Are Stock Returns Different over Weekends? A Jump Diffusion Analysis of the "Weekend Effect"

Peter Fortune

Senior Economist and Advisor to the Director of Research, Federal Reserve Bank of Boston.

The distribution of returns on common stocks is, arguably, one of the most widely studied financial market characteristics. Among the questions addressed by these studies are the following: Is the return on common stocks normally distributed, as much finance theory assumes? How many stocks should be included in a portfolio if it is to achieve most of the benefits of diversification? How has the volatility of stock returns changed over time? How is the distribution of returns affected by past returns? What influence do calendar events have on the return distribution (for example, the January effect, the Monday effect)?

The performance of stock prices during breaks in trading has also received considerable attention in recent years, especially since the advent of "circuit breakers" designed to create stability when markets are chaotic. Some studies have focused on performance surrounding periods of unscheduled trading breaks, such as trading halts in individual stocks and triggering of exchange-wide circuit breakers. These studies hope to establish whether a trading break during the "fog of battle" helps to stabilize the stock market.

Other studies look at performance around periods of scheduled trading breaks (holidays, weekends). Some of these are designed to obtain insight into whether trading breaks serve to stabilize or destabilize markets. Other studies in this genre are designed to determine whether there are financial market anomalies associated with days of the week.

This study examines the distribution of daily returns on five popular stock price indices, with a special emphasis on the difference between returns over weekends and returns over adjacent intraweek trading days. We revisit the "weekend effect" in common stock returns, focusing on two characteristics of differential returns over intraweek trading days and over weekends: the "drift" and the "volatility."

Section I of this article describes the simple diffusion model of stock returns, upon which much of modern finance rests. It also describes the jump diffusion model, upon which this paper rests. Section II summarizes the literature on the weekend effect, reporting on the evidence supporting its existence and on some possible explanations grounded in economic theory. Section III discusses the estimation methods used, while section IV discusses the data used in this paper. Section V reports the results, which are summarized just below. The article concludes with a summary.

The volatility of stock returns over weekends is much smaller than could be predicted from intraweek volatility. In short, holding stocks over weekends gives low and perhaps negative returns, but also provides relatively low risk.

We find that a jump diffusion model is superior to a simple diffusion model, and that the jump diffusion model of stock returns provides strong support for the weekend effect. For large-cap indices, like the Dow 30 and S&P 500, the normal return (or "drift") over weekends is positive but significantly less than the intraweek drift. For small-cap stock price indices, like the Russell 2000 and the Nasdaq Composite, the weekend drift is actually negative. Thus, much of the tendency for price declines over weekends is confined to stocks of small companies.[1]

The volatility of stock returns over weekends is much smaller than could be predicted from intraweek volatility. In short, holding stocks over weekends gives low and perhaps negative returns, but also provides relatively low risk. We also find that the difference between intraweek and weekend drift is smaller after October 1987 than before. Indeed, for large companies the difference disappears! This suggests that the poor performance of common stocks over weekends in the 1980s was a financial anomaly that was mitigated over time as investors incorporated it into the timing of their transactions. For example, shifting the timing of sales to Fridays and of purchases to Mondays would tend to equalize the weekend and intraweek returns on stocks. However, for small-cap stocks the 1980s pattern continued into the 1990s.

We found no change in the relative volatility of stock returns over weekends after October 1987. Thus, the observation that weekends are periods of low volatility appears to be true today as well as decades ago.

I. Modeling the Return on Common Stocks

Any analysis of stock returns rests upon a model of the evolution of those returns. The standard model of stock returns is the diffusion model. A simple diffusion model (see Box 1) assumes continuous time (that is, time is not separated into discrete intervals such as days or weeks). The rate of return at any instant of time, called the "instantaneous rate of return" is modeled as a constant, denoted by α, plus a random "surprise" with a zero mean and a constant standard error, denoted by σ and called the asset's instantaneous volatility. Formally, the diffusion model states that

$$dS/S = \alpha dt + (\sigma \sqrt{dt})\, dz \qquad (1)$$

where S is the asset's price at instant t (including any accumulated cash dividends), dt is the infinitesimally small interval of time over which the stock's return is measured, dz is a random variable with a standard normal distribution,[2] and dS/S is the instantaneous rate of return. The instantaneous return over the interval dt will be normally distributed with a mean of αdt, a variance of $\sigma^2 dt$ and a volatility of $\sigma\sqrt{dt}$.

While the theoretical model is defined in continuous time, the data we have available are measured in discrete intervals of time. Returns over discrete intervals of time are typically measured by the logarithm of the price relative over that interval. If S(t) is the price at time t (including accumulated dividends), the price relative over the subsequent T periods is defined as S(t + T)/S(t). The rate of growth over the T periods is measured by the log price relative, *ln*[(S(t + T)/S(t)], where *ln* denotes "logarithm." If the deviation of instantaneous returns from the mean of α

[1] The capitalization of a company is often measured by the market value of the company's equity. The Dow 30 and S&P 500 indices are for large companies with a median equity value of $50 billion and $6 billion, respectively. The median equity values for the Russell 2000 and Nasdaq Composite indices are $500 million and $42 million, respectively. See Fortune (1998, Table 1).

[2] A standard normal random variable is normally distributed with a zero mean and unit standard error. Any normally distributed random variable can be converted to a standard normal random variable by deducting the mean and dividing the result by the standard error. Thus, if x is normally distributed with mean μ and standard error σ, then $(x - \mu)/\sigma$ is a standard normal variable.

Box 1: The Simple Diffusion Model of Asset Prices

The simple diffusion model is a formal statement of the random walk model of stock prices, in which the price of an asset grows at a constant rate with vibrations around the trend that get smaller as time passes.

Let S(t) be the price of an asset at time t and α be the instantaneous mean rate of capital appreciation of the asset's price. Over an infinitesimally small time interval, dt, the asset's price changes by dS. If the asset is riskless, the instantaneous rate of return will be $dS/S = \alpha dt$. If the asset's return is uncertain, it will deviate from the mean return.

Suppose that dz is a normally distributed random variable with zero mean and unit variance, and that the standard deviation of the instantaneous return, also called the asset's instantaneous volatility, is denoted by σ. A simple diffusion model of the instantaneous return on the asset is

$$d\ln S = \alpha dt + (\sigma\sqrt{dt})\, dz, \qquad \text{(B1.1)}$$

where *ln* denotes a logarithm and the definition $d\ln S = dS/S$ is used. According to the simple diffusion model, the *instantaneous* return on the asset is a normally distributed random variable with mean (or "drift") α and standard deviation (or "volatility") σ.

Equation (B1.1) is a stochastic differential equation describing the probability distribution of instantaneous returns. The mathematically inclined can show that the logarithm of the "price relative," S(t + T)/S(t), over a discrete interval of time with length T, is normally distributed with mean $(\alpha - \tfrac{1}{2}\sigma^2)T$ and volatility $\sigma\sqrt{T}$; that is,

$$\ln[(S(t + T)/S(t)] \sim N[(\alpha - \tfrac{1}{2}\sigma^2)T, \sigma\sqrt{T}], \qquad \text{(B1.2)}$$

from which the distribution of the *average* return over T periods is

$$\ln[(S(t + T)/S(t)]/T \sim N[(\alpha - \tfrac{1}{2}\sigma^2), \sigma/\sqrt{T}]. \qquad \text{(B1.2')}$$

Thus, the log price relative increases in proportion to the passage of time, but the standard error of deviations around the trend increases only with the square root of elapsed time. Note that while the mean *instantaneous* return is α, the mean *average* return over T periods is $(\alpha - \tfrac{1}{2}\sigma^2)$. The deduction for volatility arises from the nonlinearity of the transformation from the price relative to its logarithm.

Equation (B1.2′) can be used to answer questions about the average return. For example, we might want to know what the probability is that the average return over 10 years will be less than the riskless rate of interest, denoted by r. We know from a simple diffusion model that $[R - (\alpha - \tfrac{1}{2}\sigma^2)]/[\sigma/\sqrt{T}]$ is a standard normal random variable, where $R = \ln[(S(t + T)/S(t)]/T$ is the average return. We want to know Prob{$R < r$}, which is equivalent to

$$\text{Prob}\{w < [r - (\alpha - \tfrac{1}{2}\sigma^2)]/[\sigma/\sqrt{T}]\}, \qquad \text{(B1.3)}$$

where w is a standard normal variable.

Suppose that for our asset we have a daily return with mean $\alpha = 0.0010$ and daily volatility $\sigma = 0.025$. Suppose also that the riskless rate is 0.0003 per day. Plugging these numbers into (B1.3) and setting T = 3650 days tells us that we want to find Prob($w < -.94$). This is easily found in tables of the standard normal distribution. The answer is Prob($w < -.94$) = 0.30; this asset has a 30 percent probability of a 10-year average return that is less than the riskless rate.

---

is normally distributed with zero mean and standard deviation σ, then the log price relative over T periods, $\ln[(S(t + T)/S(t)]$, is normally distributed with mean growth rate of $(\alpha - \tfrac{1}{2}\sigma^2)T$ and growth rate volatility $\sigma\sqrt{T}$; that is, the log price relative is normally distributed with a trend growth rate ("drift") of $(\alpha - \tfrac{1}{2}\sigma^2)T$ and a standard error ("volatility") of $\sigma\sqrt{T}$. Thus, the rate of return over a time interval of length T is described by

$$\ln[(S(t + T)/S(t)] = (\alpha - \tfrac{1}{2}\sigma^2)T + (\sigma\sqrt{T})\, dz. \qquad (2)$$

The average rate of return over T periods is the log price relative divided by T, or $\ln[(S(t + T)/S(t)]/T$. From the growth formula (2), it can be seen that the average rate of return over T periods is normally distributed with a mean rate of return of $(\alpha - \tfrac{1}{2}\sigma^2)$ and a standard deviation of $\sigma/\sqrt{T}$. The mean rate of return is positive because investors will not hold risky securities if there are no returns to compensate for risk, so we expect $\alpha > \tfrac{1}{2}\sigma^2$.

The relationship between the instantaneous rate of return at an instant of time, dS/S, and the average rate of return, *ln*[S(t + T)/S(t)], over a discrete interval of time, is often a source of confusion that requires some explanation. The mathematics of finance tells us that if the mean *instantaneous* return is α, the *average* return over T periods is ($\alpha - ½\sigma^2$). But why should the average of instantaneous returns be less than the mean instantaneous return? And why should the reduction in average return be related to the volatility?

We can think of the *log* price relative as the return associated with a given value of the price relative. The average return over an interval of time (say, a week) is the average of the returns associated with the (say, daily) price relatives that occur. The nonlinearity introduced by the transformation from price relatives to *log* price relatives ensures that the average return associated with several price relatives will be less than the return associated with the average price relative. Thus, when returns vary over time, the average return (average log price relative) will be below the return associated with the average price relative. Furthermore, the gap will widen as return volatility increases.

Demonstrating this mathematically is extremely tedious,[3] but it can be shown through a simple diagram. In Figure 1 the horizontal axis measures the price relative and the vertical axis measures the *logarithm* of that price relative. The curved line, concave from below, shows the relationship of the log price relative to the price relative. The vertical distance from any S_{t+T}/S_t on the horizontal axis to its associated $ln(S_{t+T}/S_t)$ on the curve measures the rate of return associated with that price relative. Suppose that s* is the price relative. The log price relative at s* will be x_1, and this will be the return. If exactly the same price relative applied to each instant of time, the return would be x_1 in each instant and the average return would be x_1. But only in this case of identical returns in each period will the average return be equal to the mean instantaneous return!

Now suppose that the price relative is uncertain, reflecting variation in the returns, but the average price relative remains at s*. Will the average return remain at x_1? To see that it will not, assume that there is an equal probability that the price relative will be either s_0 or s_1. The geometry of probability theory ensures that the expected log price relative will be on the line segment AB that connects the log price relative at s_0 with the log price relative at s_1. The exact position on AB will depend on the probability that s_1 materializes. Under the assumption that both outcomes are equally likely, the mean return will be at the midpoint of AB. The average price relative is still s*, but now the average return is x_2, at the midpoint of AB. The average return will be less than the original return (x_1) because of the nonlinearity: A decline in the price relative to s_0 reduces the return by more than it would be increased by a rise from s* to s_1. Thus, the introduction of uncertainty about the return has reduced the average rate of return even though the expected price relative and the mean instantaneous return have not changed.

Suppose now that the volatility of the price relative increases so that the two possible price relative outcomes are s_{00} or s_{11}, each equally likely. The expected log price relative will be at the midpoint of the red line A′B′. The expected price relative will still be s*, but the expected log price relative will fall to x_3. Thus, the greater the variability in the price relative, the lower will be the expected log price relative. The nonlinearity of the transformation from price relative to average return means that the average return will fall as variability increases. In this way, a mean instantaneous return of α becomes an average return of $\alpha - ½\sigma^2$.

Figure 1

The Log Price Relative Relationship

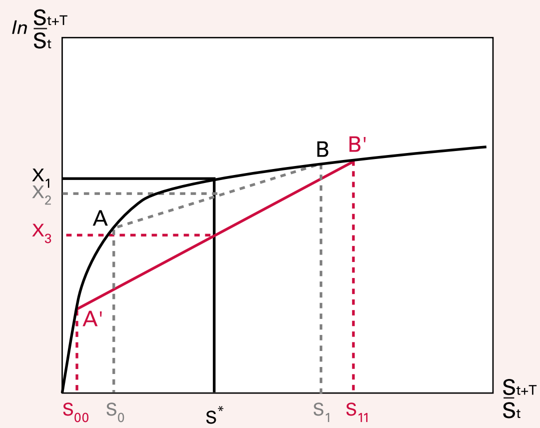

[3] Statisticians have long known that if X is a random variable with mean μ and standard error σ, then *ln* X will be a random variable with mean ($\alpha - ½\sigma^2$) and standard error σ. Thus, the use of X or *ln*X in financial analysis affects the mean, or drift, of a random variable, but not its variability around the mean.

This simple diffusion model is a two-parameter model of asset price evolution. The only parameters are the mean instantaneous growth rate, α, and the instantaneous volatility, σ. If the random component of instantaneous returns is normally distributed, and if the parameters do not change over time, forecasts of future prices, and statements about probability of pricing errors, such as confidence intervals, can easily be formed. Further, these applications can be extended to a multiple asset model if correlations among instantaneous returns are known and constant. In this case, the well-known Capital Asset Pricing Model of the structure of asset returns can be derived.

The simple diffusion model is the core model of asset price dynamics in financial theory. Unfortunately, it fails to incorporate a number of known features of stock price behavior. If returns were normally distributed, the distribution would be symmetrical around the mean, that is, the frequency of large price increases would be the same as the frequency of equally large price declines. However, it is widely known that stock returns tend to be skewed downward, that is, below-normal returns are more frequent than above-normal returns. This is particularly pronounced in the extreme tails of the distribution: Large price declines, such as the October 1987 break, are more frequent than are equally large price increases.

It is widely known that stock returns tend to be skewed downward, that is, below-normal returns are more frequent than above-normal returns.

In addition, stock returns are leptokurtic, meaning that the distribution is excessively peaked in the middle. A larger proportion of returns are in the middle of the distribution than are in the tails, and the tails are "fat," indicating a higher-than-normal frequency of big price changes.

The Jump Diffusion Model

Proposed by Press (1967) and popularized by Merton (1976) and by Cox and Ross (1976) in the context of option pricing, the jump diffusion model incorporates some of the known characteristics of stock prices that are not consistent with the simple diffusion model, such as skewness and fat tails in the returns distribution.

In recent years, papers have emerged that involve direct estimation of the parameters of this stochastic process. For example, Johnson and Schneeweis (1994) estimated a jump diffusion model to examine the impact of macroeconomic news on foreign exchange rates. Kim, Oh, and Brooks (1994) estimated the parameters of a jump diffusion model for individual stocks in the Major Market Index.

The jump diffusion model builds on the simple diffusion model. The jump diffusion model, described in Box 2, postulates that the instantaneous return is generated by a simple diffusion model with an additional source of variability in returns. This source is the jump process, in which a discrete number of shocks affect returns at any instant. Each shock is assumed to be a normally distributed random variable with a fixed mean effect, denoted by θ, and a fixed standard error, denoted by δ. The model for returns, assuming exactly x shocks, each shock having effect s_i ($i = 0, \ldots, x$), is

$$d\,ln S = \alpha dt + (\sigma\sqrt{dt})\, dz + \sum_{i=1}^{x} s_i$$

$$dz \sim N(0, 1) \qquad s_i \sim N(\theta, \delta^2). \quad (2)$$

Note that if there are no shocks ($x = 0$), the last term in (2) disappears and the simple diffusion model of instantaneous returns applies. Also, if the number of shocks is positive but fixed, the jump diffusion process and simple diffusion process are "observationally equivalent." This means that while the five parameters of a jump diffusion model might be distinguished in theory, only two parameters can be estimated, just as if the simple diffusion model applied. In short, if the number of shocks, x, is fixed, the jump diffusion process "looks like" the simple diffusion process. In this case the instantaneous return for the jump diffusion process is normally distributed, with mean $(\alpha + x\theta)$ and variance $(\sigma^2 + x\delta^2)$. Because x is fixed, we can estimate these two composite parameters but we cannot separate out the values of the components α, θ, σ, δ, and x.

The key to distinguishing between the jump diffusion model and the simple diffusion process is to allow the number of jumps experienced at any instant to vary. At some times, there are few shocks, but at other times there are many shocks. Each shock has an effect that is variable, described by a normal distribu-

Box 2: The Jump-Diffusion Model

The jump diffusion model builds on the simple diffusion model. Rather than having all variability reflected in a normally distributed "surprise," the jump diffusion model has a second source of variability in asset returns. This is the effect of a random number of "jumps," either upward or downward, in stock returns. The jump diffusion model of the log price is

$$d\mathit{ln}S = \alpha dt + (\sigma\sqrt{dt})\, dz + dv \quad dz \sim N(0,1) \quad \text{(B2.1)}$$

in which the first two terms are the simple diffusion model and dv is the variability due to the jump process.

The jump part of the model, dv, is the sum of x normally distributed "shocks." The number of shocks at any instant, x, is a random variable, distributed as a Poisson process with parameter λ. The number of shocks can range between zero and infinity, and λ is the mean number of shocks at any instant. The size of each shock, denoted by s, is normally distributed with mean θ and standard deviation δ. The mathematical description of the jump part of (B2.1) is, then,

$$dv = \sum_{i=0}^{x} s_i \quad s_i \sim N(\theta,\delta) \quad x \sim PO(\lambda) \quad x = 0, 1, 2, \ldots \quad \text{(B2.2)}$$

If the number of jumps were fixed, dv would be the sum of x normally distributed random variables, hence it would also be normally distributed. In this case, dv and dz are both normally distributed, and their sum is normally distributed; the simple diffusion model applies. Thus, it is the variability in the number of shocks that gives the jump diffusion model its power.

It can be shown that the moments for the distribution of total return, $\mathit{ln}(S_T/S_0)$, over T periods under a jump diffusion model are

Mean $(\alpha - \tfrac{1}{2}\sigma^2)T$

Standard Deviation $[\sigma^2 + \lambda(\theta^2 + \delta^2)]^{1/2}\sqrt{T}$

Skewness

$$\lambda[\theta(\theta^2 + 3\delta^2)/[\sigma^2 + \lambda(\theta^2 + \delta^2)]^{3/2}/\sqrt{T}$$

Kurtosis

$$\lambda[(3\delta^4 + 6\delta^2\theta^2 + \theta^4)/[\sigma^2 + \lambda(\theta^2 + \delta^2)]^2/T.$$

Note that if $\lambda = 0$, these parameters reduce to the simple diffusion model having zero skewness and zero kurtosis. When there are shocks, that is, when $\lambda > 0$, both skewness and kurtosis can exist.

The direction of skewness in stock returns depends solely on the mean effect of a shock. In particular, when the mean shock is negative ($\theta < 0$), the distribution of stock returns will be skewed to the left; when the mean shock is positive ($\theta > 0$), the distribution of stock returns will be skewed to the right.

Whenever shocks have a mean effect or a variable effect, the distribution of total returns will be leptokurtic, that is, the distribution will exhibit an above-normal frequency of returns around the mode.

tion with mean θ and standard error δ. The jump diffusion model postulates that x, the number of shocks in an instant, is described by a Poisson distribution, with λ, the Poisson parameter, being the expected number of jumps. Thus, the number of shocks at any instant can be as few as zero or as many as an infinite number (x = 0, 1, 2, . . .), and λ is the mean number of shocks. If λ is very low, there are few shocks on average, but occasionally there might be many shocks. If λ is very high, there are many shocks on average, but occasionally there will be instants with few shocks.

The jump diffusion model (see Box 2) has five parameters: the simple drift (α), the simple volatility (σ), the mean shock effect (θ), the standard deviation of the shock effect (δ), and the mean number of shocks (λ). The total drift and total volatility, defined to include the mean effect of jumps, are $\alpha + \lambda\theta$ and $\sqrt{\sigma^2 + \lambda\delta^2}$, respectively.

This simple extension of a diffusion process has some rich implications. The most important is that the distribution of stock returns will no longer be a normal distribution. In addition to the mean and variance of returns, which are the only characteristics of a normal distribution, there can also be skewness when θ is not zero ($\theta > 0$ implies an above-normal number of high returns, while $\theta < 0$ implies an above-normal frequency of large negative returns).

II. *The Weekend Effect*

The "weekend effect" refers to a difference in the behavior of stock returns over weekends and over intraweek trading days. This might seem vacuous at first blush, because stock markets in the United States are closed over weekends and an analysis of prices over periods with no trading seems oxymoronic. However, investors do receive and process information during periods when markets are closed, and the information processed over a weekend will affect the price at and after the opening on Monday. If the information arriving per day over weekends is of the same quantity and consequence as intraweek news, the implicit price movements over weekends should be the same as the explicit movements during the week. Thus, if the volatility over a day from close to close during a week is σ, then the volatility from close to close over a three-day interval during the week should be $\sigma\sqrt{3}$, and the same volatility should apply to a weekend from Friday's close to Monday's close. Furthermore, as we have already seen, if α is the mean daily return during intraweek trading, then the mean *daily* return over a weekend should be $(\alpha - \frac{1}{2}\sigma^2)$. In short, if a weekend were just like an intraweek day except for its length, the return volatility over weekends would be 1.73 times the daily intraweek volatility, and the daily drift over weekends would be less than the daily intraweek drift by an amount equal to $\frac{1}{2}\sigma^2$.

Table 1 (below) shows the descriptive statistics for the sample period used in this paper. For the S&P 500 (for example) the intraweek mean daily values suggest α = .000618 (.0618 percent) and σ = .008685 (.8685 percent). The implied drift over weekends is $(\alpha - \frac{1}{2}\sigma^2)$ =.00058 per day and the daily volatility on weekends should be σ = .008685. However, the weekend drift and volatility are .000115 and .005791, respectively. The mean drift over weekends is, therefore, less than during the week, and the daily volatility on weekends is less than the intraweek daily volatility.

Thus, over weekends the drift should be, and is, less than the drift found during the trading week. While this is a partial explanation of the lower returns over weekends than during the week, the weekend effect goes beyond a lower average daily return over weekends. Data from the 1970s and 1980s show that the weekend volatility is less than predicted from weekday returns data, and that the weekend drift is negative! Cross (1973) found that stock prices actually tend to decline over weekends in the three-day interval from Friday's close to Monday's close. At first this was attributed to a "Monday effect" but Rogalski (1984) found that the entire decline occurred between Friday's close and Monday's open and that the open-to-close returns on Mondays were non-negative. Har-

Table 1
Descriptive Statistics for Daily Returns
January 3, 1980 to January 29, 1999 Excluding October 1987

	S&P 500		Dow 30		Wilshire 5000		NASDAQ		Russell 2000	
	Intra-week	Week-end	Intra-week	Week-end	Intra-week	Week-end	Intra-week	Week-end	Intra-week	Week-end
Number	3,753	877	3,753	877	3,753	877	3,753	877	3,753	877
Days per year	196	138	196	138	196	138	196	138	196	138
Mean Return (Percent)										
per day	.0618	.0115	.0491	.0251	.0701	−.0034	.1084	−.0355	.0942	−.0397
per year	12.87	1.60	10.99	3.52	14.72	−.47	23.66	−4.78	20.27	−5.33
Standard Deviation (Percent)										
per day	.8685	.5791	.9035	.5984	.8015	.5420	.8483	.5974	.7118	.5028
per year	12.16	6.80	12.65	7.03	11.22	6.37	11.88	7.02	9.97	5.91
Skewness	−.224*	−1.114*	−.258*	−.926*	−.250*	−1.279*	−.498*	−1.554*	−.877*	−1.295*
Kurtosis	3.833**	6.599**	4.554**	5.810**	4.372**	7.555**	4.869**	10.437**	6.335**	7.140**

The daily returns are measured by the logarithm of relative closing prices. Weekend returns are measured as three-day log price relatives divided by 3 to convert to daily equivalents. The mean annual return and the annual standard deviation are computed as $(1 + m)^n - 1$ and $s\sqrt{n}$, where m and s are the daily mean and standard deviation and n is the number of intraweek trading days (196) or weekend days (138) in a year.

*Indicates that skewness is significantly negative at a 5 percent significance level.

**Indicates that kurtosis is significantly greater than 3.0 (the value for a normal distribution) at a 5 percent significance level.

ris (1986) further refined this, showing that prices tended to decline during the first 45 minutes of Monday trading but that early losses were recouped over the remainder of Monday. Dyl (1988), using S&P 500 futures prices, found that significant price changes are more likely to occur over weekends than during the trading week, and that price declines were more likely over weekends than intraweek. Thus, these studies created the conventional wisdom that stock prices in the United States tend to decline over weekends, an observation that became known as the "weekend effect."

The conventional wisdom has been that stock prices in the United States tend to decline over weekends, an observation that became known as the "weekend effect."

The second part of the weekend effect, the "weekend volatility effect," is that the daily volatility of returns over weekends is less than the volatility over contiguous (intraweek) trading days. French and Roll (1986) examined the descriptive statistics for returns on all common stock traded on the New York Stock Exchange (NYSE) and the American Stock Exchange (AMEX). They found that the volatility over entire weekends was only about 10 percent greater than the intraweek daily volatility. This translated to a *per day* weekend volatility well below the intraweek daily volatility. Thus, even though information might be arriving during Saturday and Sunday, either the frequency of its arrival or the volatility of returns that resulted was so low that prices behaved *as if* investors ignored any weekend information, treating Monday as if it were simply the day following Friday. French and Roll concluded that trading induces volatility and that when markets are particularly volatile a trading halt might reduce volatility.

Explanations for the Weekend Effect

The weekend drift effect has received the most attention in the literature. Miller (1988) attributes the negative returns over weekends to a shift in the *broker-investor balance* in decisions to buy and sell. During the week, Miller argues, investors are too busy to do their own research and tend to follow the recommendations of their brokers, recommendations that are skewed to the buy side. However, on weekends, investors, free from their own work as well as from brokers, do their own research and tend to reach decisions to sell. The result is a net excess supply at Monday's opening. Miller's hypothesis is supported by evidence showing that brokers do tend to make buy recommendations,[4] by evidence that odd-lot transactions tend to be net sales, and by data showing that odd-lot volume is particularly high and institutional volume is particularly low on Mondays. Thus, individual investors tend to sell on Mondays when the lack of institutional trading reduces liquidity. Ziemba (1993) provides evidence that the same phenomenon exists in Japanese stock prices.

Another explanation for the negative weekend effect is that stock prices close "too high" on Fridays or "too low" on Mondays. One variant attributes unusually high Friday closing prices to *settlement delays*. With the current T + 3 settlement schedule, settlement occurs on the third business day after the trade date. Buyers on Mondays and Tuesdays must pay during the same week (on Thursday or Friday), but buyers on Wednesday through Friday need not pay for five days because a weekend occurs before the settlement day; they get an extra two days of interest-free credit from brokers before settlement. Monday prices must be lower than Friday prices to compensate those investors who delay purchases until Monday. This hypothesis is supported by the intraweek behavior of volume and returns: Friday is the day with the greatest volume and with the most positive stock returns.

However, there are several reasons for discounting the influence of settlement delays. First, calculations of the value of the interest earned over two days gives too small a number to account for the size of the observed negative weekend drift. Second, prior to June 7, 1995, the settlement of common stocks was on T + 5, that is, buyers paid for their shares on the fifth trading day following the purchase. Thus, under the settlement practices in place at the time the weekend effect was most actively studied, it made no difference on which day the purchase was made because all purchases were paid for on the seventh business day; every settlement period included a weekend.

[4] Groth et al. (1979) found that about 87 percent of 6,000 broker recommendations were to buy, leaving only 13 percent on the sell side.

A second variant, the *dividend exclusion* hypothesis, argues that ex-dividend dates tend to cluster on Mondays and that at least part of the decline from Friday to Monday reflects the payments of dividends. Because virtually all studies of the weekend effect (including the present study) ignore dividend payments when calculating daily returns, there is a bias toward stock price decline over weekends *if* ex-dividend dates do cluster around Mondays. Any ex-dividend effects would be realized very early on Monday, after which positive returns would occur on the rest of Monday. The evidence cited above suggests some support for this pattern.

We should not ignore the hypothesis that a negative weekend drift is an anomaly that was hidden from the eyes of investors but brought to light by the data mining of academics.

According to yet another hypothesis, the *information release* hypothesis, information released during the week tends to be positive and information released over weekends tends to be negative. A firm with good news will release it quickly so investors can bid the stock price up, but bad news is an orphan, hopefully hidden from investor scrutiny by release after the Friday close. The result is that weekends are periods of absorbing bad news.

Abraham and Ikenberry (1994) find support for a *serial correlation* hypothesis, according to which Monday's price performance is conditioned on Friday's performance: A strong Friday tends to be followed by positive weekend returns; a weak Friday is followed by negative weekend returns. While only one-third of the Fridays in their sample show price decreases, these dominate the Monday results, suggesting that bad Fridays are given heavy weight in Monday trades. This observation is consistent with small investors who initiate Monday morning trades being particularly sensitive to poor performance on Fridays.

We should not ignore a final hypothesis: There is no economic rationale to justify the persistence of a negative weekend drift. Rather, it is an anomaly that was hidden from the eyes of investors but brought to light by the data mining of academics. If this is the case, the weekend effect should have disappeared after its discovery and diffusion, as investors learned to reap profits by selling short on Friday and covering their positions on Monday.

III. Estimating Jump Diffusion Models

The first estimation of a jump diffusion model, Press (1967), relied on the "method of cumulants," a method used more recently by Beckers (1981). This method involves calculating the first six moments of the sample distribution of stock returns. From these six moments, the parameters of the jump diffusion model can be estimated. While this method can be useful in developing initial estimates to be used in an iterative estimation process, it suffers from two shortcomings. First, the method is underidentified, that is, it allows several possible parameter vectors to be associated with the same distribution. Second, there are no statistical tests that would allow us to determine how reliably we can accept or reject hypotheses about the distribution of stock returns.

This study uses a more direct approach, the method of maximum likelihood. The likelihood function describing the probability of obtaining the observed sample of returns is derived from the underlying theory. This likelihood will depend upon the five jump diffusion parameters. The method of maximum likelihood, described more fully in Box 3, involves writing a computer program to calculate the likelihood associated with any specific values for the five parameters. An initial set of values is given, and the program then searches over the possible values of the parameters to find the values that maximize the likelihood. These are the maximum likelihood estimates of the jump diffusion model's parameters. We have estimated the parameters of this likelihood function using a Gauss program written by, and available from, the author.

IV. The Data

The focus of this paper is on differences in stock market returns within the week and over weekends. Because the returns on "the market" are not uniquely measured by any single stock price index, we have used five popular stock price indices: the Dow 30, the S&P 500, the Wilshire 5000, the Nasdaq Composite, and the Russell 2000. Fortune (1998) examined the

Box 3: Estimating Jump Diffusion Parameters Using Maximum Likelihood

We want to derive the likelihood function as a function of the parameters in the jump diffusion model. Let $g(R|x)$ be the probability density function for the return, conditional on the number of shocks at that instant; it is a function showing the probability that a specific value of R will be observed when there are exactly x shocks. Also let $h(x)$ be the probability density function for the number of shocks; this is a function showing the probability of each possible number of shocks. Then basic statistical theory tells us that the joint distribution of return and number of shocks is $f(R,x) = g(R|x)h(x)$ and that the marginal probability for returns will be

$$p(R) = \sum_{x=0}^{\infty} f(R, x).$$

The probability of observing a specific set of n values of R (assuming each is independently drawn from an identical distribution) will be the likelihood function

$$L(R_1, R_2, \ldots, R_n) = \prod_{i=1}^{n} p(R_i) = \prod_{i=1}^{n} \left[\sum_{x=0}^{\infty} f(R_i, x) \right]$$

$$= \prod_{i=1}^{n} \left[\sum_{x=0}^{\infty} g(R_i | x)h(x) \right], \quad \text{(B3.1)}$$

and the logarithm of the likelihood function is

$$ln(R_1, R_2, \ldots, R_n) = \sum_{i=1}^{n} ln \left[\sum_{x=0}^{\infty} g(R_i | x)h(x) \right]. \quad \text{(B3.2)}$$

In order to estimate the parameters of the jump diffusion model we must have specific functions for $g(R_i | x)$ and for $h(x)$. We assume that these are normal and Poisson, respectively. Noting that the jump diffusion model for one time period can be written

$$R = \alpha + \sigma w + \sum_{x=1}^{\infty} s_i$$

$$\text{where } w \sim N(0,1) \text{ and } s_i \sim N(\theta,\delta^2), \quad \text{(B3.3)}$$

the specific probability functions are

$$g(R_i | x) = [2\pi(\sigma^2 + x\delta^2)]^{-1/2} e^{-1/2[R-(\alpha+x\theta)]^2/(\sigma^2+x\delta^2)}$$

$$\text{and} \quad h(x) = e^{-\lambda}\lambda^x/x. \quad \text{(B3.4)}$$

Substituting these functions into the probability statement (B3.2) yields the likelihood function, which depends upon the values of the five parameters.

$$lnL(\alpha, \sigma, \lambda, \theta, \delta) = \sum_{n=1}^{N} \left\{ ln \sum_{x=0}^{\infty} \{[\mathbf{e}(-\lambda)(\lambda)^x/x!] \right.$$

$$\cdot [2\pi(\sigma^2 + x\delta^2)]^{-1/2} \quad \text{(B3.5)}$$

$$\left. \cdot\, \mathbf{e}\{-\tfrac{1}{2}\{[R - (\alpha + x\theta)]^2/(\sigma^2 + x\delta^2)\}\}\} \right\}.$$

The five parameters of interest (α', σ, λ, θ, and δ) can be directly estimated by finding the values that maximize the likelihood of occurrence of the observed sample of returns. If the sample is for 3-day breaks (as for a weekend) the per-weekend values of α, σ, and λ over weekends can be computed as 3α, $\sigma\sqrt{3}$, and 3λ, respectively.[a]

Estimation of the jump diffusion model's parameters is done using a Gauss program that is available from the author.

[a]The maximum likelihood method requires a summation over all possible values for the number of jumps at each iteration. Since the number of possible jumps is infinite, exact estimation using the method of maximum likelihood is not possible. Honore (1998) suggests a modified approach to deal with this, but we have simply summed over a large number of jumps, so that very little of the upper tail of the Poisson distribution is discarded.

properties of these five indices and found that the indices can be put into two groups. The Dow 30, the S&P 500, and the Wilshire 5000 are highly correlated, suggesting that they measure very similar segments of the market. The Nasdaq Composite and the Russell 2000 are also highly correlated with each other but are less correlated with returns from the first three indices. The returns derived from the Dow 30, the S&P 500, and the Wilshire 5000 had a beta coefficient of less than 1.0, while the Nasdaq Composite and Russell 2000 had beta coefficients above 1.0. Thus, the Dow, S&P, and Wilshire indices measure roughly the same segment of the broad market, a segment with less systematic risk than the portfolios measured by the Nasdaq Composite and Russell 2000.

Our returns data are calculated as the log price relatives for daily closing values of the five indices cited above. Thus, we use close-to-close prices of a market-capitalization-weighted index of returns rather than close-to-close or open-to-close prices for individual stocks. The data are for the period January 3, 1980 through January 29, 1999, excluding October 1987. There are 3,753 daily intraweek returns in our sample. Intraweek returns are the log relatives for closing prices on adjacent trading days (Monday and Tuesday, Tuesday and Wednesday, and so on). Weekend returns are the returns experienced over the three-day break of a normal weekend. There are 877 weekend returns in our sample period, an average of 46 in each year.[5]

Table 1 reports the descriptive statistics for our daily returns data. These support the conventional wisdom regarding the distribution of stock returns. For each of the five stock indices the daily value of the weekend return has both a lower mean and less variability than the intraweek daily returns. While weekend returns for the Dow 30 and the S&P 500 are positive (though smaller than intraweek returns), the weekend returns for the Nasdaq and Russell indices are both negative. These results are consistent with the French–Roll finding that weekends are periods of lower (perhaps negative) growth as well as of low variability. However, the larger-capitalization stocks in the S&P 500 and Dow 30 do not suffer as badly over weekends as do the small-capitalization stocks included in the Nasdaq Composite and the Russell 2000. Note that the Wilshire 5000 takes an intermediate position, with only a slight tendency to decline over weekends.

As noted above, French and Roll (1986) found that the volatility of returns over a three-day weekend was only about 10 percent greater than the daily volatility during the week; if weekends were simply three-day periods, volatility should be 73 percent greater over weekends. Our descriptive statistics support the French–Roll view that weekend volatility is especially low. For example, for the S&P 500 our sample shows a weekend volatility of 1.003 (calculated as the per day volatility times $\sqrt{3}$). This is about 15 percent greater than the intraweek daily volatility. This is true of the other four indices. Thus, we see that the volatility is especially low over weekends.

For both intraweek and weekend samples, and for all five indices, the distribution of returns is negatively skewed, showing a tendency toward larger price declines vis-à-vis large price increases. Skewness is much more negative over weekends, reflecting an increase in the chance of large price declines over weekends than during the week.

The value of the kurtosis statistic for a normal distribution is 3. Our sample shows that, for each index, the kurtosis is above the normal distribution's level, and it is higher over weekends than during the trading week. This means that the distribution of returns shows a tendency toward "peakedness," with a scarcity of returns in the middle-distance around the mean, and a greater frequency both in the extremes (the "fat tails" result) and near the mean.

Thus, our sample indicates that the returns to common stocks are not normally distributed, and that they tend both to be skewed downward and to show a high frequency of extreme values. Each of these characteristics is more pronounced over weekends than during the week, indicating that weekends are different in many respects. While something smooths out the bumps over weekends, the *mean* weekend bump is both larger and more negative than is found during the week. Weekends are safer, but news over weekends tends to arrive in bigger chunks and to be more negative.

V. The Results

The descriptive statistics in the previous section are a very rough guide to the differences between intraweek and weekend stock returns. In this section we report the results of an explicit estimation of the jump diffusion model's parameters, using the method of maximum likelihood. This estimation is done using all 4,630 returns in our sample, with the 877 weekend returns treated as single events rather than as daily returns. Because October 1987 was such an unusual month, we have excluded it from our analysis.

In order to investigate differences between intraweek and weekend parameter values, each parameter is allowed to be different over weekends than over intraweek periods. This difference is captured using

[5] The remaining six weekends in a year are "long" weekends associated with Monday holidays. These are excluded from our sample as are mid-week one-day holidays.

Table 2a

Jump Diffusion Parameters for Daily Returns[a]
Joint Estimation with Weekend Dummy Variables

January 3, 1980 to January 29, 1999, excluding October 1987

	S&P 500		Dow 30		Wilshire 5000		NASDAQ		Russell 2000	
Parameters	Intra-week	Week-end	Intra-week	Week-end	Intra-week	Week-end	Intra-week	Week-end	Intra-week	Week-end
Simple Drift (α)	.0399	.0433	.0526	.0512	.0888	.0305	.1866	.0028	.1775	−.0006
	(+3.19)*	(+.25)	(+4.78)*	(−.10)	(+9.27)*	(−5.26)*	(+16.9)*	(−15.6)*	(+24.4)*	(−20.4)*
Simple Volatility (σ)	.4023	.1937	.5841	.2104	.4490	.1723	.5070	.2089	.4400	.1663
	(+30.9)*	(−14.7)*	(+59.4)*	(−32.7)*	(+50.4)*	(−25.2)*	(+49.6)*	(−26.7)*	(+59.3)*	(−31.0)*
Jump Frequency (λ)	1.1250	.4050	.4778	.3693	.6826	.3987	.3985	.2751	.3339	.3034
	(+68.9)*	(−22.0)*	(+48.8)*	(−7.04)*	(+62.1)*	(−12.2)*	(+40.7)*	(−6.34)*	(+41.5)*	(−1.74)*
Mean Jump (θ)	.0199	−.0793	−.0007	−.0713	−.0273	−.0813	−.1851	−.1423	−.2463	−.1294
	(+1.48)	(−5.39)*	(−.40)	(−3.27)*	(−1.71)	(−2.68)*	(−11.2)*	(+2.22)*	(−12.5)*	(+5.46)*
Jump Standard Deviation (δ)	.7134	.4103	.9793	.4337	.7865	.3935	1.0204	.4935	.9121	.4015
	(+48.5)*	(−18.5)*	(62.6)*	(−41.3)*	(+50.2)*	(−25.7)*	(+70.5)*	(−32.2)*	(+47.5)*	(−33.9)*
Number of Observations	4630		4630		4630		4630		4630	

[a]Returns are measured as 100 times the daily log price relative, that is, in percent. Intraweek returns have one day between closings. Weekend returns are three-day returns, measured from close on Friday to close on Monday.

T-statistics are in parentheses. The t-statistic for the intra-week parameters is for the null hypothesis that the parameter is equal to zero. The t-statistic for the weekend parameter value is for the null hypothesis that it is different from the parameter value over intraweek intervals. An asterisk indicates rejection of the null at 5% significance.

dummy variables. Thus, if we select for attention the parameter for the mean effect of a shock, θ, then the model is estimated using a parameter defined as $\theta = \theta_0 + \theta_1 D_1$, where D_1 is a dummy variable defined as zero for intraweek observations and 1 for weekend observations. The intraweek value of the parameter is θ_0 and the weekend value is $\theta_0 + \theta_1$. The existence of a weekend effect on the parameter can be determined by assessing the statistical significance of θ_1. The same dummy variable is attached to the other four parameters (α, σ, λ, and δ).

Table 2a reports the results. The "intraweek" columns show the estimated values and the associated t-statistics for the parameters α_0, σ_0, λ_0, θ_0, and δ_0. The "weekend" columns show the parameter estimates for $\alpha_0 + \alpha_1$, $\sigma_0 + \sigma_1$, $\lambda_0 + \lambda_1$, $\theta_0 + \theta_1$, and $\delta_0 + \delta_1$, but the t-statistics for the weekend columns are for tests of whether the weekend effect is significant, that is, whether the second part of each parameter (α_1, σ_1, λ_1, θ_1, and δ_1) is significantly different from zero. If not, there is no difference between the weekend and intraweek values of the parameter.

The results support the jump diffusion model. For every index, the mean jump frequency (λ) and the standard deviation of the jump effect (δ) are both significantly greater than zero, showing that jumps do occur and that their effect is variable. For the large-cap indices (the Dow 30 and the S&P 500) and for the Wilshire 5000, the mean intraweek jump effect (θ) is not statistically significant and the coefficients have mixed signs. However, for these indices—and for the two small-cap indices—the mean jump is significantly negative over weekends. Thus, a negative weekend effect on θ is found for all indices. What is surprising is that the Nasdaq Composite and the Russell 2000 show negative values of θ during the week as well as over weekends. In fact, the small-cap indices have a smaller negative shocks size over weekends than over weekdays. It appears that the conventional wisdom that weekends are a particularly bad time to hold stocks is peculiar to the more popular large-cap indices.

The weekend results show smaller volatility over entire weekends than experienced over single weekdays. The reason is twofold. First, the weekend mean jump frequency is less than the intraweek mean frequency, a difference that is statistically significant. For example, the S&P 500 gets about 1.125 shocks on a typical weekday but only 0.405 shocks over a three-day weekend. Thus, weekends have fewer shocks than

Table 2b
Drift and Volatility of Daily Returns Joint Estimation with Weekend Dummy Variables
January 3, 1980 to January 29, 1999, excluding October 1987

Values	S&P 500 Intraweek	S&P 500 Weekend	Dow 30 Intraweek	Dow 30 Weekend	Wilshire 5000 Intraweek	Wilshire 5000 Weekend	NASDAQ Intraweek	NASDAQ Weekend	Russell 2000 Intraweek	Russell 2000 Weekend
Total Drift[a]	.0623	.0112	.0523	.0249	.0701	−.0019	.1128	−.0363	.0953	−.0399
Total Volatility[b]	.8570	.3251	.8812	.3372	.7898	.3024	.8197	.3326	.6866	.2767
Percent of Volatility Due to Jumps[c]	88.3	80.3	76.8	78.2	82.3	82.2	78.6	77.8	76.8	79.9

[a]Total drift is $(\alpha + \lambda\theta)$. The weekend total drift is for the three-day close-to-close period. Weekend total drift can be converted to a daily basis for direct comparison with weekday total drift by dividing by 3.
[b]Total volatility is $\sqrt{(\sigma^2 + \lambda\delta^2)}$. The weekend total volatility is for the three-day close-to-close period. Weekend total volatility can be converted to a daily basis for direct comparison with weekday total volatility by dividing by $\sqrt{3}$.
[c]The proportion of total volatility due to jumps is $\sqrt{(\lambda\delta^2)}/\sqrt{(\sigma^2 + \lambda\delta^2)}$.
Source: Table 2a.

do intraweek days, even though a weekend is three days long. Second, in all five indices the variability of a shock's effect is smaller over weekends.

Investors are unlikely to care about the individual parameters affecting the evolution of stock prices. Rather, they are more concerned with the mean returns and the volatility of returns. Table 2b summarizes those by calculating the "total drift" and "total volatility" implied by the parameter estimates in Table 2a. The "total drift" is defined as the mean instantaneous return in a simple diffusion model plus the expected change from jumps, or $\alpha + \lambda\theta$. It measures the normal returns that investors experience between weekdays or over weekends. The intraweek total drift is positive for all five indices, and it is smaller over weekends than during weekdays; weekends are periods of lower returns than weekdays. Whether weekends are periods of price decline depends on the index chosen: Total drift is negative on weekends for the small-cap indices (Nasdaq Composite and Russell 2000), positive for the large-cap indices (S&P 500 and Dow 30), and very small for the very broad Wilshire 5000.

The "total volatility" is the standard error of a simple diffusion model plus the contribution of jumps, or $\sqrt{(\sigma^2 + \lambda\delta^2)}$. As noted above, total volatility is lower over weekends than would be predicted from weekday data. But the differences are much more dramatic than the French–Roll study suggests. For example, the S&P 500's intraweek daily volatility is 0.8570 and the estimated weekend total volatility is only 0.3251. Weekend volatility is only about 37 percent of intraweek daily volatility, a great contrast with the French–Roll conclusion that weekend volatility was *only* 10 percent greater than intraweek volatility. In short, volatility is *much* lower over weekends than during weekdays. This is true for all five indices.

Table 2b also shows that the jump process is a very important part of stock returns. The proportion of total volatility accounted for by jumps ranges from a low of 76 percent to a high of 88 percent. Thus, at least three-fourths of stock price variability is due to jumps, strong support for a jump diffusion process.

Parameter Changes over Time

In the preceding section we treated the parameters as fixed over the entire period, with a different value over weekends than during the week. However, there are strong reasons for viewing parameters of the diffusion process as varying over time. The nature of both microeconomic and macroeconomic shocks changes over time, and this should affect the relative returns on different stock portfolios. The types of firms included in an index also change over time. For example, the S&P 500 has larger and more technology-oriented companies than it once had. Furthermore, investor behavior should lead to changes in relative returns. Also, efficient markets theory predicts that if an anomaly is identified that allows informed investors to make above-normal profits, the anomaly should disappear over time as investors take advantage of the opportunities. This will lead to a restoration of normal risk-adjusted returns. Thus, if the weekend effect of the 1980s is a true anomaly, it should disappear in the 1990s as investors change the timing of their trades, selling before weekends and buying on Monday's opening.

In order to investigate this possibility, we have split our sample into two parts: the period 1980 to September 1987 (Pre 10/87) and the period November 1987 through January 1999 (Post 10/87). A dummy variable (D_2) was formed for the Post 10/87 period; this had a zero value in the first period and a unit value in the second period. The weekend dummy variable (D_1) was also used. Each of the five parameters was constructed as (using θ as an example) $\theta = \theta_0 + \theta_1 D_1 + \theta_2 D_2 + \theta_3 (D_1 * D_2)$. Note the interaction between the weekend dummy and the Post 10/87 dummy. This allows us to estimate a separate weekend effect after October 1987 from that before October 1987.

As a result, each of the basic five parameters is now replaced by four parameters, creating a total of 20 parameters to be estimated. For each of the parameters the results shown below can be identified.

	Intraweek ($D_1 = 0$)	Weekend ($D_1 = 1$)
Pre 10/87 ($D_2 = 0$)	θ_0	$\theta_0 + \theta_1$
Post 10/87 ($D_2 = 1$)	$\theta_0 + \theta_2$	$\theta_0 + \theta_1 + \theta_2 + \theta_3$

Thus, prior to October 1987, θ_0 is the intraweek value, $\theta_0 + \theta_1$ is the weekend value, and θ_1 is the contribution of a weekend to the parameter's value. Similarly, during the week θ_0 is the value prior to October 1987, $\theta_0 + \theta_2$ is the value after October 1987, and θ_2 is the change in the parameter after October 1987.

Table 3a reports the maximum likelihood estimates of the jump diffusion model using both the weekend and Post 10/87 dummy variables. Table 3b summarizes the implications of these parameters for total drift ($\alpha + \lambda\theta$) and total volatility $\sqrt{(\sigma^2 + \lambda\delta^2)}$. Rather than discuss the individual parameters reported in Table 3a, we will focus on the summary parameters in Table 3b.

For each of the five indices, the total intraweek drift declined after October 1987, though it remained positive. Prior to October 1987 the total drift was negative on weekends, especially for small-cap stocks, a confirmation of the "weekend drift" effect found in the 1980s. However, after October 1987 the weekend drift results changed. Large-cap stocks (S&P 500 and Dow 30) exhibited about the same total drift on weekends as during the week, indicating that for them the weekend drift effect disappeared after 1987. Small-cap stocks, on the other hand, continued to experience a negative weekend drift after 1987, though it was less than prior to October 1987. For all five stock indices, the difference between intraweek and weekend drift declined after October 1987. In short, it looks as if the weekend effect disappeared for large-cap stocks, and became smaller for small-cap stocks, a result consistent with the dynamics associated with a financial anomaly: identification followed by disappearance as smart investors take advantage.

While there were substantial changes in drift after 1987, total volatility did not change much. For example, the S&P 500's total intraweek volatility was about 0.85 both before and after October 1987, but its weekend volatility was much lower (about 0.32). As found in previous sections, the total daily volatility was considerably smaller over weekends than during intraweek trading. In most cases, volatility is not very different after 1987 than before, and in all cases the weekend volatility is much less than could be predicted from the intraweek volatility.

The drift portion of a weekend effect may have been an anomaly. The weekend volatility story is persistent, suggesting stability in the intraweek and weekend differences in information arrival and information effects.

These results suggest that, particularly for large-cap stocks, the drift portion of a weekend effect may have been an anomaly with no basis in economics. Rather, it was an artifact of investor ignorance that was eliminated when investors became aware of the anomaly. The weekend volatility story is persistent, suggesting stability in the intraweek and weekend differences in information arrival and information effects.

VI. Summary

During the 1980s a flood of academic research on stock market anomalies washed over the efficient markets hypothesis. At the time, the question was whether these anomalies had an economic founda-

Table 3a

Jump Diffusion Parameters for Daily Returns[a]

Joint Estimation with Weekend and Post '87 Dummy Variables

January 3, 1980 to January 29, 1999, excluding October 1987

	S&P 500		Dow 30		Wilshire 5000		NASDAQ		Russell 2000	
Parameters	Pre 9/87	Post87	Pre 9/87	Post87	Pre 9/87	Post87	Pre 9/87	Post87	Pre 9/87	Post87
Drift (α)										
IntraWeek	.0138	.0578	.0048	.0760	.0697	.0977	.1834	.1835	.1739	.1766
	(+1.09)	(+4.19)	(+.37)	(+5.44)	(+5.93)	(+8.42)	(+16.6)	(+13.0)	(+16.3)	(+16.9)
WeekEnd	.0218	.0505	.0107	.0634	.0097	.0388	−.0354	.0357	−.0312	.0241
	(+1.69)	(+5.56)	(+.75)	(+6.38)	(+.81)	(+4.53)	(−3.72)	(+3.31)	(−3.26)	(+3.11)
Volatility (σ)										
IntraWeek	.5232	.4144	.6409	.5912	.5421	.4401	.5112	.5373	.4918	.4130
	(+38.9)	(+26.5)	(+5.0)	(+39.8)	(+42.0)	(+31.8)	(+48.0)	(+38.2)	(+45.5)	(+38.7)
WeekEnd	.2381	.1745	.3078	.1920	.2163	.1516	.2105	.2063	.1867	.1416
	(+18.9)	(+2.3)	(+26.1)	(+21.3)	(+18.5)	(+19.3)	(+25.0)	(+17.3)	(+19.0)	(+22.5)
Jump Frequency (λ)										
IntraWeek	.9446	.8479	.4673	.3398	.6461	.5158	.1916	.4531	.2430	.3550
	(+64.8)	(+48.4)	(+38.6)	(+28.1)	(+46.3)	(+33.8)	(+25.8)	(+31.1)	(+27.7)	(28.2)
WeekEnd	.4721	.3460	.2194	.2831	.4059	.3645	.1389	.3705	.2648	.4073
	(+29.1)	(+20.6)	(+18.1)	(+17.1)	(+27.6)	(+20.0)	(+14.4)	(+18.2)	(+19.4)	(+19.7)
Mean Jump (θ)										
IntraWeek	.0749	−.0136	.1344	−.1181	.0311	−.0799	−.2764	−.1957	−.1955	−.2917
	(+4.56)	(−.66)	(+9.00)	(−5.06)	(+1.70)	(−3.53)	(−14.9)	(−7.24)	(−10.4)	(−12.1)
WeekEnd	−.0885	−.0487	−.0497	−.0746	−.0973	−.0578	−.2643	−.1236	−.1420	−.1070
	(−3.76)	(−1.59)	(−1.96)	(−2.07)	(−4.18)	(−2.10)	(−.74)	(−3.28)	(−5.46)	(−3.85)
Jump Std Dev (δ)										
IntraWeek	.7132	.8049	.9162	1.1271	.7400	.8946	1.0906	1.0882	.9411	.9341
	(+42.5)	(+37.8)	(+47.6)	(+45.2)	(+4.1)	(+37.2)	(+58.1)	(+42.6)	(+46.6)	(+37.0)
WeekEnd	.3319	.4540	.3294	.5120	.3508	.4165	.4515	.4767	.3955	.3667
	(+14.4)	(+17.3)	(+13.0)	(+15.1)	(+13.8)	(+14.1)	(+16.3)	(+12.9)	(+14.8)	(+12.5)

[a]Returns are measured as 100 times the daily log price relative, that is, in percent. Intraweek returns have one day between closings. Weekend returns are three-day returns, measured from close on Friday to close on Monday. T-statistics are in parentheses. All t-statistics are for the null hypothesis that the parameter is equal to zero.

tion, so they did not reflect opportunities for profit and could be expected to be persistent, or whether they were real anomalies arising from investor ignorance that would be eliminated as investors became aware of them and took advantage of profitable opportunities.

This paper revisits the weekend effect, interpreting the conventional wisdom as having two parts. First, stock returns are normally positive during the week but are negative over weekends. There are a number of hypotheses suggesting that this arose from variation across the week in the risks and rewards of investing in stocks. If, on the other hand, this pattern is simply a financial anomaly, we would expect investors to eliminate it over time by shifting their sales toward Fridays and their purchases toward Mondays. Second, stock returns are less volatile over weekends than they should be, given intraweek volatility. This pattern has been interpreted by some as suggesting that trading itself promotes volatility, with the implication that trading halts can be a stabilizing force when markets are chaotic.

We have revisited these weekend effects by directly estimating the parameters of the probability distribution of stock returns. These propositions are tested using daily close-to-close data for the S&P 500 from January 1980 through January 1999. Unlike previous studies that were based on descriptive statistics from a large number of individual stocks, this paper is based on the explicit estimation of the parameters of

Table 3b
Drift and Volatility of Daily Returns
Joint Estimation with Weekend and Post '87 Dummy Variables
January 3, 1980 to January 29, 1999, Excluding October 1987

	S&P 500		Dow 30		Wilshire 5000		NASDAQ		Russell 2000	
Values	Pre 9/87	Post87	Pre 9/87	Post87	Pre 9/87	Post87	Pre 9/87	Post87	Pre 9/87	Post87
Total Drift[a]										
IntraWeek	.0846	.0462	.0677	.0359	.0898	.0565	.1305	.0949	.1264	.0730
Weekend	−.0200	.0336	−.0002	.0422	−.0298	.0177	−.0721	−.0101	−.0688	−.0194
Total Volatility[b]										
Intraweek	.8684	.8491	.8961	.8838	.8048	.7788	.6994	.9083	.6760	.6931
Weekend	.3297	.3190	.3443	.3333	.3110	.2936	.2695	.3560	.2762	.2735
Percent of Volatility Due to Jumps[c]										
Intraweek	79.8	87.3	69.9	74.3	73.9	82.5	68.2	80.6	68.6	80.3
Weekend	69.2	83.7	44.8	81.7	71.9	85.6	62.4	81.5	73.7	85.6

[a]Total drift is $(\alpha + \lambda\theta)$. The weekend total drift is for the three-day close-to-close period. Weekend total drift can be converted to a daily basis for direct comparison with weekday total drift by dividing by 3.
[b]Total volatility is $\sqrt{(\sigma^2 + \lambda\delta^2)}$. The weekend total volatility is for the three-day close-to-close period. Weekend total volatility can be converted to a daily basis for direct comparison with weekday total volatility by dividing by $\sqrt{3}$.
[c]The proportion of total volatility due to jumps is $\sqrt{(\lambda\delta^2)}/\sqrt{(\sigma^2 + \lambda\delta^2)}$.
Source: Parameter estimates in Table 3a.

the distribution of returns on stock indices. We have used five indices: two large-cap indices (the S&P 500 and the Dow 30), two small-cap indices (the Nasdaq Composite and the Russell 2000), and one very broad index (the Wilshire 5000).

Our results are derived from maximum likelihood estimation of the five parameters of a jump diffusion model of stock returns. The five parameters are the instantaneous mean return and volatility of a simple diffusion process (α and σ, respectively) and the three parameters of the jump process: the mean frequency of a shock (λ), the mean size of a shock's effect on returns (θ), and the standard deviation, or volatility, of a shock's effect on returns (δ). The results suggest strong support for the jump diffusion model. Not only are the instantaneous mean return and volatility estimates sensible, but the parameters associated with the jump part of the process are also sensible and significant.

Our analysis of parameter estimates (Tables 2a and 2b) strongly supports the jump diffusion model of stock returns: At least 75 percent of total volatility is attributable to the jump part of the diffusion model. We also find strong support for the weekend volatility effect. If weekends were the same as weekdays, volatility would be about 73 percent greater over three-day weekend periods than during weekdays. But we find that volatility over an entire three-day weekend is only about 40 percent of the volatility during a single intraweek trading day. The lower weekend volatility occurs for two reasons. First, the arrival rate of new shocks (λ) is much lower over weekends than over weekdays: We estimate for the S&P 500 that there are about 1.2 shocks per day during the week, but only about 0.4 shocks per three-day weekend. Second, the standard error of the impact of a shock on a weekend is only about half the weekday standard error. In short, weekends have less new information, and the variability of the effect of a bit of new information is smaller over weekends. These results are of the same order for all five stock indices.

The weekend drift effect is more mixed. For all five indices, the normal return, or drift, over weekends is smaller than during weekdays. However, the major finding of the 1980s literature—that weekend returns are actually negative—is not found for indices containing large-capitalization stocks (the S&P 500 and the Dow 30). The negative weekend drift is a phenomenon confined primarily to small-cap stocks, such as those in the Nasdaq Composite and the Russell 2000, but these indices also show a negative drift during the week!

Thus, weekends do provide smaller returns than weekdays, but they also have smaller return volatility.

The mystery of a negative weekend drift seems to have resolved itself for large-cap stocks, but it remains for small-cap stocks. For all five stock indices, the difference between weekday and weekend performance has narrowed as time has passed, suggesting that the weekend drift effect is a financial anomaly that will ultimately correct itself.

For all five stock indices, the difference between weekday and weekend performance has narrowed as time has passed, suggesting that the weekend drift effect is a financial anomaly that will ultimately correct itself.

Table 3b reports the drift and volatility results when the estimates are split into Pre-October '87 and Post-October '87 periods. The volatility differences between weekdays and weekends that were reported above remain; that is, the ratio of weekend-to-weekday volatility is about the same after 1987 as before. But the weekend drift effects are different after 1987. In particular, after 1987 the weekend drift is about the same as the weekday drift for the large-cap indices, and, though still negative, the weekend drift is closer to the weekday drift for small-cap stocks.

This is consistent with a diminution or elimination of the weekend effect in the 1990s, and it suggests that the 1980s efforts to provide an economic rationale for negative returns over weekends in terms of economics might have ignored the possibility that investors are typically ignorant of subtleties about stock returns. Once made aware of anomalies through the data mining of financial economists, investors behave in ways that eliminate the newfound anomaly.

The sharp decrease in volatility over weekends is consistent with the view that active trading actually increases volatility, so that a close in trading will be consistent with a reduction in volatility. However, important differences between trading halts under circuit breakers and trading halts for weekends suggest that a volatility decline in the latter might not occur in the former. A weekend is a scheduled event, while a circuit-breaker halt is an unscheduled event, resulting from chaos in markets and excessive price variations. The scheduled event of a weekend might simply reduce the rate of new information flow, creating a decline in volatility, while a sudden halt in trading from a circuit breaker might eliminate all information flow from price discovery, creating an environment that elicits the volatility it is designed to mitigate.

References

Abraham, Abraham and David Ikenberry. 1994. "The Individual Investor and the Weekend Effect." *Journal of Financial and Quantitative Analysis*, 29, June, pp. 263–77.

Beckers, Stan. 1981. "A Note on Estimating the Parameters of the Diffusion-Jump Model of Stock Returns." *Journal of Financial and Quantitative Analysis*, 16, March, pp. 127–38.

Cox, John, and Stephen Ross. 1976. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics*, 2, pp. 145–66.

Cross, Frank. 1973. "The Behavior of Stock Prices on Fridays and Mondays." *Financial Analysts Journal*, November/December, pp. 67–69.

Dyl, Edward. 1988. "A Possible Explanation of the Weekend Effect." *Financial Analysts Journal*, May/June, pp. 83–84.

Fortune, Peter. 1998. "A Primer on Stock Price Indices." Federal Reserve Bank of Boston, *New England Economic Review*, November/December, pp. 25–40.

French, Kenneth and Richard Roll. 1986. "Stock Price Variances: The Arrival of Information and the Reaction of Traders." *Journal of Financial Economics*, 17, September, pp. 5–26.

Groth, J., W. Lewellen, G. Schlarbaum, and R. Lease. 1979. "An Analysis of Brokerage House Security Recommendations." *Financial Analysts Journal*, pp. 32–40.

Harris, Lawrence. 1986. "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns." *Journal of Financial Economics*, 16, May, pp. 99–117.

Honore, Peter. 1998. "Pitfalls in Estimating Jump Diffusion Models." The Aarhus School of Business, Denmark, mimeo, January 24.

Johnson, Gordon and Thomas Schneeweiss. 1994. "Jump Diffusion Processes in the Foreign Exchange Markets and the Release of Macroeconomic News." *Computational Economics* 7, pp. 309–29.

Kim, Myung-Jug, Young-Ho Oh, and Robert Brooks. 1994. "Are Jumps in Stock Returns Diversifiable? Evidence and Implications for Option Pricing." *Journal of Financial and Quantitative Analysis*, 29, December, pp. 609–31.

Merton, Robert. 1976. "Option Pricing When Underlying Stock Returns Are Discontinuous." *Journal of Financial Economics*, 3, January–March, pp. 125–44.

Miller, Edward. 1988. "Why a Weekend Effect?" *Journal of Portfolio Management*, 14, Summer, pp. 42–48.

Press, John. 1967. "A Compound Events Model for Security Prices." *Journal of Business*, 40, July, pp. 317–35.

Rogalski, Richard. 1984. "New Findings Regarding Day of the Week Returns Over Trading and Non-Trading Periods." *Journal of Finance*, 39, pp. 1603–14.

Ziemba, William. 1993. "Comment on 'Why a Weekend Effect?'" *Journal of Portfolio Management*, Winter, pp. 93–99.