Minute by Minute: Efficiency, Normality, and Randomness in Intradaily Asset Prices

Feinstone, Lauren J.

Working Paper No. 15
June 1985

Minute by Minute: Efficiency, Normality, and Randomness in Intradaily Asset Prices

Lauren J. Feinstone *

University of Rochester

and Rochester Center for Economic Research

June 1985

ABSTRACT

In this study we test the efficiency of asset markets at intervals as short as thirty seconds. We also describe the properties of a simple new stochastic process as a potential model of the behavior of asset prices and test it on intradaily Deutsche Mark futures prices. According to this process, asset prices are constant between economically relevant events, which occur at the random times generated by a Poisson process. At the moments of these events, prices jump to new values; the size of the jump is drawn from a normal distribution. Tests of this process indicate that it cannot be rejected for almost all days in the sample.

* This paper is part of my doctoral dissertation undertaken at the University of Chicago. I especially wish to thank my adviser, Jacob A. Frenkel, for his boundless inspiration and support. I also thank Lester G. Telser and James J. Heckman for their incisive comments and suggestions. Any remaining errors are my own.

1. Introduction

A key issue in the study of asset markets has been the speed at which market prices reflect new information, or "news." Many studies have investigated the efficiency of a multitude of asset markets using daily, weekly, and monthly data, with the conclusion that they efficiently incorporate publicly available information. Only a few studies have pursued the question of market efficiency within the day (see, for example, Wasserfallen and Zimmerman, 1981), despite the fact that it is only possible to see if the release of news has an immediate impact on market prices if one is very precise about the timing.

A simultaneous and related line of finance literature has pursued the stochastic process generating asset prices. Here again there are few studies using intradaily data (see, for example, Niederhoffer and Osborne, 1966), despite the obvious advantage that can be gained by using continuous time asset price series.

In this study, we propose to extend the study of the efficiency of asset markets to intradaily data, envisioning a floor trader or an individual who follows the market during the day. Although this study is univariate in that only one price series is examined, the data set contains explicit information about the times of trades. This allows us to perform an improved version of a filter rule test, autocorrelation tests at intervals as short as 30 seconds, and particularly extensive tests concerning the price process. It also sheds light on one previously unexamined aspect of the nature of news: namely that the market knows as little about when the news will occur as it knows about the content of the news.

In order to carefully confront the greater precision afforded by an intradaily data set, we explore an intuitively simple stochastic process as a model of the behavior of asset prices. This process is an attempt to resuscitate the properties of normality and continuity that appealed to Bachelier (1900) and other economists because of the likelihood of eventually providing underlying models. While it is an example of the much larger class of subordinated processes proposed by Clark (1973), it is an economically appealing variant because it explicitly incorporates the effect of news. According to this process, new economically relevant events occur randomly over time. When the information about the event is revealed, prices jump instantaneously to reflect that information; between events, since there are no economic forces impinging on prices, prices are constant.

A formal model is provided by postulating that events are generated by a Poisson process, while the price jumps at the times of the events are drawn from a normal distribution. Hence the name for this process -- the compound Poisson-normal process, or CPN. Both assumptions of the CPN process are tested below; we cannot reject either the hypothesis that the timing of events is Poisson or the hypothesis that the price jumps at the times of occurrence are normally distributed. The evidence for normality is particularly strong -- on only one day in the month-long sample is normality rejected.

That the jump process just described is a random walk that is continuous only almost always is a satisfactory feature of this model, as it will also be shown that the data are inconsistent with the hypothesis that the price process is a random walk that is continuous everywhere. That is, the (continuous) Wiener process is rejected as generating the observed prices, because we find that price changes over intervals of uniform length are not normally distributed, while the evidence strongly supports the hypothesis that prices are not autocorrelated.

Some properties of stochastic price processes can have major implications for the behavior of price-taking market participants, although sometimes the properties of greatest economic interest are not the properties that are directly tested. For example, if an asset's price is a Wiener process, a market participant can be perfectly sure (with probability 1) that any stop or limit order he places will not be passed through before it can be executed, because prices move continuously almost always. According to a jump process such as the CPN, stop and limit orders will occasionally be ineffective in preventing large capital losses. Knowledge of the parameters of the process would allow us to work out the probability of such occurrences. Only by testing both the autocorrelation properties and the distributional properties can we distinguish between these possibilites.

Finally, it has been noticed that rates of return in many asset markets have sample distributions that have been leptokurtic (sharply peaked and fat-tailed) relative to a normal distribution with the sample mean and variance. This is at odds with the Wiener hypothesis, and has led to the study of the stable Paretian distribution (e.g. Mandelbrot, 1963 and Fama, 1963). However, the stable Paretian distribution has infinite variances of prices over any finite time interval and generates sample paths that are almost all discontinuous at almost every moment of time. These implications for the behavior of prices are at the opposite extreme from the properties that drew Bachelier to study the Wiener process. The returns generated by the CPN process, which has finite variances and is continuous except at discrete moments, are leptokurtic. Moreover, the extent of the leptokurtosis diminishes as the length of the sampling interval increases, a fact which has also been observed.

The CPN process requires the estimation of extremely few parameters (in this study, only two are required; the maximum would be three). Although this is an appealing feature of the CPN process, the nature of intradaily data makes difficult the comparison of the CPN process to alternatives which require more parameters to be estimated. Thus, in this study the chief alternative is the Wiener process.

The data set used in this study consists of Deutsche Mark futures prices provided by the Chicago Mercantile Exchange, of which the International Monetary Market (IMM) is a division. Futures prices have the advantage of representing open outcry bidding, unaffected by the role of market specialists as in the stock market. In addition, unlike for forward prices, the data are actual sales prices rather than bid and ask spreads. Finally, the question of market efficiency and the role of news has been particularly vigorously pursued in the literature on exchange rates (Frenkel, 1981).

In section 2 of this paper we briefly explore the CPN process, discussing its properties and deriving the various testable density functions that it implies. We consider the special case of its application to intradaily data and to futures prices, both of which require minor modifications. In section 3 we provide the empirical tests of market efficiency, of the Wiener process, and of the CPN process. The conclusions are in section 4.

2. The CPN Process

2.1 Randomly Timed Events

Suppose that economically relevant events occur at random times. If the market is efficient at evaluating this information, we will observe prices changing instantly at the revelation of these events. The randomness of the timing of events or of price jumps can be captured by positing that the number of events per unit time interval is an independent drawing from a Poisson distribution. That is, let

$x(t)$ = the number of events prior to time t, $x(t)\in\{0,1,2,\ldots\}$ $\forall t$.

Then $\{x(t)\}$ is a Poisson process if

(1) $$pr\{x(s)-x(t)=j\} = (1/j!)\, e^{-c(s-t)} c^j (s-t)^j, \quad s>t, \ j\in\{0,1,2\ldots\}.$$

Here c represents the average number of events per minute. The Poisson process is purely random in the sense that, although knowledge of c allows an individual to predict that the number of events in a time interval of length s-t is $c\cdot(s-t)$, the individual still will not be able to predict when in the interval the events are most likely to occur.

If $\{x(t)\}$ is a Poisson process, then for any fixed interval of observation, $h>0$, the first differences of $\{x(t)\}$, i.e. the number of events in each interval of length h, will be stationary independent random variables with the Poisson density function

(2) $$f(j) = (1/j!)e^{-ch}(ch)^j, \qquad j\in\{0,1,2,\ldots\},$$

where ch = the expected number of events per interval of length h.

Suppose that at the times of the events, price jumps of size p are drawn from an underlying normal distribution with mean ν and variance τ^2. Normality of price movements is appealing because the central limit theorem provides the possibility of supplying underlying models with large numbers of agents (see, for example, Telser, 1981). It is fundamental because of its formal simplicity and its long history of study as a distribution relevant to asset markets (e.g. Bachelier, 1900). Nevertheless, it is an independent assumption from the assumption that the timing of events is randomly generated by a Poisson process.

We will refer to the combined process generating prices that is described here as the compound Poisson-normal process, or CPN. It is a particularly simple example of a subordinated process (Clark 1973). Not only does it have only three parameters (c, ν, and τ^2), but the driving process, the Poisson, is both formally easy to work with and intuitively easy to motivate. Although some economically relevant events (such as weekly money supply announcements) do not occur randomly over time, most events (war in Iran, greater production in a particular steel factory, changes in the demand for money) are unpredictable in timing as well as in magnitude. Such is the nature of news.

With sufficiently detailed data these two hypotheses about the timing of events and the distribution of price jumps can be independently tested. In particular, an intradaily data set such as a ticker tape would contain the times at which prices moved, allowing computation of the number of such movements per minute, as well as the size of each movement. Moreover, it is possible to jointly test both hypotheses with any ordinary time series of asset prices observed only at intervals of uniform length. Such a test would allow direct comparison with other processes that have been suggested for asset price movements. To do this joint test, we need the distribution of changes in prices (or rates of return) that would be generated by observing the CPN process at intervals of uniform length. We turn next to the derivation of that distribution.

2.2 The Distribution of Price Changes

Suppose a particular interval of length h contains exactly $z_h \geq 0$ events, i.e. prices jump exactly z_h times in that time interval. At each such event, the price jump p_j, $j=1,\ldots,z_h$, is drawn independently from the distribution $N(0,\tau^2)$. We restrict the mean to zero both for simplicity and because it is an interesting economic hypothesis that we will later test -- none of the results in this section depend on this restriction. Now the total price change for that interval is

$$p_h = \begin{cases} \sum_{j=1}^{z_h} p_j & \text{if } z_h > 0 \\ 0 & \text{if } z_h = 0. \end{cases}$$

Next, recall that z_h is itself a random variable whose distribution function is Poisson (with density function $f(z_h)$ given by equation (2)). That is, price changes p_h are sums of normally distributed random variables, where the number of terms in the sum is distributed Poisson. Since the sum of j random variables distributed $N(0,\tau^2)$ is distributed $N(0,j\tau^2)$, we can write the probability that a price change is $\leq p_h$ as the product of the probability that there is one jump per interval times the probability that a normal density function with variance $1\tau^2$ generates a value $\leq p_h$, plus the product of the probability of two jumps per interval times the probability that a normal density function with variance $2\tau^2$ generates a value $\leq p_h$, etc. In the case that $p_h \geq 0$, we must also add the discrete probability that there are no events in the interval at all, which according to the Poisson density function is e^{-ch}. This can be summarized in a density function with the following form:

$$(3) \quad \varphi(p_h) = \begin{cases} \sum_{j=1}^{\infty} \left[\frac{1}{\tau\sqrt{2\pi j}} \exp\left\{ - \frac{p_h^2}{2j\tau^2} \right\} \right] \left[(1/j!) e^{-ch} (ch)^j \right] & \text{if } p_h \neq 0 \\ e^{-ch} + \sum_{j=1}^{\infty} \left[\frac{1}{\tau\sqrt{2\pi j}} \right] \left[(1/j!) e^{-ch} (ch)^j \right] & \text{if } p_h = 0. \end{cases}$$

It is because of equation (3) that we call this process a compound Poisson-normal. A compound distribution, or mixture, is the distribution that results when a parameter of some distribution is itself a random variable. In our case, the variance of the normal distribution, $z_h\tau^2$, is a itself a random variable whose distribution is Poisson. Using obvious notation, we will sometimes refer to the above process as the CPN(ν,ch,τ^2); in the above derivation we restricted ν=0.

Properties of the CPN distribution function and $\varphi(p_h)$ are derived and discussed in Feinstone (1984). In particular, it can be shown that its mean is ν, its variance is $ch\tau^2$, and its kurtosis is

$$\gamma_2 = \frac{3(ch+1)}{ch} - 3 = \frac{3}{ch} > 0.$$

That is, it is leptokurtic, or sharply-peaked and fat-tailed relative to a normal distribution with the same mean and variance. Moreover, like the well-known Wiener process, the CPN process is a continuous time random walk: because of the independent draws from the underlying Poisson process, it is infinitely divisible, which is to say that increments in prices are independent over any disjoint time intervals. Unlike the Wiener process, however, the sample paths followed by prices under a CPN process are frequently discontinuous, exhibiting discrete jumps at events.

In the remainder of the paper, we present a variety of tests of the CPN process using an intradaily data set which contains sale prices and the times of sales. We test the infinite divisibility of the process using autocorrelation tests and filter rules; as is well-known, these tests simultaneously bear on the efficiency of the asset market in question. We separately test to see if the number of events per time interval is consistent with the Poisson process, and if the rates of return at the times of the events are normally distributed. We also perform a joint test to see whether the returns are distributed according to the CPN density function $\varphi(p_h)$. Finally, we also test the alternative hypothesis that the intradaily prices are generated by the Wiener process, according to which the returns or price changes observed at intervals of uniform length h would be distributed normally with mean μ and variance $\sigma^2 h$. Because of the large number of parameters involved, we are unable to test several other alternative hypotheses that have been presented elsewhere, such as the stable Paretian process (Mandelbrot, 1963) or the compound ln-normal-normal process (Clark, 1973).

Before we can turn to the empirical evidence, we must discuss how we account for two features of the particular data set, which has been supplied by the Chicago Mercantile Exchange and contains data from the futures market in the Deutsche Mark (DM). The first of these features is the effect of rounding the data to the nearest point, which is quite constraining on asset prices within a single day. The second concerns the appropriate way to measure asset prices in a futures market.

2.3 The Effect of Rounding the Price Data

Consider first the price changes or rates of return generated by a Wiener process. According to the Wiener process, for any interval h>0, the net change in price over that interval, p_h, is drawn independently from a normal distribution with constant mean, and variance proportional to h. If these price changes are observed only after being rounded to the nearest point, then the observed price changes p_h are drawn from a distribution with discrete density function

$$(4) \qquad \bar{g}_h(p_h) = \int_{p_h-.5}^{p_h+.5} g_h(x)\,dx \quad , h>0,\ p_h \in \{0,\pm 1,\pm 2,\ldots\},$$

where $g_h(x)$ is the normal density function with mean μ and variance $\sigma^2 h$.

In the case of the CPN process, it is price jumps at the time of occurrence rather than the accumulated changes in prices over an interval that are normally distributed. In the particular data set provided by the CME, trades that give rise to no price change are not recorded. Because of the effect of rounding, price jumps within half a point of 0 are thus truncated from the sample. Thus, under the CPN hypothesis, we expect observed price jumps, p, to be drawn from the truncated and rounded distribution function

$$(5) \qquad \bar{g}(p) = \left[\int_{p-.5}^{p+.5} g(x)dx\right] \Big/ \left[1 - \int_{-.5}^{.5} g(x)dx\right] \quad , \quad p \in \{\pm 1,\pm 2,\ldots\},$$

where g(x) is the normal density function with mean ν and variance τ^2, and where the denominator acts to spread the unobserved probability weight associated with the interval (-.5,.5) proportionately over all observable values of p.

Finally, the density function that results when price changes over intervals of length h are generated by the CPN density function $\varphi(p_h)$ but are observed only after rounding is

$$(6) \quad \bar{\varphi}(p_h) = \begin{cases} \displaystyle\sum_{j=1}^{\infty} \left[\int_{p_h-.5}^{p_h+.5} \frac{1}{\tau\sqrt{2\pi j}} \exp\left\{-\frac{x^2}{2j\tau^2}\right\}dx\right] \left[(1/j!)e^{-ch}(ch)^j\right], p_h=\pm 1,\pm 2,\ldots \\ \displaystyle e^{-ch} + \sum_{j=1}^{\infty} \left[\int_{-.5}^{.5} \frac{1}{\tau\sqrt{2\pi j}} \exp\left\{-\frac{x^2}{2j\tau^2}\right\}dx\right] \left[(1/j!)e^{-ch}(ch)^j\right], \ p_h=0. \end{cases}$$

2.4 Asset Prices in the Futures Market

In this section we discuss the ways to appropriately measure asset prices or rates of return using intradaily futures market data.

In the absence of transactions costs or collateral, the buyer of a futures or forward contract in a commodity does not pay for the commodity until the maturity date, T, of the contract.[1] Just as the seller has no claim on the buyer's money until date T, the buyer has no claim on the commodity until date T. However, at any time after t, when the future is bought, and prior to T, the buyer has a current claim on the contract itself -- that is, on the right (and obligation) to take delivery of the commodity in exchange for the agreed upon price when date T arrives. This distinction is economically non-trivial: it is the difference between a commodity and an asset, and allows us to derive the market price of the asset associated with the commodity's future price.

If an individual buys a futures or forward contract for a unit of the commodity deliverable at date $T>t$, we call the price at which the exchange will take place at time T the "commodity futures prices," and denote it $F^T(t)$. (The word "commodity will hereafter frequently be suppressed.) If the futures price happens to rise at a later time $t_1<T$, the buyer will find that he can sell his contract in the market at a price equal to the difference $F^T(t_1)-F^T(t)$, because a buyer at time t_1 is indifferent between the prospect of surrendering $F^T(t_1)$ dollars at time T for the commodity, and the prospect of surrendering $F^T(t_1)-F^T(t)$ dollars at time t_1 and $F^T(t)$ at time T. (In

---

[1] It is convenient to consider foreign currency to be a commodity when discussing futures markets in foreign exchange. This classification is purely taxonomic and does not preclude the economic function of money as an asset.

theory, at time t_1 the agent's asset would only be worth the discounted present value of the change in the futures price. In practice, traders in futures agree to neglect discounting. This is enforced by the Clearing House of the futures market, which daily transfers the full face value of the price change from the margin accounts of the holders of contracts on which there has been a capital loss to the accounts of the holders of contracts on which there has been a capital gain. This is one potential source of the small price difference between futures and forward contracts, since the latter must be discounted to be resold.) Thus the individual who bought the contract at time t is in possession of an asset; the asset price at time t_1 is

$$p(t_1) = F^T(t_1) - F^T(t) \quad . \tag{7}$$

Note that the value of the asset consists entirely of accumulated capital gains and losses.

It is possible to extend the uniformity of contracts by a process known as "marking." Two individuals buying (or selling) futures contracts at different times, t_0 and t_1, such that $F^T(t_0) \neq F^T(t_1)$, are holding assets that differ in asset price as well as futures price. If the contracts are identical in size and delivery date, however, an accounting intermediary can perform the following service at time t_1: for each party to the earlier, t_0, contract, it can issue a new contract at futures price $F^T(t_1)$, simultaneously requiring the surrender of the old contract along with payment made by the party who incurred the capital loss (for example, the seller, if $F^T(t_1)$) $F^T(t_0)$,) to the party who received the capital gain. Thus at time t_1, all outstanding contracts are made identical not only in size and maturity date but in futures and asset prices as well.

This extended uniformity can be achieved amongst futures contracts, but not amongst forward contracts. Forward contracts in the organized forward market are not necessarily identical in size (although blocks traded in round numbers are common), or maturity date (except on the first date of purchase). The day after a forward contract with maturity date T is sold, the organized market no longer trades contracts of that maturity date, but of date T+1. A forward contract can be resold after the purchase date, but not without direct contact between the parties involved. In contrast, the futures market continues to trade contracts of maturity date T until date T arrives, and all contracts are for identical quantities of the commodity.

The Clearing House of the futures market is able to perform the accounting of marking by requiring margin deposits for all parties involved in transactions. While in practice it "marks" its customers to the market only at the end of each day, it may equivalently be viewed as recrediting and redebiting the accounts of its trading members with each price change.[2] Thus a buyer need not ask if the seller with whom he is matched on the trading floor bought his contract ten minutes or ten days ago. All contracts on the market are identical at each moment.

A consequence of this is that asset prices in the futures market are equal to zero at all times except at the moments when prices actually change. Thus, rates of return do not exist. The implications of this for models of asset price behavior have been explored by Black (1976). In general the changes in commodity futures prices rather than their rates of return can be

---

[2] Black (1976) described the daily rewriting of the contracts, but did not note that the Clearing House actually provides this service continuously. He wrote that this "applies only to the end of the day, after the futures contract has been rewritten. During the day, the futures contract may have a positive or negative value, and its value will be equal to the value of the corresponding forward." Black (1976), p.170.

expected to exhibit the same sort of behavior as rates of return on unmarked assets. Thus, in the following sections we measure asset prices in the futures market as the changes, rather than the percentage changes, in the futures prices.

3. Empirical Evidence

3.1 The Data Set

The data set used throughout this study consists of intradaily futures prices on the Deutsche Mark (DM) as traded at the IMM during July 1977 for maturity in September 1977. The sale price data are reported with the times of the sales within ten second intervals. Only sales which give rise to a change from the last recorded price are reported, and the number of contracts traded are not reported. That is, the data are not what are called "transactions data," even though the data set is unusually detailed in its timing information.

In Table 1. we present a summary of the data. The average price of a DM during the month was $.44, giving a contract (125,000 DM) an average value of $55,000. Prices are reported in points, with one point equaling $.0001. Thus, a one point change in the price of a DM represents $12.50 in capital loss or gain for each contract owned. Of the 19 trading days in the month, only the first day had a volume of fewer than 200 contracts traded; the average number of contracts exchanging hands on a given day was 839. Each trading day lasts 4.5 hours, from 8:48 a.m. to 1:18 p.m. (Chicago time); the prevailing futures price moved, or jumped, on average 69 times per day, or about once every four minutes. Despite the close match between settlement prices and the averages of daily prices sampled at thirty second intervals, it is seen that prices exhibit wide swings during the day.

Price movements on two fairly typical days are plotted in Figures 1 and 2. The data points are connected for ease of viewing. For some purposes in this study we create data sets consisting of the set of prices that would be observed if observations were taken at intervals of uniform length h (in particular, h = thirty seconds and h = three minutes). The price at any time of observation is taken to be the last recorded price in the raw data set, which does not contain observations on trades that give rise to no price change, even though such trades frequently occur. The first differences of these recreated price time series will be called price "changes," while the set of changes in prices at the times of occurrence will be called price "jumps."

3.2 Tests of Efficiency and Infinite Divisibility

In this section we report on autocorrelation tests to see if the data are consistent with the economic hypothesis that the market is efficient in its use of the information contained in past prices. Because the data set contains timing information, we are also able to perform filter rule tests that improve on previous filter rule tests of market efficiency. All these tests can equivalently be viewed as testing the formal hypothesis that the process generating the data is infinitely divisible or has independent increments.

The autocorrelation tests can be performed for any fixed observation interval h. Two particular time intervals have been chosen for analysis here: h = 3 minutes and h = 30 seconds. Longer intervals were examined with a concomitant reduction in sample size. The results did not differ.

In Tables 2 and 3 we present the autocorrelation functions obtained by recreating the daily price change series that would have been observed by sampling at these intervals. As may be seen, few autocorrelation coefficients are significantly different from zero, and we cannot reject the hypothesis that prices are a random walk for any observation interval $h>0$. Even where such non-zero coefficients occur, no outstanding patterns are evident, although some regularity does appear on days 18, 19, and 28. To use this regularity, the chartist must be able to predict these days or recognize them early in the day. A chartist with no transactions costs who not only receives price information with a delay, but also takes the time to analyze its conformation to patterns such as "heads and shoulders," would profit only by coincidence in this market. Nevertheless, the tests most appropriate for testing the chartist's hypothesis are filter rule tests conducted on all data points (as opposed to only those sampled at fixed time intervals).

We now consider these filter rules. A day trader -- one who holds a net non-zero position only during the day -- can use the information contained in all prices. If price rises are correlated, he can profit by buying at the start of an upswing, and selling as close to the peak as possible. We cannot study a trader who uses information other than that contained in the price series itself to judge the troughs and peaks. However, we can model a trader who is a univariate chartist.

Alexander (1961) considered the following sort of filter rule. Suppose that a trader chooses an arbitrary number of points (or an arbitrary percentage), n. If the price rises by n points, he buys a contract and waits. When the price then falls by n points from a subsequent peak (the peak could be the price at which he bought if he is unlucky), he sells out, and also sells another contract. He then holds this short position until the price rises by n points, buys, and continues this process. At the end of the day, he closes out his position.

An n-point filter like that described above will generate many transactions if n is small. This will result in high brokerage fees, and more physical exhaustion. (Anyone who has seen the floor of an exchange knows that continuous-time trading is a strenuous activity.) On the other hand, if n is high, the turning points may be missed completely.

Alexander's filter rule has been criticized on the grounds that a trader is unlikely to always be able to trade at the price that triggers his desire to trade. Not only does time elapse before his bid or offer is presented, but his bid or offer is new information to the other floor traders about the quantities available, and may itself affect the price.

We can consider a more sophisticated n-point filter or "trigger" rule. Suppose the price rises by n points, triggering a desire to buy at that price. Let the trader immediately bid at that price. Several things can happen. One is that an offer is tendered at his bid price, and he does buy at that price. In this data set, such an action is not marked because the price would not change. We assume arbitrarily that if 60 seconds pass with no price jump after the trigger then the trader was able to conclude his deal at the trigger price.

Possibly, however, all offers are higher than his bid. In this case, though the trader was right that the price would rise, he cannot take the whole rise as a capital gain. We assume that he immediately takes the lowest of the offers, so that the next observation of a higher price (lower, if he desired to sell) is the price at which he is assumed to have concluded his deal. That is, if we observe that within 60 seconds of the n-point trigger the price jumped in the same direction, we assume that the trader traded at that price rather than the trigger price.

Finally, it is possible that someone is willing to sell at a price lower than the trader's bid because the market price has actually fallen in the interim. Since the bid is outstanding, the seller gladly accepts it; since the market price is now less than that, he immediately buys at the lower price, reaping his capital gain.[3] That is, if we observe that within 60 seconds of the n-point trigger the price jumped in the opposite direction, we assume that the trader did actually trade at the trigger price.

In Table 4 we present the results of 2-point and 3-point Alexander-type filter rules, and the results of the 2-point trigger rule described above. Note that a 3-point filter rule would have generated a loss for the trader for the month even in the absence of transactions costs. The smaller 2-point filter does generate positive returns for the month -- at $12.50 per point, a 2-point filter would have resulted in $525 per contract; a trader dealing in 1 million DM would have made $2100 for the month. However, his 264 round trip transactions would have generated substantial transactions costs. The costs for a trader who was not a member of the exchange would have been prohibitive: a minimum of $20 per round trip transaction in commissions, plus $1.50 per round trip as the fee to the Clearing House. A member trading on the floor on his own account would pay no commission and only $.50 per round trip for the

---

3 In actuality, our trader can cancel his outstanding bid, which he would like to do if his information is the same as the seller's. Whether he is successful or not depends on whether he or the seller is quicker. In any event, our test strictly deals with a univariate chartist who is not concerned with the seller's non-price information.

Clearing House fee. We may also assume that as a member, he could hold his margin deposit as Treasury Bills, so that the margin does not generate a substantial interest loss to him. However, the lost interest on the price of his seat on the exchange must be counted as a cost: in July 1977 interest rates were about 5.5% and a seat on the IMM cost about $50,000.[4] Thus, the costs to a member trading four contracts for the month were about $757, just over a third of his potential gain -- if he could actually achieve his desired price on each of his actions.

Finally, under the 2-point trigger rule, the trader executes the same number of transactions as under the 2-point filter rule, but at prices sufficiently less favorable for him to incur a loss for the month even without transactions costs.

It appears that univariate information is insufficient to generate consistent profits in this market. However, we cannot rule out the possibility that outside information or news that reaches the traders on the floor before it reaches the public can create a profit opportunity. A floor trader who receives information at the same time as outsiders also has time to profit because he need not take the time to place a phone call to his broker. Also, the price transmission delay from the floor can give floor traders an edge over outsiders.

3.3 The Price Process

We turn now to the distributions of events, price jumps (measured at the time of occurrence), and price changes (measured at intervals of uniform

---

4 Seat prices and Clearing House fees for July 1977 were provided by Matt Jackson and the research department of the CME.

length h). The tests in this section consist primarily of χ^2 goodness-of-fit tests, which are appropriate for hypotheses about discrete distributions. Because the prices are rounded to the nearest point, the price data are discrete; the data on events per interval (i.e. the number of price jumps per interval) are naturally discrete. Thus, tests of the Kolmogorov-Smirnov type cannot be used.

The χ^2 statistic is computed as

$$X^2 = \sum_{j=1}^{k} \frac{(z_j - m_j)^2}{m_j} \sim \chi^2(k-r-1)$$

where k = the number of cells into which the density function is broken,

z_j = the number of observations per cell,

m_j = the expected number of observations per cell, based on the hypothesized density subject to any estimated parameters, and the total number of observations in the sample, N,

r = the number of estimated parameters of the hypothesized density.

If the parameters are estimated using the (ungrouped) microdata rather than using the z_j, then X^2 is bounded between a $\chi^2(k-r-1)$ and a $\chi^2(k-1)$. That is, critical values from the $\chi^2(k-r-1)$ will produce a more stringent test than necessary. All parameters in this paper are estimated from microdata, and all of the χ^2 statistics are bounded below by the χ^2 distribution with the reported degrees of freedom.

For the χ^2 test to be reliable it is desirable for each expectation, m_j, to be greater than 5 or 6, and that k-r-1 be greater than 2. For h smaller than 3 minutes, the expectations and degrees of freedom conditions were generally difficult to meet. To remedy this, in some cases we use a related test suggested in Heckman (1985). Heckman derives an alternative statistic, G, whose distribution is $\chi^2(k-1)$ whenever the parameters are maximum likelihood estimates taken from the microdata. For given k, this statistic is numerically greater than the traditional goodness-of-fit statistic. Instead of reducing the degrees of freedom, this test adds a function of the covariance matrix of the parameters to the traditional χ^2 statistic. This test is particularly useful when the number of estimated parameters equals or barely exceeds the number of cells. It is also useful when the traditional χ^2 statistic is ambiguous -- that is, when it exceeds the critical value of the $\chi^2(k-r-1)$ but falls short of the critical value of the $\chi^2(k-1)$. These conditions will be met in the test of the joint CPN hypothesis.

We turn first to the distribution of events. A test of equation (2) based on the 3 minute sample is given in Table 5. The estimated parameter is allowed to vary by day; the data are then pooled. As can be seen, the null hypothesis that the daily timing of events, or price jumps, is a Poisson process is rejected at the 5% level on only four days.

That these days are all days for which the random walk hypothesis for prices at 3 minute intervals was rejected (or almost so for day 28) is a curious feature. It suggests that certain days have characteristic patterns of both price movements and their timing that may be recognizable to floor traders or those market participants who choose to spend time watching the intradaily market. It is not clear whether these characteristics signal profit opportunities. Although a trader engaging in a 2-point trigger rule would have done very well for the month had he traded only on these four days, the price and timing patterns must be recognizable in advance in order to use them. One possibility is that these four days were days of heavy Federal Reserve or Bundesbank intervention in spot exchange markets. Central banks attempt to hide their currency purchases by spreading them out over the day; the times of their trades may not be statistically random and might be anticipated by the traders.

On the basis of the pooled sample we reject the hypothesis that the parameter does not vary over the course of the month. (We do not difference across days -- we simply omit those observations from the sample.) Finally, in order to rely on these tests we must satisfy ourselves that the observations on the first differences of $\{x(t)\}$ are independent. Autocorrelation functions computed for the differenced series failed to reject this hypothesis.

We next consider the distributions of the price changes and jumps. In order to proceed with the tests of equation (5) (the truncated and rounded density consistent with price jumps being normally distributed, and hence generated by a CPN process) and equation (4) (the rounded density consistent with price changes being normally distributed, and hence generated by a Wiener process), we must first provide maximum likelihood estimates of the various parameters. We restrict the means of both densities to be 0, leaving us with univariate estimation problems for the variances τ^2 and σ^2.

It is not obvious that the maximum likelihood estimators of either τ^2 or σ^2 can be computed from simple sample statistics such as the sample variances of the price changes. In fact it is clear that the sample variance of the price jumps will always be larger than τ^2 under the truncated and rounded CPN hypothesis, because the observed sample will contain no price jumps of size zero. It is less obvious whether the sample variances of the price changes, appropriately corrected for N rather than N-1 degrees of freedom, are perhaps the maximum likelihood estimators of $\sigma^2 h$. However, the range of observed price changes in the data set is small as well as discrete, so it is not obvious that we can treat the sample as if it is large.

The estimates of the maximum likelihood estimates of σ^2 and τ^2 and the details of the estimation procedure are described in the Appendix.

In Table 6 we present various tests of normality. Looking first at the studentized ranges, we see that the hypothesis that the price changes are normally distributed is rejected on all but two days (at the 5% level) for the 3 minute sample and on all days at 30 second intervals. The test done by pooling the data for all days indicates a rejection for both observation intervals. On the other hand, the studentized range is in fact consistent with the hypothesis that the price jumps are normally distributed for each day independently and for the pooled sample. However, because of the truncation near 0 discussed above, the sample variances are too large and thus these studentized ranges artificially small. Note that the studentized range test does not restrict the mean to be zero.

The χ^2 statistics to test the hypothesis that the price changes are generated by an underlying Wiener process with parameters $\mu=0$ and $\sigma^2=\hat{\sigma}^2$, which are observed only after rounding, are performed only for the 3 minute sample. This is because the degrees of freedom condition was violated on almost all days for the 30 second sample. For the 3 minute sample, this null hypothesis is rejected independently for each day at the 5% level, and also for the pooled sample.

Two χ^2 tests are performed on the price jumps. The first tests the hypothesis that price jumps are drawn from an underlying distribution that is N(0,1), and then rounded and truncated. There is no particular reason to restrict the variance of the underlying distribution to equal one, but this restriction does improve the degrees of freedom condition because it requires no estimated parameters. As can be seen, this null hypothesis is rejected on days 11, 13, 27, 28, and 29, and for the pooled sample. However, we note that this rejection is not inconsistent with the more general hypothesis where τ^2 is not necessarily equal to one, because these are in fact days when $\hat{\tau}^2$ is very far from one.

The second test allows the variance of the underlying normal distribution to be determined by the sample, i.e. it tests for an underlying distribution that is $N(0,\hat{\tau}^2)$, which is observed after rounding and truncating.[5] The null hypothesis cannot be rejected at the 5% level for any day except day 27; it is not rejected at the 1% level on that day. The pooled sample is not consistent with the null hypothesis. This can be viewed as a rejection that the variance τ^2 is constant for the entire month.

The G statistic proposed by Heckman was also calculated for the $N(0,\hat{\tau}^2)$. The inference was the same as under the χ^2 for all days and for the pooled sample.

A χ^2 test based on equation (6) using price changes at intervals of length h would jointly test the hypotheses that the timing of events is Poisson and that the jumps are normally distributed. Each of these hypotheses has been tested independently above. Thus, such a test does not yield independent evidence about the CPN hypothesis; it is an alternative to the previous tests. It does, however, present a useful summary of the evidence on

---

[5] Note that it is not the case that the χ^2 statistic must improve when τ is chosen by the sample. First, the number of degrees of freedom change. Second, an increase in τ causes increases in the expectations of some cells and decreases in other cells, with an ambiguous effect on the combined statistic.

the CPN hypothesis, as well as providing a fairly stringent test that can be rejected for any number of reasons.

In order to perform a test based on $\bar{\varphi}(p_h)$ we must have maximum likelihood estimates of c and τ^2 based on the null hypothesis. The maximum likelihood estimates of τ^2 took account of the fact that true events which resulted in very small price changes are not recorded. The estimates of the speed of events $\hat{c}$ (Table 5) did not take this into account and thus are not appropriate for use with $\hat{\tau}^2$. There is no directly comparable method for finding an estimator of c that includes an adjustment for events which produce price jumps of less than half a point. However, an indirect approach is available.

Note that the variance of p_h under the CPN hypothesis is the same as the variance of the p_h if they are normally distributed. The latter has already been estimated: $3\hat{\sigma}^2$ is the maximum likelihood estimate of the variance of prices at 3 minute intervals under the hypothesis that they are normally distributed. According to the CPN hypothesis that variance also equals $3c\tau^2$. Thus we can estimate the speed of occurrence inclusive of unmeasured events, $\bar{c}$, by

$$3\hat{\bar{c}} = 3\hat{\sigma}^2/\hat{\tau}^2 \quad . \tag{8}$$

Note that this method assumes that $\bar{c}$ and τ^2 are independent.

In Table 7 we present the estimates of $\bar{c}$ as well as the χ^2 and G statistics to test the joint hypothesis for the 3 minute sample. Because two parameters were estimated, it was very difficult to satisfy the degrees of freedom condition for the χ^2 statistic, despite the fact that the improved fit tended to increase the number of cells for which the expectations condition was satisfied. For completeness, these statistics are presented even when there is only one degree of freedom, but we will not draw any inference from them. The G statistics have satisfactory degrees of freedom in all cases.

Of the twelve days with χ^2 statistics of 2 or more degrees of freedom, six surpass the 5% critical value of the $\chi^2(k-r-1)$. However, for only two of these days is the null hypothesis rejected if the G statistic is used. Note that on the other four days, the χ^2 statistic is *below* the 5% critical value of the $\chi^2(k-1)$ distribution, so that these four days are precisely those for which the traditional χ^2 test is ambiguous.

Both of the two days for which the CPN hypothesis is rejected are days when the Poisson distribution for events was rejected on the basis of the non-joint tests (Table 5). Several other days are consistent with the joint hypothesis despite the fact that the non-joint hypotheses were rejected. It is reassuring that in no case for which the χ^2 test unambiguously failed to reject the joint hypothesis did the G test result in a rejection. The joint CPN hypothesis is rejected for the pooled sample, as it had been for each of the non-joint tests.

Measures of the sample kurtosis are also presented in Table 7. A sample kurtosis that is higher than $3/3\hat{\bar{c}}$ roughly indicates a sample distribution that is sharply peaked and fat-tailed relative to the CPN.

The statistics of Table 7 represent a striking improvement of fit of the CPN relative to the Wiener hypothesis, particularly at the central peak. The distribution $N(0,3\hat{\sigma}^2)$ of Table 6 underpredicted the number of observations with no price change by an average of 19 observations for the daily samples, so that on average 19/89=21% of each daily χ^2 statistic is due to error at the peak. It underpredicted the pooled sample's peak by over 500 observations.

In contrast, the distribution CPN$(0,3\hat{\bar{c}},\tau^2)$ underpredicted the daily peak by an average of 4.2 observations, and by 136 observations for the pooled sample.

The two rejections of the joint CPN hypothesis in Table 7 are still associated with tails too fat and peaks too high relative to the expectations. However, because we performed independent tests of the Poisson and normal components of the CPN, we are able to go further. In particular, because on the days when the joint hypothesis was rejected, price jumps appeared to be normally distributed, it appears to be more likely that the kurtosis problem is due to our inability to adequately model the timing of events than to an innate non-normality of price movements.

4.1 Conclusions

Our tests of market efficiency indicate that the futures market in foreign exchange functions almost instantaneously to absorb new price information. Both autocorrelation tests and trigger rule tests indicate that a chartist would have a difficult time making a profit from this market, and even a floor trader may not be able to profit.

The compound Poisson-normal process was empirically contrasted with the Wiener process. We could not reject the hypothesis that the distribution of daily events was Poisson, or the hypothesis that the distribution of price jumps at the time they occurred was normal, although in both cases we reject the hypothesis that the parameters of the distributions were constant. The Wiener process was rejected at the usual significance levels.

A joint test of the CPN process was provided by the derivation of the price changes generated by this process. The implied distribution of price changes is leptokurtic. The CPN process cannot be rejected by the evidence, and results in a dramatically better fit than the normal distribution, which is rejected. Rejections of the joint hypothesis are accompanied by even greater kurtosis than that already implied by the CPN hypothesis, and by rejections that the events associated with those days were distributed Poisson. This suggests that we must delve deeper into the timing of economic responses to shocks before we will unravel the kurtosis question.

APPENDIX

By restricting $\mu=0$, we can estimate σ^2 in the following manner: using the 3 minute sample of price changes, choose the $\hat{\sigma}$ that maximizes the likelihood function resulting from density $\bar{g}_h(p_h)$ with h=3 and $\mu=0$. This is a univariate problem because of the restriction, and can be done by a grid search. Similarly, the maximum likelihood estimate of τ^2 is found by choosing the $\hat{\tau}$ that maximizes the likelihood function resulting from density $\bar{g}(p)$ with $\nu=0$.

In Tables A.1 and A.2 we present the sample means, ranges, and a variety of estimated variances. Note first that the means for price changes over both 3 minute and 30 second intervals, and for the price jumps, are close to zero for all days and for the pooled samples. In every case a T-test failed to reject the hypothesis that the mean is zero. Thus, the restrictions of $\mu,\nu=0$ are not unreasonable.

The maximum likelihood estimates $3\hat{\sigma}^2$ and $.5\hat{\sigma}^2$ are all very close to the corrected sample variances (i.e. the maximum likelihood sample variances if the underlying distributions were unrounded normal distributions). This is very reassuring, but was by no means a necessary result in a small sample.

As expected, the corrected sample variances for the price jumps are uniformly greater than the maximum likelihood estimates, $\hat{\tau}^2$.

TABLE A1

SAMPLE MEANS AND RANGES OF PRICES

Day	Price Changes				Price Jumps	
	3 Minutes		30 Seconds			
	$\hat{\mu}$	Range	$\hat{\mu}$	Range	$\hat{\gamma}$	Range
1	-0.01	-2 to 3	0.00	-2 to 3	0.00	-2 to 3
5	-0.01	-2 to 2	0.00	-2 to 2	-0.02	-2 to 3
6	0.04	-2 to 6	0.01	-2 to 6	0.14	-2 to 6
7	-0.01	-3 to 3	0.00	-4 to 3	-0.03	-4 to 3
8	0.21	-2 to 7	0.03	-2 to 4	0.38	-2 to 4
11	0.16	-2 to 5	0.03	-3 to 5	0.27	-3 to 5
12	0.13	-5 to 4	0.02	-3 to 4	0.10	-3 to 4
13	-0.11	-4 to 5	-0.02	-4 to 4	-0.14	-4 to 4
15	0.34	-2 to 3	0.01	-2 to 3	0.08	-2 to 3
18	0.35	-3 to 6	0.06	-2 to 4	0.22	-3 to 3
19	-0.06	-6 to 3	-0.01	-3 to 2	-0.08	-3 to 3
20	-0.13	-4 to 5	-0.02	-4 to 4	-0.12	-3 to 4
21	-0.06	-4 to 3	-0.01	-4 to 3	-0.10	-3 to 3
22	0.09	-2 to 3	0.02	-2 to 3	0.18	-2 to 3
25	0.08	-3 to 3	0.01	-2 to 3	0.09	-2 to 3
26	0.03	-4 to 3	0.01	-4 to 3	0.06	-4 to 3
27	0.20	-4 to 9	0.03	-4 to 4	0.19	-4 to 6
28	-0.27	-8 to 7	-0.04	-4 to 4	-0.24	-4 to 4
29	-0.01	-6 to 9	0.01	-5 to 6	-0.01	-4 to 6
pooled	0.03	-8 to 9	0.01	-5 to 6	0.04	-4 to 6

NOTES: There are 89 observations per day in the 3-minute sample, 535 observations per day in the 30-second sample, and the number of observations per day in the sample of price jumps is one less than the number of events on that day.

TABLE A2

ESTIMATED VARIANCES OF PRICES

Day	Price Changes					Price Jumps	
		3 Minutes		30 Seconds			
	$\hat{\sigma}^2$	V^{2a}	$3\hat{\sigma}^2$	V^{2a}	$.5\hat{\sigma}^2$	V^{2a}	$\hat{\tau}^2$
1	0.11	0.37	0.34	0.07	0.06	2.55	1.69
5	0.16	0.55	0.48	0.12	0.08	1.63	0.88
6	0.27	0.85	0.81	0.15	0.14	2.75	1.90
7	0.21	0.71	0.64	0.13	0.11	1.97	1.19
8	0.38	1.16	1.14	0.17	0.19	2.01	1.37
11	0.43	1.35	1.30	0.25	0.22	3.47	2.59
12	0.76	2.34	2.28	0.33	0.38	2.00	1.23
13	0.62	1.92	1.85	0.36	0.31	2.79	1.93
15	0.21	0.71	0.64	0.13	0.11	1.87	1.10
18	0.78	2.32	2.34	0.37	0.39	1.78	1.06
19	0.64	2.01	1.93	0.25	0.32	1.71	0.96
20	0.64	2.00	1.93	0.29	0.32	1.97	1.21
21	0.38	1.22	1.14	0.20	0.19	2.09	1.32
22	0.26	0.85	0.77	0.20	0.13	1.50	0.81
25	0.27	0.88	0.81	0.16	0.14	1.41	0.69
26	0.45	1.43	1.35	0.28	0.22	1.96	1.19
27	0.92	2.79	2.76	0.45	0.46	3.15	2.28
28	1.25	3.75	3.76	0.54	0.63	3.10	2.25
29	1.13	3.47	3.39	0.51	0.56	3.93	2.96
pooled	0.52	1.63	1.56	0.26	0.26	2.32	1.51

NOTES: There are 89 observations per day in the 3-minute sample, 535 observations per day in the 30-second sample, and the number of observations per day in the sample of price jumps is one less than the number of events on that day.

[a]The corrected sample variance is defined by $V^2 = (1/N)\cdot$(sum of squared deviations from the sample mean), where N = the number of observations.

REFERENCES

Alexander, Sidney S. "Price Movements in Speculative Markets: Trends or Random Walks." Industrial Management Review 2 (May 1961): 7-26

Bachelier, L. J. B. A. "Theorie de la speculation." Annales Scientifiques de l'Ecole Normale Superiente 3, No. 1018 Paris: Gauthier-Villars, 1900. Translated in The Random Character of Stock Market Prices. Edited by Paul H. Cootner. Cambridge: The M.I.T. Press, 1964

Black, Fischer. "The Pricing of Commodity Contracts." Journal of Financial Economics 3 (January/March 1976): 167-179

Clark, Peter K. "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices." Econometrica vol. 41, No. 1 (January 1973):135-155

Fama, Eugene F. "Mandelbrot and the Stable Paretian Hypothesis." Journal of Business 36 (October 1963): 420-429

Feinstone, Lauren J. "Efficiency and Normality in the Futures Market in Foreign Exchange." University of Rochester Working Paper, no. 84-10: October 1984

Frenkel, Jacob A. "Flexible Exchange Rates, Prices, and the Role of 'News': Lessons from the 1970s." Journal of Political Economy 89 (August 1981): 665-705

Heckman, James J. "The x^2 Goodness of Fit Statistic for Models with Parameters Estimated from Microdata." Econometrica 52 (November 1984): 1543-1547

Mandelbrot, Benoit. "The Variation of Certain Speculative Prices." Journal of Business 36 (October 1963): 394-419

Niederhoffer, Victor and Osborne, M. F. M. "Market Making and Reversal on the Stock Exchange." Journal of the American Statistical Association 61 (December 1966): 897-916

Telser, Lester G. "Margins and Futures Contracts." The Journal of Futures Markets 1, 1981

Wasserfallen, Walter and Zimmerman, Heinz. "The Behavior of Intra-Daily Exchange Rates." Paper prepared for the Annual Meeting of the European Finance Association in Schevningen, Holland, July 1981

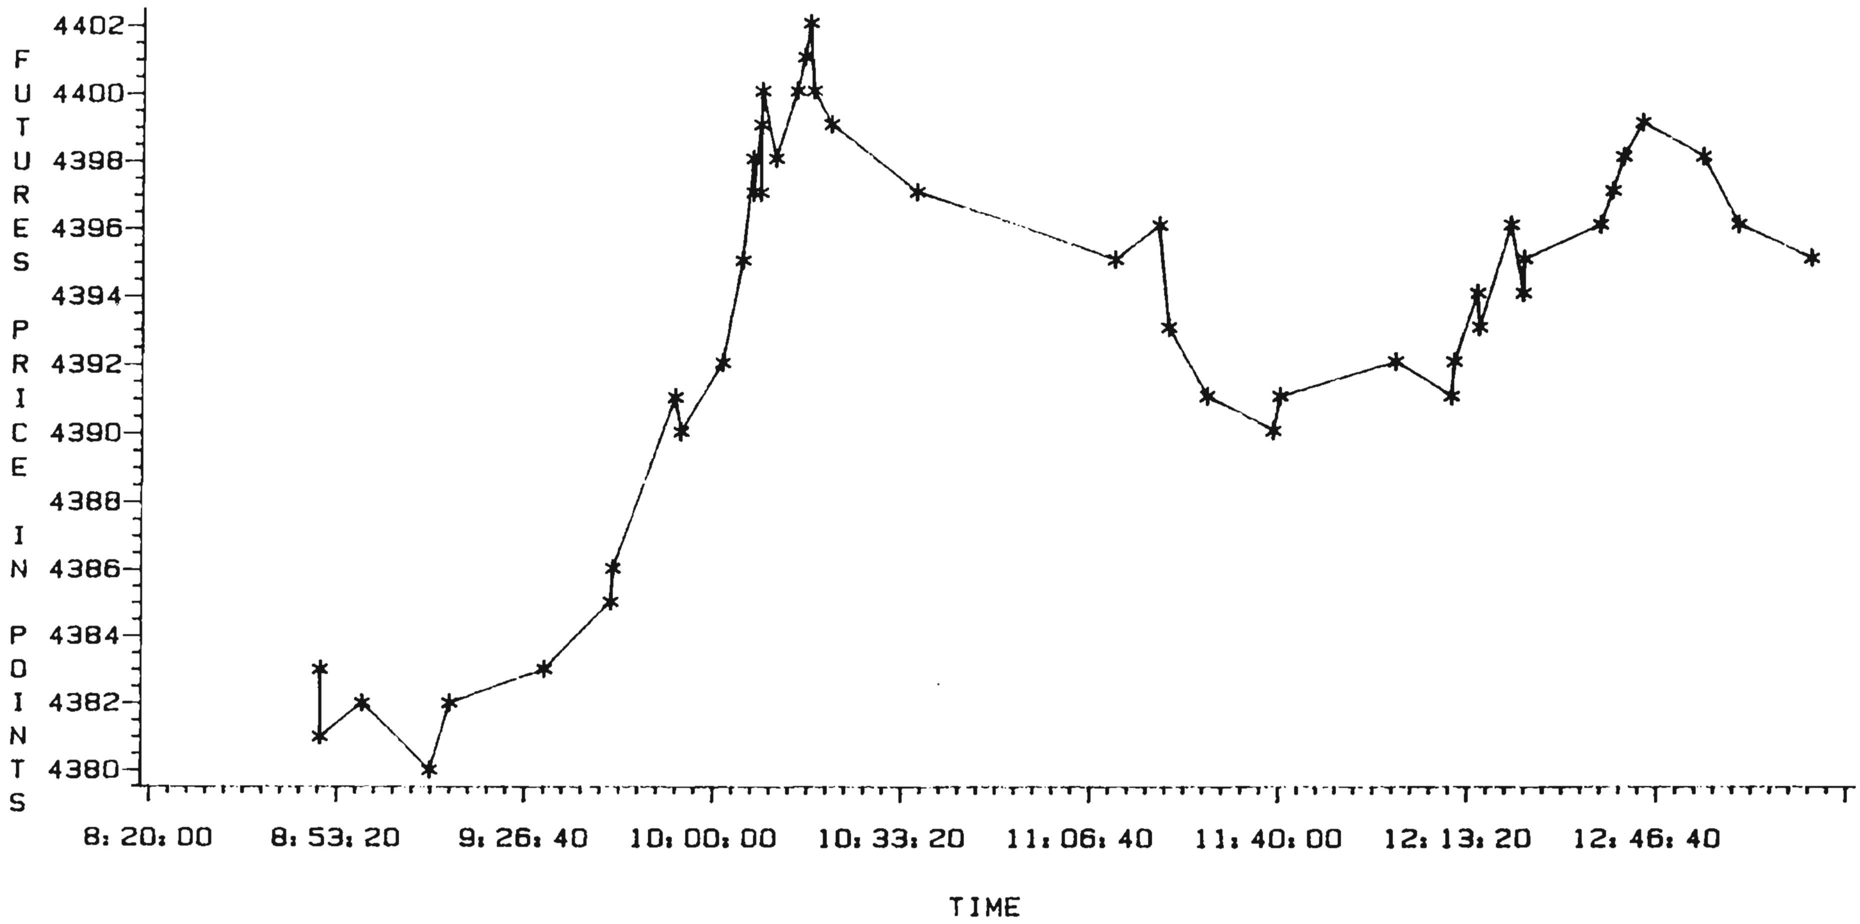

Fig. 1.--September DM futures prices on July 11, 1977

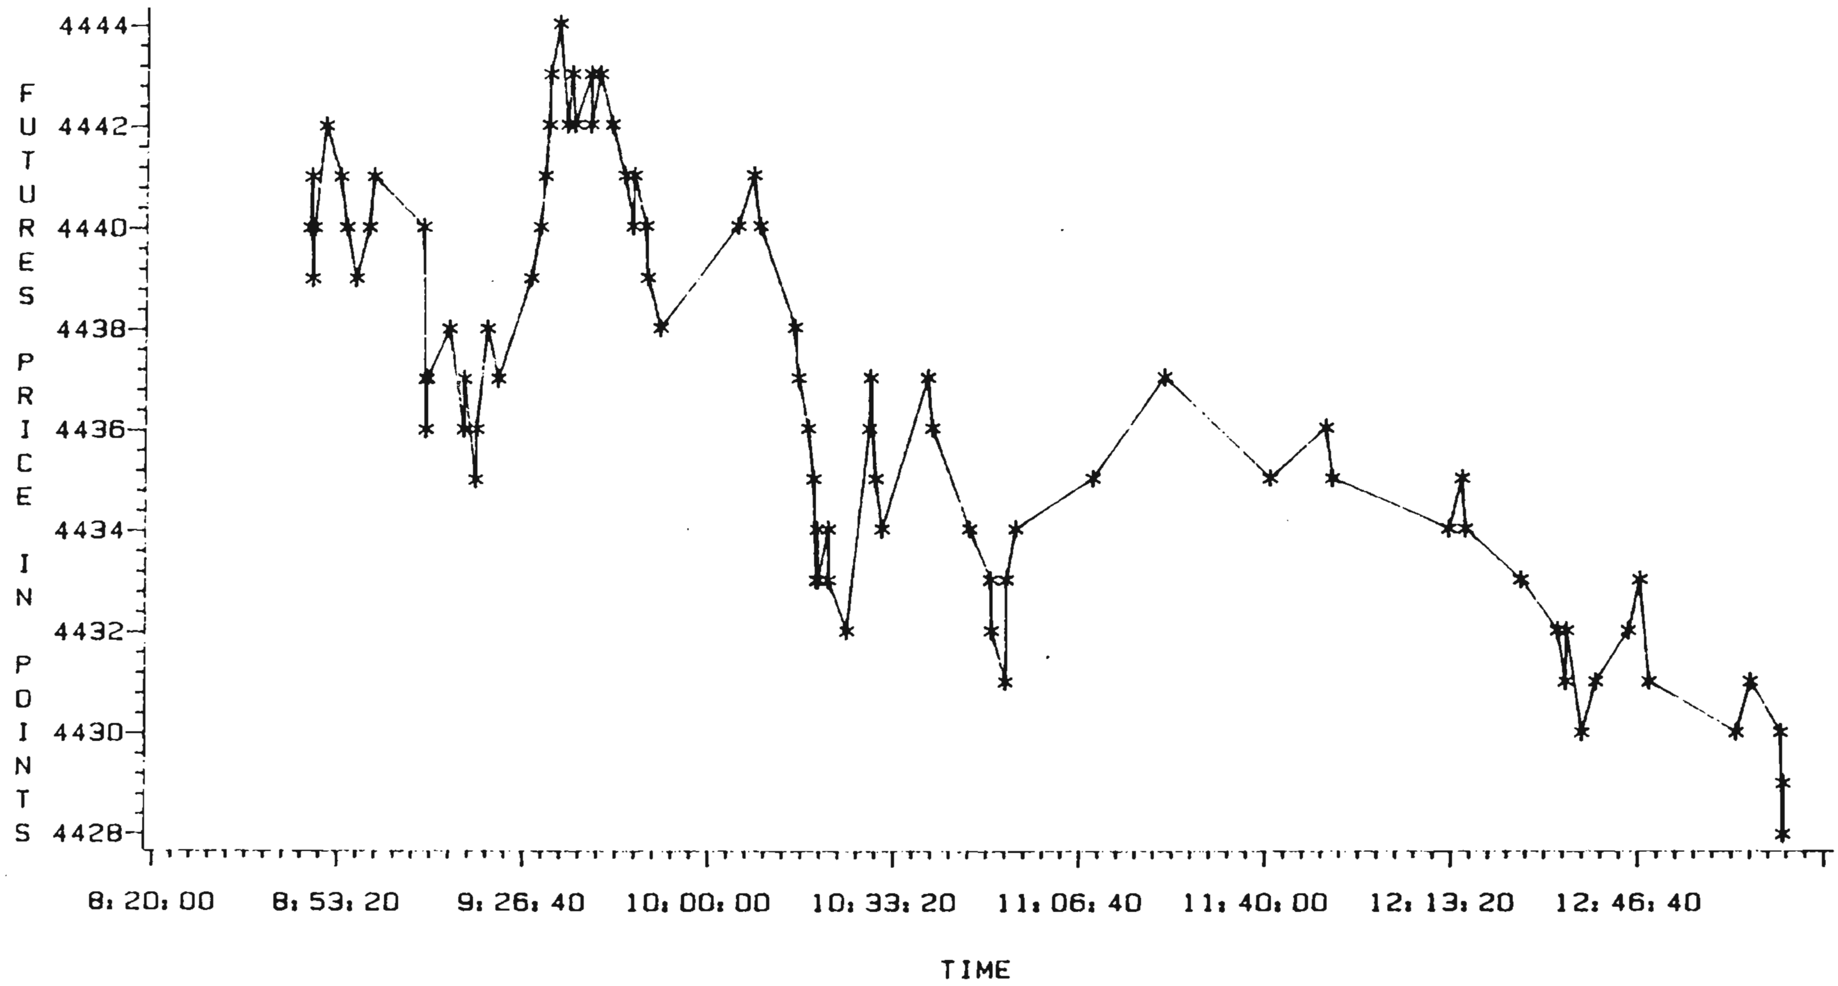

Fig. 2.--September DM futures prices on July 20, 1977

TABLE 1

SEPTEMBER 1977 DM FUTURES DURING JULY 1977

July 1977 Trading Date	Daily Volume[a]	Number of Price Jumps	Settlement Price[a]	Average Price[b]	Standard Deviation of Price[b]
1	187	23	4294	4294	1.44
5	400	47	4338	4338	2.03
6	376	36	4338	4336	2.64
7	291	37	4354	4354	2.47
8	400	46	4365	4357	6.76
11	512	45	4395	4392	6.01
12	1456	88	4395	4384	5.41
13	996	85	4396	4402	4.25
15	554	40	4397	4393	2.33
18	1689	140	4433	4421	10.42
19	1294	101	4426	4427	2.48
20	880	87	4429	4436	3.40
21	716	49	4426	4429	2.90
22	1133	66	4447	4445	3.96
25	682	58	4464	4460	3.01
26	1202	88	4459	4463	2.52
27	1126	97	4443	4439	3.43
28	1267	106	4393	4404	12.13
29	777	72	4385	4380	5.14
Total	15,938	1,311			
Average	839	69	4399	4398	

[a]Source: 1977-1978 Yearbook, International Monetary Market, Chicago Mercantile Exchange. Prices are reported in points/DM; 1 point = $.0001; 1 contract = 125,000 DM.

[b]Average prices and their standard deviations are computed from the data set resulting from observations taken at 30-second intervals. The number of observations is 535 per day.

TABLE 2

AUTOCORRELATION FUNCTIONS OF PRICE CHANGES AT THREE MINUTE INTERVALS

Day	Lag 1	2	3	4	5	6	7	8	9	10	11	12	$\tilde{Q}_{12}$	$\tilde{Q}_{24}$	$\tilde{Q}_{36}$
1	.09	.00	.12	.03	.03	.03	-.21	-.09	.00	-.15	-.30*	-.06	19.37	25.56	37.10
5	.18	-.22*	-.18	.04	.04	.09	.00	-.09	.06	.13	.13	-.07	16.64	23.87	33.11
6	-.12	-.12	.00	-.15	-.02	.13	.07	-.11	.05	-.02	.17	.02	11.41	36.64	46.34
7	.10	-.13	.06	.05	.00	.00	-.03	-.03	-.05	.00	-.03	-.10	4.60	25.15	33.88
8	.03	-.08	.09	-.04	-.16	-.12	-.02	-.08	-.27*	-.08	.28*	.01	22.12	39.62	44.35
11	.01	.11	.20	.13	.23*	-.08	.14	.08	-.04	-.10	-.05	.04	16.26	33.06	44.21
12	-.04	.01	.10	-.14	.02	.03	-.07	.07	.20	-.04	-.07	.10	9.74	33.10	43.75
13	.08	.18	-.01	.10	-.11	-.22*	-.14	-.09	-.10	-.08	-.03	.12	16.31	33.35	48.05
15	-.07	-.07	.06	.11	-.11	-.03	-.05	.05	.03	.13	-.13	.00	7.65	14.66	34.20
18	.21	.11	-.22	-.29*	.09	.08	.07	.16	-.03	.04	-.14	-.15	26.59	41.79	45.72
19	.09	-.10	-.05	-.28*	.11	.01	-.13	.10	-.12	-.17	-.09	-.18	21.68	37.10	51.96
20	-.09	-.04	-.17	-.01	.05	-.12	-.18	-.09	.19	.07	.09	.02	14.18	29.05	43.69
21	-.01	-.05	.04	-.06	-.10	-.19	.13	.00	-.13	.01	-.02	-.05	8.83	21.25	33.06
22	.00	.11	-.04	-.02	.19	-.02	.16	-.15	-.18	.00	-.17	-.10	16.78	22.99	29.10
25	.13	.07	-.03	-.16	-.12	.01	-.08	.11	.08	-.19	-.20	-.07	16.63	25.57	34.66
26	-.04	-.05	-.06	.12	.12	.06	-.09	-.02	.07	.15	-.17	-.03	10.46	23.09	38.21
27	.04	.09	.06	.09	-.11	.02	-.10	.03	-.13	-.03	.13	.04	7.94	16.65	27.40
28	.33*	.14	-.01	-.24*	-.11	-.09	-.04	.00	.05	.06	-.04	-.07	20.54	33.64	47.07
29	.05	.05	.05	.08	-.06	.07	.02	-.07	-.11	-.16	.01	-.05	6.72	23.43	27.96
Sum													274.45	539.57	743.82

NOTES: There are 88 observations per day. The large lag standard error is .11 in all cases. Autocorrelation coefficients significantly different from 0 at the 5 percent level have been marked with asterisks. Critical values for $\tilde{Q}_m$ are

m	12	24	36	19·12	19·24	19·36
5%	21.03	36.42	50.96	264.06	506.66	745.87
2.5%	23.04	39.36	54.40	271.23	516.58	757.89

TABLE 3

AUTOCORRELATION FUNCTIONS OF PRICE CHANGES AT THIRTY SECOND INTERVALS

Day	Lag 1	2	3	4	5	6	7	8	9	10	11	12	$\tilde{Q}_{12}$	$\tilde{Q}_{24}$	$\tilde{Q}_{36}$
1	.00	-.03	-.07	.07	.00	.03	.00	.00	.00	.00	-.03	.00	6.73	19.75	22.17
5	.03	-.03	-.10*	-.02	.03	.00	.08*	-.07	.02	-.03	.03	-.03	14.73	35.74	38.16
6	.00	.00	-.02	-.01	.01	-.02	-.10*	.01	-.01	.00	-.05	-.02	7.64	14.74	22.53
7	.01	.01	-.06	.01	.01	-.01	.03	.00	.00	.04	-.01	.00	3.61	14.13	15.37
8	.11*	-.03	-.05	-.02	.05	-.01	.05	-.02	-.05	-.04	.06	-.01	15.73	24.99	37.56
11	-.01	-.04	.09*	-.03	-.05	.05	.02	-.05	-.03	.01	.06	.01	12.53	43.38	61.51
12	-.02	.07	-.01	.04	.03	-.05	-.01	.02	.00	-.05	.00	-.02	7.36	14.85	37.04
13	-.04	.00	.00	.04	.05	-.03	.06	.00	-.04	-.03	.01	.09*	11.40	27.70	37.21
15	-.02	.00	.00	-.10*	.04	.07	.03	-.06	-.03	-.07	.01	.01	14.70	28.50	40.83
18	.01	.03	.08*	-.02	.00	.01	.10*	.04	-.05	.06	.03	-.02	14.49	32.25	48.07
19	.01	.12*	.03	.00	.08*	.05	.01	.01	.01	-.05	.05	-.03	16.49	29.95	42.61
20	.02	.01	-.01	-.05	.02	-.03	.02	.01	-.09*	.02	-.01	.00	7.34	23.58	39.53
21	.02	-.01	-.01	-.03	.04	.03	.01	-.05	-.04	-.04	-.04	-.01	6.21	16.98	33.94
22	.00	.01	-.06	-.05	.01	-.06	.03	.02	.03	.07	.04	.01	10.08	25.57	40.51
25	-.02	-.02	.10*	-.06	.13*	.01	-.02	-.05	.01	.06	.03	-.09*	25.50	39.96	61.66
26	-.02	-.02	.01	-.09*	-.03	.05	-.02	.02	.00	-.05	-.01	.06	10.43	28.11	37.59
27	-.05	.12*	.02	.01	.00	.02	.00	.00	.00	-.02	-.04	.05	11.89	28.51	38.74
28	.03	.01	.04	.09*	.03	.15*	.05	.00	-.02	.04	.08*	.01	24.28	38.12	51.29
29	.05	-.01	.00	.07	-.01	.02	.06	.02	.02	-.08*	.01	.01	10.27	16.59	35.12
SUM													231.41	503.40	741.44

NOTES: There are approximately 530 observations per day. The number of observations differ slightly by day. Autocorrelation coefficients significantly different from 0 at the 5 percent level have been marked with asterisks: Critical values for $\tilde{Q}_m$ are

m	12	24	36	19·12	19·24	19·36
5%	21.03	36.42	50.96	264.06	506.66	745.87
2.5%	23.04	39.36	54.40	271.23	516.58	757.89

TABLE 4

FILTER AND TRIGGER RULES

Day	3-point Filter		2-point Filter		2-point Trigger
	No. of Round-trip Transactions	Net Gain (in points)	No. of Round-trip Transactions	Net Gain (in points)	Net Gain (in points)
1	3	- 9	6	-10	-10
5	3	4	9	- 1	- 2
6	3	-12	13	-29	-30
7	4	- 7	8	- 1	- 1
8	5	8	7	8	0
11	4	20	10	15	15
12	13	-14	17	- 6	- 9
13	10	10	18	12	12
15	6	- 2	8	4	3
18	15	3	19	18	8
19	15	-15	18	3	0
20	12	- 9	17	- 7	- 8
21	8	- 2	9	5	0
22	8	3	17	-10	-17
25	9	-14	14	- 3	- 3
26	11	- 4	13	3	2
27	17	-14	23	6	- 9
28	12	20	21	22	11
29	9	19	17	13	6
Total	167	-15	264	42	-32

TABLE 5

GOODNESS-OF-FIT TESTS FOR EVENTS

Day	Number of Events: Per Day	Per 3 Minutes ($3\hat{c}$)	Range	$\chi^2(k-2)$: Poisson with Parameter $3\hat{c}$
1	23	0.26	0 - 4	0.69(1)
5	47	0.52	0 - 3	3.17(2)
6	36	0.40	0 - 3	3.57(2)
7	37	0.41	0 - 3	1.75(2)
8	46	0.51	0 - 3	8.23(2)[a]
11	45	0.50	0 - 5	1.48(2)
12	88	0.98	0 - 5	3.73(3)
13	85	0.94	0 - 5	3.51(3)
15	40	0.44	0 - 3	0.38(2)
18	140	1.56	0 - 8	21.59(4)[b]
19	101	1.12	0 - 5	19.15(3)[b]
20	87	0.97	0 - 5	4.05(3)
21	49	0.54	0 - 3	3.64(2)
22	66	0.73	0 - 5	1.22(2)
25	58	0.64	0 - 4	5.45(2)
26	88	0.98	0 - 6	2.52(3)
27	97	1.08	0 -10	2.89(3)
28	106	1.18	0 - 8	10.32(3)[a]
29	72	0.80	0 - 5	5.19(2)
pooled	1311	0.77	0 -10	264.87(4)[b]

NOTES: There are 90 3-minute intervals on each day, and 1,710 for the month.

[a]The 5 percent critical values are 3.83(1), 5.99(2), 7.81(3), 9.49(4).

[b]The 1 percent critical values are 6.63(1), 9.21(2), 11.30(3), 13.30(4).

TABLE 6

TESTS OF NORMALITY OF PRICES

Day	Price Changes			Price Jumps		
	3 Minutes		30 Sec.			
	s.r.[a]	$\chi^2(k-2):N(0,3\hat{\sigma}^2)$[b]	s.r.[c]	s.r.[d]	$\chi^2(k-1):N(0,1)$[b]	$\chi^2(k-2):N(0,\hat{\tau}^2)$[b]
1	8.17	22.07(1)	18.27	3.06	3.04(2)	0.29(1)
5	5.36	12.82(1)	11.65	3.87	0.60(3)	0.15(2)
6	8.62	38.93(2)	20.42	4.75	1.32(2)	2.21(2)
7	7.09	31.03(2)	19.76	4.92	1.09(2)	2.69(2)
8	8.32	52.44(3)	14.58	4.18	3.73(3)	3.97(2)
11	6.00	57.00(3)	15.99	4.25	21.89(3)	1.75(2)
12	5.85	22.21(2)	12.17	4.92	0.52(3)	2.79(2)
13	6.46	33.77(2)	13.31	4.76	8.28(3)	3.09(2)
15	5.91	21.97(2)	14.12	3.61	0.06(3)	0.32(2)
18	5.88	25.33(3)	9.80	4.48	3.77(3)	3.76(2)
19	6.32	27.87(3)	9.91	4.57	1.49(3)	1.23(2)
20	6.32	30.42(3)	14.81	4.96	1.86(3)	1.44(2)
21	6.30	35.25(3)	15.79	4.11	1.44(3)	2.36(2)
22	5.41	18.06(2)	11.24	4.05	2.75(3)	1.34(2)
25	6.35	24.30(2)	12.46	4.17	2.51(3)	0.57(2)
26	5.83	10.67(3)	13.25	4.97	2.59(3)	3.46(2)
27	7.74	25.84(5)	11.94	5.60	7.94(3)	11.28(4)
28	7.70	56.06(5)	10.88	4.52	25.25(3)	2.62(4)
29	8.00	74.71(5)	15.33	5.01	26.45(3)	2.75(4)
pooled	13.30	769.89(5)	21.51	6.56	158.95(5)	36.13(4)

NOTES: There are 89 observations per day in the 3-minute sample, 535 observations per day in the 30-second sample, and the number of observations per day in the sample of price jumps is one less than the number of events on that day.

[a]The critical values for the daily samples are 5%: 5.81 and 1%: 6.26.

[b]The critical values are 5%: 5.99(2), 7.81(3), 9.49(4), 11.10(5), and 1%: 9.21(2), 11.30(3), 13.30(4), 15.10(5).

[c]The critical values for the daily samples are 5%: 7.15 and 1%: 7.62.

[d]The critical values vary by day.

TABLE 7

JOINT TESTS OF THE COMPOUND POISSON-NORMAL PROCESS ON PRICE CHANGES OBSERVED AT THREE MINUTE INTERVALS

Day	$3\hat{c}$	CPN$(0,3\hat{c},\hat{\tau}^2)$ $\chi^2(k-3)$	CPN$(0,3\hat{c},\hat{\tau}^2)$ G(k-1)	Sample Kurtosis[a]	$3/3\hat{c}$ = Kurtosis of CPN$(0,3\hat{c},\hat{\tau}^2)$[a]
1	0.20	[0.27(1)]	0.61(2)	10.19	15.00
5	0.54	[2.50(1)]	0.78(2)	2.61	5.56
6	0.43	[1.72(1)]	0.06(2)	19.35	6.98
7	0.54	[7.27(1)]	1.57(2)	3.84	5.56
8	0.84	7.38(2)[b]	7.43(4)	17.78	3.57
11	0.50	1.78(2)	1.88(4)	6.06	6.00
12	1.85	4.79(2)	4.94(4)	2.65	1.62
13	0.96	6.39(2)[b]	6.43(4)	2.67	3.13
15	0.58	[0.51(1)]	0.15(2)	4.49	5.17
18	2.21	11.54(2)[c]	11.54(4)[c]	2.69	1.36
19	2.01	8.64(2)[b]	8.73(4)	3.29	1.49
20	1.60	7.24(2)[b]	7.35(4)	2.56	1.88
21	0.87	2.38(2)	2.42(4)	2.79	3.45
22	0.96	[1.50(1)]	0.73(2)	1.38	3.13
25	1.18	[3.95(1)]	4.71(2)	2.83	2.54
26	1.13	0.73(2)	0.75(4)	1.78	2.65
27	1.21	2.22(4)	2.25(6)	8.22	2.48
28	1.67	16.11(4)[c]	16.29(6)[b]	4.11	1.80
29	1.14	9.38(4)	9.50(6)	7.28	2.63
pooled	1.03	50.86(6)[c]	50.87(8)[c]	7.56	2.91

NOTES: There are 89 observations per day. [Bracketed statistics violate the degrees of freedom condition and may not have distributions that are well-approximated by a $\chi^2(1)$.]

[a]Positive values of the kurtosis indicate a distribution that is leptokurtic (sharply peaked) relative to a normal distribution with the same mean and variance.

[b]The 5% critical values are: 3.84(1), 5.99(2), 7.81(3), 9.49(4), 11.10(5), 12.60(6), 15.50(8).

[c]The 1% critical values are: 6.63(1), 9.21(2), 11.30(3), 13.3(4) 15.10(5), 16.80(6), 20.10(8).

To order copies of the above papers complete the attached invoice and return to Christine Massaro, W. Allen Wallis Institute of Political Economy, RCER, 109B Harkness Hall, University of Rochester, Rochester, NY 14627. Three (3) papers per year will be provided free of charge as requested below. Each additional paper will require a $5.00 service fee which must be enclosed with your order. For your convenience an invoice is provided below in order that you may request payment from your institution as necessary. Please make your check payable to the **Rochester Center for Economic Research**. **Checks must be drawn from a U.S. bank and in U.S. dollars**.

W. Allen Wallis Institute for Political Economy

Rochester Center for Economic Research, Working Paper Series

OFFICIAL INVOICE

Requestor's Name ______________________________

Requestor's Address ______________________________

____

____

____

Please send me the following papers free of charge (**Limit: 3 free per year**).

WP# ______ WP# ______ WP# ______

I understand there is a $5.00 fee for each additional paper. Enclosed is my check or money order in the amount of $__________. Please send me the following papers.

WP# ______ WP# ______ WP# ______

WP# ______ WP# ______ WP# ______

WP# ______ WP# ______ WP# ______

WP# ______ WP# ______ WP# ______