
THE CLUSTERING OF EXTREME MOVEMENTS: STOCK PRICES AND THE WEATHER

Atanu Saha^a, Burton G. Malkiel^b and Alex Grecu^c

A robust finding in this paper is that extreme movements in stock prices and temperature are usually preceded by large average daily movements during the preceding three-day period. This suggests that investors might fashion a market timing strategy, switching from stocks to cash in advance of predicted extreme negative stock returns. In fact, we have been able to simulate market timing strategies that are successful in avoiding nearly eighty percent of the negative extreme move days, yielding a significantly lower volatility of returns.



1 Introduction

There is considerable empirical evidence suggesting that the random walk model for changes in stock prices can be rejected and that the distribution of stock returns is distinctly non-normal. Researchers have observed that the key deviation from normality stems from the existence of “fat tails”: there are

a significantly higher number of extreme values—both negative and positive—than would be found in a normal distribution. This paper deals with the stock-market returns that reside in the fat tails. Its main objective, however, is not to present further evidence on higher-than-normal frequency of days with large changes. Instead, we focus on examining the duration between one extreme-move day and the next. We use duration-model analysis to examine the factors that are associated with the onset of days with unusually large changes in the Dow Jones Industrial Index for the years 1900 through 2006. We undertake a similar analysis in examining the duration between days with extreme movements in temperature using New York Central Park data for the years 1901 through 2006. We find striking similarity in the patterns of extreme change in the Dow

^aManaging Director, AlixPartners, 9 West 57th Street, Suite 3420, New York, NY 10019; USA. Tel.: 212 297 6322; e-mail: asaha@alixpartners.com

^bProfessor, Princeton University, 26 Prospect Avenue, Princeton, NJ 08544; USA. Tel.: 609 258 6445; e-mail: bmalkiel@princeton.edu

^cManager, Huron Consulting Group, 1120 Avenue of the Americas, 8th floor, New York, NY 10036; USA. Tel.: 212 785 1313; e-mail: agrecu@huronconsultinggroup.com

and in New York temperature. We also find that the existence of extreme-move days is at least partially predictable. Finally, we ask whether the predictability of extreme-value changes in the stock market can be employed to develop a useful portfolio-timing strategy.

We make no attempt to propose a stochastic process for stock returns or temperature changes. The modest goal of this paper would be accomplished if it triggers further inquiry into the question whether there exists a stochastic process common to many random phenomena, including stock market returns.

2 The relevant literature

The idea that no useful regularities exist in security prices that would enable investors to earn “abnormal” profits is one that goes back for more than a century. While Paul Samuelson (1965) is often associated with the general idea that stock prices change randomly, the random walk thesis dates back at least to the time of Bachelier (1900). Bachelier’s doctoral dissertation developed a theory of stochastic processes that was applied to prices of French government bonds. Bachelier found that price changes looked very much like a random walk process. His work lay dormant for almost 60 years until it was discovered by Paul Samuelson.

Essentially Samuelson argued that economists should expect price changes to be random. The profit seeking behavior of investors should eliminate any predictable movement in stock prices. Samuelson defines “properly intercepted prices” as prices that at every date $t \leq T$ are based on all available information at Φ_t , including all past price realizations for the security. If the security has a single payoff X_t , assumed to be a random variable, then for all $t \leq T$:

$$P_t = E(X_t/\Phi_t).$$

Samuelson then proved that prices will fluctuate randomly since for all $t \leq T$:

$$P_t = E(P_{t+1}/\Phi_t), \quad \text{or} \quad E(\Delta P_{t+1}/\Phi_t) = 0.$$

If prices are properly anticipated, then all useful information contained in past price series will be incorporated into current prices. Prices will follow a martingale and successive price changes will be uncorrelated.

Much of the subsequent empirical work on stock prices has found that the stock market does not meet the conditions for a random walk. Lo and MacKinlay, in their book *A Non-Random Walk Down Wall Street*, are able to reject the random walk hypothesis. They define X_t as the log of P_t and assume prices have the recursive relationship

$$X_t = \mu + X_{t-1} + \varepsilon_t,$$

where μ is an arbitrary drift parameter and ε_t is the random disturbance term. The traditional random-walk hypothesis restricts the ε_t s to be independently and identically distributed normal random variables with a constant variance. Lo and MacKinlay are able to reject the random-walk hypothesis for weekly stock returns (of both indices and individual stocks) using a simple volatility-based specification test. Not only do they report significant positive serial correlation for weekly and monthly holding period returns, they further conclude that mean reverting models of Shiller and Perron (1985) and Fama and French (1987) cannot account for the departures of returns from random walk.

There is considerable evidence of several non-random patterns in the movement of stock prices. For example, Keim and Stambaugh (1985), Cross (1973), French (1980), Gibbons and Hess (1981), and Rogalski (1984) find evidence of day-of-the-week and weekend effects on stock returns, while Bremer and Sweeny (1991) find extremely large negative 10-day returns are followed on average by larger-than-expected positive rates of return over the

following days. But none of this work appears to contradict the efficient-market hypothesis (EMH). Whatever predictable patterns that do exist do not appear to be large enough (or dependable enough) to permit an investor to make abnormal returns after paying transactions costs. Therefore, the evidence against the random walk hypothesis is not inconsistent with the practical implications of EMH.

Another predictable pattern in the behavior of stock prices is the subject of this paper. Work by Osborne (1963), Alexander (1964), and Mandelbrot (1963) finds that the occurrence of transactions in a given stock is not independent of the past history of trades in that stock. There is a clustering of activity. Trading tends to come in “bursts.” If recent trading activity is heavy, it is highly likely to continue to be heavy. Similarly, large price changes are likely to be followed by large changes in prices; as Mandelbrot (1963) writes, “...large changes tend to be followed by large changes—of either sign—and small changes by small changes...”.

In this context, the ARCH model (Engel, 1982) and its generalization to the GARCH (Bollerslev, 1986) specifications are relevant. These models set conditional variance equal to a constant plus a weighted average of past squared residuals. These models (as well as their numerous generalizations) have been proposed to explain two fundamental characteristics of stock returns: the presence of a surprising large numbers of extreme values and the fact that both the extremes and quiet periods are clustered in time. Despite the ubiquitous presence of GARCH models, however, many financial empiricists have observed that volatility responds asymmetrically to the nature of news—volatility tends to rise in response to “bad news” and to fall in response to “good news,” contrary to the predictions of GARCH models (Nelson, 1991). GARCH models assume that only the magnitude and not the sign of unanticipated excess returns determines future

volatility. Nelson (1991) recognized that large price declines forecast greater volatility than similarly large price increases. Since then numerous studies have found evidence of nonlinearity, asymmetry, and long memory properties of volatility (Engel, 2004).

Our paper complements this body of literature. To our knowledge ours is the only study that examines the temporal properties of “fat tails.” We analyze the time duration between extreme daily movements in the Dow Jones Index. Consistent with Nelson’s findings, we find that today’s extreme return is indeed a powerful predictor of the next extreme value day. Additionally, days with extreme returns are preceded by large moves, both for negative and positive returns, although the relationship is asymmetric. Similarly, there is a temporal asymmetry—extreme value days are preceded by but are not followed by large moves days. Interestingly, we find a remarkably similar pattern in extreme changes in New York temperature (measured by the percent deviations from daily normal temperature). We undertake various tests to confirm the robustness of the results.

Finally, we explore whether the negative extreme value days are predictable enough to devise a useful asset allocation portfolio strategy. We ask whether an investor could unwind a long position in the stock market (as measured by the Dow Index) when the model predicts a high likelihood of a negative extreme-move day and temporarily hold cash. After a fixed period over which the portfolio remains unwound, the investor is assumed to re-initiate a long position in the Dow Index. We demonstrate that this portfolio strategy does not outperform a buy and hold investment in the Index when one allows for reasonable transaction costs. However, this asset allocation strategy is successful in avoiding nearly eighty percent of the extreme negative move days, yielding a significantly lower volatility of returns.

3 The data and descriptive statistics

In this section, we describe the data used in the analysis of daily stock-market returns and NY temperature changes. We also provide some simple statistics on the magnitude and frequency of the extreme-move days to motivate the duration analysis that follows in the next section.

We employed daily data on the Dow Jones Industrial Index for all trading days in the period between January 1, 1900 and December 31, 2006. Panel A

of Table 1A provides some relevant statistics about the daily returns, measured in terms of logarithmic changes (i.e., $\ln(I_t/I_{t-1})$, where I_t is the value of the Index on day t).

Our sample consists of 26,884 daily returns. For this sample, the average daily return is 0.0195%, and the standard deviation (σ) of daily returns is 1.131%, which translates to an annualized¹ mean of 7.02%

Table 1A Dow Jones daily stock returns compared with draws from a normal distribution.

<i>Panel A: Comparison of Frequencies of Days</i>					
Dow Jones daily returns: 1900–2006			Random draws from a normal distribution		
# of Sigmas from the mean	Number of observations	Percentage of total	# of Sigmas from the Mean	Number of observations	Percentage of total
Greater than 1	5,203	19.4%	Greater than 1	8,568	31.9%
Greater than 2	1,217	4.5%	Greater than 2	1,208	4.5%
Greater than 3	427	1.6%	Greater than 3	70	0.3%
Greater than 4	183	0.7%	Greater than 4	1	0.0%
Greater than 5	82	0.3%	Greater than 5	0	0.0%
Greater than 6	52	0.2%	Greater than 6	0	0.0%
Sub-Total	7,164		Sub-Total	9,847	
Total Trading Days	26,884		Total Observations	26,884	

<i>Panel B: Time Interval Between 3-Sigma Plus Days</i>					
Dow Jones daily returns: 1900–2006			Random draws from a normal distribution		
Time interval in days	Number of observations	Percentage of total	Time interval in days	Number of observations	Percentage of total
1	76	17.8%	1	0	0.0%
2	57	13.4%	2	0	0.0%
3 to 5	96	22.5%	3 to 5	0	0.0%
6 to 25	93	21.8%	6 to 25	4	5.8%
Subtotal	322	75.6%	Sub total	4	5.8%
Greater than 25	104		Greater than 25	65	
Total	426		Total	69	

Note: A “3-Sigma Plus day” is one where the daily absolute return of Dow Jones Average is greater than or equal to 3-sigmas from the mean daily returns.

and sigma of 21.46%. To provide a benchmark, we generated 26,884 pseudo returns by drawing random numbers from a normal distribution with the same mean and sigma as in the sample of returns of the Dow Jones Index.

Panel A of Table 1A shows that in about nearly eighty percent of the trading days—80.6% to be precise—the Dow's daily change was less than one sigma away from its mean daily return. In the random draws from the normal distribution only 68.1% of the observations were less than one sigma away from the mean. However, the contrast between the two distributions (Dow returns and draws from a normal distribution) is most pronounced for days where the returns are three or more sigmas away from the mean. Extreme value days are far more prevalent for the Dow than for the normal distribution. These findings are consistent with the evidence of “fat tails” noted in numerous prior research studies.

In Panel B of Table 1A, we examine the daily time interval between one extreme-move day to the next. Here we define an “extreme-move” day as one where the absolute value of the daily return is more than three sigmas away from the mean (which we will call a “ $3\sigma+$ day”). Interestingly, of the 426 time intervals between the $3\sigma+$ days, 322 had intervals less than or equal to 25 days; and more than 31% of the $3\sigma+$ days occurred within two days of each other! This clustering of extreme-move days for the Dow is very different from what is found from random draws from a normal distribution. If the stochastic process underlying the Dow's return were normal, one would have observed less than 6% of extreme-move days to have a time interval of less than or equal to 25 days; by contrast, for the Dow this percentage is 76%.

We further explore the phenomenon of the clustering of extreme-move days for the Dow in Panel A of Table 1B. The average duration in days (i.e. the time interval) between one $3\sigma+$ day and the next

for the entire period, 1990–2006, was 178 days (the median was 116 days). Excluding the 1930s, this average rises to 194 days (the median is 126 days). Of the 427 $3\sigma+$ days, more than half occurred in 1930s. Indeed, in the 1930s, 9% of all trading days were characterized by returns greater than three sigmas from the mean. By contrast, the 1950s were years of calm: less than 1/4 of 1 percent of trading days were ones with extreme moves.

The clustering of $3\sigma+$ days over the decades is illustrated in Figure 1. The figure suggests several interesting observations. First, there seems to be a considerable clustering of extreme-change days. We note that the greatest clustering took place during the late 1920s and early 1930s, reflecting the boom and bust of the stock market and the depression that followed. Another clustering is apparent around the recession (depression) of economic activity in the late thirties and the uncertainty leading up to World War II. The most recent clustering occurred during the late 1990s and early 2000s associated with the high-tech internet bubble, the busting of the bubble, and the recession of the early 2000s. Note the existence of a few outliers of extremely large one-day declines in the stock market, of which the crash of October 19, 1987 stands out. Finally, note that the picture of large positive and negative changes is generally symmetric, with the exception of the aforementioned highly unusual steep declines, and that there is no evidence of increasing volatility over time. Indeed, there have been large gaps of several years during which no three-sigma events occurred.

We also gathered data on the daily temperature in New York's Central Park² for the years 1901 to 2006. We then defined the temperature's daily “change” as the percent deviation from the “normal” temperature for that day. In calculating the daily normal temperature, we used the 106-year average for any given day. Our sample consists of 38,655 observations of daily percent deviations from normal temperature. The standard deviation (σ) for the

Table 1B Extreme-move days for the Dow Jones and New York temperature: Number of days by decade.

Panel A				Panel B			
Dow Jones daily returns: 1900–2006				New York daily average temp's % Deviation from normal: 1901–2006			
Decade	Number of 3 -sigma plus days	% of trading days	Average duration between 3-sigma plus days	Decade	Number of 3-sigma plus days	% of Calendar days	Average duration between 3-sigma plus days
1900s	37	1.47%	62	1900s	11	0.34%	292
1910s	28	1.16%	93	1910s	30	0.82%	121
1920s	43	1.72%	59	1920s	13	0.36%	279
1930s	225	8.94%	11	1930s	27	0.74%	124
1940s	17	0.67%	135	1940s	15	0.41%	264
1950s	4	0.16%	567	1950s	13	0.36%	228
1960s	5	0.20%	309	1960s	20	0.55%	181
1970s	11	0.43%	343	1970s	23	0.63%	189
1980s	24	0.94%	116	1980s	21	0.57%	175
1990s	11	0.44%	207	1990s	30	0.83%	108
2000s	22	1.25%	51	2000s	17	0.66%	171
Total	427			Total	220		
Average		1.58%	178	Average		0.57%	194
Average w/o 1930s		0.84%	194				

Note: A “3-Sigma Plus day” is one where the daily absolute return of Dow or the absolute deviation from normal temperature is greater than or equal to 3-sigmas.

sample of daily temperature changes is 1.46%. We then defined an “extreme-move” day as one where the percent deviation in temperature exceeded three sigmas away from the mean. As we did for the stock market, we call these $3\sigma+$ days.

Panel B of Table 1B contains data on the occurrence of and the duration between $3\sigma+$ days for New York temperatures. Comparisons of Panels A and B reveal some interesting differences and similarities. Unlike the Dow, extreme-move days for NY temperature are more evenly distributed across the decades. For the years 1901–2006, there were a total of 220 extreme movement days in temperature. While temperature change extremes occurred only about half as often as extreme movements for the Dow, such movements are far more comparable

if the 1930s are excluded for the Dow. For NY temperature, the average duration in days between one $3\sigma+$ day to the next is 194 days, which is higher than the 178 days figure for the Dow. However, for the Dow, if one excludes the 1930s, the average duration between $3\sigma+$ days becomes 194, exactly the same as the average duration between extreme value days for the NY temperature.

4 Duration analysis

In this section, we explore the factors that may explain the onset of an extreme-move day. In particular, we examine whether extreme-value days—both for the Dow and the NY temperature—are preceded and followed by larger-than-average

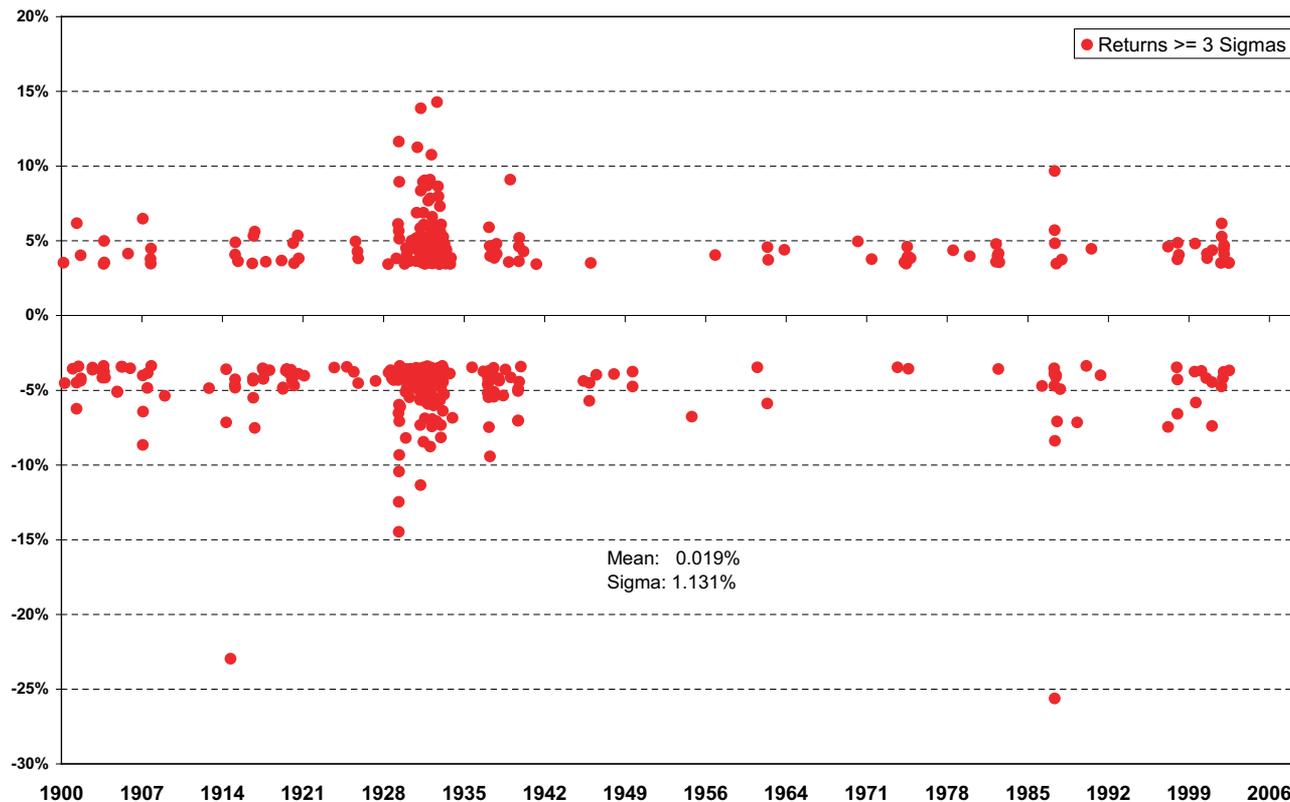


Figure 1 The Clustering of extreme movements in stock prices: Daily returns of Dow Jones index on large-move days, 1900–2006.

changes. We begin with a simple regression analysis of the time between extreme value days and then undertake a more formal duration analysis using the Cox Proportional Hazard model.

The key explanatory variables we consider are: (i) the average (of the absolute value of returns for the Dow and the percent deviations from normal for temperature) daily move in the three-day interval immediately *preceding* a $3\sigma+$ day; (ii) the average daily move in the three-day interval immediately *after* a $3\sigma+$ day; and (iii) the percentage change, in absolute terms, on a $3\sigma+$ day.³

The important statistics (mean and standard deviation) for these three variables are presented in Table 2. As a benchmark for comparison we also present the statistics for the absolute value of returns and the daily deviation for the entire sample of

26,884 observations for the Dow and 38,655 for the NY temperature. The first row of this table shows that the mean absolute value of the Dow's daily returns is 0.743% for the entire sample and the standard deviation is 0.853%. These figures differ from the ones presented in Table 2 because here we report the mean and standard deviation of the *absolute* value of the returns.

The next three rows of Table 2 contain summary statistics on the three explanatory variables. The mean and standard deviation for these three variables are based on 427 observations for Dow and 220 observations for NY temperature, which corresponds to the number of $3\sigma+$ days for each.

Comparison of the figures in the first row to the ones in rows two and three of Table 2 reveals that

Table 2 Descriptive statistics for extreme-move days for stock prices and NY temperature.

	Dow Jones daily returns: 1900–2006		New York daily average temp's deviation from normal: 1901–2006	
	Mean	Standard deviation	Mean	Standard deviation
Absolute value of daily return/daily change	0.743%	0.853%	1.138%	0.909%
	(N = 26,884)		(N = 38,655)	
Average daily move 3-days before a 3 sigma-plus day	2.105%	1.741%	2.563%	1.107%
Average daily move 3-days after a 3 sigma-plus day	2.122%	1.602%	2.441%	1.102%
Average daily move on the day of a 3-sigma plus day	4.952%	2.166%	4.984%	0.556%
Number of observations	427		220	

the average daily movement in the three-day interval preceding and following a $3\sigma+$ day is markedly different from the typical daily movement. For example, the average daily return in the three-days before a large-move days is 2.11%, which is almost three times larger than the average daily return in the sample. The same holds true for the average daily return during the three-day period after a $3\sigma+$ day. The mean absolute return on all $3\sigma+$ days is 4.96%, almost seven times larger than the average for all days. This difference is to be expected because of the 427 days with $3\sigma+$ returns, there are a sizeable number of days with returns as much as six or seven sigmas away from the mean.

Turning now to the summary statistics on the explanatory variables for the NY temperature, we find a remarkable similarity with the Dow's moves. The mean daily deviations in the three days before and after a $3\sigma+$ day is considerably larger (approximately 2 1/2 times) than the average daily deviations from normal temperature for all days. Interestingly, the means of the variable "daily move on the

day of a $3\sigma+$ day," are virtually identical for the NY temperature and the Dow; 4.98% and 4.95%, respectively.

The summary statistics in Table 2 suggest that these variables may have some explanatory power to predict the onset of extreme-move days. However, before undertaking a formal duration analysis, it may be useful to examine the time interval between $3\sigma+$ days using a simple least squares regression. The results of two least squares analyses, one for the Dow and the other for NY temperature, are presented in Table 3. The dependent variable in each regression is $\ln(T)$, where T denotes the time interval in days from one extreme movement day to the next. The explanatory variables are the ones discussed in the preceding paragraphs.

Qualitatively, the results for the Dow and the NY temperature are very similar. In both cases, the coefficient of the variable "average daily move 3-days before a $3\sigma+$ day" is negative and highly significant (with t -statistics exceeding 10), whereas the

Table 3 Least squares analysis of the time interval between 3 sigma-plus days.

	Dow Jones daily returns		New York daily average temp's deviation from normal	
	Estimated coefficient	<i>T</i> -stat	Estimated coefficient	<i>T</i> -stat
Three-days before	-54.87	-12.20***	-123.73	-10.04***
Three-days after	-8.21	-1.64	5.21	0.41
The day of	-6.37	-1.77	-68.14	-2.65***
Intercept	3.72	19.54***	9.84	8.22***
Number of obs	427		220	
Adjusted <i>R</i> -squared	0.327		0.362	

Note: The dependent variable in each regression is the log of the duration in days between 3-sigma-plus days.

Definitions:

For the Dow, the sigma is based on the standard deviation of daily log returns.

For New York Temp, the sigma is based on the standard deviation of daily % deviation from normal (= 106-year average) temperature.

The 3-Days Before (after): for the Dow, it is the average absolute return in the three days prior (after) to a 3-sigma-plus day; for NY temp, it is the average absolute % deviation three days prior (after).

The Day of: For the Dow it is the absolute value of return, for NY Temp it is the absolute % deviation on a 3-sigma-plus day.

***denotes statistically significant at 99% level of confidence.

coefficient of the variable “average daily move 3-days after” is not. Thus, we find that extreme value days are preceded by, but not necessarily followed by, changes that are relatively large.

The coefficient of the variable “daily move on a $3\sigma +$ day” is negative for both the Dow and the NY temperature, but not significant at the 95% level for Dow. The negative coefficient of this variable suggests that the magnitude of the change on a $3\sigma +$ day is itself negatively correlated with the time interval between extreme-move days.

4.1 Duration analysis using the cox proportional hazard model

In Table 4A, we present the results of the duration analysis using the Cox Proportional Hazard model. We use the semi-parametric Cox model rather than a parametric model such as the Weibull or logistic

models because the Cox model does not impose a particular shape to the hazard function.

In comparing the estimated coefficients of the explanatory variables in the least squares regression and the Cox model, one should note that the signs of the coefficients are expected to be positive rather than negative. Positive estimated coefficients imply that these variables shorten the duration, i.e., “hasten” the arrival of the next $3\sigma +$ day. For example, consistent with the regression findings, the estimated coefficient of the variable “average daily move 3-days before a $3\sigma +$ day” is positive and significant at the 99% confidence level in the Cox model both for the Dow and the NY temperature. Similarly, the estimated coefficient of the variable “daily move on a $3\sigma +$ day” is also positive and statistically significant, both for the Dow and the NY temperature. However, the coefficient of the variable, “average daily move 3 days after a $3\sigma +$ day” is not

Table 4A Estimation of time between extreme-move days using the Cox duration model.

	Dow Jones: 1900–2006					
			Excluding the 1930s		NY Temp: 1901–2006	
	Coeff-Est	Z-stat	Coeff-Est	Z-stat	Coeff-Est	Z-stat
<i>All 3-Sigma Plus Days</i>						
3-Days Before	32.74	13.25***	29.80	8.11***	45.60	6.69***
3-Days After	4.49	1.42	1.21	0.27	−2.03	−0.31
That Day	6.26	2.51**	7.27	2.09**	43.98	3.38***
Number of Obs		427	202		220	
<i>Only Positive 3-Sigma Plus Days</i>						
3-Days Before	89.38	9.97***	115.94	6.10***	36.46	3.76***
3-Days After	23.53	3.07***	15.98	1.13	9.48	1.06
That Day	10.72	2.50**	21.57	2.16**	47.19	2.08**
Number of Obs		178	80		115	
<i>Only Negative 3-Sigma Plus Days</i>						
3-Days Before	46.35	9.84***	41.13	6.18***	45.75	4.48***
3-Days After	3.77	0.89	2.64	0.48	−6.34	−0.59
That Day	9.07	3.45***	10.98	3.25***	39.78	2.31**
Number of Obs		249	122		105	

Note: See Table 3 for definition of the variables.

*** denotes statistically significant at 99% level of confidence.

** denotes statistically significant at 95% level of confidence.

statistically significant for both the Dow and the NY temperature.

For the Dow, we re-estimated the Cox model by excluding the 1930s (which leaves us with 202 $3\sigma+$ days). The results stay virtually unchanged: the signs and significance of the estimated coefficients are identical to those using the entire sample period 1990–2006.

To explore Nelson's (1991) observation that there is an asymmetry in the market's response to "good news" and "bad news," we next considered the two subsets—positive and negative $3\sigma+$ days—separately. In examining these subsets, we define the explanatory variables slightly differently than before. For example, for the set of only positive

$3\sigma+$ days, the variable "average daily move 3-days before a $3\sigma+$ day" now measures the average of only the positive returns or temperature changes in the three-day intervals. The converse applies to the explanatory variables for the analysis of negative $3\sigma+$ days: the average of only the negative moves in the three-day intervals is measured.

The results are consistent across sub-samples. The coefficients of the variables "average daily moves three days before" and "the day of" for both negative and positive $3\sigma+$ days, remain statistically significant for the Dow (with and without the 1930s included) and for the NY temperature. Similarly, the coefficient of the variable "three day after" is insignificant in all cases except for positive $3\sigma+$ days for the Dow. However, even in this case, the

coefficient estimate becomes insignificant when the 1930s are excluded from the estimation.

Although the estimation results for positive and negative $3\sigma+$ days are qualitatively similar in terms of the signs and significance of the estimated coefficients, there are some differences, particularly for the Dow. The coefficient of the variable “average daily move 3-days before” is substantially larger for positive $3\sigma+$ days than the negative ones for the Dow. This result holds true even when one excludes the 1930s. This implies that the positive returns in the three-day period before a positive $3\sigma+$ day “hasten” the arrival of the next extreme, positive-move day much more than does the negative returns preceding a negative $3\sigma+$ day.

However, a similar asymmetry is not observed for positive and negative large-movement days in the NY temperature. In fact, the coefficient of the variable “average daily move three days before” for the positive $3\sigma+$ days is slightly smaller than the coefficient of this variable for negative extreme-move days.

Despite these minor dissimilarities, the nine sets of duration model results contained in Table 4A are remarkably consistent. For both the Dow Jones and the NY temperature, larger-than-average moves three days prior and on the day of a $3\sigma+$ day “hasten” the arrival of the next extreme-move day.

4.2 Validation of the robustness of results

We have undertaken a large number of variations of the Cox duration model to verify that our results are robust. For example, (a) we have considered the average moves during different windows of time (e.g., three-day, five-day, seven-day windows) before and after a $3\sigma+$ day; (b) we have excluded various sub-periods, including the years 1928–1935 for the Dow; (c) we have excluded observations where the difference between two consecutive trading days exceed four days as was the case during the

World Wars or 9-11-2001 for the Dow; (d) we have excluded all observations with moves greater than or equal to 5-sigmas to examine the impact of outliers; (e) we have included a variety of dummy variables for decades, for months of the year, for periods of recession, etc. While we do not report the results of all these robustness checks here, they are available on request. They confirm that the key findings of this paper are remarkably robust. In Table 4B, however, we provide a few examples of these robustness checks.

As shown in this table, we re-estimated the duration model for the Dow with two sub-periods: 1900–1950 and 1951–2006. The daily σ for the first sub-period is a fair bit higher since it includes the 1930s: 1.33% versus 0.91%. Despite this difference, the explanatory power of the model and the qualitative results across the two sub-periods are very similar. The coefficient of the variable “average daily move 3-days before” continues to be positive and statistically significant at the 99% level of confidence.

In this table, we also report results when a dummy variable for recession quarters is included in the model. The “recession” dummy variable takes a value of one for all trading days in the quarters that the National Bureau of Economic Research⁴ has identified as being recessionary. The recession dummy is not statistically significant in either sub-period, 1900–1950 or 1951–2006; and the sign and significance of the “three days before” variable’s coefficient estimate remain unchanged.

The remarkable consistency of the results—particularly the finding that extreme-move days are preceded by but not necessarily followed by periods of higher volatility—across two seemingly unrelated random phenomena, stock-market returns and temperature changes, seem to suggest the existence of a common stochastic process. We do not propose such a process here. We believe our findings simply lend further support to the arguments of financial economists and mathematicians who

Table 4B Results of Cox duration model estimation: Pre and Post 1950s.

	Coeff-Est	Z-stat	Coeff-Est	Z-stat
<i>Dow Jones: 1900–1950 (Sigma = 1.334%)</i>				
3-Days before	38.18	10.11***	37.53	9.83***
3-Days after	-2.60	-0.57	-3.25	-0.71
That day	7.50	2.19**	7.41	2.18**
Recession			0.19	1.26
Number of Obs	223		223	
<i>Dow Jones: 1951-2006 (Sigma = 0.909%)</i>				
3-Days before	28.07	6.40***	28.08	6.39***
3-Days after	6.92	1.10	6.91	1.10
That day	0.59	0.11	0.61	0.11
Recession			0.01	0.05
Number of Obs	154		154	

Note: See Table 3 for definition of the variables.

*** denotes statistically significant at 99% level of confidence.

** denotes statistically significant at 95% level of confidence.

have contended that many complex random processes, including stock-market returns, cannot be explained by simplistic processes such as the random walk. The challenge, however, lies in uncovering and formulating the stochastic processes that can be shown to possess a reliable degree of predictability.⁵

5 Predictability and the returns of a portfolio strategy

In this section, we explore the predictive power of our duration-model results and examine whether they can be used to produce a useful market timing strategy. In particular, we examine whether our model can be used to predict an extreme, negative-move day in the stock market where the investor could shift a portfolio's asset allocation from stocks to cash. Of the three explanatory variables used in the duration analysis, only one, "the average daily move three days before a $3\sigma+$ day," has the potential to be used as a predictor, because the value of other explanatory variables is not known until *after* a large-move has already occurred.

Our proposed asset allocation strategies are based on our empirical finding that negative extreme-move days are typically preceded by days when equity returns are negative and unusually large in absolute value.⁶ Based on this finding, we use the following rules to construct asset-allocation strategies:

- From the beginning of our sample period, the portfolio is 100 percent invested in the Dow index. That is, the portfolio's daily return is assumed to be the same as the Dow's.
- If a day's return is negative, and the average absolute value of the preceding three day's negative returns exceeds a set threshold, then the portfolio is unwound and reinvested in cash.
- The portfolio stays unwound for a set period of time over which the assets earn a rate of return equivalent to the risk-free rate.
- After the set period of time, the investor is assumed to re-initiate a long position in the Dow index.
- The portfolio incurs a fixed transaction cost, both for unwinding and re-initiating the portfolio.

We define the “set threshold” in (b) above, as the average of the absolute values of daily returns of all non- 3σ days with a negative return in the entire sample. We call this threshold the “trigger condition to unwind.”

We have considered various periods of time over which the portfolio remains unwound, including 2, 3, 5, 7, and 10 days. On each day in this interval during which the portfolio remains unwound it is assumed to earn the daily risk-free rate of 0.0103%. This is based on the average annual Treasury bill rate of 3.7% for the years 1926 through 2006.⁷

As a measure for transaction cost, we have used 5 basis points each way—that is, the portfolio incurs a cost of 0.05% on the day it is unwound, and another 0.05% when it is re-initiated. We have used two sample periods in the portfolio analysis: the

entire sample of the 107 years, and the sub-sample comprising the years 1936–2006. The results of our analysis are presented in Table 5.

When the entire 107-year-period is considered, the portfolio is found to outperform the Dow’s actual annual average returns when the unwound period is at least two days. The percentage of all trading days during which the portfolio stays invested in cash varies from 12% to 34%, depending on the length of the unwound period.

Two findings warrant special attention: First, the portfolio strategy yields a considerably lower volatility of returns than the actual Dow. While the annualized standard deviation for the actual Dow index is 21.46%, the portfolio’s annual σ is almost 400 basis points lower at 16.85%. Not unexpectedly, the portfolio’s σ diminishes considerably when

Table 5 Performance statistics for alternative trading strategies based on extreme-move days.

	Annual mean return (%)	Annual sigma (%)	Sharpe ratio	% Days	% of negative 3-sigma-plus days avoided
<i>All data (1900–2006)</i>					
Actual	7.01	21.46	0.15		
<i>Hypothetical Portfolio</i>					
Unwound period 2 days	6.59	16.85	0.17	12%	54%
Unwound period 3 days	7.63	16.05	0.25	16%	65%
Unwound period 5 days	8.62	14.66	0.34	22%	75%
Unwound period 7 days	8.42	13.61	0.35	28%	80%
Unwound period 10 days	8.28	12.53	0.37	34%	86%
<i>1936–2006 Only</i>					
Actual	9.45	17.76	0.32		
<i>Hypothetical Portfolio</i>					
Unwound period 2 days	7.42	14.57	0.26	12%	47%
Unwound period 3 days	7.04	14.03	0.24	16%	55%
Unwound period 5 days	7.53	12.94	0.30	23%	67%
Unwound period 7 days	7.45	11.98	0.31	29%	74%
Unwound period 10 days	7.07	11.03	0.31	36%	79%

Note: “% Days” is the percent of the trading days during which the portfolio remains invested in cash as opposed to being invested in Dow.

the portfolio stays unwound for a longer period of time. In this context, it is informative to consider the Sharpe ratio (where we divide the difference between the mean return and the risk-free annual rate of 3.7% by the sigma of the return). In all cases, the portfolio yields a higher Sharpe ratio than the actual Dow's returns. In fact, for the unwound period of ten days, the Sharpe ratio for the portfolio (0.37) is more than twice as large as the actual Dow's ratio (0.15).

Second, our portfolio strategy, with a rather simple trigger condition for unwinding, seems to be fairly effective in avoiding negative $3\sigma+$ days. Even with an unwound period of only two days, the strategy avoids 54% of the negative $3\sigma+$ days. When the unwound period increases to 10 days, this percentage jumps to 86%, implying that the trigger condition is fairly accurate in predicting the next negative extreme-move day.⁸

In Table 5, we also report the results of the portfolio analysis using data only for the period 1936–2006. For this period, the Dow's actual mean annualized return is 9.45% (as opposed to 7.01% for the entire period of 1900–2006) and annualized sigma is 17.76% (as opposed to 21.46% for 1900–2006). Now, however, the portfolio does not outperform the Dow. In fact, at its best—when the unwound period is 5 days—the mean return is 7.53%, about 200 basis points worse than Dow's actual performance. It is not surprising that the asset-allocation strategy has a more negative effect during the past 70 years. The stock market has enjoyed a particularly strong long-term uptrend following the great depression and the buy-and-hold investor has generally benefited from being totally invested in all periods. But the portfolio's volatility (σ) continues to be substantially smaller than Dow's actual volatility. Even with a two-day unwound period, the portfolio's annualized σ is 14.57% more than 250 basis points lower than Dow's. The portfolio's Sharpe ratio is roughly comparable to the actual

Dow's when the unwound period is 5 days or longer.

The percentage of trading days affected by the portfolio strategy is virtually identical regardless of which period is considered, all 107 years or the last 70 years. However, that is not the case for the number of negative $3\sigma+$ days avoided. Consistent with the finding of lower returns for the sub-period 1936–2006, we find that the percentage of avoided negative $3\sigma+$ days is slightly smaller as well. With a two-day unwound period, this percentage is 47%, which rises to 79% when a ten-day unwound period is applied.

While there is strong empirical evidence of a relationship between an extreme-move day and its preceding three-day's average return, we find that this relationship is not predictable enough to devise a dependable asset-allocation (market-timing) strategy to outperform the market. However, if the objective of the trading strategy is to reduce the portfolio's volatility, then the strategy seems to be reasonably effective: from 50% to 80% of the negative extreme-move days are avoided by adopting the strategy suggested by our duration model results.

6 Concluding comments

There has been considerable empirical work confirming that stock returns are not normally distributed. Moreover, we know that stock prices do not follow a strict random-walk process and may, to some extent, be predictable. The paper has examined the temporal properties of the "fat tails" of the stock-return distributions. We found that there has been a clustering of days when the stock market has been subject to "extreme movements," which we define as daily stock returns that are three or more standard deviations away from the mean return. Moreover, a similar clustering of extreme movements is shown to exist in the daily time series of temperatures recorded in New York City.

Utilizing a duration model we found that extreme movements can, to some extent, be predicted on the basis of unusually large movements in preceding periods. One particularly robust finding is that extreme movements in stock prices are usually preceded by large average daily movements during the preceding three trading days. A similar finding was reported for daily NY City temperatures.

The strong results suggest that it may be possible to fashion an asset-allocation strategy, whereby the investor would switch from stocks to cash in advance of predicted extreme negative returns in the stock market. We found, however, that it was not possible to use the predictability that exists to improve an investor's long-run return over the return that would be earned by the buy-and-hold investor who simply stayed fully invested in the stock market. The market-timing strategy that we simulated was, however, useful in reducing the volatility of returns. Because the timing strategy does permit the investor to avoid between half and more than three quarters of the trading days in which the market suffers substantial declines, the suggested strategy does substantially lower the volatility of the portfolio's returns over time.

We have shown that a duration model analysis can shed considerable light on the factors that are associated with unusually large changes in the stock market and the weather. The task for the future is to better understand the underlying stochastic processes common to many random phenomena.

Notes

- ¹ Throughout the paper, when annualizing daily returns we have assumed a 360-day calendar year.
- ² We procured this data set from Weather Source.
- ³ We examine not only a three-day interval *after* a $3\sigma+$ day, but also *before* to explore whether a large move day can be predicted by "pre-shocks"; by definition "after-shocks" cannot have any predictive power. Later in this paper (see Section 5), we explicitly explore this predictive relationship by examining an asset allocation strategy.

- ⁴ The NBER study can be accessed electronically at www.nber.org/cycels.html.
- ⁵ In this context, the writings of Benoit Mandelbrot, who has long argued for a fractal explanation of various disparate and diverse stochastic phenomena, deserve special attention. In his latest book, "The Misbehavior of Markets: A Fractal View of Risk, Ruin, and Reward" he elaborates on his fractal view of various random phenomena, including the behavior of the financial markets.
- ⁶ To illustrate the point, consider one of the worst negative move days in Dow's history, October 19, 1987, when the Dow fell by 27% (in log returns). This day was preceded by fairly large negative move days: in the three days prior to this day, the Dow fell on average by 3.7% per day and cumulatively by nearly 11%. Yet, in the three day-period following October 19, 1987 the Dow's moves were relatively uneventful.
- ⁷ Ibbotson Associates, 2007 Yearbook, Chicago, IL.
- ⁸ The reason that returns are not improved considerably even though so many negative days are avoided is that many extreme positive days are eliminated as well. As is well known, a substantial part of the generous returns earned by long-run equity investors comes from the infrequent number of days when the market rallies sharply.

References

- Alexander, S. S. (1964). *Price Movements in Speculative Markets: Trends or Random Walks*. 338–372.
- Bachelier, L. (1964). "Theory of Speculation." *Annales Scientifiques de l'École Normale Supérieure*, Sup (3), 1018 (1900), [English translation by A.J. Boness in Cootner (1964)], 17–78.
- Bollerslev, T. (1986). "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31, 307–327.
- Bremer, M. and Sweeney, R. J. (1991). "The Reversal of Large Stock-Price Decreases." *The Journal of Finance* 46, 747–754.
- Cootner, P. H. ed. (1964). *The Random Character of Stock Prices*. The MIT Press, Cambridge, MA.
- Cross, F. (1973). "Price Movements on Fridays and Mondays." *Financial Analysts Journal* 29, 67–79.
- Engel, R. (1982). "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica* 50, 987–1007.
- Engel, R. (2004). "Risk and Volatility: Econometric Models and Financial Practice." *The American Economic Review* 94, 405–420.

- French, K. R., Schwert, G. W., and Stambaugh, R. F. (1987). "Expected Stock Returns and Volatility." *Journal of Financial Economics* 19, 3–29.
- French, K. (1980). "Stock Returns and the Weekend Effect." *Journal of Financial Economics* 8, 55–69.
- Gibbons, M. R. and Hess, P. J. (1981). "Day of the Week Effects and Asset Returns." *Journal of Business* 54, 569–596.
- Keim, D. B. and Stambaugh, R. F. (1985). "A Further Investigation of The Weekend Effect in Stock Returns." *The Journal of Finance* 39, 819–837.
- Lo, A. and Craig MacKinlay, A. (1988). "Stock Market Prices do not Follow Random Walks: Evidence from a Simple Specification Test." *The Review of Financial Studies* 1, 41–66.
- Lo, A. and Craig MacKinlay, A. (2001). *A Non-Random Walk Down Wall Street*. Princeton University Press.
- Mandelbrot, B. (1963). "The Variation of Certain Speculative Prices." *Journal of Business* 36, 394–419.
- Mandelbrot, B. and Hudson, R. (2004). *The Misbehavior of Markets: A Fractal View of Risk, Ruin, and Reward*, Perseus Publishing.
- Nelson, D. B. (1991). "Conditional Heteroskedasticity in Asset Returns, A New Approach." *Econometrica* 59, 347–370.
- Osborne, M. F. M. (1959). "Brownian Motion in the Stock Market." *Operations Research* 7, 145–173.
- Rogalski, R. (1984). "New Findings Regarding Day of the Week Returns Over Trading And Non-Trading Periods: A Note." *The Journal of Finance* 39, 1603–1614.
- Samuelson, P. A. (1965). "A Proof That Properly Anticipated Prices Fluctuate Randomly." *Industrial Management Review* 6, 41–49.
- Shiller, R. J. and Perron, P. (1985). "Testing the Random Walk Hypothesis: Power Versus Frequency of Observation." *Economics Letters* 18, 381–386.

Keywords: Volatility clustering; duration analysis; portfolio strategy