prices since they indicate an increase in volatility and the agent is risk averse. This "volatility feedback effect" dampens the price effects of positive dividend news and exaggerates the price effects of negative dividend news.

While the model of Campbell and Hentschel generates negative skew, it does not give news-free jumps: prices are a continuous function of the dividend shock. One point of our model is that if some traders believe that *prices* follow a GARCH process, and if there are rational traders with a coordinating device, then there can be news-free crashes.

Barsky and De Long [6] suggest that excess volatility is driven by investor uncertainty about the dividend growth rate. They assume that investors form their expectations of future growth rates adaptively, as a geometrically weighted average of past growth rates. In periods of sustained dividend growth, investors become more optimistic about future growth, which can cause large changes in the price-earnings ratio. This model does not explain skewed returns or news-free jumps.

---

[12] In response to Pindyck [43], Poterba and Summers [44] produced evidence that volatility changes are not persistent enough to effect stock prices much. They model volatility as an AR(1) process. However, Chou [14] subsequently found much stronger persistence using GARCH, a more flexible specification.

[13] Campbell and Hentschel [13] assume that the log change in dividends in period $t+1$, η_{t+1}, is normal with mean zero and conditional variance σ_t^2, where σ_t^2 follows the Quadratic GARCH process $\sigma_t^2 = \omega + \alpha(\eta_t - b)^2 + \beta\sigma_{t-1}^2$. This reduces to GARCH if $b = 0$. If $b > 0$, negative dividend shocks lead to higher volatility than positive shocks of the same size, which can capture the leverage effects first posited by Black [8].

3 Adaptive Volatility Predictions

Naive traders' expectations of volatility are formed in the following way. Let p_t be the stock price in period t and let $r_t = \frac{p_t - p_{t-1}}{p_{t-1}}$ be the period-t capital gains return. The naive traders' estimate V_t of the next period's volatility r_{t+1}^2 is a weighted sum of the past realized and expected volatilities:

$$V_t = \alpha r_{t-1}^2 + \beta V_{t-1} \tag{1}$$

In his popular textbook, Sharpe [48, p. 441] recommends predicting volatility as a weighted sum of past volatilities, with more weight on more recent squared returns. Equation (1) is an example of this: by substituting repeatedly for V on the right hand side, one can express V_t as a geometric weighted sum of past squared returns:

$$V_t = \alpha \sum_{i=1}^{\infty} \beta^{i-1} r_{t-i}^2 \tag{2}$$

One implication of equation (2) is that the naive traders react to price changes with a lag. V_t, the naive traders' prediction of period $t+1$ volatility, depends on the realized volatilities in periods $t-1$, $t-2$, and so on, but not on the realized volatility in period t. This is consistent with evidence from Shiller [49, p. 388] that the average investor heard about the crash of October 19, 1987 at 1:56 Eastern Time, 2 hours before the market closed and, based on intraday price charts in Kleidon and Whaley [34], well after the major portion of the decline. The effect of this assumption is to given the rational traders an incentive to sell when they see the crash signal. They sell so as to get a jump on the naive traders, who react only in the following period.

4 The Model

Time is discrete and the horizon is infinite: $t = 0, 1, 2, ...$ In each period there are young and old agents, each present in measure one. Agents live two periods. Each generation consists of a measure μ of rational traders and $1 - \mu$ of naive traders. Agents work only when young, earning a fixed amount L of labor income, and consume only when old. There are two assets: one ("stocks"), in fixed unit supply, pays an i.i.d. dividend of $\delta_t \sim N(\overline{\delta}, \sigma^2)$ per share in each period. The other ("bonds"), in infinitely elastic supply at the price 1, pays interest of r each period.

In each period t, a common signal θ_t is first observed by all young rational traders. Young naive traders do not observe θ_t (or at least do not attach any significance to it). Following this, young traders of both types submit their demand functions: the quantity of shares they wish to buy at each possible price. These demand functions can be interpreted as a schedule of limit orders. Old agents then receive a dividend δ_t (which is not seen by the young agents before they submit their demand curves[14]) and transfer their shares at the market-clearing price p_t, which is determined by the condition that the demand of young traders equal the fixed unit supply of old agents.

The sole role of the common signal θ_t is to let young rational traders bid down the stock price in a coordinated fashion. Young traders in the prior period must not observe θ_t, lest they bid the price down in anticipation. Thus, the signal can be any information that comes out after young agents submit their demand curves in the prior period. For example, it may be the prior period's dividend, δ_{t-1}.

If a young agent buys x_t shares of stock in period t, costing her $p_t x_t$, her wealth in period $t+1$ is

$$W_{t+1} = x_t\left(p_{t+1} + \delta_{t+1}\right) + (L - p_t x_t)(1+r)$$

Young agents maximize expected utility $EU(W_{t+1}) = E\left[-e^{-\lambda W_{t+1}}\right]$, where λ is the coefficient of absolute risk aversion.

If p_{t+1} is normally distributed, a rational agent's demand for stocks, x_t^R, equals the expected difference between stock and bond returns, $E(p_{t+1} - p_t) + \overline{\delta} - rp_t$, divided by

[14]This timing guarantees that p_t and δ_t (which are both received by old agents in period t) are independent. Hence, the dividend risk cannot be reduced by building an equilibrium in which there is a negative correlation between the price and the dividend that old agents receive. This is important since the dividend risk is what creates limits to arbitrage in the model - limits that are both realistic and essential for our results. All the equilibria we study would remain equilibria if young traders observed δ_t before submitting their demand curves.

λ times the variance of the total stock return $p_{t+1} + \delta_{t+1}$:[15]

$$x_t^R = \frac{E(p_{t+1} - p_t) + \overline{\delta} - rp_t}{\lambda\left(\mathrm{Var}(p_{t+1}) + \sigma^2\right)} \tag{3}$$

where the expectation and variance are conditioned on all information available at time t (including p_t). In the simulations, we use this approximation to rational traders' demands. As explained in section 6, this has a negligible effect on our results, since nearly all the risk for the rational agents comes from the normally distributed dividend shocks.

The naive traders know that $\delta_t \sim N(\overline{\delta}, \sigma^2)$ and believe that the capital gains return $\frac{p_{t+1}-p_t}{p_t}$ is normally distributed[16] with mean R_t and variance V_t. Substituting these into (3), we obtain the demand x_t^N of a naive trader:[17]

$$x_t^N = \frac{R_t p_t + \overline{\delta} - rp_t}{\lambda\left(V_t p_t^2 + \sigma^2\right)} \tag{4}$$

---

[15]Since δ_{t+1} and p_{t+1} are independent (see note 14), the variance of the total stock return equals $\mathrm{Var}(p_{t+1})+\sigma^2$. The agent buys a number x of shares to maximize $E\left\{\exp\left[-\lambda\left(x\rho + (L - xp)(1+r)\right)\right]\right\}$, where $\rho = p_{t+1} + \delta_{t+1}$. If $\rho\tilde{\ } N(\overline{\rho}, \sigma_\rho^2)$, this equals

$$\int_{y=-\infty}^{+\infty} \exp\left[-\lambda\left(xy + (L - xp)(1+r)\right)\right] \frac{1}{2\pi\sigma_\rho} \exp\left[-(y-\overline{\rho})^2/2\sigma_\rho^2\right] dy$$

Completing the square and integrating, this becomes

$$\exp\left[-\lambda(L - xp)(1+r) - \left(\overline{\rho}\lambda x - \sigma_\rho^2\left(\lambda x\right)^2/2\right)\right]$$

The first order condition with respect to x is

$$0 = \lambda p(1+r) - \overline{\rho}\lambda + \sigma_\rho^2\lambda^2 x \Longrightarrow x = \frac{\overline{\rho} - (1+r)p}{\sigma_\rho^2\lambda}$$

and the result follows by substituting $\overline{\rho} = Ep_{t+1} + \overline{\delta}$, $\sigma_\rho^2 = \mathrm{Var}_t\, p_{t+1} + \sigma^2$.

[16]Some functional form for the distribution of the capital gains return must be assumed for the naive traders since they estimate only its mean and variance. We assume a normal distribution because it gives rise to the tractable form for the demand function given in (4).

[17]We use this formulation to make naive traders' beliefs match the GARCH specification, which predicts the volatility of capital gain *returns* rather than the volatility of the capital gain $p_{t+1} - p_t$ itself. The results are qualitatively the same with either formulation.

The condition for market clearing in each period t is

$$\mu x_t^R + (1-\mu)x_t^N = 1 \tag{5}$$

where x_t^R and x_t^N are given above. With the exception of section 8.2, we will assume that naive traders have random walk beliefs ($R_t \equiv 0$).

To rule out bubbles, we assume that agents have an arbitrarily small margin requirement $m > 0$. That is, an agent who wants to invest x in stocks must have at least mx in labor income. This implies that the stock price can never exceed L/m, where L is the aggregate amount of labor income.

We define an equilibrium as follows. Consider a pricing function $f(t, \theta_t, h_t) \rightarrow p_t$, where θ_t is the signal observed by young rational agents in period t and h_t is the history of prices and signals through time $t-1$, $(p_0, \theta_0, ..., p_{t-1}, \theta_{t-1})$. This pricing function is an *equilibrium* if $p_t = f(t, \theta_t, h_t)$ satisfies (5), where $E[p_{t+1}] = E[f(t+1, \theta_{t+1}, h_{t+1}) \mid \theta_t, h_t, p_t]$ and where $\text{Var}(p_{t+1})$ is defined analogously.

5 Theoretical Results

The following two theoretical results assume that naive traders have random walk beliefs. The results do not rely on the normal approximation to rational agents' demands. We first show that there is only one constant-price equilibrium:

Proposition 1 *There is only one constant-price equilibrium. In this equilibrium, the price always equals* $\overline{p} = \frac{1}{r}\left[\overline{\delta} - \lambda\sigma^2\right]$.

Proof. *Since the price is constant,* $V_t \equiv 0$ *and* $Var_t(p_{t+1}) \equiv 0$*, so the normal approximation of rational trader demand is exact. By setting* $V_t = R_t = E[p_{t+1} - p_t] = \text{Var}(p_{t+1}) = 0$ *in (5), one finds that* $\overline{p}$ *must equal* $\frac{1}{r}\left[\overline{\delta} - \lambda\sigma^2\right]$.

We call $\overline{p}$ the *variance-free price*. Proposition 2 shows that under random walk beliefs, even if the price is not a constant, it can never exceed the variance-free price. This places a tight limit on the size of any frenzy: it cannot exceed the difference between the variance-free price and the precrash price (which is set to 0.5% in the simulations).

Proposition 2 *In any equilibrium, the stock price can never exceed* $\overline{p} = \frac{1}{r}\left[\overline{\delta} - \lambda\sigma^2\right]$.

Proof. *Appendix A.*

The intuition is that since there can be no bubbles, there must be a maximum price $p^{\max}$ that can ever be attained in equilibrium. When the price is $p^{\max}$, rational traders must expect the next period's return to be zero or negative. Naive traders, by assumption, expect it to be zero. Agents' expected returns in the constant price equilibrium are at least as optimistic as this; furthermore, agents in the constant-price equilibrium expect zero volatility. So if the price is $p^{\max}$, stocks cannot offer a more attractive return distribution to either type of agent than in the constant-price equilibrium. But then no agent will ever be willing to pay more than the price in that equilibrium, which is $\overline{p}$; hence, $p^{\max} \leq \overline{p}$.

6 Simulations

One aim of the simulations is to determine the relative importance of frenzies vs. crashes. We thus focus on equilibria in which there is initially a constant probability q of a crash, in which prices fall by a proportion $c \in [0, 1]$, and the *same* constant probability q of a frenzy, in which prices rise by a proportion $c' \in [0, 1]$.[18] Crashes are more important than frenzies insofar as $c > c'$.[19]

After either event occurs, the signal θ in every subsequent period is ignored by the rational agents. This approach is adopted because it yields a unique equilibrium price

[18]This can be captured in the model by the existence of two sets S_C and S_F of possible values of the signal θ_t, each occurring with probability q, such that there is a crash if $\theta_t \in S_C$ and a frenzy if $\theta_t \in S_F$.

[19]Another potential approach might be to assume that frenzies and crashes are of the same size but arrive at different frequencies. However, this approach does not work, since crashes and frenzies must generally be of specific sizes to be consistent with equilibrium. For example, if naive traders believe that prices follow a random walk ($R_t \equiv 0$), a 20% crash can be sustained for some proportion μ of rational traders. But by Proposition 2, the price can never exceed $\overline{p}$, so frenzies cannot exceed the difference between the variance-free price $\overline{p}$ and the precrash price (which is set to 1% in the simulations).

path following a crash and a unique (though different) path following a frenzy.[20] One implication of this approach, however, is that frenzy signals are accompanied by the good news that there will be no crashes in the future. This good news causes prices to rise, which should not be confused with a frenzy. We minimize this effect by focusing on equilibria in which crashes happen very infrequently, so that the crash risk has a very small (0.5%) effect on the precrash price level.[21] Further details are found in Appendix B.

The parameters are calibrated to be as realistic as possible given the stylized nature of the model. Let T equal the length of one period as a fraction of a year. For crashes the appropriate trading period is one day, so we set $T = 1/250$, where 250 is the approximate number of trading days in a typical year. We take the annual interest rate to be 5%, so $r = e^{.05T} - 1$. We normalized the expected annual dividend to 1, so the expected dividend in one day, $\overline{\delta}$, equals T.

In our simulations the parameter σ^2 proxies for all of the noncrash risk that accompanies investment in the stock market. This includes not only dividend risk but also risk due to possible noise trading, monetary shocks, etc.[22] We thus calibrate σ^2 so that it captures the empirical volatility in the stock market. We use the daily closing S&P Composite Index from January 3, 1928 to October 16, 1987, the last trading day before the 1987 crash.[23] The daily noncrash volatility in the index over this period is

[20]Since agents ignore any subsequent signals θ, any randomization must be independent across agents. There is a continuum of rational agents, so each rational agent's demand has an infinitesimal effect on total demand. Thus, the price is nonrandom by the law of large numbers.

[21]We refer to the period before (after) a crash or frenzy as the "precrash" ("postcrash") period for brevity.

[22]In addition, our dividends are i.i.d. while actual dividends are positively correlated over time, so even if we were concerned only with modelling dividend risk, σ^2 would have to be higher than the empirical dividend volatility in order to capture the price risk resulting from dividend shocks.

[23]The data are taken from the Yahoo Finance web page. Similar results, available from the author, are obtained using the Dow Jones Industrial Average.

$V^{nc} = 0.00011638$.[24] In computing this we first remove any 1-day returns below -5% to obtain the noncrash risk present in the market. The computation of σ^2 from V^{nc} is straightforward (Appendix C).

Since traders have constant absolute risk aversion, the coefficient of relative risk aversion is not a useful metric. We thus calibrate λ, the coefficient of absolute risk aversion, so that the precrash equity premium equals its historic average of 6% (Kocherlakota [35]). We verify in section 7 that crashes can occur for much lower levels of λ as well. This is important since the historic equity premium may entail implausibly high levels of risk aversion (Mehra and Prescott [40]).

We use the approximation to rational agents' demand that assumes returns are normally distributed (equation (3)). The use of this approximation has miniscule effects on the simulation. On and after the day of the crash or frenzy, the true variance of p_{t+1} is zero, so the return risk comes solely from the normally distributed dividend shock. Prior to the crash or frenzy, the only price risk comes from the chance of a crash, which we assume is small enough to reduce prices by only 0.5%. In contrast, the dividend risk reduces prices by over 50%. Thus, the normal dividend shock dominates the return distribution in this period as well.

Since the proportion μ of rational traders is unobservable, our approach is to specify a crash size (e.g., 20%) and then let μ be that proportion for which such a crash can be sustained. The estimated μ is always less than 5%. This is lower than the 1/3 of investors whom Shiller [49, p. 394] reports as having relied on technical analysis in the 1987 crash. However, the traders who correctly forecasted the crash must have been a very small subset of technical analysts. Out of Jack Schwager's 32 interviews with successful traders [46, 47], there were only two (Stanley Druckenmiller and Paul Tudor Jones) who claim to have known prior to October 19, 1987 that a crash would occur on that specific day.

The naive traders predict volatility according to equation (1). To estimate the parameters α and β, we fit Bollerslev's GARCH model [10] to daily closing S&P Composite

[24] The day-t return is $\frac{p_t - p_{t-1}}{p_{t-1}}$.

Index levels from 1/3/28 to 10/16/87.[25] The specification we fit is as follows:

$$\begin{aligned} r_{t+1} &= a + \varepsilon_{t+1} \\ \varepsilon_{t+1} &\sim N(0, V_t) \\ V_t &= \alpha_0 + \alpha\varepsilon_{t-1}^2 + \beta V_{t-1} \end{aligned} \tag{6}$$

where $r_{t+1} = \frac{p_{t+1}-p_t}{p_t}$ and V_t is the conditional variance of r_{t+1}.

The parameter estimates are given in Table 1.

Variable	**Parameter**	**t Value**
a	0.000482	9.24
α_0	1.0154×10^{-6}	18.77
α	0.0903	34.53
β	0.9034	343.67

Table 1: Parameter estimates from the GARCH model.

In the simulations, we set a to zero since the theoretical model is stationary. We also set α_0 to zero since noncrash price risk in the model is already proxied by σ^2.

[25]In GARCH, future volatility depends on negative and positive returns symmetrically. Empirical evidence suggests that volatility rises more following market declines. The Exponential GARCH model of Nelson [42] captures this by weighting negative returns more in the volatility estimate. We do not study Exponential GARCH beliefs since they would only increase the potential for crashes and the tendency for skewed returns.

Figure 2 shows a simulation of a 20% crash (marked with solid diamonds) and the corresponding frenzy (marked with open circles).

Figure 2: Crash and Frenzy Prices

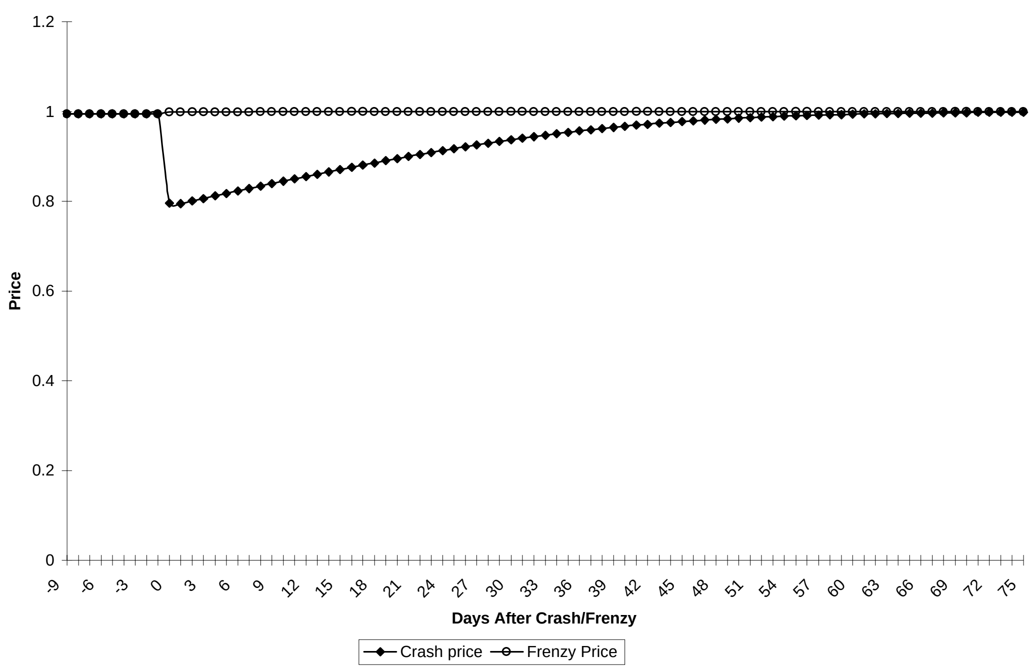

The parameters from this simulation appear in Table 2.

Parameter	**Value**
λ (CARA)	0.22696
μ (proportion rational)	3.83%
c (crash proportion)	20%
c' (frenzy proportion)	0.42%
Average wait to crash/frenzy	41 years

Table 2: Simulation parameters and estimates.

The average wait until a crash or frenzy is 41 years. The proportion μ of rational traders that makes a 20% crash possible is 3.83%. Returns are skewed: while prices fall by 20% in a crash, they rise by only 0.42% in a frenzy. This follows from random-walk beliefs: by Proposition 2, prices cannot rise above the variance-free price $\overline{p}$ in a frenzy, so the frenzy cannot exceed 0.5% of $\overline{p}$.

Figure 3 shows the share of stock owned by rational traders during the crash. (The share owned by the naive traders is one minus this number.) For reference, the path of prices is also shown.[26]

Figure 3: Prices and Rational Agent Stock Ownership in a Crash

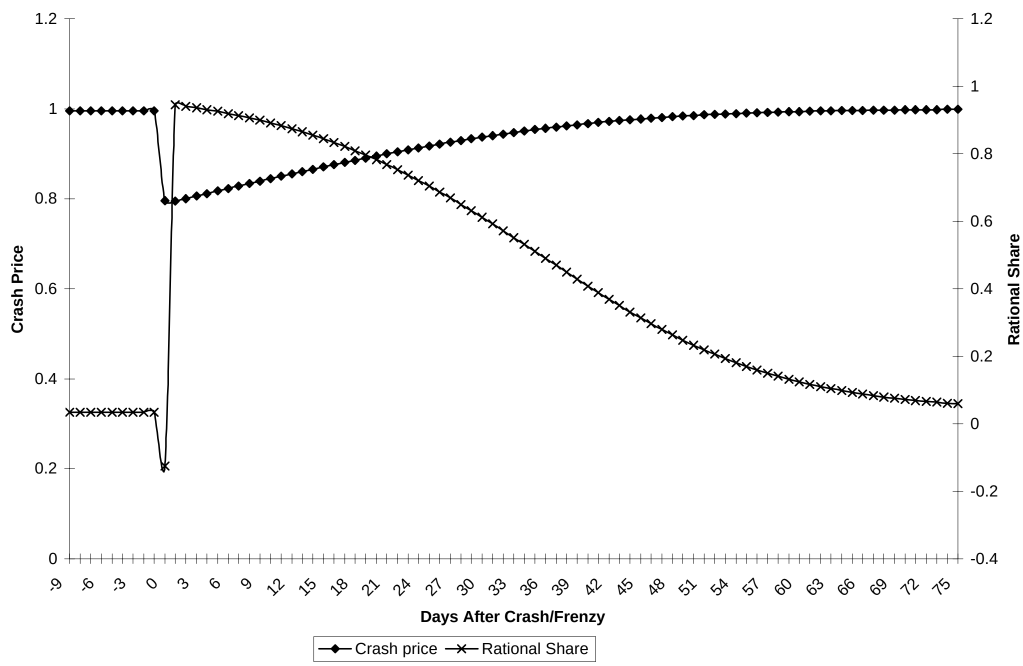

On the day of the crash, rational traders sell about 16% of the market short. This causes a 20% decrease in price. The following day they reverse course, buying nearly 100% of the naive traders' shares. In contrast, the total stock and stock index trading volume on October 19, 1987 amounted to only about 1% of total stock market value (Gennotte and Leland [28, p. 999]). However, Gennotte and Leland [28] estimate that about 98% of stock is held by investors who do *not* trade to alter their overall exposure to the stock market, in contrast to the rational and naive traders in our model. These investors would sit pat during a crash. Thus, the stock sold by naive traders probably corresponds to no more than 2% of the U.S. stock market. Since some who sold in

[26] The share of stock owned by rational traders is marked with crosses and corresponds to the vertical axis on the right side. The price path is marked with solid diamonds and corresponds to the left side axis.

response to the crash probably did so the day after, our story is consistent with the empirical trading volume of 1% on October 19.

Figure 4 shows the consumption of old agents, according to type.[27]

Figure 4: Current Consumption of the Old in a Crash

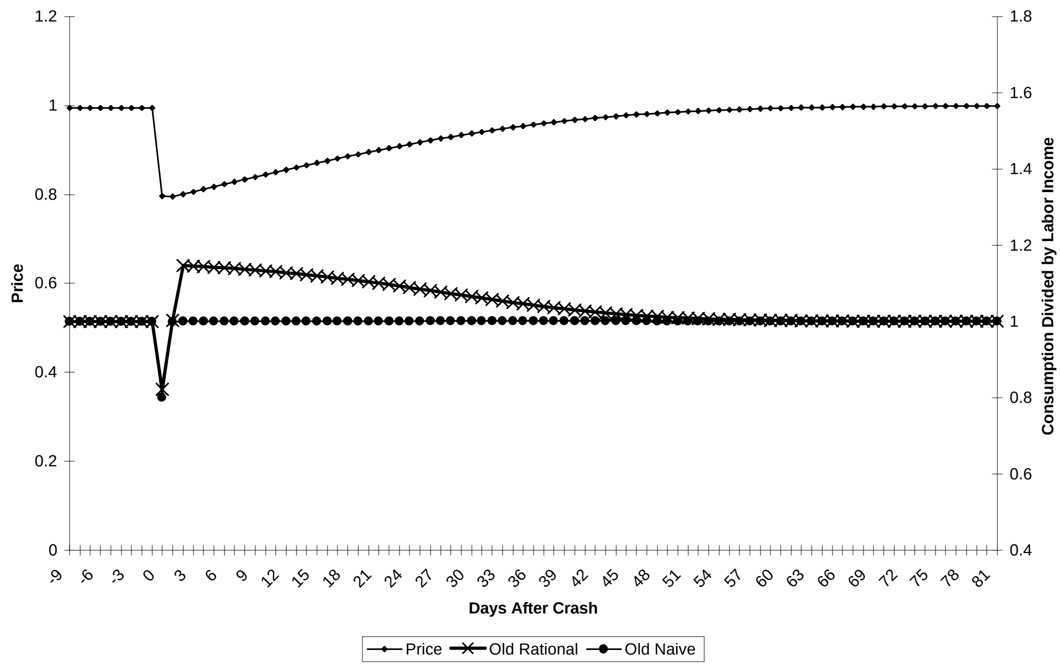

The reason old rational agents don't do better on the crash day is that the crash signal is unexpected, so they cannot lighten up on stocks in anticipation.[28] The day after the crash, the naive traders exit en masse, causing prices to fall again slightly. On this day, old rational traders (who sold short on the crash day) do about as well as naive traders. After this, however, rational traders do much better than naive traders. Unlike

[27]Labor income is normalized to 1. The current consumption of old rational traders is marked by crosses, while the consumption of old naive traders is marked with solid circles. These curves virtually coincide until the second day after the crash. The consumption levels correspond to the right side axis. For reference, the stock price is also shown, above; it corresponds to the left side axis.

[28]This might not be true if each rational agent had a truly private signal - e.g., the results of her own technical analysis - that was correlated with other agents' current and future signals (Jackson and Peck [33]). In this case, some agents might come to believe before the crash day that a crash is likely. Indeed, traders' fluctuating assessments of the likelihood of a major correction may help explain the large price swings that occurred in the week that preceded the 1987 crash.

naive traders, they know the high volatility of the crash is a short-lived phenomenon. They also know that as the naive traders calm down and reenter the market, prices will rise. Knowing these facts, they hold large long positions and receive large capital gains as the price gradually returns to its fundamental value.

7 Sensitivity Analysis

The level of risk aversion λ is chosen to replicate the historic equity premium of 6% (Kocherlakota [35]). However, the historic equity premium is hard to reconcile with other estimates of risk aversion (Mehra and Prescott [40]), so in this section we check what happens if agents are less risk averse. We also examine which crash sizes can be sustained for all possible proportions μ of rational agents

Figure 5: Equilibrium Crash Sizes by Proportion of Rational Agents for Four Different Degress of Risk Aversion

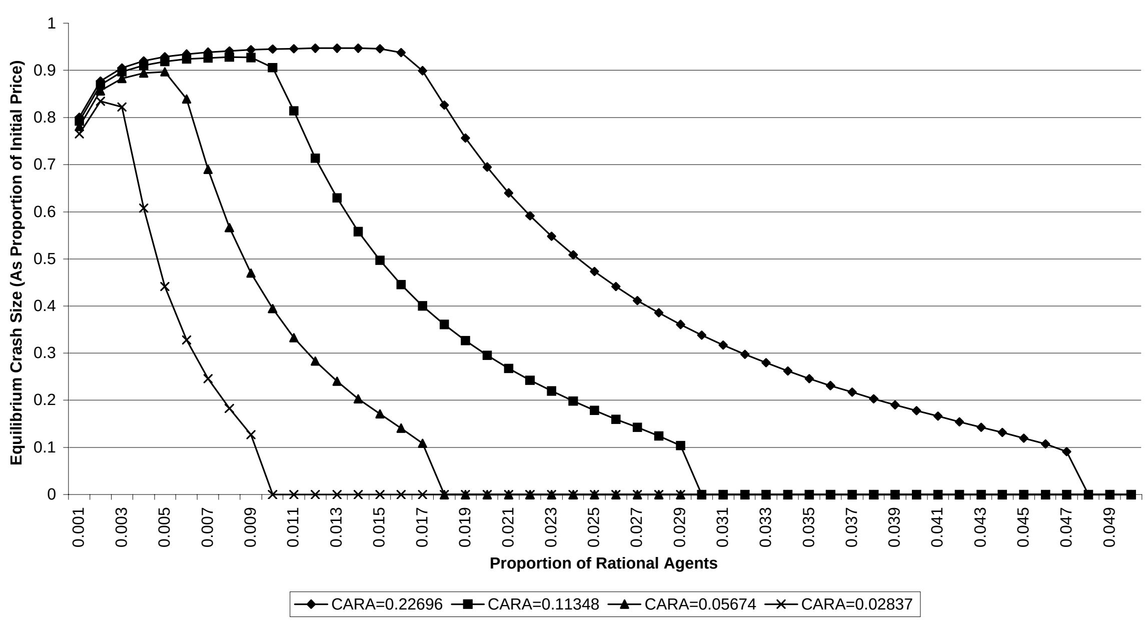

Figure 5 shows the crash sizes that occur in equilibrium for different proportions μ and for four different values of λ. (λ is denoted "CARA" in the Figure.) The vertical axis gives the crash size as a proportion of the precrash price. The proportion of rational traders, μ, appears on the horizontal axis. There are four curves, each corresponding to a different coefficient of absolute risk aversion. The highest curve uses the original coefficient λ from Table 2. The other three curves correspond to coefficients of absolute risk aversion that are a half, a quarter, and an eighth of this value. At these levels of risk aversion, crashes cannot be sustained if the proportion of rational traders exceeds 5%. Thus, values of μ above 0.05 are not depicted.

As shown in Figure 5, the crash size is increasing in the level of risk aversion and decreasing in the proportion of rational agents. This is intuitive. The day after the crash, young naive traders are unwilling to buy stock at almost any price. Their demand elasticity is virtually zero since their demand is proportional to the reciprocal of their variance estimate, which is now very high. Thus, the price is determined almost entirely by the demand of young rational traders. As their proportion shrinks or they become more risk averse, prices have to fall further the day after the crash to clear the market. The anticipation of this leads rational traders to bid the price down further on the day of the crash, so the crash is larger.

8 Naive Traders' Two Mistakes

The naive traders make two mistakes that cause them to sell after the crash. The first is that they overestimate the postcrash variance. The second is that they do not anticipate that the gradual reentry of naive traders to the stock market will lift prices. Instead, they continue to believe that stock prices follow a random walk - neither rising nor falling on average.

8.1 Mistake #1: Overestimating the Variance

The first mistake - overestimating the variance - is essential *within the class of equilibria we simulate.* If the naive traders know that the postcrash variance will be zero and have random walk beliefs, then there can be no crash, as the following proposition shows:

Proposition 3 *Consider the class of simulated equilibria, in which there is at most one crash. If naive traders know the true variance but have random-walk beliefs, there is* *no crash: the price must equal $\overline{p}$ on the day the "crash signal" is observed and on all subsequent days.*

Proof. *Appendix A*

The intuition is that on the day of the crash, the price will never again be as low. Thus, on this day, both types of traders know that the further price variance is zero and both believe that the next period's capital gain will be nonnegative. But the price $\overline{p}$ is computed on the assumption that traders expect no variance and zero capital gains. Since beliefs on the crash day are at least as optimistic as this, the price on the crash day cannot be lower than $\overline{p}$.

What happens if we do not restrict the form of the equilibrium? This remains an open question. For example, the original model (with random walk beliefs and GARCH expectations of variance) undoubtedly has other equilibria in which the variance of the capital gains return is not zero after the crash - e.g., because of a risk of subsequent crashes. We do not know whether this sort of equilibria would also exist if naive traders knew the true variance but had random walk beliefs.

8.2 Mistake #2: Random Walk Beliefs

Popular books and best-selling finance texts teach that since markets are efficient, stock prices must follow a random walk. However, this is not quite correct: the efficient markets paradigm implies that when stocks are riskier, their excess return should be higher. If naive traders understood this, then they would have less incentive to sell in periods of high volatility - such as after a crash. In this section we assume that the naive traders' expected return R_t, rather than being identically zero, is an increasing function of their volatility prediction V_t. Crashes can still occur for reasonable parameter values.

Our approach is to estimate the dependence of expected returns on volatility *empirically*, as we did for the equation governing volatility. Engle, Lilien, and Robins [22] show that if stocks are in fixed supply and agents have constant absolute risk aversion, the relation between V_t and R_t has the form

$$R_t = b_0 + b_1\sqrt{V_t} \tag{7}$$

where b_0 and b_1 are constants.[29] We estimate the following model empirically using the daily closing S&P Composite Index from January 3, 1928 to October 16, 1987:

$$
\begin{aligned}
r_{t+1} &= R_t + \varepsilon_{t+1} \\
R_t &= b_0 + b_1 \sqrt{V_t} \\
\varepsilon_{t+1} &\sim N(0, V_t) \\
V_t &= \alpha_0 + \alpha \varepsilon_{t-1}^2 + \beta V_{t-1}
\end{aligned}
$$

The parameter estimates are shown in Table 3.

GARCH Parameter Estimates		
Variable	**Parameter**	**t Value**
b_0	0.0000937	0.61
b_1	0.0545	2.60
α_0	1.023×10^{-6}	18.69
α	0.0904	34.33
β	0.9031	339.66

Table 3: Parameter estimates from GARCH model, using S&P Composite Index returns from 1/3/28-10/16/87.

The simulation requires a few minor modifications. In actual market crashes there are two sources of volatility: the crash itself and noncrash risk that comes from earnings surprises, liquidity shocks, and so on. Since the relation between V and R is nonlinear, we must add an estimate V^{nc} of this noncrash risk to V. We use the estimate $V^{\text{nc}} = 0.00011638$ from Section 6, which captures the average noncrash risk in the market. In addition, we subtract $b_0 + b_1\sqrt{V^{\text{nc}}}$ from the return estimate so that the naive traders expect zero capital gains on average if the crash-related volatility V is zero. This reflects the fact that dividends are i.i.d. and guarantees that, as before, the price will converge

[29] They show that it has the form $R_t = b_0 + b_1\sqrt{c + V_t}$ but when this model is estimated, the point estimate for c is 0, so we use the simpler model instead.

to $\overline{p}$ as V gradually shrinks to zero after the crash. Thus, naive traders predict the mean of the capital gains return as $R_t = b_1\left(\sqrt{V_t + V^{\text{nc}}} - \sqrt{V^{\text{nc}}}\right)$. We also set α_0 to zero as before.

Crashes can still occur. Figure 6 shows a simulation of a 20% crash. The path of prices is shown, together with the share of stocks owned by rational traders (as in Figure 3).

Figure 6: Prices and Rational Agent Stock Ownership in a Crash when Naive Traders Believe that Expected Returns are Increasing in Volatility

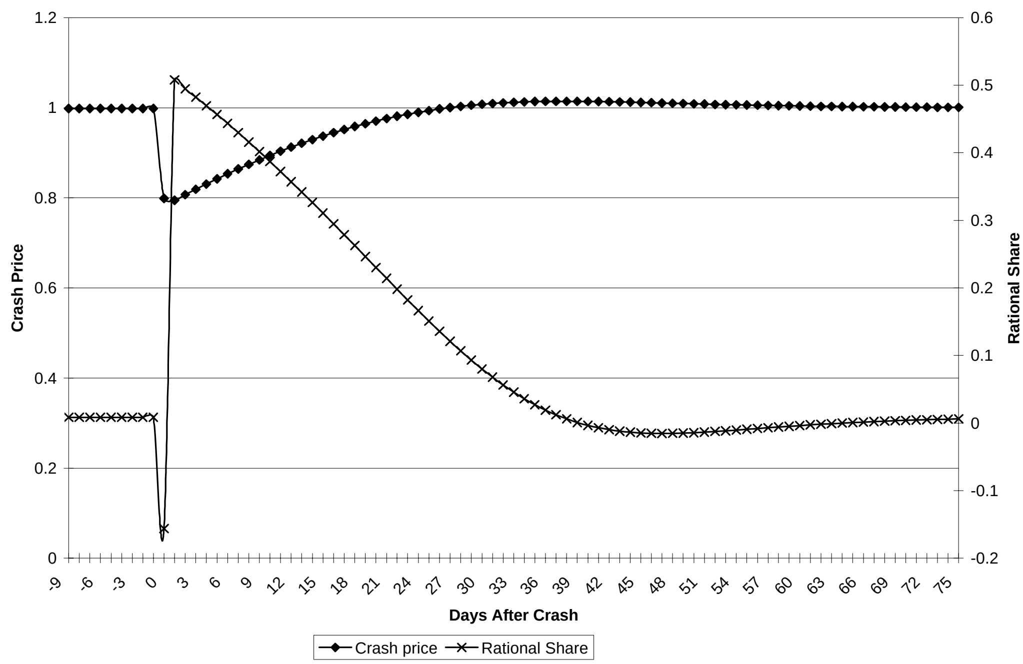

Since R_t is increasing in V_t, naive traders expect positive returns after the crash. This does counteract their decision to sell. However, their overestimate of postcrash volatility is large enough by itself to drive them out of the market following a 20% crash. Consequently, a crash remains a self-fulfilling prophecy for the rational traders.

One difference from the base model is that after prices recover, they actually rise very slightly above the variance-free price (here normalized to 1). As long as their prediction of volatility is positive, naive agents expect positive capital gains. For the parameters we calibrate, this makes them willing to pay more than the variance free price when predicted volatility is small.

9 Discussion

In the simulation of section 6, prices quickly return to their precrash level. 62 trading days after the crash, the entire 20% price decline has been erased. After the 1987 crash, it took 256 trading days for the S&P 500 to return to its precrash level. One reason may be that our naive traders believe that returns are normally distributed. In contrast, one conclusion investors may have drawn from the 1987 crash is that returns are highly skewed and 20% crashes can occur. This is supported by evidence that option prices reflected crash fears only *after* the 1987 crash (Jackwerth and Rubinstein [32]). Thus, the crash led traders not only to expect more volatility in the following 60 days, but also to fear more crashes in the years to come. This may have led to a long-lasting decrease in their willingness to buy stock, thus depressing stock prices for much longer than our model predicts. Adding this dimension to traders' beliefs would be a realistic addition to our model that we do not pursue in the current paper.

In this paper we also do not address the issue of how rational traders end up interpreting the same signals in the same way. Interviews in Schwager [46, 47] suggest a dynamic process in which some traders discern new patterns in recent price behavior and start using these patterns to govern their trades. Others hear about the new technical signals and start to follow them. Even if the initial patterns were spurious, the signals can retain their value since they now act as a coordinating device. When finally too many traders learn about the signals, they become less valuable,[30] and a search for new technical signals begins. This topic merits further study.

A Proofs

Proof of Proposition 2. Since returns may not be normal, we must use the exact exponential utility function for rational trader demand to prove this result. A trader's

[30] This is consistent with our finding that crashes can occur only if less than 5% of active traders see the technical signal.

demand D in period t, given the price p_t, solves

$$\max_D \int_p \int_\delta -e^{-\lambda[(1+r)W+D(p+\delta-(1+r)p_t)]} d\Phi(\delta) dF(p) \tag{8}$$

where p represents the next period's price, which has some c.d.f. F, and δ the next period's dividend, which has the normal c.d.f. Φ. Since $\Phi'(\delta) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(\delta-\overline{\delta})^2/2\sigma^2}$, we drop leading constants and complete the square to show that (8) is equivalent to:

$$\begin{aligned} & \max_D \int_p -e^{-\lambda D(p-(1+r)p_t)} \left(\frac{1}{\sqrt{2\pi}\sigma} \int_\delta e^{-\lambda D\delta-(\delta-\overline{\delta})^2/2\sigma^2} d\delta \right) dF(p) \\ = \; & \max_D \int_p -e^{-\lambda D(p+\overline{\delta}-(1+r)p_t)+(D\lambda\sigma)^2/2} dF(p) \end{aligned}$$

The derivative of this integral with respect to D is:

$$e^{-\lambda D(\overline{\delta}-(1+r)p_t)+(D\lambda\sigma)^2/2} \int_p \left(\lambda(p+\overline{\delta}-(1+r)p_t) - D\lambda^2\sigma^2 \right) e^{-\lambda Dp} dF(p) \tag{9}$$

Since the leading exponential term is positive, the sign of this expression equals the sign of the integral.

Let $p^{\max}$ be the supremum of the set of prices reached with positive probability. Since the stock price can never exceed L/m, $p^{\max} < \infty$. Thus, for any $\varepsilon > 0$ there is a time t, a history h_t, and a signal θ_t such that p_t is within ε of $p^{\max}$. There are now two cases.

1. If $x_t^R < 0$, then by (5), $x_t^N > 1$. But $x_t^N \leq \frac{\overline{\delta}-rp_t}{\lambda\sigma^2}$ by (4)[31] so $p_t \leq \overline{p}$, whence $p^{\max} \leq \overline{p} + \varepsilon$.

2. If $x_t^R \geq 0$, there are two subcases. (a) If $p_t \leq \overline{p}$ then $p^{\max} \leq \overline{p} + \varepsilon$ as in case 1. (b) If $p_t > \overline{p}$ then $x_t^N \leq \frac{\overline{\delta}-rp_t}{\lambda\sigma^2} < \frac{\overline{\delta}-r\overline{p}}{\lambda\sigma^2} = 1$ so

$$\begin{aligned} 1 &= \mu x_t^R + (1-\mu)x_t^N < \mu x_t^R + (1-\mu) \\ &\implies x_t^R > 1 \end{aligned}$$

Thus, letting $D = x_t^R$ in (9),

$$\lambda(p+\overline{\delta}-(1+r)p_t) - D\lambda^2\sigma^2 < \lambda(p+\overline{\delta}-(1+r)p_t) - \lambda^2\sigma^2$$

[31] Naive traders believe that the capital gains return is normally distributed, so the normal approximation always holds exactly for them.

But $p - p_t \leq p^{\max} - p_t \leq \varepsilon$, so

$$\lambda(p + \overline{\delta} - (1+r)p_t) - \lambda^2\sigma^2 \leq \lambda(\varepsilon + \overline{\delta} - rp_t) - \lambda^2\sigma^2$$

However, the integrand in (9) must be nonnegative for some values of p or rational traders will set $x_t^R = -\infty$. Thus,

$$\begin{aligned} 0 &< \lambda(\varepsilon + \overline{\delta} - rp_t) - \lambda^2\sigma^2 \\ &\Longrightarrow p_t < \overline{p} + \frac{\varepsilon}{r} \end{aligned}$$

so that $p^{\max} \leq \overline{p} + \frac{\varepsilon}{r} + \varepsilon$.

In all cases, $p^{\max} < \overline{p} + k\varepsilon$ for a fixed constant k. Since this is true for all ε, $p^{\max} \leq \overline{p}$. Q.E.D.

Proof of Proposition 3: Since old traders can dispose of their shares freely, the price of stock is never negative. Accordingly, let $p^{\min} \in [0, \overline{p}]$ be the infimum of prices ever reached in this class of equilibria on or after the period in which the crash or frenzy signal is observed. For any $\varepsilon > 0$ there is a price p_t, observed in some equilibrium, that is within ε of $p^{\min}$. Since all traders know the variance of the capital gains return is zero, the normal approximation is exact. By (5), since $p_{t+1} - p_t \geq -\varepsilon$,

$$\begin{aligned} 1 &\geq \mu\frac{-\varepsilon + \overline{\delta} - rp_t}{\lambda\sigma^2} + (1-\mu)\frac{\overline{\delta} - rp_t}{\lambda\sigma^2} \\ &\Longrightarrow p_t \geq \overline{p} - \frac{\mu\varepsilon}{r} \end{aligned}$$

Thus, $p^{\min} \geq p_t - \varepsilon \geq \overline{p} - \frac{\mu\varepsilon}{r} - \varepsilon$. Since this is true for any $\varepsilon > 0$, $p^{\min} \geq \overline{p}$. By Proposition 2, the price must equal $\overline{p}$ both in and after the period in which the crash or frenzy signal is observed. Q.E.D.

B Simulated Equilibria

In the simulations, the price is assumed constant in the precrash period. It is determined in the following way. We use the normal approximation to rational traders' demands.[32] Since the precrash price is constant, naive traders' initial predictions of the mean and variance of returns are both zero ($V_t = R_t = 0$ in any precrash period t).[33] For the rational traders, $E[p_{t+1} - p_t]$ equals $(qc' - qc)p^{PC} \overset{d}{=} Ap^{PC}$ and $\text{Var}(p_{t+1}) = (qc'^2 + qc^2 - (qc' - qc)^2)\left(p^{PC}\right)^2 \overset{d}{=} B\left(p^{PC}\right)^2$. The precrash price p^{PC} is the solution to (5) with these values substituted:

$$1 = \mu \frac{Ap^{PC} + \overline{\delta} - rp^{PC}}{\lambda\left(B\left[p^{PC}\right]^2 + \sigma^2\right)} + (1 - \mu)\frac{\overline{\delta} - rp^{PC}}{\lambda\sigma^2} \tag{10}$$

The postcrash price path is computed as follows. Suppose that a crash or frenzy occurs at some time $t' \leq t$. Since the continuation path of prices is unique, $\text{Var}(p_{t+1}) = 0$ and $E(p_{t+1} - p_t) = p_{t+1} - p_t$. Substituting these into (5) and solving for p_{t+1}, we obtain

$$p_{t+1} = p_t(1 + r) - \overline{\delta} + \frac{\lambda\sigma^2}{\mu}\left[1 - (1 - \mu)\frac{R_t p_t + \overline{\delta} - rp_t}{\lambda\left(V_t p_t^2 + \sigma^2\right)}\right] \tag{11}$$

We use this to solve forwards for the price path and impose the requirement that the price converges to $\overline{p}$.

The algorithm for computing the simulation is as follows. We first compute the variance-free price $\overline{p}$ as described in section 6. The precrash price p^{PC} is set to be 0.5% less than $\overline{p}$. We then find crash and frenzy sizes c, c' that lead to convergent price paths if the price is initially p^{PC}. Finally, given these crash and frenzy sizes, we find a per-period crash/frenzy probability q such that p^{PC} satisfies (10).

[32]This approximation is exact on and after the day of the crash/frenzy since the variance of the next day's price is zero and all risk is due to the normally distributed dividend. It is inexact only before the crash; however, this creates only a small distortion since crashes are very unlikely, so overall risk is dominated by the normally distributed dividend.

[33]This is necessary since if naive traders observe a constant price, they see zero empirical volatility and a zero return. Thus, their adaptive estimates V_t and R_t must converge to zero over time unless they are already zero. Since changes in V_t and R_t will cause changes in the market-clearing price, they must be identically zero in the precrash, constant-price period.

C Calibration

We compute σ^2 from $V^{\mathrm{nc}} = 0.00011638$ (section 6) as follows. V^{nc} is the variance per dollar invested. The per-share variance is thus $\overline{p}^2 V^{\mathrm{nc}}$, where the price is $\overline{p}$ in a noncrash (constant-price) world. If σ^2 is to capture this risk, we must have

$$\sigma^2 = \overline{p}^2 V^{\mathrm{nc}} \tag{12}$$

Moreover, if σ^2 represents the only risk to the investor (i.e., if the price is constant), the stock price must equal

$$\overline{p} = \frac{1}{r}\left(\overline{\delta} - \lambda\sigma^2\right) \tag{13}$$

Equations (12) and (13) are solved to yield

$$\overline{p} = \frac{1}{2\lambda V^{\mathrm{nc}}}\left(-r + \sqrt{\left(r^2 + 4\lambda V^{\mathrm{nc}}\overline{\delta}\right)}\right) \tag{14}$$

σ^2 is computed by substituting (14) into (12). The equations for p and σ^2 depend on λ. We choose λ so that the equity premium equals 6% per year . Since we assume an the annual risk-free interest rate of 5%, this is the λ that gives a $\overline{p}$ for which

$$\overline{\delta}/\overline{p} = e^{.11T} - 1 \tag{15}$$

References

[1] Aït-Sahalia, Yacine, Yubo Wang, and Francis Yared. 2000. "Do Options Markets Correctly Assess the Probabilities of Movement of the Underlying Asset?" Forthcoming, *Journal of Econometrics*.

[2] Bachelier, L.J.B.A. 1900. *Theorie de la Speculation.* Paris: Gauthier-Villars, 1900.

[3] Barberis, Nicholas, Ming Huang, and Tano Santos. 1999. "Prospect Theory and Asset Prices." Forthcoming, *Quarterly Journal of Economics*.

[4] Barberis, Nicholas, Andrei Shleifer, and Robert Vishny. 1998. "A Model of Investor Sentiment." Mimeo.

[5] Barlevy, Gadi, and Pietro Veronesi. 2000. "Rational Panics and Stock Market Crashes." Mimeo.

[6] Barsky, Robert B., and J. Bradford De Long. 1993. "Why Does the Stock Market Fluctuate?" *Quarterly Journal of Economics* 108:291-311.

[7] Bates, David S. 1991. "The Crash of '87: Was It Expected? The Evidence from Options Markets." *Journal of Finance* XLVI:1009-1044.

[8] Black, F. 1976. "Studies of Stock Market Volatility Changes." 1976 *Proceedings of the American Statistical Association, Business and Economics Statistics Section*, 177-181.

[9] Board of Governors of the Federal Reserve System. March 9, 2001. *Flow of Funds Accounts of the United States: Annual Flows and Outstandings, 1985-1994.*

[10] Bollerslev, Tim. 1986. "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31:307-328.

[11] Brealey, Richard A., and Steward C. Myers. 1988. Principles of Corporate Finance: International Edition. Singapore: McGraw-Hill.

[12] Brock, William, Josef Lakonishok, and Blake LeBaron. 1992. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *Journal of Finance* XLVII:1731-1764.

[13] Campbell, John Y., and Ludger Hentschel. 1992. "No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns." *Journal of Financial Economics* 31:281–331.

[14] Chou, Ray Yeutien. 1988. "Volatility Persistence and Stock Valuations: Some Empirical Evidence Using GARCH." *Journal of Applied Econometrics* 3:279-294.

[15] Christie, A.A. 1982. "The Stochastic Behavior of Common Stock Variances: Value, Leverage, and Interest Rate Effects." *Journal of Financial Economics* 10:407-432.

[16] Cutler, David M., James M. Poterba, and Lawrence H. Summers. 1989. "What Moves Stock Prices?" *Journal of Portfolio Management* 15:4-12.

[17] De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldmann. 1990. "Positive Feedback Investment Strategies and Destabilizing Rational Expectations." *Journal of Finance* 45:379-395.

[18] Donaldon, R. Glen, and Harold Y. Kim. 1993. "Price Barriers in the Dow Jones Industrial Average." *Journal of Financial and Quantititative Analysis* 28:313-330.

[19] Duffie, Darrell. 1996. *Dynamic Asset Pricing Theory.* New Jersey: Princeton University Press.

[20] Engle, Robert F. 1982. "Autoregressive Conditional Heterskedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica* 50:987-1008.

[21] Engle, Robert F., and Tim Bollerslev. 1986. "Modelling the Persistence of Conditional Variances." *Econometric Reviews* 5:1-50.

[22] Engle, Robert F., David Lilien, and Russel P. Robins. 1987. "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model." *Econometrica* 55:391-407.

[23] Engle, Robert F., and V.K. Ng. 1993. "Measuring and Testing the Impact of News on Volatility." *Journal of Finance* 48:1749-1778.

[24] Fama, Eugene F. 1965. "The Behavior of Stock-Market Prices." *Journal of Business* 38:34-105.

[25] Frankel, Jeffrey A., and Kenneth A. Froot. 1990. "Chartists, Fundamentalists, and Trading in the Foreign Exchange Market." *American Economic Review Papers and Proceedings* 80:181-185.

[26] French, Kenneth R., G. William Schwert, and Robert F. Stambaugh. 1987. "Expected Stock Returns and Volatility." *Journal of Financial Economics* 19:3-29.

[27] Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein. 1992. "Herd on the Street: Informational Inefficiencies in a Market with Short-Term Speculation." *Journal of Finance* XLVII:1461-1484.

[28] Gennotte, Gerard, and Hayne Leland. 1990. "Market Liquidity, Hedging, and Crashes." *American Economic Review* 80:999-1021.

[29] Grossman, Sanford J., and Zhongquan Zhou. 1996. "Equilibrium Analysis of Portfolio Insurance." *Journal of Finance* LI:1379-1403.

[30] Hansen, Lars Peter, and Thomas J. Sargent. 2000. "Wanting Robustness in Macroeconomics." Mimeo.

[31] Hong, Harrison, and Jeremy C. Stein. 1999. "Differences of Opinion, Rational Arbitrage and Market Crashes." Mimeo.

[32] Jackwerth, Jens C., and Mark Rubinstein. 1996. "Recovering Probability Distributions from Option Prices." *Journal of Finance* 51:1611-1631.

[33] Jackson, Matthew O., and James Peck. 1991. "Speculation and Price Fluctuations with Private, Extrinsic Signals." *Journal of Economic Theory* 55:274-95.

[34] Kleidon, Allan W., and Robert E. Whaley. July 1992. "One Market? Stocks, Futures, and Options During October 1987." *Journal of Finance Papers and Proceedings*, pp. 851-877.

[35] Kocherlakota, Narayana R. 1996. "The Equity Premium: It's Still a Puzzle." *Journal of Economic Literature* XXXIV:42-71.

[36] MacKinlay, A. Craig, and Krishna Ramaswamy. 1988. "Index-Futures Arbitrage and the Behavior of Stock Index Futures Prices." *Review of Financial Studies* 1:137-158.

[37] Malkiel, Burton. 1979. "The Capital Formation Problem in the United States." *Journal of Finance* 34:291-306.

[38] Malkiel, Burton. 1985. *A Random Walk Down Wall Street, 4th ed.* New York: Norton.

[39] Mandelbrot, Benoit. 1963. "The Variation of Certain Speculative Prices." *Journal of Business* 36:394-419.

[40] Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15:145-61.

[41] Muth, John. 1960. "Optimal Properties of Exponentially Weighted Forecasts." *Journal of the American Statistical Association* 55:299-306.

[42] Nelson, Daniel B. 1991. "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica* 59:347-370.

[43] Pindyck, Robert S. 1984. "Risk, Inflation, and the Stock Market." *American Economic Review* 74:335-351.

[44] Poterba, James M., and Lawrence H. Summers. 1986. "The Persistence of Volatility and Stock Market Fluctuations." *American Economic Review* 76:1142-1151.

[45] Presidential Task Force on Market Mechanisms (Brady Commission). January 1988. *Report.* Washington: U.S. Government Printing Office.

[46] Schwager, Jack D. 1989. *Market Wizards: Interviews with Top Traders.* New York: Harper-Collins.

[47] Schwager, Jack D. 1992. *The New Market Wizards: Conversations with America's Top Traders.* New York: Harper-Collins.

[48] Sharpe. William F. 1981. *Investments.* New Jersey: Prentice-Hall.

[49] Shiller, Robert J. 1989. *Market Volatility.* Cambridge, Mass.: MIT Press.

[50] U.S. Department of Commerce, Bureau of Economic Analysis. August 2000. *National Income and Product Accounts.*

[51] U.S. Securities and Exchange Commission, Trading and Exchange Division. 1947. *A Report on Stock Trading on the New York Stock Exchange on September 3, 1946.* Washington: Securities and Exchange Commission.