ON THE NATURE OF THE STOCK MARKET: SIMULATIONS AND EXPERIMENTS

by

Hendrik J. Blok

B.Sc., University of British Columbia, 1993
M.Sc., University of British Columbia, 1995

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE STUDIES

(Department of Physics and Astronomy)

We accept this dissertation as conforming
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

November 2000

Chapter 6

Experiments with a hypothetical portfolio

This chapter is somewhat of a departure from the rest of the thesis. It describes some experimental results obtained purely to satiate my own curiosity. As such, the experiments are less rigorous than they could be (in particular, the dataset is quite small) and the results should not be taken too seriously. However, these experiments may yield valuable insight for the reader because they describe a real-world application of some of the theory discussed in prior chapters.

6.1 Motivation

The agents in Chapter 3 trade using a *fixed investment strategy* (FIS) which states that they should keep a fixed fraction of their capital in stock and the remainder as cash, to minimize risk and maximize returns. As was discussed, the theory underlying it has two important assumptions: (1) that there are no costs associated with trading, and (2) that moments of the return distribution higher than two (in particular, the kurtosis) are negligible on short timescales.

I was curious how well FIS would work in a real market environment where these assumptions may not hold so I constructed a hypothetical portfolio to track real stocks. Sandbox Entertainment (`http://www.sandbox.net/business/`) provides an online simulated stock market called "PortfolioTRAC" which gives users an imaginary bankroll of $100,000 and allows them to invest it in stocks listed on the major American markets. The simulation uses real trading prices and allows daily trades. Although it requires a few other idealizations, it is quite thorough and supports such complexities as short positions, limit and stop orders, broker fees, and daily interest on cash. Note that trades are only processed once per day (after

Symbol	Company	Price	Shares	Value
AAPL	Apple Comp Inc	$41.25	244	$10,065.00
AMD	Adv Micro Device	$28.00	345	$9,660.00
AU	Anglogold Ltd	$19.56	514	$10,055.13
CHV	Chevron Corp	$82.06	120	$9,847.50
EK	Eastman Kodak Co	$71.19	141	$10,037.44
IMNX	Immunex Corp	$116.50	82	$9,553.00
MSFT	Microsoft Corp	$141.00	69	$9,729.00
NSCP	Netscape Comm	$63.25	157	$9,930.25
RG	Rogers Comm	$8.56	1139	$9,752.69
	Cash			$11,004.89
	Total			$99,634.89

Table 6.1: Initial holdings of a hypothetical portfolio on January 4, 1999.

closing) so limit and stop orders execute based on closing prices.

I began with $100,000 on January 4, 1999 and decided to divide my assets uniformly among 9 stocks and one cash account. Table 6.1 lists my initial portfolio after commissions have been accounted for. The FIS goal was to keep one tenth of my total capital in each stock.

6.2 Choice of companies

My choice of companies was not completely random: I chose Apple Computers (AAPL), Advanced Micro Devices (AMD) and Netscape Communications (NSCP) because they were all "underdogs" in their respective industries and would have to be innovative and aggressive to survive. Similarly, I chose Eastman Kodak (EK) because, although currently a large, stable company, I expected the emerging digital camera technology to threaten its dominance and I wanted to see how it fared. I chose the cable company Rogers Communications (RG) because I was interested in the newly available cable modem technology which they were investing in. Microsoft Corporation (MSFT) seemed a low-risk choice to balance my high-risk portfolio. My focus to this point had been in the high technology sector so I determined to diversify: in the petroleum sector I chose Chevron Corporation (CHV) for its apparent low-risk and because it was the most recent gas station I had visited, and I chose Anglogold (AU) as a gold stock simply because of its catchy symbol. I couldn't think of a last company I was interested in so I let my wife choose Immunex Corporation (IMNX) from the biomedical sector. Although these choices were biased by my own interests it was hoped that they would prove sufficiently representative to test the performance of the fixed investment strategy. (As will be seen, this portfolio correlated strongly with the Nasdaq composite index.)

6.3 Friction

The derivation of FIS in Chapter 3 and other sources [56, 57] neglected commissions; they considered a completely fluid portfolio, capable of adjusting instantly to infinitesimal price changes. This market simulation was more realistic, with commissions which were handled as follows: in each trade a \$39.95 charge was levied for the first 1000 shares traded (bought or sold) and \$0.04 per share over 1000. Obviously, it would be unprofitable to trade on every minuscule price fluctuation so a *friction* f was introduced. Orders are not placed until a stock's price p exceeds a threshold as given by Eqs. 3.20–3.21.

6.3.1 Minimum friction

Under some particular conditions it is possible to estimate how large the friction needs to be for profitable trading in a commission-enabled market. To calculate the necessary scale of the friction consider an imaginary scenario: we begin with a total capital of w divided uniformly between a cash account and $N-1$ stocks, so the ideal investment fraction is $i = 1/N$. The scenario consists of

1. a single stock's price moving to a trade limit (buy or sell),
2. the stock being rebalanced (traded),
3. returning to its original price, and
4. being rebalanced again.

In this scenario we assume all the other stocks are unchanged, each maintaining a value iw. We are interested in what the minimum friction f_{min} can be such that we don't lose any money given an absolute transaction cost T.

We begin with the fluctuating stock at its ideal price

$$p^* = \frac{iw}{s}, \tag{6.1}$$

where s is the number of shares held of the stock.

So the limit prices are

$$p_{\pm} = p^*(1+f)^{\pm 1}, \tag{6.2}$$

where p_+ is the sell limit and p_- is the buy limit.

If the stock moves to one of the limits while all others remain constant, then our wealth (before trading) will become

$$w_\pm = (1-i)w + sp_\pm \tag{6.3}$$

and the quantity to be traded, from Eq. 3.15, will be

$$\begin{aligned} \Delta s_\pm &= s^*(p_\pm) - s & (6.4)\\ &= (1-i)s\left[(1+f)^{\mp 1} - 1\right] & (6.5)\end{aligned}$$

maintaining the same notation (the upper symbol of $\pm$ and $\mp$ indicates an initial rise in price, and the lower indicates an initial drop).

The trade also changes our cash holdings by

$$\begin{aligned} \Delta c_\pm &= -\Delta s_\pm p_\pm - T & (6.6)\\ &= -\Delta s_\pm p^*(1+f)^{\pm 1} - T & (6.7)\end{aligned}$$

where T is the transaction cost (in dollars).

Now we assume the stock's price returns to its original value p^* and we trade to recover our original portfolio $\Delta s'_\pm = -\Delta s_\pm$ (for simplicity), yielding another change in cash

$$\Delta c'_\pm = \Delta s_\pm p^* - T \tag{6.8}$$

so the net change is

$$\Delta\hat{c} = \Delta s_\pm p^*\left[1 - (1+f)^{\pm 1}\right] - 2T. \tag{6.9}$$

Inserting the computation for $\Delta s_\pm$ gives a net change

$$\begin{aligned} \Delta\hat{c} &= (1-i)iw\left[(1+f)^{\mp 1} - 1\right]\left[1-(1+f)^{\pm 1}\right] - 2T & (6.10)\\ &= (1-i)iw\left[f + \frac{1}{1+f} - 1\right] - 2T, & (6.11)\end{aligned}$$

regardless of whether the stock price rose then fell or fell and then recovered.

After some algebra we find the condition requiring a profit $\Delta\hat{c} > 0$ holds when $f > f_{min}$ where

$$f_{min} = \frac{TN^2}{w(N-1)}\left[1 + \sqrt{1 + \frac{2w(N-1)}{TN^2}}\right]. \tag{6.12}$$

which simplifies to

$$f_{min} \approx N\sqrt{\frac{2T}{w(N-1)}} \tag{6.13}$$

Date	Event
January 4, 1999	Experiment started
March 23, 1999	Takeover: 1 NSCP $\rightarrow$ 0.9 AOL
March 26, 1999	Stock split: 2-for-1 IMNX
March 29, 1999	Stock split: 2-for-1 MSFT
August 27, 1999	Stock split: 2-for-1 IMNX
November 11, 1999	Stock split: 2-for-1 AOL
December 10, 1999	Tolerance changed from 10% to $2f_{min}$
March 21, 2000	Stock split: 3-for-1 IMNX
April 14, 2000	NASDAQ correction
May 12, 2000	Experiment ended

Table 6.2: Events relating to the hypothetical portfolio which occurred during the course of the experiment.

in the limit $w \gg T$.

This informal derivation is only meant to set a scale for the minimum friction, it is not meant to be rigorous. A more detailed calculation may be possible by assuming each stock's price moves as geometric Brownian motion but the derivation would be cumbersome and the benefit dubious.

When I began trading with my hypothetical portfolio in the beginning of 1999 I arbitrarily chose $f = 10\%$, a fortuitous choice, as it turns out, because anything less than $f_{min} = 9.90\%$ might have been a losing strategy.

Since Eq. 6.12 sets the break-even friction it is best to set the actual friction somewhat higher. Once the minimum friction for my portfolio had been estimated (December 1999) I chose a dynamic value of $f = 2f_{min}$.

6.4 FIS Experimental results

In this section, the results of using the fixed investment strategy (with friction) on a hypothetical portfolio will be discussed.

6.4.1 Events

The experiment began on January 4, 1999 and ran until May 12, 2000 for a total of 343 trading days. The portfolio was rebalanced faithfully, as needed almost every day (excepting a few rare and brief vacations). Note that the simulation only executed trades after closing so intra-day trading was not supported and the trading price was always the stock's closing price. Also note that the simulation did not constrain orders to be in round lots and most orders, in fact, were odd sizes.

A list of important events occurring over the course of the experiment is shown in Table 6.2. The majority of events consist of stock splits, a division of the shares owned by each shareholder of a company such the stake held by each is unchanged. For example, a 2-for-1 stock split means each share is split into two, each worth half its original value. Although theoretically a stock split should not affect an investor's capital in a company, stock splits are considered good news and often drive the stock's price up both before and after the split. The main reason is that a split lowers the price of a stock and thereby makes it accessible to more potential investors—increasing demand.

Another interesting event which occurred during the experiment run was the takeover of Netscape Communications by America Online in March, 1999. AOL purchased Netscape for roughly \$4 billion and each share of Netscape stock was converted to 0.9 shares of AOL. Takeovers tend to engender a great deal of speculation which can precipitate large fluctuations in the stock's value.

On the book-keeping side, the only change in methodology was the move from a constant friction of $f = 10\%$ to a floating value of $f = 2f_{min}$, as given by Eq. 6.12, on December 10, 1999 (giving $f = 15\%$ at the time). The main consequence was a somewhat decreased trading frequency.

The most exciting event was the correction in the high-technology sector (which dominates my portfolio) in the week of April 14, 2000, evinced by an over-35% drop in the Nasdaq Composite index from its all-time high only weeks earlier [84]. This is a particularly fortunate occurrence because it tests the ability of the FIS to handle drawdowns. Market-wide fluctuations of this magnitude are rare but an important consideration when devising a trading strategy. This aspect of the experiment will be discussed below.

6.4.2 Performance

The final state of the portfolio at the end of the experiment is shown in Table 6.3. In this section the performance of the fixed investment strategy will be evaluated. As a control, a simple "Buy-and-Hold" Strategy (BHS) with the initial portfolio shown in Table 6.1 held fixed, will be contrasted with the fixed investment strategy (FIS).

Fig. 6.1, which shows the evolution of total capital for both strategies, on the surface seems to indicate BHS outperforms FIS. BHS reaches a high of \$237,000, a full 14% higher than the maximum achieved with FIS. Also, BHS maintained a higher capital on 292 of the 343 days (85%) the market was open. Evidently, FIS does not perform well in real-world applications.

However, a closer inspection suggests FIS should not be discarded too rashly.

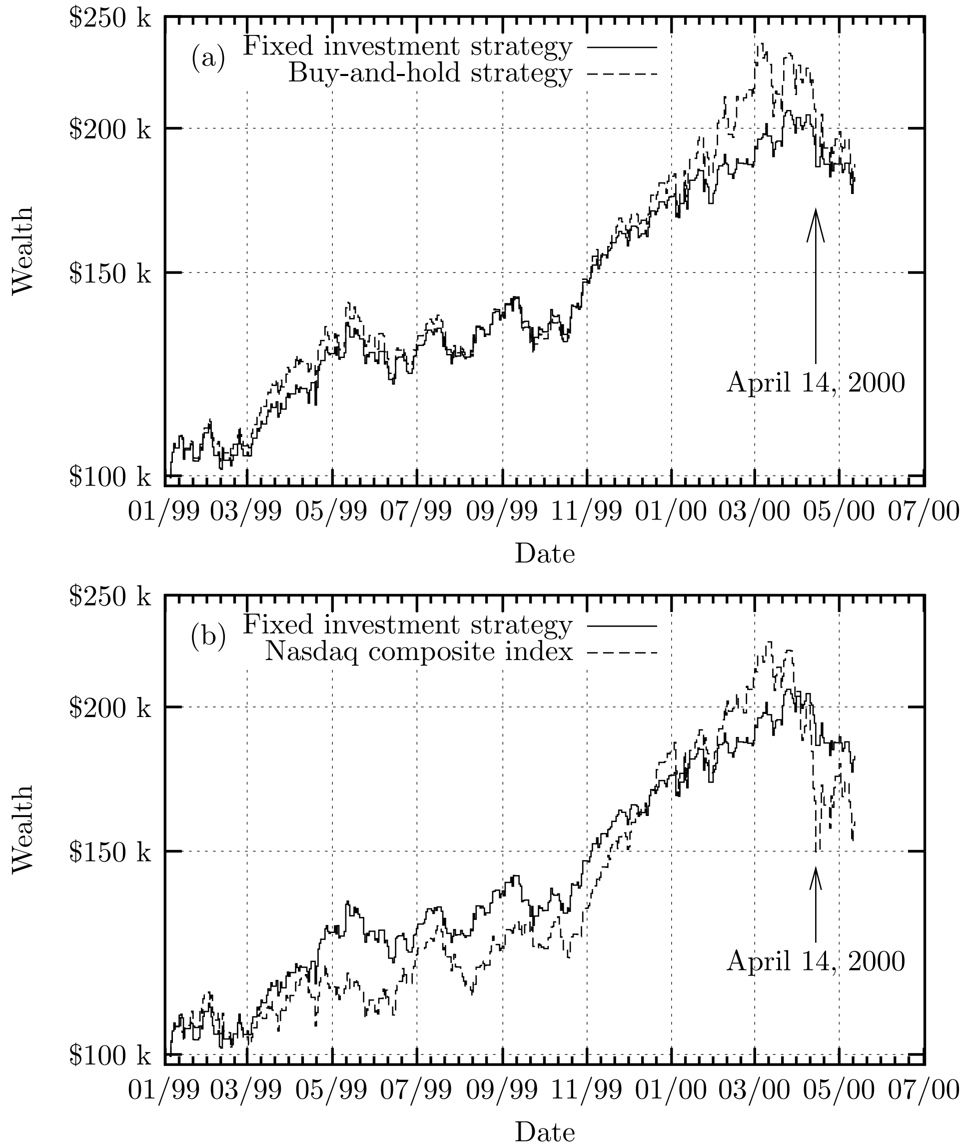

Figure 6.1: Historical wealth using FIS versus (a) the Buy-and-Hold strategy and (b) the Nasdaq Composite Index over the same interval (rescaled to be equal at the start of the experiment).

Symbol	Company	Price	Shares	Value
AAPL	Apple Comp Inc	$107.63	159	$17,112.38
AMD	Adv Micro Device	$85.69	216	$18,508.50
AOL	America Online Inc	$55.38	325	17,996.88
AU	Anglogold Ltd	$20.06	965	$19,360.31
CHV	Chevron Corp	$94.25	190	$17,907.50
EK	Eastman Kodak Co	$56.13	310	$17,398.75
IMNX	Immunex Corp	$35.00	456	$15,960.00
MSFT	Microsoft Corp	$68.81	269	$18,510.56
RG	Rogers Comm	$25.56	732	$18,711.75
	Cash			$18,792.39
	Total			$180,259.02

Table 6.3: Final holdings of a hypothetical portfolio on May 12, 2000.

For example, consider the market correction on and around the week of April 14, 2000. The Nasdaq Composite peaked at 5,049 points on March 10 and fell to a low of 3,321 on April 14, a drop of 34%. The Buy-and-hold strategy fared somewhat better, dropping to $180,000 for a drawdown of 25%. But the fixed investment strategy suffered the smallest decrease—down only 15% to $176,000, finishing with almost the same value as BHS. (It should be noted that Maslov and Zhang [85] demonstrated that the FIS is the most aggressive possible strategy that keeps the *risk*—measured as the expected drawdown from the maximum—bounded.)

This suggests FIS is less susceptible to large fluctuations. By rebalancing the portfolio, one moves capital out of stocks which may be overvalued and into safer companies which may be more resilient to perturbations. In this sense, FIS reduces risk.

This can be seen in Fig. 6.2 which demonstrates that BHS is more prone to large fluctuations (both positive and negative). To test whether these histograms are compatible with the Gaussian hypothesis [46] the first four moments of the distributions are calculated in Table 6.4. The moments of the daily returns of the Nasdaq Composite index and each of the stocks in the portfolio over the same interval are also shown.

The means for all the distributions are all small, negligible in comparison to the standard deviations (which had an average of 3% daily). The skewness, given by

$$\mathrm{Skew}(\{r_i\}) = \frac{1}{N} \sum_i \left[\frac{r_i - \bar{r}}{\sigma_r} \right]^3 \tag{6.14}$$

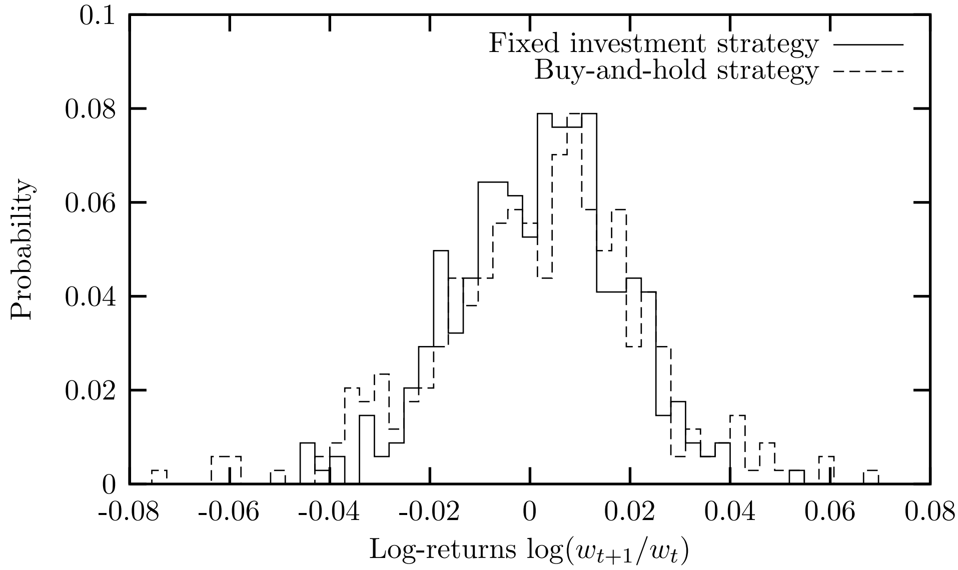

Figure 6.2: Histograms of log-returns of capital $r_{t+1} = \log(w_{t+1}/w_t)$ for both strategies. Notice BHS exhibits more large fluctuations (fatter tails) than FIS.

Symbol	Mean	Std. Dev.	Skewness	Kurtosis
FIS	0.00178	0.016	-0.18	0.12
BHS	0.00183	0.021	-0.23	0.82
Nasdaq	0.00119	0.022	-0.58	1.65
AAPL	0.00274	0.040	0.04	0.54
AMD	0.00342	0.046	0.23	2.76
AOL	0.00134	0.044	0.33	1.32
AU	0.00042	0.029	0.60	3.93
CHV	0.00050	0.020	0.39	0.95
EK	-0.00062	0.019	0.25	5.40
IMNX	0.00375	0.064	0.29	1.80
MSFT	-0.00010	0.029	-0.90	5.18
RG	0.00361	0.034	0.66	0.80
Averages	0.00166	0.032	0.08	2.11
Significance			0.21	0.53

Table 6.4: First four moments of the distribution of log-returns for each stock, the two trading strategies under review and the Nasdaq Composite Index. The skewness characterizes the asymmetry of the distribution and the kurtosis indicates the presence of outliers. The average skewness is not found to be significant but the kurtosis is.

where σ_r is the standard deviation of the returns, indicates the degree of asymmetry in the distribution. The sign of the skewness indicates which tail of the distribution contains more outliers. The skewness uncertainty for a (symmetric) Gaussian-distributed sampling of N points is $\sqrt{15/N}$ [20, Ch. 14], thereby setting a scale for deciding if a particular skewness is significant or not. Interestingly, Table 6.4 indicates that individual stocks tend to be skewed positively but the portfolios and the Nasdaq index are negatively skewed. This suggests that negative movements tend to be correlated between stocks but positive movements are not.

The (excess) kurtosis is a measure of the spread of the distribution. A positive kurtosis indicates the presence of many outliers or "fat tails". The kurtosis is defined as

$$\mathrm{Kurt}(\{r_i\}) = \left\{ \frac{1}{N} \sum_i \left[\frac{r_i - \bar{r}}{\sigma_r} \right]^4 \right\} - 3 \tag{6.15}$$

where 3 is subtracted in order to fix the kurtosis at zero for a normal distribution. The standard deviation of the kurtosis from a dataset of size N sampled from a Gaussian is $\sqrt{96/N}$ so the Gaussian hypothesis is rejected if the kurtosis if found to be greater.

The table shows that almost every calculated kurtosis is significant, indicating fat tails for these daily return distributions. The only exception is the fixed investment strategy which only has a kurtosis of 0.12 (versus a significance level of 0.53). This further confirms the hypothesis that FIS reduces risk: by rebalancing the portfolio the frequency of large fluctuations is reduced. (Another favourable consequence is that the return distribution appears to converge to a Gaussian, as was assumed in the derivation of FIS.)

Nevertheless, it is undeniable that BHS outperformed FIS in the experiment. But given the atypical trend seen in the portfolios, this conclusion may not be generalizable. Both portfolios realized almost a 100% growth over the first year, a gain which can hardly be expected to be repeated often (except, perhaps, during other speculative bubbles). A more typical realization may have proven FIS superior. By more typical is meant a smaller trend, relative to the scale of the fluctuations. As can be seen in Fig. 6.1 FIS performs best in the presence of fluctuations and slowly loses ground against BHS in the presence of an upwards trend.

Even if this experiment doesn't demonstrate that FIS is optimal it still shows that it is a reasonable investment strategy, and a suitable choice for the agents in DSEM (Chapter 3). The striking feature of FIS is that it reduces the kurtosis (risk of large events) and thereby misses large downturns (and upturns) in the market.

6.5 Log-periodic precursors

In this section the results of another experiment performed, using the same hypothetical portfolio, will be examined. At issue is whether there exists a reliable method to forecast imminent crashes in the market. First, some background theory is necessary.

6.5.1 Scale invariance

Scale invariance is a property of some systems such that a change of scale in a parameter $x' = \lambda x$ only has the effect of changing the scale of some observable $F' = F/\mu$, such that

$$F(x) = \mu F(\lambda x). \tag{6.16}$$

The above scaling relation has a power law solution $F(x) = Cx^z$ where

$$z = -\frac{\log \mu}{\log \lambda} \tag{6.17}$$

and C is an arbitrary constant.

The important point to notice is that the scalings along both axes are related by $\mu = \lambda^{-z}$. No matter how much the control parameter x is scaled by (even infinitesimally, $\lambda \to 1$), it is always possible to rescale the observable so that it is invariant. This is known as *continuous* scale invariance.

6.5.2 Discrete scale invariance and complex exponents

In contrast, *discrete* scale invariance only allows fixed-size rescalings of the parameter. To see how this comes about we begin by substituting the solution $F(x) = Cx^z$ into the renormalization equation (Eq. 6.16),

$$Cx^z = \mu C \lambda^z x^z \tag{6.18}$$

$$\Rightarrow 1 = \mu \lambda^z. \tag{6.19}$$

Now notice that $1 = e^{2\pi i n}$ for any integer n. Applying this and taking the logarithm of both sides gives

$$2\pi i n = \log \mu + z \log \lambda \tag{6.20}$$

which has the solution

$$z = -\frac{\log \mu}{\log \lambda} + i\frac{2\pi n}{\log \lambda}. \tag{6.21}$$

For the scaling relation to hold z must be a constant, which can only hold when $n = 0$ (which allows λ to take on any value, recovering continuous scale invariance) or when λ is some fixed constant, the *preferred* scaling ratio. Hence, the only invariant transformations are the discrete rescalings $x' = \lambda x$ with corresponding scalings in the observable $F' = F/\mu$ (for some fixed μ).

6.5.3 Log-periodic precursors

So far this might all look like mathematical trickery to the reader but the theory does have testable consequences. If we use the notation $z_n = \alpha + i\omega_n$ with

$$\alpha = -\frac{\log \mu}{\log \lambda} \tag{6.22}$$

$$\omega_n = \frac{2\pi n}{\log \lambda} \tag{6.23}$$

then $F_n(x) = C_n x^\alpha x^{i\omega_n}$ is a solution to the scaling relation for each n and the general solution is the linear combination over all integers n,

$$\begin{aligned} F(x) &= x^\alpha \sum_n C_n x^{i\omega_n} && (6.24) \\ &= x^\alpha \sum_n C_n \exp(i\omega_n \log x) && (6.25) \\ &= x^\alpha \left[C_0 + e^{i\omega \log x} \sum_{n \neq 0} C_n \exp(i\omega(n-1)\log x) \right] && (6.26) \end{aligned}$$

where we have defined $\omega = 2\pi/\log \lambda$, for convenience.

The final form of $F(x)$ indicates that the function has a periodic component with angular frequency ω. Expanding the periodic component as a Fourier series gives, to first order,

$$F(x) \approx x^\alpha \left[C_0 + C_1' \cos(\omega \log x + \phi) \right] \tag{6.27}$$

where ϕ is an unknown phase constant.

This argument, a variation of those presented in Refs. [13, 86, 87], concludes that discrete scale invariance leads naturally to log-periodic (in x) corrections to the scaling function F.

6.5.4 Critical points

Near a critical point many properties of a system exhibit power law scaling relations as described above. Therefore they are prime test-cases for the existence of complex exponents characterized by log-periodic precursors.

Seismicity, studied in the context of critical phenomena, have been successfully modeled as self-organizing (with the build up of stress) to a critical point in time t_c characterized by an earthquake, a sudden release of energy [13,88,89]. In the neighbourhood of the critical time (small $|t_c - t|$) the stress exhibits classic power laws seen in critical phenomena. One important goal in seismology is forecasting the time of occurrence t_c of large earthquakes. It has been argued that log-periodic fluctuations are present both before (foreshocks) and after (aftershocks) large events and that the precursors improve earthquake forecasts considerably [86, 90, 91].

The premise is that the rate of change of the regional strain ϵ exhibits critical scaling near the critical point,

$$\dot{\epsilon} = F(|t_c - t|) \tag{6.28}$$

so that the strain (a measurable quantity) obeys

$$\boxed{\epsilon = A + |t_c - t|^{\alpha+1} \left[B + C\cos(\omega \ln|t_c - t| + \phi)\right] \qquad (6.29)}$$

in the vicinity of t_c. The curve is fit to known data by tuning the seven model parameters (A, B, C, t_c, α, ω, and ϕ) and the forecast of t_c is read off from the best fit to the data. This method has significantly improved precision over curve fits neglecting log-periodicity ($C = 0$), validating the adoption of the three extra parameters.

6.5.5 Application to financial time series

The same group of researchers who developed the concept of log-periodic precursors in seismology have recently turned their attention to the stock market, arguing that market crashes should be predictable by the same methodology [1, 84, 92].

Johansen et al. [92] construct a theory for price fluctuations with the risk of crash such that price series obeys precisely the relationship given in Eq. 6.29. The basic argument is that stock prices enjoy exponential growth but with some risk of crash. As time progresses the risk accumulates and the exponential growth rate increases to compensate for the risk (to remunerate rational investors for their risk). At some point in time the risk diverges and a critical point emerges.

The fundamental component of the theory is that the instantaneous risk of crash (which they call the "hazard rate") is assumed to obey a scaling relation like Eq. 6.16 with a control parameter $t_c - t$. Since it is related to the rate of change of the price, Eq. 6.29 arises.

In theory then, financial crashes should be predictable by curve fitting to the price series. In practice, though, this is an extremely difficult task: while searching through the seven-dimensional parameter space for the optimum fit one often gets stuck in local optima, missing the global one. This complaint has been raised against the theory [93] and is acknowledged by Johansen et al. [84].

Another problem with the research is that the experiments are all performed on known crashes *after they have occurred*! This introduces two problems: Firstly—with no disrespect intended—it may bias the results. If one knows there was a crash at such-and-such a time it would be very difficult to be satisfied with a curve-fit which made no such prediction. One would probably suspect the parameters were stuck in a local minimum and tweak them. This is perfectly natural but without foreknowledge one might have accepted the results without prejudice.

Secondly, all the curve fits were performed around well-established crashes. It would be as useful to test the theory during other periods when no crashes occur in order to test for "false positives." If the theory predicts too many crashes when none actually occur it is of no use.

In order to avoid these pitfalls I conducted a "blind" experiment to test the ability of Eq. 6.29 to forecast crashes. As I was already running my FIS experiment it was convenient to use it as my input data for the curve fit. Instead of forecasting a crash in a single stock, then, I was attempting to forecast a crash in a portfolio of nine stocks (and one cash account). But this is not seen as problematic since the other studies used composite market indices instead of individual stocks, as well [1, 84, 92, 93].

6.5.6 Experimental design

The experiment consisted of collecting portfolio wealth data w_t and fitting the curve given by Eq. 6.29 with $\epsilon = w$. The experiment began on February 14, 2000 and ran through May 12, 2000 but the dataset used was the entire historical set from the FIS experiment (which began on January 4, 1999).

The dataset consisted of sets of date-wealth pairs which were only collected on days when a trade was executed. At the beginning of the experiment this consisted of 113 points which grew to 134 points by the close of the experiment.

The fitting over the seven model parameters was performed using Microsoft Excel's *Solver Add-in* which uses the Generalized Reduced Gradient (GRG2) nonlinear optimization technique [94]. The GRG2 method is suitable for problems involving up to 200 variables and 100 constraints. The optimization condition was the minimization of the sum of the squared deviations (χ^2 nonlinear least-squares fitting).

The fit was performed on a logarithmic price scale on the basis that it is the relative (fractional) fluctuations in capital which are fundamental, not the absolute variations. Fitting on a linear scale, then, would significantly bias the curve to fit better at greater wealths at the expense of the fit at lesser wealths. So the fit consisted of minimizing

$$\chi^2 = \sum_t [\ln w(t) - \ln w_t]^2 \tag{6.30}$$

where w_t is the actual wealth at time t and $w(t)$ is the fitting function as given by Eq. 6.29.

The *Solver* routine did not provide a measure of the quality of the fit or an estimate of the fitted parameters' uncertainties but an estimate of the quality is provided by the χ^2 measure itself. If the fit is of high quality then the data should be randomly distributed around the curve with a total χ^2 variance proportional to $N-7$ [20, Ch. 15]. Hence, the ratio $(N-7)/\chi^2$ should be independent of the number of data points N acquired. A small value indicates a large χ^2 variance and a poor fit, while a large value indicates a good fit. Hence, the *quality* of the fit Q, defined as

$$Q \equiv \frac{N-7}{\chi^2}, \tag{6.31}$$

is a dimensionless (strictly positive) quality which increases as the fit gets better. Using Q allows us to compare fits at different times t with different amounts of data N. Note that the parameter Q is only useful so far as ordering the fits: if $Q_i > Q_j$ for fits i and j (possibly at different times) then i is a better fit—more likely to explain the data and with more meaningful parameter values (in particular, the forecasted crash date t_c).

Every day of the experiment a new data point was recorded (if a trade had been executed) and then the curve was refit to the dataset generating a new forecast for the next market crash t_c. The forecasted date of the crash, date the forecast had been generated and the quality of fit Q were then recorded. A new forecast was made everyday, even in the absence of new data, because the critical time t_c was constrained to occur in the future.

Nonlinear curve fitting is basically a parameter space exploration which depends crucially on the initial choice of parameters. The initial parameter set, at the initiation of the experiment, was chosen by first trying to establish a good power law fit (with $C = 0$) and then refitting over all seven parameters. Subsequent fits all began with parameter values that were produced by the last fit with one important exception: the critical point t_c was always initialized to be the current day. The motivation was to avoid getting stuck in sub-optimal solutions at later times and miss an impending crash. It was preferable to impose a bias towards imminent events. It is still possible to converge to sub-optimal solutions but it was decided that false positives were preferable to false negatives.

It is important to stress that all the forecasts from this experiment were true predictions, tabulated as the experiment progressed for analysis later. The

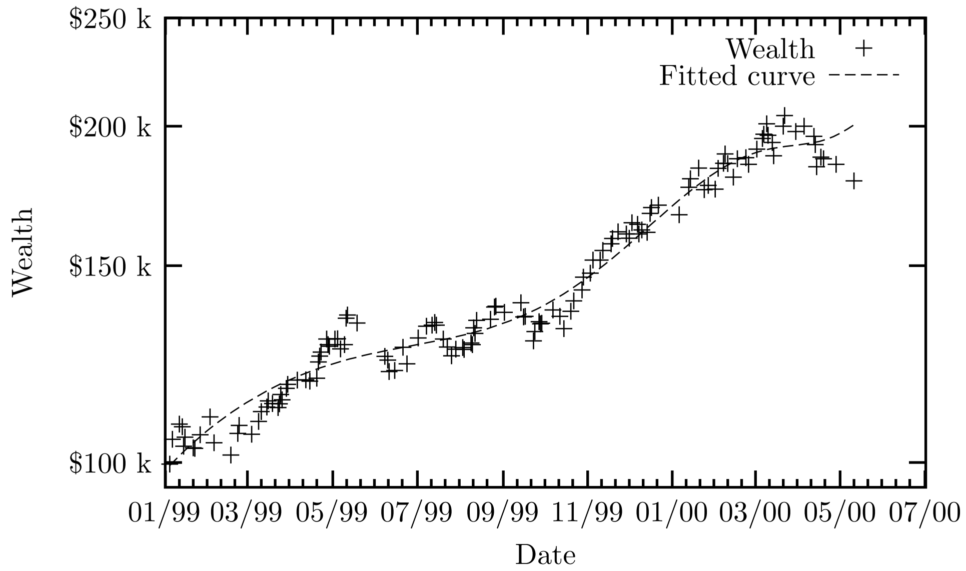

Figure 6.3: Sample fit of Eq. 6.29 to portfolio wealth on May 12, 2000. The best fit parameters indicate a crash is anticipated on or around t_c =July 4, 2000.

calculations were not performed after-the-fact so the results are not biased by foreknowledge.

6.5.7 Results

Each day, a fit of Eq. 6.29 to the portfolio wealth data was performed giving a fitted curve similar to the one shown in Fig. 6.3 and the value of t_c the fitting procedure converged upon was interpreted as a forecast of the time of the next crash.

The experiment ran for 63 (week-)days and predicted a remarkable 30 crashes in that period. Obviously, the theory predicts too many false positives. However, it may still have some merit if the false positives have some correlation with returns.

The dates of forecasted crashes and their actual returns (fractional change of wealth) are plotted in Fig. 6.4. Notice the increased numbers of forecasts for a crash in early April, in agreement with the observed decline on the 14th. However, besides that, it is difficult to discern any pattern from the graph so some further statistical analysis is in order.

We want to determine if there is a statistically significant signal in the returns on the days the market was forecasted to crash versus other days so the dataset is split into two: "Forecasted" and "Not Forecasted." The mean returns and standard deviations were computed for both data sets (and the entire dataset) as shown in

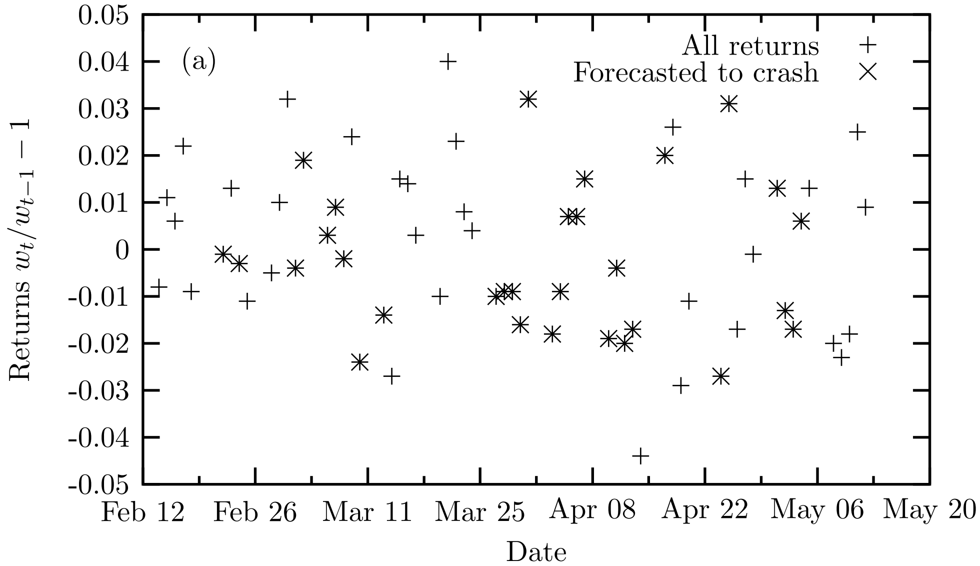

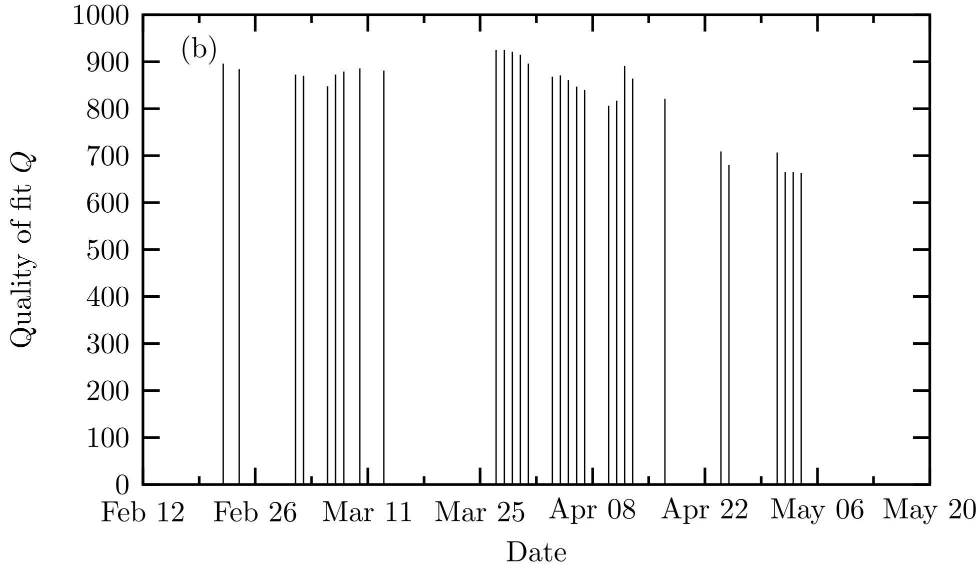

Figure 6.4: Daily wealth returns ($w_t/w_{t-1} - 1$) are shown along with the dates forecasted to crash in (a). The qualities of the curve fits corresponding to the forecasted crashes, which suggest the reliability of the predictions, are shown in (b).

Data set	Mean return	Std. Dev.
All data	0.0%	1.8%
Forecasted	−0.2%	1.5%
Not Forecasted	0.2%	1.9%

Table 6.5: Average values and standard deviations of the daily portfolio returns ($w_t/w_{t-1} - 1$) for all data and separately for days a crash was forecasted and not forecasted.

Data set	Mean return	Std. Dev.
Forecasted	−0.4%	1.4%
Not Forecasted	0.5%	1.9%

Table 6.6: Same as Table 6.5 except only including data up until the observed decline on April 14.

Table 6.5. There does appear to be a small deviation between returns on forecasted days versus not-forecasted days but the deviation is insignificant when compared to each dataset's standard deviation.

The likelihood that the two datasets come from the same underlying distribution can be calculated using the Kolmogorov-Smirnov (K-S) test which compares the cumulative probability distribution of the two samples [20, Ch. 14]. The calculation estimates a 24% chance that the underlying distributions are the same, which is still statistically significant so the log-periodic precursor prediction method is not conclusive.

Recall that there was a (fortuitous) correction in the markets on and around April 14 as indicated by Fig. 6.1(b). It would be interesting to know whether this forecasting method "saw it coming." Interestingly, Johansen and Sornette submitted a paper to the LANL preprint archive on April 16 claiming to have predicted, as early as March 10, a major event between March 31 and May 2 [84].

To test whether the crash around April 14 was predictable the data from Fig. 6.4 are reused neglecting everything after April 14. (Notice the quality of the fit declined markedly after April 14, suggesting the reliability of the later predictions is dubious.) The average returns and their standard deviations for both "Forecasted" and "Not forecasted" dates is again shown in Table 6.6, with a somewhat more significant difference betweens the means (the standard deviations are almost unchanged). Applying the K-S test now yields a much less significant 11.7% likelihood that the datasets are samples from the same underlying distribution.

In conclusion, it appears that there may be some value in interpreting market crashes as critical phenomena with log-periodic precursors but the predictive advantage of doing so is limited. The main difficulty lies in fitting seven nonlinear model parameters to a given dataset—often the fitting algorithm converges to a suboptimal solution, thereby forecasting an erroneous crash date t_c.

6.5.8 Universality of scaling ratio

In this section an open problem in the theory of log-periodic precursors will be presented.

It has been observed that the scaling ratio λ in Eq. 6.29 seems to be universal, almost always converging to a value near $\lambda \approx 2.5 - 3.0$ [1, 84, 90–92]. (Note this corresponds to a universal log-periodic frequency $\omega = 2\pi/\ln\lambda \approx 5.5 - 7.0$.) The emergence of a universal scaling ratio has come as a surprise to researchers [84, 92] since it describes some natural hierarchy within the specific system of interest and is not expected to be general.

Another peculiarity is that log-periodic fluctuations occur in some systems which do not have an obvious discrete scale invariance, such as the stock markets. In the derivation of log-periodicity, discrete scale invariance was a fundamental ingredient, without which it did not emerge. Why then might markets, which are not suspected to have any discrete scale invariant structures, exhibit log-periodicity?

It is my belief that these two idiosyncrasies are tied together: with the lack of a preferred scaling ratio a natural ratio is chosen, Euler's constant, $\lambda = e \approx 2.72$. The log-periodic frequency is then $\omega = 2\pi \approx 6.28$, in agreement with observation. A mechanism that might produce this preferred scaling ratio is unknown and this issue is only discussed here to generate interest in the problem. The discovery of a mechanism whereby ω is fixed would be a great boon to forecasting because this parameter is one of the most problematic for the optimization routine. (Incidentally, fixing $\omega = 2\pi$, the next crash (in the hypothetical portfolio) is forecasted to occur in the third week of October, 2000.)

6.6 Summary

In this chapter two experiments were performed with a hypothetical portfolio of nine stocks. In the first experiment it was observed that the fixed investment strategy (FIS) performs sufficiently well to justify its application in the Decentralized Stock Exchange Model (DSEM). Although it underperformed when compared to a trivial "Buy-and-hold" strategy, this is attributed to the strong upward trend in the portfolio over the course of the experiment. In each case when the climb was interrupted the FIS managed to "catch up to" and surpass the Buy-and-hold strategy, only to lag behind again when the trend re-emerged. The FIS also had the favourable property that it significantly reduced the kurtosis of the distribution of returns, essentially taming the largest fluctuations. This may be relevant to derivative pricing theory [65] which assumes Gaussian-distributed increments with no excess kurtosis.

The second experiment tested a method for forecasting financial crashes. The method relies on log-periodic oscillations in the price series which accelerate as the time of the crash approaches. The data suggest that log-periodic precursors probably do exist but they offer little, if any, prediction advantage because the method requires solving an optimization problem involving seven nonlinear parameters. Thus, the optimization procedure tends to get stuck in local, sub-optimal regions of the parameter landscape, frequently producing false-positive forecasts. It would be interesting to discover whether stochastic optimization techniques, such as simulated annealing [61], could provide better forecasts.