

Stock Prices and Volume ¹

A. Ronald Gallant
North Carolina State University

Peter E. Rossi
University of Chicago

George Tauchen
Duke University

January 1990
Revised August 1990

¹Corresponding author: George Tauchen, Department of Economics, Duke University, Durham, NC 27706, (919) 684-3705. This material is based upon work supported by the National Science Foundation under Grants No. SES-8808015 and SES-8810357, North Carolina Agricultural Experiment Station Project NCO-6134, the PAMS Foundation, and the Graduate School of Business, the University of Chicago.

Abstract

We undertake a comprehensive investigation of price and volume co-movement using daily New York Stock Exchange data from 1928 to 1987. We adjust the data to take into account well-known calendar effects and long-run trends. To describe the process, we use a seminon-parametric estimate of the joint density of current price change and volume conditional on past price changes and volume. Four empirical regularities are found: 1) positive correlation between conditional volatility and volume, 2) large price movements are followed by high volume, 3) conditioning on lagged volume substantially attenuates the “leverage” effect, and 4) after conditioning on lagged volume, there is a positive risk/return relation.

1 Introduction

1.1 Background and Motivation

The recent history of the stock market has been characterized by sharp downward price movements accompanied by high volume and associated with increased future volatility. On Black Monday II (October 19, 1987), the S&P 500 composite index plunged 22.9 per cent on the second highest volume ever recorded (604 million shares). On the day after the crash of '87, the S&P 500 index rose by 5.2 per cent on the highest volume ever recorded of 608 million shares. Two days after the crash on October 21st, the S&P 500 index rose 8.7 percent (the 7th highest one day increase in the period from 1928 to 1989) with the trading of 450 million shares. In 1989, the seven per cent drop on October 13th was accompanied by a fifty per cent increase in volume and followed by heavy trading on Monday, October 16th, of two and a half times the normal volume. These events of the late eighties suggest strong interrelationships among the sign and magnitude of price movements, the volatility of prices, and the trading volume.

Studies of volatility dynamics examine the relationship between large price movements and increased volatility. For the (univariate) S&P composite price index data, French, Schwert, and Stambaugh (1987) employ the GARCH specification developed by Bollerslev (1986). Engle and Bollerslev (1986) extend the GARCH model to include a unit root in the variance evolution term to accommodate the observed strong persistence in conditional variances. Nelson (1989a,b) introduces an "exponential" GARCH model which overcomes some problems with positivity restrictions and symmetry in the conditional variance function associated with standard GARCH models. Efforts to explore the determination of the risk premium for stocks have employed a variety of ARCH-M specifications (see French, Schwert and Stambaugh, 1987; Bollerslev, Engle, and Wooldridge, 1988; and Nelson, 1989a). Pagan and Hong (1989) and Harvey (1989) use nonparametric techniques to study the risk premium. The presence of ARCH-like variance shifts is thus well-documented, though there is considerable disagreement about which variant of a parametric ARCH model is most appropriate to describe the price change process.

In this study, we investigate the joint dynamics of price changes and volume on the stock

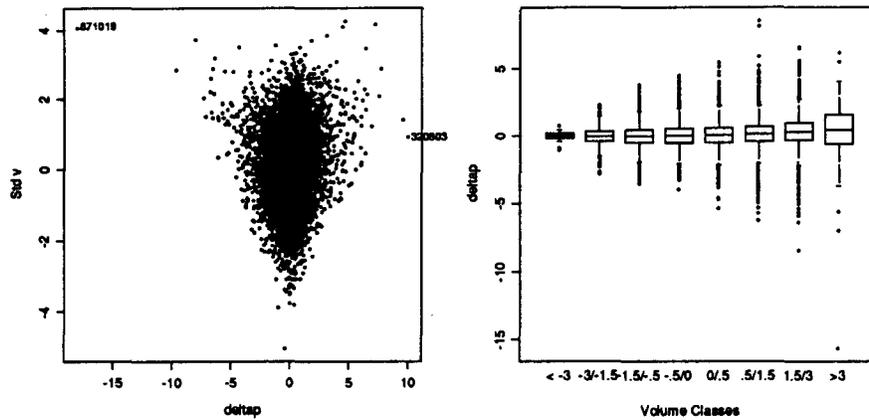


Figure 1: Scatter Plots

market. We use daily data on the S&P composite index and total New York Stock Exchange trading volume from 1928 to 1987 which is a bivariate time series of 16,127 observations.

Previous empirical work on the price and volume relationship has focused primarily on the contemporaneous relationship between price changes and volume. Transactions level, hourly, daily, weekly, and monthly data on individual stocks, futures, and stock price indices have been used to document a positive correlation between the absolute value of stock prices changes and volume (see Karpoff, 1987, and Tauchen and Pitts, 1983, for a summaries of this literature). Foster and Viswanathan (1990) use transactions level data to examine within-day price/volume relations; Mulherin and Gerety (1988) use both hourly and daily volume and returns data to document the relationship between the magnitude of price changes and volume, as well as patterns in volume by time of day and week for the period from 1900 to 1987. Lamoureux and Lastrapes (1989) enter volume directly into the GARCH variance equation in their analysis of individual stock returns data. Schwert (1989) uses monthly aggregates of daily data and finds a positive relationship between estimated volatility and current and lagged volume growth rates in linear distributed lag and VAR models. With the exception of Mulherin and Gerety (1988), most empirical studies using daily or within-day data examine relatively short time periods of between three and five years.

For our data set, the contemporaneous relationship between price movements and volume is summarized in Figure 1. The figure presents a scatterplot of adjusted price changes versus standardized volume as well as boxplots of the distribution of Δp_t for various volume ranges. The scatterplot shows that, for the most part, large price movements are associated

with unusually high volume. The boxplots demonstrate that the dispersion of the distribution of Δp_t (the height of the box is the interquartile range) increases uniformly as the volume increases. The patterns in the figure are generally consistent with existing findings on the contemporaneous positive correlation between the magnitude of price movements and volume.

Generally speaking, the empirical work on price/volume relations tends to be very data-based and not guided by rigorous equilibrium models of market behavior. The models are more statistical than economic in character, and typically neither the optimization problem facing agents nor the information structure is fully specified. The intrinsic difficulties of specifying plausible, rigorous, and empirically implementable models of volume and prices are the reasons for the informal modeling approaches commonly used. For nontrivial trading volume to emerge in equilibrium, an economic model needs to incorporate heterogeneous agents and incomplete markets, both of which are substantive complications of the familiar representative agent asset pricing models.

Recently, some interesting theoretical work investigates these complications. For example, Admati and Pfleiderer (1988, 1989) explore the implications for within-day and weekend volume and price movements of a model comprised of informed traders and liquidity traders. Huffman (1987) presents a capital growth model with overlapping generations that yields a contemporaneous volume/price relationship. More recently, Huffman (1988) and Ketterer and Marcet (1989) examine trading volume and welfare issues in various economies comprised of heterogeneous infinitely-lived agents facing limited trading opportunities. Existing models, however, do not confront the data in its full complexity and have not evolved sufficiently to guide the specification of an empirical model of daily stock market data. For instance, there seems to be no model with dynamically optimizing heterogeneous agents that can jointly account for major stylized facts — serially correlated volatility, contemporaneous volume volatility correlation, and excess kurtosis of price changes — each of which we discuss below.

In this paper, we undertake an empirical investigation of the dynamic interrelationships among price and volume movements on the stock market. Our work is motivated in part by the recent events on the stock market, which suggest that more can be learned about the

market, and in particular about volatility, by studying prices in conjunction with volume, instead of prices alone. It is also motivated by an objective of providing a full set of stylized facts that theoretical work will ultimately have to confront. Because of the limitations of existing theory, the empirical work is not organized around the specification and testing of a particular model or class of models. Instead, the empirical effort is mainly data-based. We begin with an informal graphical look at the data and then proceed ultimately to the estimation and interpretation of a seminonparametric model of the conditional joint density of market price changes and volume in Section 4. The conditional density is the fundamental statistical object of interest, as it embodies all of the information about the probabilistic structure of the data.

1.2 Objectives

The investigation has four objectives: (1) to analyze the relationships between contemporaneous volume and volatility in an estimation context that explicitly accounts for conditional heteroskedasticity and other forms of conditional heterogeneity; (2) to characterize the intertemporal relationships among prices, volatility, and volume; (3) to examine the leverage effect and the extent to which conditioning on past volume reduces or increases the symmetry of the conditional variance function; and (4) to determine what, if any, relationship there is between the conditional mean and variances of price changes.

Many of the issues implicit in these objectives relate to features of the time series probability distribution of the price change and volume data, and not to the signs and magnitudes of specific parameters. Hence, as noted above, we employ a seminonparametric approach to estimate this distribution directly, and we use kernel-based methods to corroborate major findings.

The remainder of the paper is organized as follows. In Section 2, we describe the data sources and the adjustments made to remove systematic calendar and trend effects from the location and scale of the price change and volume series. In Section 3, we review the seminonparametric approach to modeling nonlinear time series and undertake the estimation, which involves a specification search and diagnostic checking procedures. In Section 4, we examine various features of the fitted SNP density in order to address the basic research

goals described above. In Section 5, we summarize our findings.

2 Data Sources and Adjustments

The basic raw data consist of the daily closing value of the Standard and Poor's composite stock index and the daily volume of shares traded on the New York Stock Exchange. Price index data for the period from 1928 to 1985 were generously supplied to us by R. Stambaugh. We extended the price data through 1987. The volume data is from the Standard and Poor's *Security Price Index Record* (various years). The *Security Price Record* appears to be only source of a long time series on daily market volume.

The Standard and Poor's composite price index is a value-weighted arithmetic index of prices of common stocks, most of which are traded on the NYSE. In the period before March 1, 1957, the S&P composite index was made up of 90 stocks. On March 1, 1957, the index was broadened to include 500 stocks. In July 1976, S&P added a group of financial stocks to the 500 composite index. Some of these financial stocks are traded Over-the-Counter so that in recent years the 500 has included a few non-NYSE stocks.

The raw price index series, P_t , is differenced in the logs to create the price movement series, $100(\log P_t - \log P_{t-1})$, and is plotted in the top panel of Figure 2. It is immediately obvious that there is a "U"-shaped pattern in the volatility of the raw price change series. In the early thirties and the late eighties the volatility is very high, while in the middle part of the sample the volatility is low. We do not expect to explain or model very long-run shifts in the volatility of price changes. We decided, therefore, to allow for a quadratic trend in the variance of the price changes in order to focus our modeling efforts on the shorter run pattern of conditional heteroskedasticity.

Many authors have noted systematic calendar effects in both the mean and variance of price movements. Rozeff and Kinney (1976) report a January seasonal in stock market index returns, i.e. mean returns are higher in January. Keim (1983) refined this analysis of the January seasonal by studying the magnitude of the seasonal for various size-based portfolios of stocks. Keim finds that most of the seasonal is associated with the returns on small stocks in January. Constantinides (1984) points out that tax-related trading might occur around the turn of the year but that some sort of irrationality on the part of investors would be

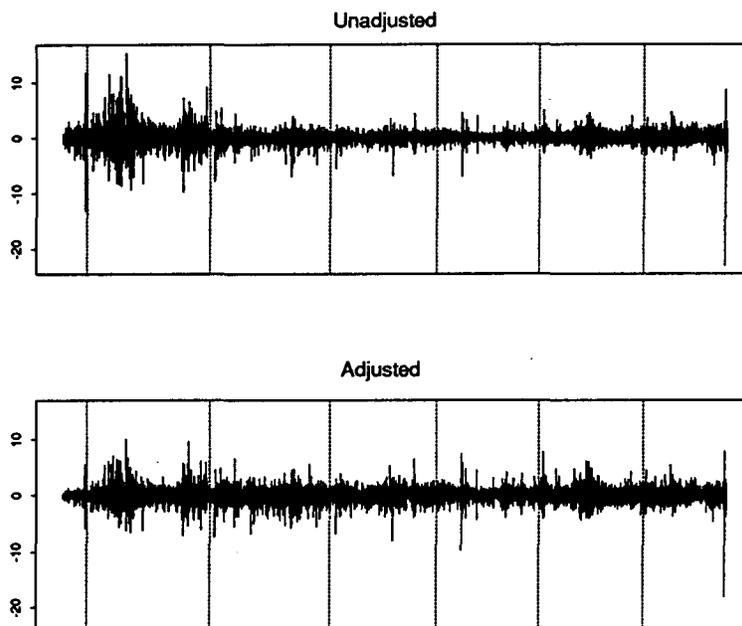


Figure 2: Price Change Data

required to induce systematic shifts in the mean of stock returns. Thus, we might expect to see a January/December seasonal in the volume series even in the absence of a mean effect on prices. French (1980) notes a weekend effect in stock returns with lower than average returns on Monday. French and Roll (1986) study the variance of stock returns around weekends and exchange holidays and document shifts in the variance associated with these non-trading periods. Ariel (1988) has uncovered evidence of an intra-month pattern of higher returns in the first half of the month. Recently, Schwert (1990a) also finds evidence of monthly and daily seasonals in means and standard deviations of returns.

In order to adjust for these documented shifts in both the mean and variance of the price and volume series, we performed a two stage adjustment process in which systematic effects are first removed from the mean and then from the variance. We use the following set of dummy and time trend variables in the adjustment regressions to capture these systematic effects:

1. Day of the week dummies (one for each day, Tuesday through Saturday).
2. Dummy variables for each number of non-trading days preceding the current trading day (dummies for each of 1, 2, 3, and 4 non-trading days since the preceding trading day). These “gap” variables capture the effects of holidays and weekends. The

distribution of these gaps in the trading record are as follows:

- gap of one non-trading day: 1339
 - gap of two non-trading days: 1873
 - gap of three non-trading days: 223
 - gap of four non-trading days: 5
3. Dummy variables for months of March, April, May, June, July, August, September, October, and November.
 4. Dummy variables for each week of December and January. These variables are designed to accommodate the well-known “January” effect in both the mean and variance of prices and volume.
 5. Dummy variables for each year, 1941-1945.
 6. t, t^2 , time trend variables. (Note: these variables are not included in the mean regressions for the price change.)

This list of variables is generally self-explanatory, though we should elaborate on a few points concerning the “gap” variables in the second group. If trading occurred on the preceding day, then there is no gap in the trading record and no dummy is included; there are 12,686 such days. The Bank Holiday of 1933 is associated with a gap of eleven days over which the increase in the raw S&P index is the largest close-to-close movement in the entire data set. No dummy is included for this single eleven-day gap because doing so would, in effect, replace the largest upward change in the price index with the unconditional mean of the price changes, which in our view would not accurately reflect what transpired over the Bank Holiday. Finally, in both the adjustments for the mean and variance of volume, the coefficients of the four gap variables are constrained to lie along a line. Without such constraints, the adjustment process itself appears to create some very extreme and implausible values in the adjusted volume process, particularly for the five days in which $\text{gap}=4$.

To perform the adjustment, we first regress $100(\log P_t - \log P_{t-1})$ or $\log(V_t)$ (V_t represents volume) on the set of adjustment variables:

$$w = x'\beta + u \quad (\text{mean equation})$$

Here w is the series to be adjusted and x contains the adjustment regressors. The least squares residuals are taken from the mean equation to construct a variance equation:

$$\log(u^2) = x'\gamma + \epsilon \quad (\text{variance equation})$$

This regression is used to standardize the residuals from the mean equation and then a final linear transformation is performed to calculate adjusted w :

$$w_{adj} = a + b\left(\frac{\hat{u}}{\exp(x'\hat{\gamma}/2)}\right),$$

where a and b are chosen so that the sample means and variances of w and w_{adj} are the same. The linear transformation makes the units of measurement of adjusted and unadjusted data the same, which facilitates interpretation of our empirical results. In what follows, Δp_t denotes adjusted $100(\log P_t - \log P_{t-1})$ and v_t denotes adjusted $\log(V_t)$.

Table 1 shows the estimated coefficients in the mean and variance adjustment equations for the price movements series. The patterns confirm the well-known day-of-week and January effects. Monday has a lower mean return than any other day of the week along with a lower variance than any other weekday. Price changes are higher in the last week of December and the first week of January. The variance is higher in the first week of January. The effect of wartime seems to be confined to a reduction in the variance.

The Δp_t series is plotted in the bottom panel of Figure 2. The adjustments have considerably homogenized the series, allowing us to focus on the day-to-day dynamic structure under an assumption of stationarity.

We do not regard the Δp_t series as a market returns series since the S&P index is not adjusted for dividend payout. If dividend payout is lumpy and the payout has an appreciable effect on the index (due to groups of stocks going ex-dividend together), then the lumpy dividends can create yet another possible systematic calendar effect. In order to investigate the lumpiness of dividend payout, we obtained daily data on the total dividend payout of

	Location			Variance		
	Coefficient	Std Dev	p-value	Coefficient	Std Dev	p-value
<i>Day of the week</i>						
Monday	—	—	—	—	—	—
Tuesday	0.115	0.065	0.079	0.349	0.136	0.010
Wednesday	0.167	0.066	0.011	0.294	0.137	0.032
Thursday	0.122	0.066	0.065	0.171	0.138	0.215
Friday	0.142	0.065	0.030	0.179	0.137	0.190
Saturday	0.220	0.071	0.002	-0.742	0.149	0.000
<i>Number of days since the preceding trading day</i>						
GAP1	-0.125	0.059	0.035	0.385	0.124	0.002
GAP2	-0.117	0.068	0.083	0.517	0.142	0.000
GAP3	-0.257	0.083	0.002	0.443	0.174	0.012
GAP4	0.439	0.519	0.398	0.617	1.078	0.567
<i>Month or week</i>						
Jan 1- 7	0.206	0.078	0.008	0.294	0.163	0.071
Jan 8-14	0.033	0.072	0.652	0.070	0.150	0.640
Jan 15-21	-0.003	0.072	0.970	-0.177	0.150	0.239
Jan 22-31	0.077	0.063	0.225	-0.122	0.131	0.354
February	—	—	—	—	—	—
March	0.016	0.045	0.722	0.026	0.094	0.782
April	0.053	0.046	0.245	0.036	0.095	0.708
May	-0.032	0.045	0.486	0.093	0.095	0.325
June	0.060	0.046	0.190	0.117	0.095	0.217
July	0.086	0.046	0.060	0.078	0.095	0.410
August	0.064	0.045	0.156	0.029	0.094	0.757
September	0.060	0.046	0.191	0.357	0.096	0.000
October	0.000	0.045	0.989	0.239	0.094	0.011
November	0.031	0.047	0.502	0.430	0.097	0.000
Dec 1- 7	0.094	0.072	0.195	0.137	0.150	0.362
Dec 8-14	-0.078	0.072	0.283	0.013	0.150	0.929
Dec 15-21	0.016	0.072	0.826	-0.111	0.150	0.459
Dec 22-31	0.201	0.067	0.003	-0.183	0.140	0.187
<i>Year</i>						
1941	-0.094	0.068	0.1621	-0.528	0.141	0.000
1942	0.008	0.068	0.8992	-0.338	0.141	0.017
1943	0.028	0.068	0.6798	-0.586	0.141	0.000
1944	0.016	0.068	0.8163	-0.907	0.142	0.000
1945	0.075	0.069	0.3020	-0.373	0.145	0.010
<i>Trend</i>						
(t/16127)	—	—	—	-8.678	0.270	0.000
(t/16127) ²	—	—	—	7.112	0.260	0.000

Table 1: Adjustment Regressions for Differenced Log Price, 16,127 observations

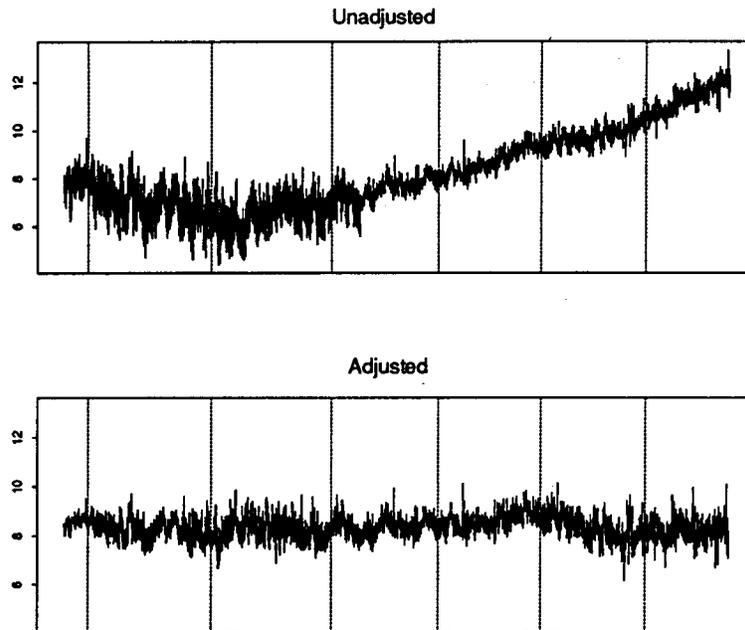


Figure 3: Volume Data

the S&P 100 index in the period 1979 to 1987. (These data are used in Harvey and Whaley, 1989, and we thank the authors for allowing us access to this data.) Our analysis, which is available upon request, indicates that dividends are lumpy with payouts concentrated at certain times of each quarter. In spite of the dividend lumpiness, the S&P index itself does not show detectable movements in times of high dividend payouts. We therefore do not regard the failure to adjust for dividends as an important factor in modelling the daily S&P price index. Schwert (1990a) also finds that volatility estimates are not influenced appreciably by dividends.

The top panel of Figure 3 shows the unadjusted log volume series. The series exhibits a clear trend in level as might be expected. We experimented with transforming the volume series into a turnover series by dividing the volume by measures of the number of outstanding shares. However, plots revealed that the turnover series has a U-shaped pattern with very high turnover in the late twenties and the late eighties. The pattern suggests that division by the number of outstanding shares is an inadequate detrending strategy. We thus decided to include a quadratic trend in both the mean and variance equation for volume along with the same dummy variables to account for calendar and wartime effects as were used in adjusting the price change series.

	Location			Variance		
	Coefficient	Std Dev	p-value	Coefficient	Std Dev	p-value
<i>Day of the week</i>						
Monday	—	—	—	—	—	—
Tuesday	0.035	0.022	0.109	0.292	0.110	0.008
Wednesday	0.065	0.022	0.003	0.263	0.113	0.020
Thursday	0.058	0.022	0.009	0.317	0.114	0.005
Friday	0.023	0.022	0.307	0.339	0.113	0.003
Saturday	-0.776	0.025	0.000	0.574	0.127	0.000
<i>Number of days since the preceding trading day</i>						
GAP1	-0.069	0.022	0.001	0.545	0.110	0.000
GAP2	-0.008	0.020	0.697	0.374	0.104	0.000
GAP3	0.053	0.026	0.043	0.204	0.134	0.129
GAP4	0.115	0.036	0.002	0.033	0.185	0.858
<i>Month or week</i>						
Jan 1- 7	0.040	0.029	0.169	-0.045	0.148	0.750
Jan 8-14	0.077	0.027	0.004	-0.074	0.136	0.587
Jan 15-21	0.021	0.027	0.424	-0.019	0.136	0.890
Jan 22-31	0.025	0.023	0.278	-0.074	0.119	0.532
February	—	—	—	—	—	—
March	-0.025	0.017	0.143	-0.047	0.085	0.582
April	-0.010	0.017	0.548	-0.106	0.086	0.219
May	-0.063	0.017	0.000	0.045	0.086	0.602
June	-0.114	0.017	0.000	-0.053	0.086	0.540
July	-0.134	0.017	0.000	-0.204	0.086	0.018
August	-0.211	0.017	0.000	-0.059	0.086	0.493
September	-0.067	0.017	0.000	-0.073	0.087	0.399
October	-0.029	0.017	0.085	0.035	0.086	0.680
November	0.022	0.017	0.198	-0.064	0.088	0.470
Dec 1- 7	0.021	0.027	0.433	-0.051	0.137	0.710
Dec 8-14	0.060	0.027	0.024	-0.051	0.137	0.711
Dec 15-21	0.055	0.027	0.041	-0.219	0.137	0.110
Dec 22-31	0.028	0.025	0.254	0.018	0.126	0.890
<i>Year</i>						
1941	-0.779	0.025	0.000	-0.414	0.128	0.001
1942	-1.058	0.025	0.000	-0.436	0.128	0.000
1943	-0.266	0.025	0.000	-0.181	0.128	0.159
1944	-0.311	0.025	0.000	-0.418	0.129	0.001
1945	0.080	0.026	0.002	-0.666	0.131	0.000
<i>Trend</i>						
($t/16127$)	-5.117	0.048	0.000	-3.651	0.244	0.000
($t/16127$) ²	9.577	0.046	0.000	1.486	0.235	0.000

Table 2: Adjustment Regressions for Log Volume, 16,127 observations

As seen from Table 2, volume is lower on Monday and Saturday and there are pronounced seasonal patterns by month of the year, with lower volume in the summer months. In the war years, the level of volume was much lower than normal. The adjusted log of volume series given in the bottom panel of Figure 3 shows relatively homogeneous variation around a mean level.

It is important to note that our quadratic detrending of price volatility and the volume series is best viewed as a band limited filter which passes everything except extremely low-frequency behavior. We certainly do not suggest that these trends can be extrapolated. Officer (1973) and Schwert (1989) conclude that great depression is simply an unusual event with two to three times higher volatility than any other period since 1870. Schwert (1990b) suggests that the crash of 1987 is also characterized by unusually high volatility. An alternative to detrending would be to introduce dummy variables for the depression and the 87 crash period. But evidence in the data that there has been a gradual increase in volatility since the early 1970's persuaded us that a "u" shaped quadratic trend is a more reasonable procedure.

The adjustment procedures are designed to remove long-run trend and systematic calendar effects, which are well-documented. We take care to make only those adjustments for which there is strong evidence either in the data or from previous work and avoided making a large number of arbitrary adjustments. In our view, Figure 2 and Figure 3 completely summarize the characteristics of the adjustments. The adjusted series exhibit homogeneous behavior and can reasonably be modeled by a stationary time series model.

3 Conditional Density Estimation

We use the nonparametric estimation strategy proposed by Gallant and Tauchen (1989, 1990). Their SNP approach, which is explained below, has the advantage of giving reasonably smooth density estimates even in high dimensions. It is a series expansion whose leading term can be chosen to be a particularly successful parametric model and whose higher order terms accommodate deviations from the parametric model. We use kernel method to corroborate our SNP estimates.

3.1 Semiparametric (SNP) Estimators

The method is based on the notion that a Hermite expansion can be used as a general purpose nonparametric estimator of a density function. Letting z denote an M -vector, the particular Hermite expansion employed has the form $h(z) \propto [P(z)]^2 \phi(z)$ where $P(z)$ denotes a multivariate polynomial of degree K_z and $\phi(z)$ denotes the density function of the (multivariate) Gaussian distribution with mean zero and the identity matrix as its variance-covariance matrix. The constant of proportionality is the divisor $\int [P(z)]^2 \phi(z) dz$ which makes $f(z)$ integrate to one. Because of this division, the density is a homogeneous function of the coefficients of the polynomial $P(z)$ and these coefficients can only be determined to within a scalar multiple. To achieve a unique representation, the constant term of the polynomial part is put to one.

The location/scale shift $y = Rz + \mu$, where R is an upper triangular matrix and μ is an M -vector, followed by a change of variables leads to a parameterization that is easy to interpret $f(y|\theta) \propto \{P[R^{-1}(y - \mu)]\}^2 \{\phi[R^{-1}(y - \mu)]/|\det(R)|\}$. Because $\{\phi[R^{-1}(y - \mu)]/|\det(R)|\}$ is the density function of the M -dimensional multivariate Gaussian distribution with mean μ and variance-covariance matrix $\Sigma = RR'$, and because the leading term of the polynomial part is one, the leading term of the entire expansion is the multivariate Gaussian density function; denote it by $n_M(y|\mu, \Sigma)$. When K_z is put to zero, one gets $n_M(y|\mu, \Sigma)$ exactly. When K_z is positive, one gets a Gaussian density whose shape is modified due to multiplication by a polynomial in the normalized error $z = R^{-1}(y - \mu)$. The shape modifications thus achieved are rich enough to accurately approximate densities from a large class that includes densities with fat t -like tails, densities with tails that are thinner than Gaussian, and skewed densities.

The parameters θ of $f(y|\theta)$ are made up of the coefficients of the polynomial $P(z)$ plus μ and R and are estimated by maximum likelihood. Equivalent to maximum likelihood but more stable numerically is to estimate θ in a sample of size n by minimizing $s_n(\theta) = (-1/n) \sum_{t=1}^n \ln[f(y_t|\theta)]$. If the number of parameters p_θ grows with the sample size n , the true density and various features of it such as derivatives and moments are estimated consistently. Because the method is parametric yet has nonparametric properties, it is termed semiparametric to suggest that it lies halfway between parametric and non-

parametric procedures.

This basic approach is adapted to the estimation of the conditional density of a multiple time series that has a Markovian structure as follows. By a Markovian structure, one means that the conditional density of the M -vector y_t given the entire past y_{t-1}, y_{t-2}, \dots depends only on L lags from the past. For notational convenience, we collect these lags together in a single vector denoted as x_{t-1} which has length $M \cdot L$. As above, a density is obtained by a location/scale shift $y_t = Rz_t + \mu_x$ off a sequence of normalized errors $\{z_t\}$. μ_x is a linear function of x_{t-1} making the leading term of the expansion $n_M(y|\mu_x, \Sigma)$ which is a Gaussian vector autoregression or Gaussian VAR. In time series analysis, the z_t are usually referred to as linear innovations. In order to permit the innovations to be conditionally heterogeneous, the coefficients of the polynomial $P(z)$ are, themselves, polynomials of degree K_x in x_{t-1} . This polynomial is denoted as $P(z, x)$. When $K_x = 0$, the $\{z_t\}$ are homogenous, as the conditional density of z_t does not depend upon x_{t-1} . When $K_x > 0$, the $\{z_t\}$ are conditionally heterogeneous.

To keep K_x small when the data exhibit marked conditional heteroskedasticity, the leading term of the expansion can be put to a Gaussian ARCH rather than a Gaussian VAR. This is done by letting R be a linear function of the absolute values of (the elements of) L_r of the lagged y_t , centered and scaled to have mean zero and identity covariance matrix. The classical ARCH (Engle, 1982) has Σ_x depending on a linear function of squared lagged residuals. The SNP version of ARCH is more akin to the suggestions of Nelson (1989a) and Davidian and Carroll (1987). Denoting this function by R_x and letting $\Sigma_x = R_x R_x'$, the form of the conditional density becomes $f(y|x, \theta) \propto [P(z, x)]^2 n_M(y|\mu_x, \Sigma_x)$ where $z = R_x^{-1}(y - \mu_x)$ and θ denotes the coefficients of the polynomial $P(z, x)$ and the Gaussian ARCH $n_M(y|\mu_x, \Sigma_x)$ collected together. The parameters are estimated by minimizing $s_n(\theta) = (-1/n) \sum_{t=1}^n \ln[f(y_t|x_{t-1}, \theta)]$.

Hereafter, we shall distinguish between the total number of lags under consideration, which is L , the number of lags in the x part of the polynomial $P(z, x)$, which we denote by L_p , and the number of lags in Σ_x , which is L_r . The vector x has length $M \cdot L$ where $L = \max(L_r, L_p)$.

Large values of M can generate a large number of interactions (cross product terms) for

even modest settings of degree K_z ; similarly, for $M \cdot L_p$ and K_x . Accordingly, we introduce two additional tuning parameters, I_z and I_x , to represent filtering out of these high order interactions. $I_z = 0$ means no interactions are suppressed, $I_z = 1$ means the highest order interactions are suppressed, namely those of degree $K_z - 1$. In general, a positive I_z means all interactions of order $K_z - I_z$ and larger are suppressed; similarly for $K_x - I_x$.

In summary, L_r and L_p determine the location/scale shift $y = R_x z_t + u_x$ and hence determine the nature of the leading term of the expansion. The number of lags in the location shift u_x is the overall lag length L which is the maximum of L_r and L_p . The number of lags in the scale shift R_x is L_r . The number of lags that go into the x part of the polynomial $P(z, x)$ is L_p . The parameters K_z and K_x determine the degree of $P(z, x)$ and hence the nature of the innovation process $\{z_t\}$. I_z and I_x determine filters that suppress interactions when set to positive values.

Putting certain of the tuning parameters to zero implies sharp restrictions on the process $\{y_t\}$, the more interesting of which are:

Parameter setting	Characterization of $\{y_t\}$
$L_r = 0, L_p = 0, K_z = 0, K_x = 0$	iid Gaussian
$L_r = 0, L_p > 0, K_z = 0, K_x = 0$	Gaussian VAR
$L_r = 0, L_p > 0, K_z > 0, K_x = 0$	non-Gaussian VAR, homogeneous innovations
$L_r > 0, L_p = 0, K_z = 0, K_x = 0$	Gaussian ARCH
$L_r > 0, L_p > 0, K_z > 0, K_x = 0$	non-Gaussian ARCH, homogeneous innovations

3.2 Model Selection

We used the model selection strategy suggested by Gallant, Hsieh, and Tauchen (1989) and Gallant, Hansen, and Tauchen (1990). The Schwarz criterion (Schwarz, 1978; Potscher, 1989) is used to move along an upward expansion path until an adequate model is determined. This Schwarz-preferred model is then subjected to a battery of specification tests. These tests can indicate that further expansion of the model is necessary. For data from financial markets, experience suggests that this strategy will inevitably select $L_r \geq L_p$ and $K_z \geq 4$. Thus, specifications that violated these conditions were excluded from consideration *a priori*.

Specification tests are conducted for each fit from scaled residuals $\{\hat{u}_t\}$ which are cal-

culated as follows. By computing analytically the moments of the estimated conditional density, the estimated conditional mean $\hat{\mathcal{E}}(y|x_{t-1})$ and variance $\hat{V}ar(y|x_{t-1})$ are obtained at each $x_{t-1} = (y_{t-L}, \dots, y_{t-1})$ in the sample. Using these, a scaled residual is computed as $\hat{u}_t = [\hat{V}ar(y|x_{t-1})]^{-1/2}[y_t - \hat{\mathcal{E}}(y|x_{t-1})]$ where $[\hat{V}ar(y|x_{t-1})]^{-1/2}$ denotes the inverse of the Cholesky factor of the conditional variance.

We conduct diagnostic tests for predictability in both the scaled residuals and the squares of the scaled residuals. Predictability of the scaled residuals would suggest inadequacies in the conditional mean estimate implied by the fitted density, and thus below such tests are termed mean tests. Similarly, predictability of the squared scaled residuals would suggest inadequacies in the implied estimate of the conditional variance, and thus such tests are termed variance tests. For both mean and variance, we conduct two types of tests for predictability, one of which is sensitive to short-term misspecification while the other is sensitive to long-term misspecification.

For the conditional mean, the short-term diagnostic test is a test for the significance of a regression of scaled residuals on linear, quadratic, and cubic terms on twenty lagged values of the elements of the series. The long-term test is a test for the significance of a regression of scaled residuals on annual dummies to check for a failure to capture long-term trends. For the conditional variance, the tests are the same with the squares of the scaled residuals as the dependent variable in these regressions. The significance test is the F -test when the residuals are from the univariate price series and is the Wilk's test when the residuals are from the bivariate price and volume series. It should be noted that because of the "Durbin effects" of prefitting discussed in Newey (1985) and Tauchen (1985), the p-values could be somewhat inaccurate, even asymptotically.

For each of the specifications considered, the settings of the tuning parameters L_r , L_x , K_z , I_z , K_x , I_x , the number of parameters p_θ that they imply, the value of the minimized objective function $s_n(\hat{\theta})$, Schwarz's criterion, and the battery of diagnostic tests are reported in Table 3 for the univariate price series $y_t = \Delta p_t$ and in Table 4 for the bivariate price and volume series $y_t = (\Delta p_t, v_t)$. All reported values are comparable as the same number of leading observations (27) were set aside to provide the initial lags in every fit. The net sample size is 16,100 observations.

										Specification tests							
										Annual dummies				Twenty lag cubic			
										Mean		Variance		Mean		Variance	
L_r	L_p	K_x	I_x	K_x	I_x	p_θ	Obj	Schwarz		F	p-val	F	p-val	F	p-val	F	p-val
Raw data										1.14	.2147	20.77	.0000	9.27	.0001	45.11	.0001
Adjusted data										1.32	.0510	11.24	.0000	10.68	.0001	30.21	.0001
2	2	0	0	0	0	6	1.3471	1.3489	1.23	.1140	10.63	.0000	2.24	.0001	17.00	.0001	
4	4	0	0	0	0	10	1.3219	1.3249	1.16	.1859	6.96	.0001	1.46	.0114	8.37	.0001	
6	6	0	0	0	0	14	1.3082	1.3124	1.19	.1565	4.62	.0001	1.27	.0778	4.73	.0001	
6	6	4	0	0	0	18	1.2811	1.2865	1.21	.1274	4.35	.0001	1.25	.0922	4.55	.0001	
6	6	6	0	0	0	20	1.2805	1.2865	1.16	.1929	4.31	.0001	1.31	.0543	4.54	.0001	
6	6	4	0	1	0	48	1.2728	1.2872	1.14	.2167	4.59	.0001	1.06	.3520	2.48	.0001	
8	8	0	0	0	0	18	1.3032	1.3086	1.14	.2169	3.88	.0001	1.09	.2906	4.34	.0001	
8	8	4	0	0	0	22	1.2778	1.2844	1.17	.1726	3.64	.0001	1.09	.2911	4.19	.0001	
8	8	4	0	1	0	62	1.2687	1.2873	1.11	.2633	3.97	.0001	0.93	.6378	1.70	.0006	
10	10	0	0	0	0	22	1.2981	1.3047	1.12	.2431	3.45	.0001	1.06	.3558	4.22	.0001	
10	10	4	0	0	0	26	1.2743	1.2821	1.17	.1708	3.26	.0001	1.06	.3539	4.08	.0001	
10	10	4	0	1	0	76	1.2642	1.2870	1.09	.2946	3.63	.0001	0.84	.8114	1.22	.1185	
12	12	0	0	0	0	26	1.2945	1.3023	1.09	.2982	2.88	.0001	1.01	.4489	3.74	.0001	
12	12	4	0	0	0	30	1.2714	1.2804	1.16	.1898	2.70	.0001	1.00	.4688	3.66	.0001	
12	12	4	0	1	0	90	1.2610	1.2880	1.07	.3360	3.09	.0001	0.81	.8595	0.82	.8420	
16	16	0	0	0	0	34	1.2917	1.3019	1.06	.3504	2.44	.0001	0.96	.5635	3.75	.0001	
16	16	4	0	0	0	38	1.2688	1.2802	1.12	.2471	2.26	.0001	0.95	.5749	3.75	.0001	
16	2	4	0	1	0	34	1.2634	1.2737	1.18	.1642	2.37	.0001	1.19	.1513	1.87	.0001	
16	4	4	0	1	0	46	1.2607	1.2745	1.10	.2816	2.46	.0001	0.96	.5617	1.34	.0404	
16	6	4	0	1	0	58	1.2597	1.2771	1.11	.2545	2.50	.0001	0.89	.7154	1.09	.2952	
16	8	4	0	1	0	70	1.2588	1.2799	1.12	.2465	2.56	.0001	0.83	.8258	0.83	.8256	
16	16	4	0	1	0	118	1.2571	1.2925	1.11	.2680	2.52	.0001	0.70	.9599	0.77	.9003	
20	20	0	0	0	0	42	1.2893	1.3019	1.10	.2718	2.16	.0001	0.90	.6877	3.58	.0001	
20	20	4	0	0	0	46	1.2671	1.2809	1.15	.2060	2.00	.0001	0.89	.7035	3.58	.0001	

Table 3: Univariate Series: Optimized Likelihood and Residual Diagnostics

First consider Table 3. The Schwarz criterion is computed as $s_n(\hat{\theta}) + (1/2)(p_\theta/n) \ln(n)$ with small values of the criterion preferred. The criterion rewards good fits as represented by small $s_n(\hat{\theta})$ but uses the term $(1/2)(p_\theta/n) \ln(n)$ to penalize good fits gotten by means of excessively rich parameterizations. The criterion is conservative in that it selects sparser parameterizations than the Akaike information criterion which uses the penalty term p_θ/n in place of $(1/2)(p_\theta/n) \ln(n)$. Schwarz is also conservative in the sense that it is at the high end of the permissible range of penalty terms in certain model selection settings (Potscher, 1989). Of the models in Table 3, the Schwarz preferred model has $L_r = 16$, $L_p = 2$, $K_x = 4$, $I_x = 0$, $K_x = 1$, $I_x = 0$ with $p_\theta = 34$. The short-term variance diagnostic indicates that there is short-term conditional heterogeneity of some sort that is not accounted for by the Schwarz preferred model but is adequately approximated if the lag on the polynomial part of the model is moved from $L_p = 2$ to $L_p = 6$, thus increasing p_θ to 58 (278 observations per parameter). The long-term variance diagnostic indicates that there is heterogeneity of some sort associated with long-term trends in variance that are not removed by the adjustments described previously nor adequately approximated by any of the SNP models. We return to

										Specification tests							
										Annual dummies				Ten lag cubic			
										Mean		Variance		Mean		Variance	
L_r	L_p	K_z	I_z	K_x	I_x	p_θ	Obj	Schwarz		W	p-val	W	p-val	W	p-val	W	p-val
Raw data										.057	.0000	.037	.0000	.026	.0000	.016	.0000
Adjusted data										.617	.0000	.577	.0000	.244	.0000	.210	.0000
2	2	0	0	0	0	25	2.1083	2.1158		.974	.0001	.928	.0000	.900	.0000	.930	.0000
3	3	0	0	0	0	35	2.0707	2.0812		.980	.0001	.934	.0000	.933	.0000	.951	.0000
4	4	0	0	0	0	45	2.0519	2.0655		.984	.0001	.942	.0000	.950	.0000	.960	.0001
4	4	4	0	0	0	59	1.9187	1.9365		.987	.0001	.942	.0000	.941	.0000	.958	.0000
4	4	4	0	1	0	179	1.8624	1.9162		.985	.0001	.950	.0000	.963	.0001	.977	.0001
6	4	0	0	0	0	65	2.0293	2.0489		.987	.0001	.949	.0001	.963	.0001	.971	.0001
8	4	0	0	0	0	85	2.0177	2.0432		.988	.0001	.950	.0001	.969	.0001	.971	.0001
10	4	0	0	0	0	105	2.0085	2.0400		.989	.0001	.951	.0001	.975	.0001	.971	.0001
12	4	0	0	0	0	125	2.0022	2.0397		.990	.0023	.952	.0001	.976	.0001	.971	.0001
12	4	4	2	0	0	134	1.9110	1.9512		.978	.0001	.948	.0001	.958	.0001	.953	.0001
12	4	4	2	1	0	214	1.8634	1.9277		.985	.0001	.954	.0001	.983	.0001	.981	.0001
12	4	4	1	1	0	232	1.8260	1.8957		.989	.0005	.954	.0001	.983	.0001	.985	.0001
12	4	4	0	1	0	259	1.8180	1.8958		.989	.0005	.955	.0001	.983	.0001	.986	.0001
12	4	4	2	2	1	294	1.8358	1.9241		.985	.0001	.958	.0001	.991	.0962	.987	.0001
12	6	4	2	1	0	254	1.8516	1.9279		.986	.0001	.956	.0001	.984	.0001	.976	.0001
16	4	0	0	0	0	165	1.9950	2.0445		.990	.0109	.954	.0001	.976	.0001	.971	.0001
16	4	4	2	0	0	174	1.9058	1.9581		.979	.0001	.950	.0001	.957	.0001	.953	.0001
16	4	4	2	1	0	254	1.8577	1.9340		.985	.0001	.954	.0001	.984	.0001	.978	.0001
16	4	4	2	2	1	334	1.8275	1.9278		.985	.0001	.959	.0001	.991	.0961	.988	.0001
16	4	4	1	2	1	368	1.7848	1.8953		.988	.0001	.961	.0001	.992	.1674	.990	.0027
16	4	4	0	2	1	419	1.7735	1.8994		.989	.0006	.960	.0001	.992	.1620	.989	.0005

Table 4: Bivariate Series: Optimized Likelihood and Residual Diagnostics

this point below.

Similar considerations applied to Table 4 have $L_r = 16$, $L_p = 4$, $K_z = 4$, $I_z = 1$, $K_x = 2$, $I_x = 1$ with $p_\theta = 368$ (88 observations per parameter) as the Schwarz preferred model for the bivariate price and volume process. The short-term diagnostics do not suggest movement away from the Schwarz preferred model.

Both the Schwarz preferred univariate and bivariate fitted models fail the regression-based diagnostic for long-term variance shifts. This suggests that our leading ARCH-like term may be insufficient to capture the degree of persistence in variance found in the price data. It may also be the case that the rejections are caused by very small but statistically significant departures found with our extremely large datasets. Figure 4 compares the monthly range of the price change data with monthly averages of our estimated conditional standard deviations computed from the bivariate fit. Our fitted standard deviations display an extremely high degree of persistence which is characteristic of the data. Hence, the omitted heterogeneity detected by the long-term variance diagnostics is probably very slight and should not affect our examination of the short-term price and volume dynamics. If extreme long-term persistence in variance is required for long-term dynamics, it would be possible

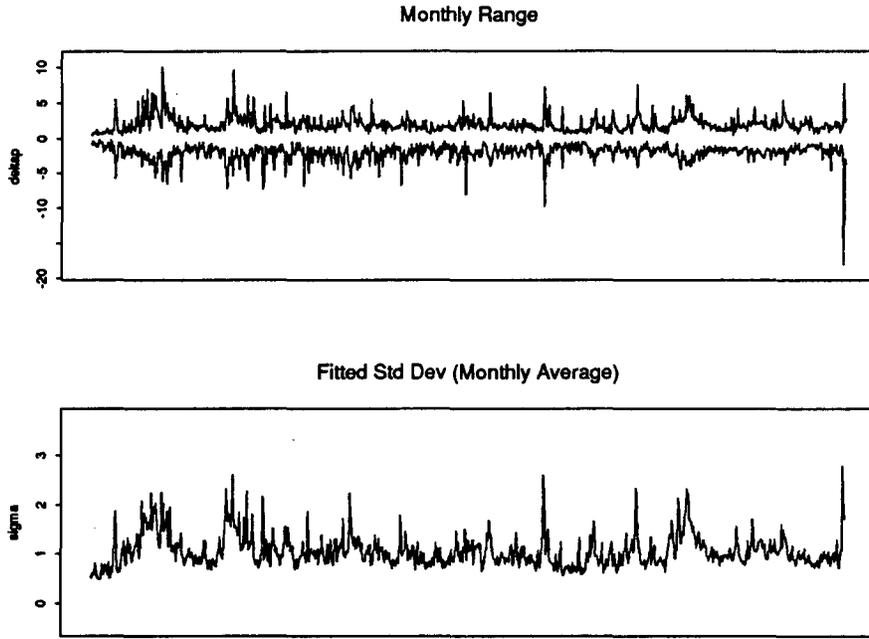


Figure 4: Monthly Volatility

to build in a unit root in variance in a similar manner to the integrated GARCH model literature (c.f. Engle and Bollerslev, 1986).

3.3 Kernel Estimators

We use Gaussian kernel estimates of the conditional mean and variance as a means of cross-checking seminonparametric estimates. A Gaussian kernel estimate of the mean conditional on an x comprised of L lags is computed as

$$\hat{\mathcal{E}}(y|x) = \frac{\sum_{t=1}^n y_t n_{ML}(x_{t-1}|x, h\Sigma)}{\sum_{t=1}^n n_{ML}(x_{t-1}|x, h\Sigma)}$$

where $n_{ML}(x_{t-1}|x, \Sigma)$ denotes the multivariate Gaussian density function of dimension $M \cdot L$, mean x , and variance-covariance matrix Σ evaluated at $x_{t-1} = (y_{t-L}, \dots, y_{t-1})$; note that y_t has length M and x_{t-1} has length $M \cdot L$. The tuning parameters are L, Σ , and h . For all estimates we took Σ to be diagonal with estimates of the unconditional variance of x_{t-1} along the diagonal. Similarly, a kernel estimate of the variance-covariance matrix conditional on an x comprised of L lags is computed as

$$\hat{V}ar(y|x) = \frac{\sum_{t=1}^n [y_t - \hat{\mathcal{E}}(y|x)][y_t - \hat{\mathcal{E}}(y|x)]' n_{ML}(x_{t-1}|x, h\Sigma)}{\sum_{t=1}^n n_{ML}(x_{t-1}|x, h\Sigma)}.$$

As seen from the form of the estimator, a kernel estimate is a locally weighted average. Points x_{t-1} that are far from x in the metric $\rho(x, x_{t-1}) = n_{ML}(x_{t-1}|x, h\Sigma)$ have little influence.

Because of this, as one moves x to the fringes where data is sparse, estimates become very erratic. SNP estimates are not as erratic as kernel estimates in the fringes as the leading terms of the expansion smooth and interpolate between points. Well outside the data, the leading terms dominate completely. Inside the mass of the data, the SNP estimates appear to behave qualitatively like kernel estimates, as seen later. Kernel estimates and SNP estimates both differ in kind from parametric estimates in that aberrant observations tend to have only a local effect, whereas they have a global effect on parametric estimates.

Following Robinson (1983) we selected the tuning parameters visually, trying for smoothness within the data and disregarding erratic behavior in the fringes. There is some trade-off between L and h ; larger L can be compensated by larger h . We put L to 4 so as to correspond to the SNP value $L_p = 4$ determined above. However, one could equally as well put L to 5 or 6 and make a compensating change to h and get approximately the same estimates as we report. These values for L all imply some smudging of variance estimates, because L should be much larger to control for conditional heteroskedasticity, as seen above. Unfortunately, kernel estimates break down completely at large settings of L , and this smudging just has to be accepted. For mean estimates, the selected values were $h = 15$ for the univariate price process and $h = 8$ for the bivariate price and volume process. For variance estimates the values were $h = 7$ and $h = 4$, respectively.

4 Empirical Findings

The introduction raises a number of economic issues concerning the properties of the price process and the relationship between the price process and volume. Some issues pertain to the predictability of price movements, the nature of the relationship between price volatility and volume, and the shape characteristics of the probability density of Δp_t . Others concern asymmetry of the conditional variance function (the leverage effect) and the relationship between the risk premium and conditional price volatility.

The seminonparametric estimate of the one-step ahead, bivariate, conditional density $f(\Delta p_t, v_t \mid \Delta p_{t-1:16}, v_{t-1:16})$, embodies the sample information on these issues. Being a nonlinear function of thirty four variables, though, the fitted conditional density is difficult to describe. Our reporting strategy is to examine various statistical features of the density —

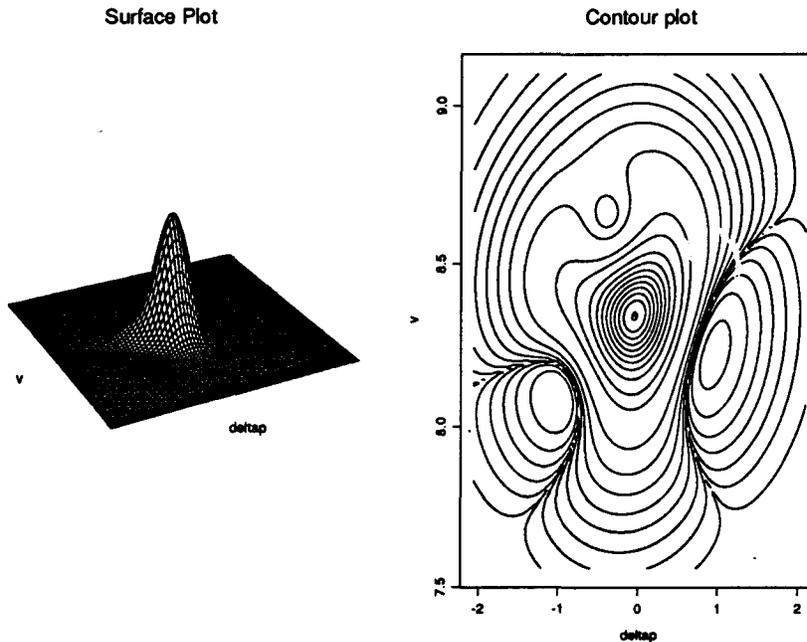


Figure 5: Bivariate Conditional Density

marginals, low order moments, and conditional moment functions — and to interpret these features in view of the economic issues raised before. Such a reporting strategy is naturally graphically oriented.

4.1 The Conditional Density at the Mean

Figure 5 shows the bivariate, conditional density of $(\Delta y_t, v_t)$ given that all lags in the conditioning set are put to their unconditional means, which is denoted as $f(\Delta p_t, v_t | \Delta p_{t-1:16} = mean, v_{t-1:16} = mean)$. The surface plot in the left-hand panel suggests that over most of its support the fitted density is quite smooth. There is some roughness as indicated in the contour plot in the right-hand panel. (By roughness we mean oscillations in the fitted density which occur when the SNP estimator attempts to fit small clumps of isolated data points.) In this plot, we highlight roughness by choosing contours associated with very low density values. From this plot, the roughness is seen to be well out in the tails. All told, the SNP density estimation procedure achieves a high degree of smoothing of the empirical distribution of price change and volume. We should, however, remain alert to the roughness in the tails, as it could heavily influence features of the density that depend strongly on extreme tail behavior.

Figure 6 is a plot of the marginal, conditional density of Δp_t given that all lags in the

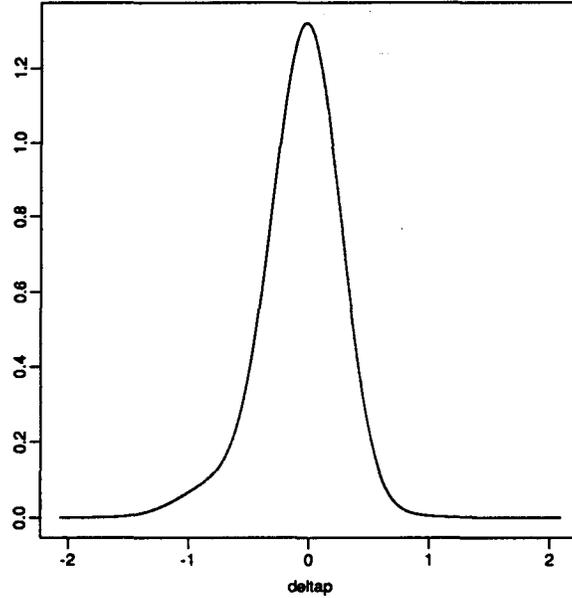


Figure 6: Marginal of the Bivariate Conditional Density

conditioning set are put to their unconditional means, which is computed as

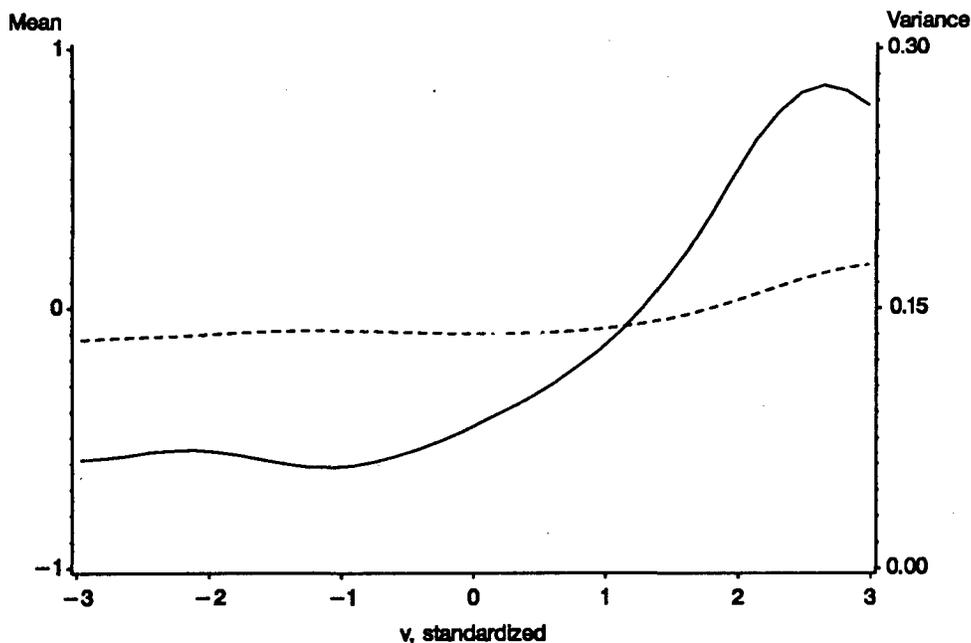
$$\begin{aligned}
 & f_{\Delta p}(\Delta p_t | \Delta p_{t-1:16} = \text{mean}, v_{t-1:16} = \text{mean}) \\
 & = \int f(\Delta p_t, v_t | \Delta p_{t-1:16} = \text{mean}, v_{t-1:16} = \text{mean}) dv_t
 \end{aligned}$$

The density is slightly skewed to the left. It assumes the classic shape for financial data — peaked near zero and thick in the extreme tails relative to the Gaussian density. The excess kurtosis for this density is 4.14, versus 11.22 for the unconditional kurtosis of the $\{\Delta p_t\}$ series. Thus, conditioning on past prices and volume removes much, but not all, of the excess kurtosis.

4.2 Contemporaneous Conditional Price/Volume Relationships

The contemporaneous relationships between price movements and volume are revealed by looking at the conditional mean and variances of Δp_t given v_t along slices of the bivariate $(\Delta p_t, v_t)$ density shown in Figure 5.

Figure 7 shows the first two moments of Δp_t conditional on v_t , with all lagged values of Δp_t and v_t set to their unconditional means. These are the mean and variance of the univariate density obtained by slicing the bivariate density shown in Figure 5 along a line through $(0, v_t)$ parallel to the Δp_t axis. The horizontal axis is in standardized units, $(v_t -$



v is standardized with $\mu = 8.35$, $\sigma = 0.124$
 Dashed: Conditional Mean Solid: Conditional Variance

Figure 7: Bivariate SNP Estimate of the Conditional Moments of Δp_t Given v_t

μ)/ σ , where μ and σ are the moments of the marginal, conditional density of v_t , $f_v(v_t | \Delta p_{t-1:16} = \text{mean}, v_{t-1:16} = \text{mean})$. The range of the horizontal axis extends for three standard deviations on either side of μ . Outside that range, the moment functions become oscillatory. We determined that the estimated moments of Δp_t given such large v_t were adversely affected by the roughness seen in the extreme tails in Figure 5 and were therefore unreliable.

Interestingly, Figure 7 shows that the direction of the daily change in the stock market is unrelated to contemporaneous volume. The market is as likely to fall or rise on heavy volume as it is on light volume, at least over the range of the data in which we can reliably estimate the contemporaneous conditional moment functions.

On the other hand, volatility is related to contemporaneous volume. Days with high volume are associated with high price volatility. The contemporaneous conditional variance function for Δp_t shown in Figure 7 is a nonparametric conditional analogue of the function shown in Figure 1 of Tauchen and Pitts (1983, p. 502) for Treasury Bill futures. Their estimate was obtained from a fitted lognormal-normal parametric mixing model which did not take account of conditional heteroskedasticity. Still, the Tauchen-Pitts plot possesses the same convex shape as the variance function in Figure 7.

4.3 The Conditional Moment Structure of Δp_t

We now examine those features of the density related to the conditional mean and variance properties of Δp_t given past Δp_t . We are mainly interested in the symmetry of the conditional variance function, as this relates to the leverage effect discussed in the introduction. In view of our previous findings regarding the contemporaneous volume-volatility relationship, we are also interested in understanding how lagged volume modifies the conditional variance function.

The upper panel of Figure 8 shows the conditional mean and variance of Δp_t as a function of Δp_{t-1} in standardized units. Specifically, the figure shows

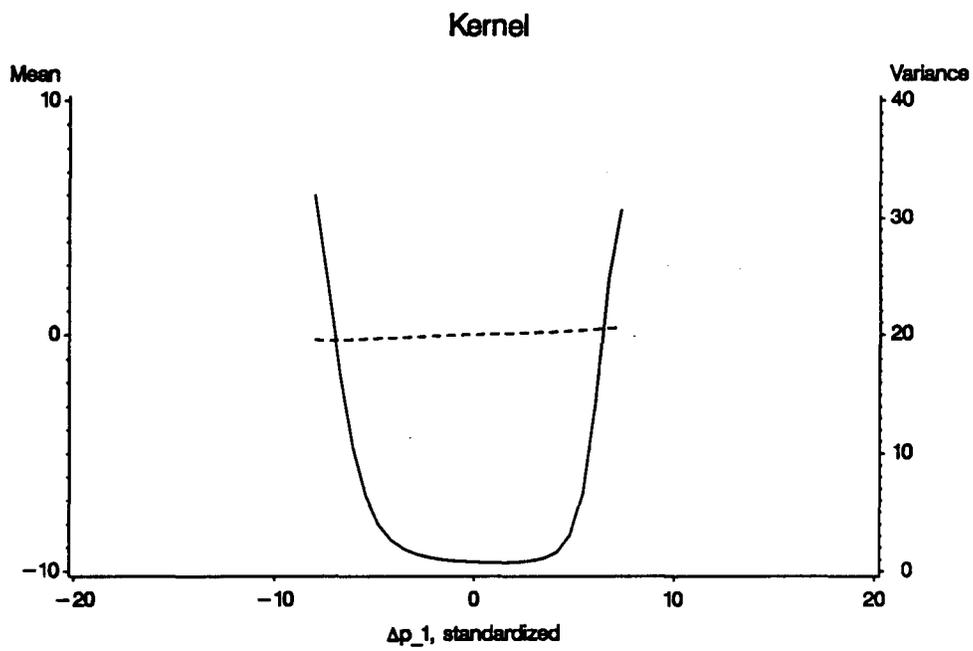
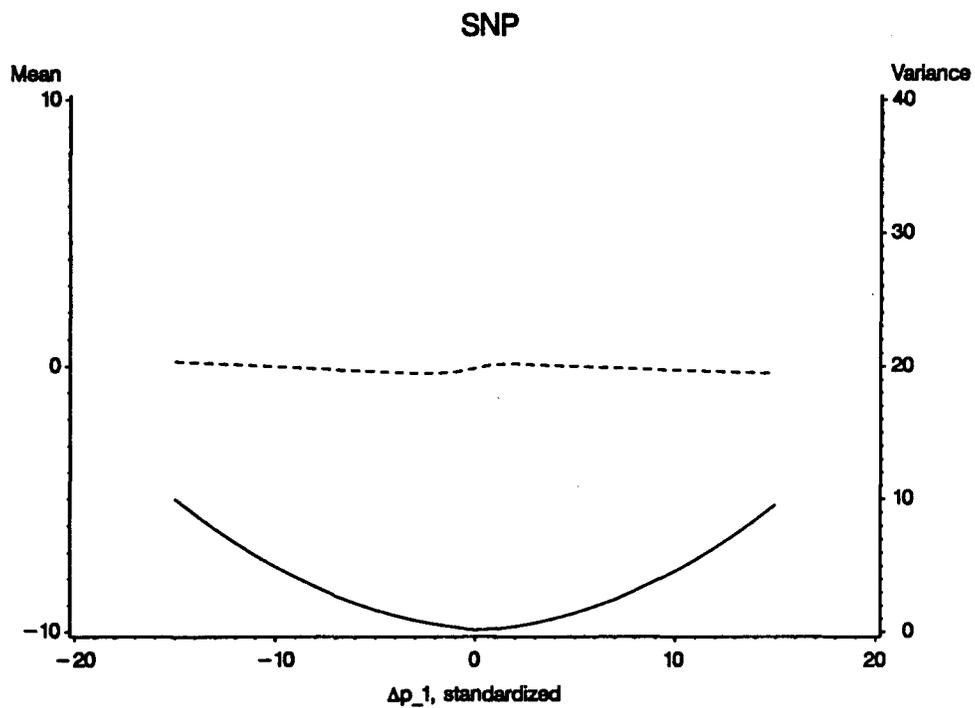
$$\mathcal{E}(\Delta p_t | \Delta p_{t-1} = \mu_{\Delta p} + u\sigma_{\Delta p}, \Delta p_{t-2:16} = \text{mean}, v_{t-1:16} = \text{mean})$$

$$\text{Var}(\Delta p_t | \Delta p_{t-1} = \mu_{\Delta p} + u\sigma_{\Delta p}, \Delta p_{t-2:16} = \text{mean}, v_{t-1:16} = \text{mean})$$

where u varies between ± 15 unconditional standard deviations of Δp . Since $\mu_{\Delta p} \doteq 0$ and $\sigma_{\Delta p} \doteq 1.15$, the range along the horizontal axis corresponds to movements in Δp_{t-1} over a band slightly wider than -15 percent to 15 percent.

A linear analysis (VAR model) reveals only a very modest amount of predictability in the price change series. The first order autocorrelation of the adjusted $\{\Delta p_t\}$ series is .129; the unadjusted series has first order autocorrelation of .065. This low level of autocorrelation is to be expected in a value-weighted index such as the S&P composite index. For example, the weekly return on the CRSP value-weighted index computed by McCulloch and Rossi (1989) has a first order autocorrelation of .089 over the period from 1963 to 1987; we might expect the autocorrelation coefficient of the S&P index to be higher due to the thinner trading in the period 1928 to 1964. (See Lo and MacKinlay, 1988, for a discussion of the effects of non-synchronous trading.) The conditional mean function displayed in the upper panel of Figure 8 is an essentially constant function of Δp_{t-1} . This indicates that no further nonlinear predictability is detected in the fitted SNP model.

The conditional variance function shown in the upper panel of Figure 8 clearly displays the sort of conditional heteroskedasticity found in many ARCH applications. Even though the fitted SNP model does not impose symmetry on the conditional variance function as the traditional ARCH and GARCH models do, the estimated function is symmetric. In order to



Δp_1 is standardized with $\mu = 0.016$, $\sigma = 1.15$
 Dashed: Conditional Mean Solid: Conditional Variance

Figure 8: Bivariate Estimate of the Conditional Moments of Δp_t Given Δp_{t-1}

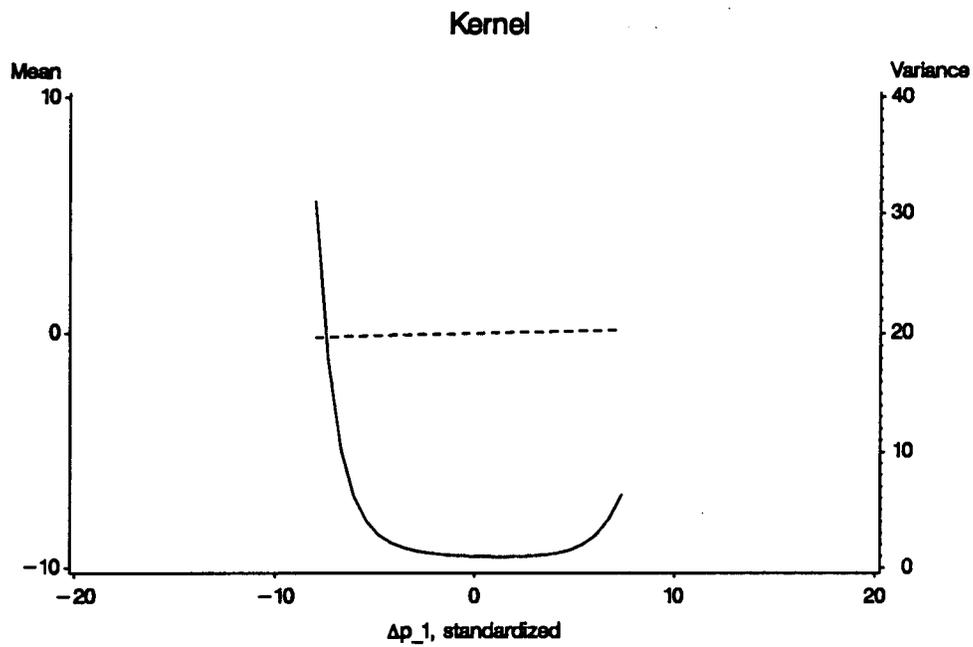
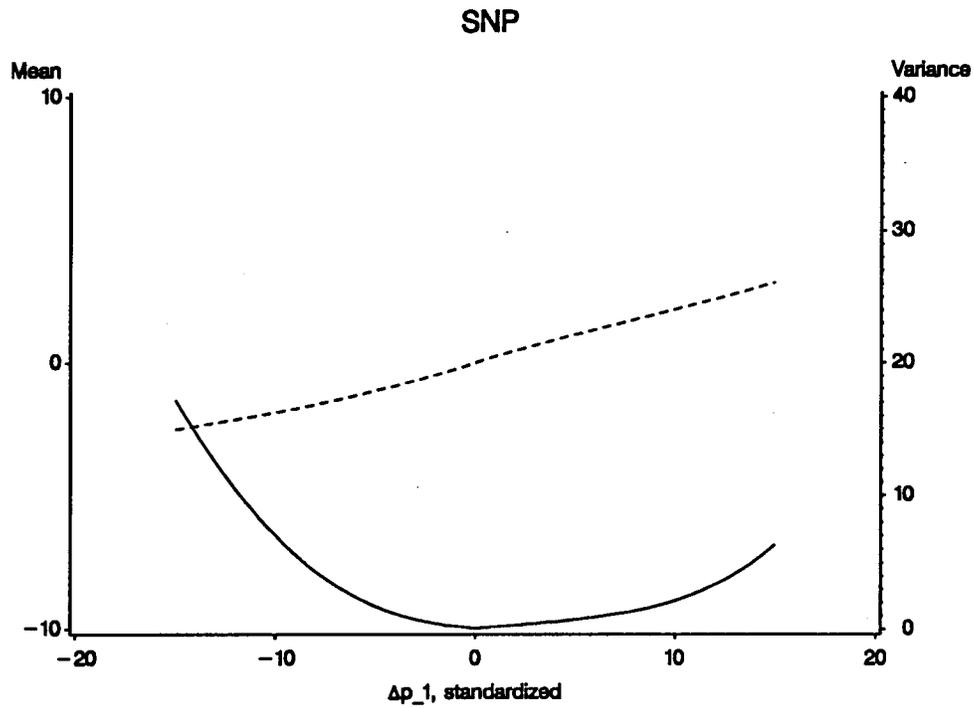
insure that these findings are indeed representative of the data and not an artifact of the SNP approach, we used kernel methods to estimate the conditional variance and mean functions. (See Section 3 above for a discussion of kernel methods.) The kernel-based functions are displayed in the lower panel of Figure 8. The kernel-based estimates give an independent confirmation of the SNP results. The kernel estimates are not shown for Δp_{t-1} exceeding five standard deviations. Beyond that range, the data are exceedingly sparse and the method of local averaging used by the kernel estimator gives estimates of the conditional moment functions that are so variable as to be of little use.

The evidence on symmetry of the conditional variance function is interesting in view of the findings of Nelson (1989a,b), Pagan and Schwert (1989), and others who find evidence of asymmetry in the conditional variance function.

This asymmetry in the variance function has been dubbed the “leverage” effect after early work by Black (1976) and Christie (1982) in which changes in the equity value of a firm affect the riskiness of the firms equity. However, tests of the leverage hypothesis by French, Schwert, and Stambaugh (1987) and by Schwert (1989 and 1990b) suggest that financial leverage could not be responsible for asymmetries of the magnitude found by Nelson (1989a). Nonetheless, common parlance is “leverage effect” in reference to the asymmetry.

The chief difference between our estimation and that of these other papers is that we model a joint price and volume process, while the other studies examine a marginal price process. This difference suggests that introducing volume into the analysis is responsible for producing the symmetry seen in Figure 8.

We can confirm this conjecture. When we fit the univariate price change process $\{\Delta p_t\}$ alone, we also uncover evidence of asymmetry. The fact that we can reproduce the findings of others using only the price data is seen in Figure 9. The upper panel of Figures 9 shows the conditional mean, $\mathcal{E}(\Delta p_t \mid \Delta p_{t-1} = \mu_{\Delta p} + u\sigma_{\Delta p}, \Delta p_{t-2:16} = mean)$, and variance, $Var(\Delta p_t \mid \Delta p_{t-1} = \mu_{\Delta p} + u\sigma_{\Delta p}, \Delta p_{t-2:16} = mean)$, functions of u that are computed from the preferred SNP fit to the univariate price change process $\{\Delta p_t\}$. (This estimation is summarized in the discussion of Table 3.) The conditional variance function is higher on the left than on the right, which is consistent with previous findings on leverage. The lower panel of Figure 9 is the kernel-based counterpart to the upper panel of Figure 9, and likewise provides independent,



Δp_1 is standardized with $\mu = 0.016$, $\sigma = 1.15$
 Dashed: Conditional Mean Solid: Conditional Variance

Figure 9: Univariate Estimate of the Conditional Moments of Δp_t Given Δp_{t-1}

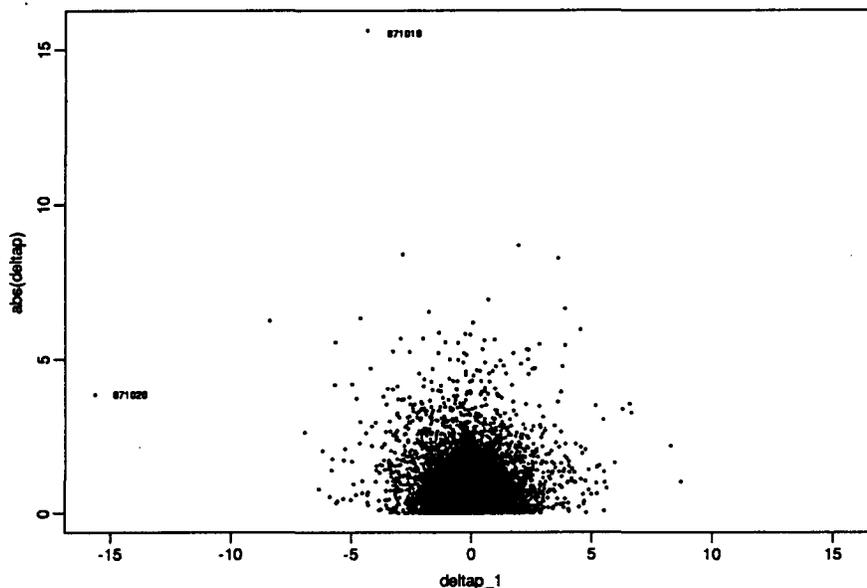


Figure 10: Volatility Scatterplot

corroborative evidence on the asymmetry obtained with only price data.

Figure 10 helps reconcile this disparate evidence on the characteristics of the conditional variance function. Figure 10 is a scatter plot of $|\Delta p_t|$ versus Δp_{t-1} , which is the cloud of points that the various models are attempting to fit. Overall, the cloud appears to be asymmetric and shows a leverage effect, in the sense of being oriented towards the northwest instead of towards the vertical. This visual interpretation, though, also appears to be heavily influenced by a few extreme events, as the central part of the cloud appears more symmetric.

Figure 10 suggests that introducing volume into the analysis mean has the effect of downweighting extreme observations. The SNP and kernel estimates of the conditional variance function show that in the center of the data cloud the conditional variance function is symmetric. It should be emphasized that neither the SNP nor the kernel variance function is constrained to be symmetric, as is the standard GARCH or ARCH model. Evidently, observations corresponding to large price movements appear more extreme when considered relative to the marginal distribution of price alone than when considered relative to the bivariate price/volume distribution. The extent to which large price movements appear less extreme in the bivariate data set is perhaps apparent in Figure 1 above. Thus when fitting the univariate data alone, the estimation procedures attempt to accommodate these more

extreme outliers by introducing asymmetry into the variance function estimate. This view of the estimated asymmetry as a by-product of a small number of outlying observations is bolstered by the fact that the asymmetry disappears when industry or size portfolios are used (see Braun, Nelson and Sunier, 1990) or when additional variables are introduced (see Glosten, Jagannathan and Runkle, 1989).

4.4 Dynamic Price/Volume Relationships

To describe fully the dynamic relationship between volume and price changes, it is important to assess how much lagged volume contributes to the prediction of the future distribution of price changes. We performed a series of regressions in which lagged volume variables were added to the one step-ahead predictors of the conditional mean and variance of the price series. The one-step ahead predictors used are the conditional mean and variance of the SNP estimate of the univariate, conditional density. Thus, these predictors depend only on past prices and not on volume. In the first equation, we regressed Δp_t on $\mathcal{E}(\Delta p_t | \Delta p_{t-1:16})$ to establish a baseline measure of goodness of fit and then added 20 lags of v_t , v_t^2 and v_t^3 . Likewise, In the second equation, we regressed the squared univariate residual, $[\Delta p_t - \mathcal{E}(\Delta p_t | \Delta p_{t-1:16})]^2$, on the conditional variance computed from the univariate fit and then added the functions of lagged volume series. These regressions and associated statistics are shown below:

	R^2	F test on inclusion of volume variables
<i>Predictability in the Mean:</i>		
p_t on $\mathcal{E}(\Delta p_t \Delta p_{t-1:16})$.0221	F(60,16038)=1.936
With volume variables included	.0291	$p\text{-val} = 10^{-5}$
<i>Predictability in the Second Moment:</i>		
$[\Delta p_t - \mathcal{E}(\Delta p_t \Delta p_{t-1:16})]^2$ on $Var(\Delta p_t \Delta p_{t-1:16})$.1034	F(60,16038)=2.273
With volume variables included	.1109	$p\text{-val} = 10^{-8}$

The F statistic reported is the joint test for the inclusion of the 60 volume variables. It should be noted that the F statistic is not adjusted for the conditional heteroskedasticity that is undoubtedly present in the residuals from these regressions. The F is only suggestive, and

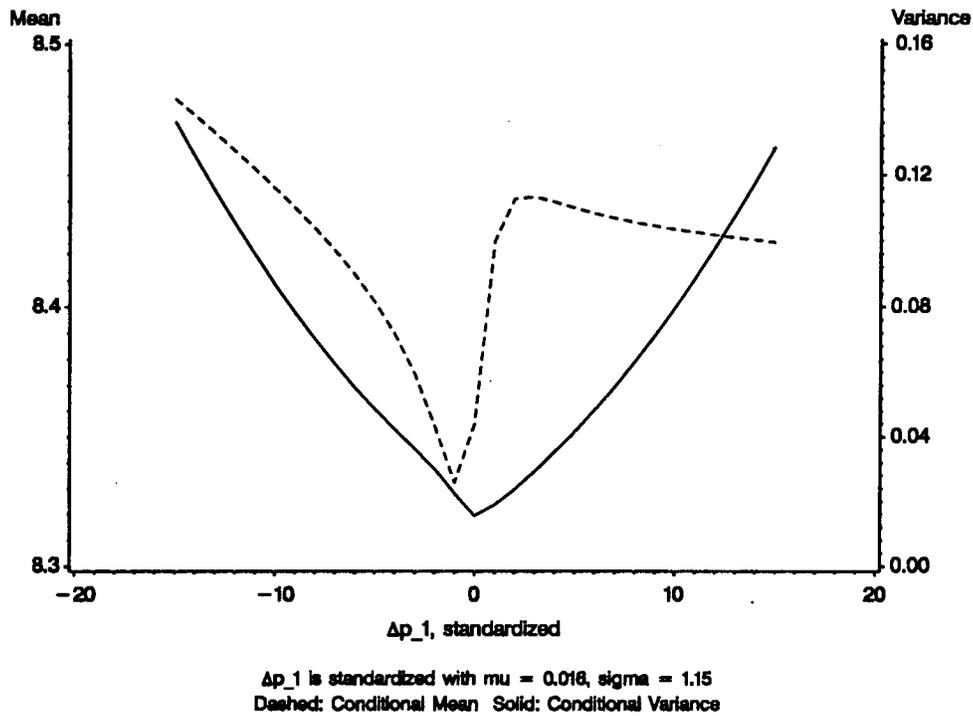


Figure 11: Bivariate SNP Estimate of the Conditional Moments of v_t Given Δp_{t-1}

the best strategy is to view the R^2 s as estimates of the corresponding population measures of fit.

Three conclusions emerge from these regressions: 1) As to be expected, there is a great deal of noise in daily data; 2) Volatility is surprisingly predictable; and 3) Volume contributes slightly to the prediction of both the mean and the variance of the price change series. Item 3) is consistent with Schwert (1989), who finds a significant incremental effect from adding volume to a linear VAR model for volatility.

We examine further the effect that price volatility has on the volume in Figure 11. The figure displays the conditional mean and variance functions of v_t as a function of Δp_{t-1} , in standardized units. The figure shows

$$\mathcal{E}(v_t | \Delta p_{t-1} = \mu_{\Delta p} + u\sigma_{\Delta p}, \Delta p_{t-2:16} = \text{mean}, v_{t-1:16} = \text{mean})$$

$$\text{Var}(v_t | \Delta p_{t-1} = \mu_{\Delta p} + u\sigma_{\Delta p}, \Delta p_{t-2:16} = \text{mean}, v_{t-1:16} = \text{mean})$$

as u varies between -15 and 15 standard deviations. The figure suggests that large price movements lead to increases in both the mean and variability of the volume. Both functions are fairly symmetric, indicating market declines have the same effect on subsequent volume as market increases.

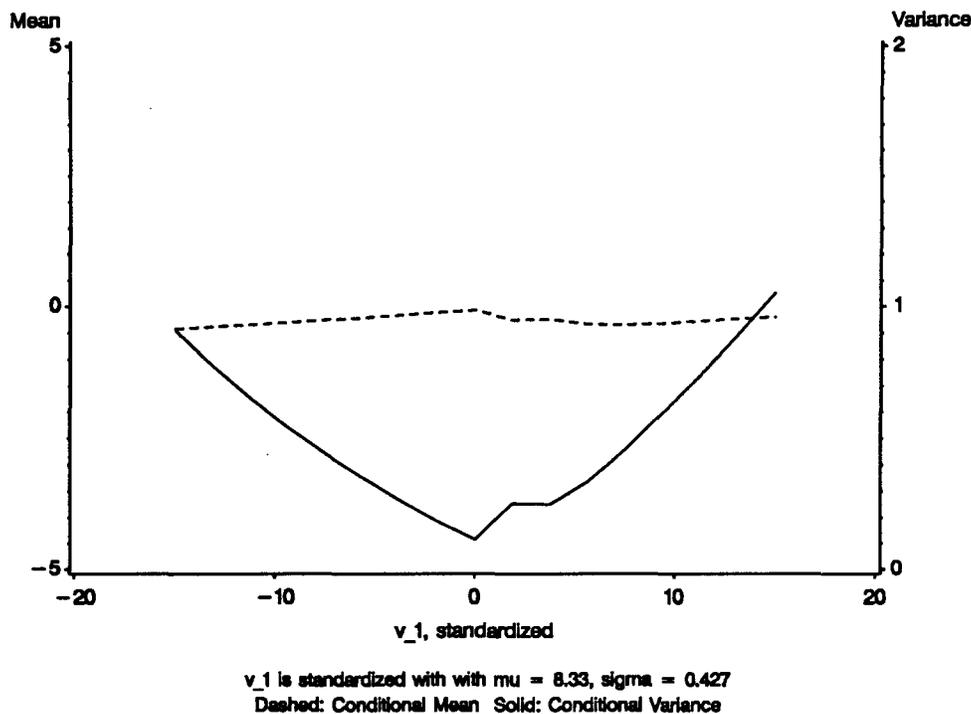


Figure 12: Bivariate SNP Estimate of the Conditional Moments of Δp_t Given v_{t-1}

Figure 12 shows the effect that lagged volume has on current price changes. Both abnormally high and low volumes are associated with increased future price volatility, while the conditional mean of price change is constant across a very wide range of lagged volume levels.

4.5 The Risk Premium and Conditional Price Volatility

The final feature of the density we examine is the relationship between the conditional mean and variance of Δp_t . Motivating this effort is recent empirical work aimed at measuring the relationship between risk premiums on financial assets and the conditional second moments of returns. Bollerslev, Engle, and Wooldridge (1988), French, Schwert, and Stambaugh (1987), and Nelson (1989a), use ARCH-in-mean specifications to relate risk premiums to conditional second moments. Much of this effort is directed towards measurement of a hypothesized monotone increasing relationship between the risk premium on the market return and its own conditional variance.

The existence of such a relationship, though, has been the subject of debate on both empirical grounds (Pagan and Hong, 1989) and theoretical grounds (Backus and Gregory, 1988). This debate is perhaps not surprising given that, in general, equilibrium asset pricing models relate the conditional means of asset returns to generalized notions of a marginal

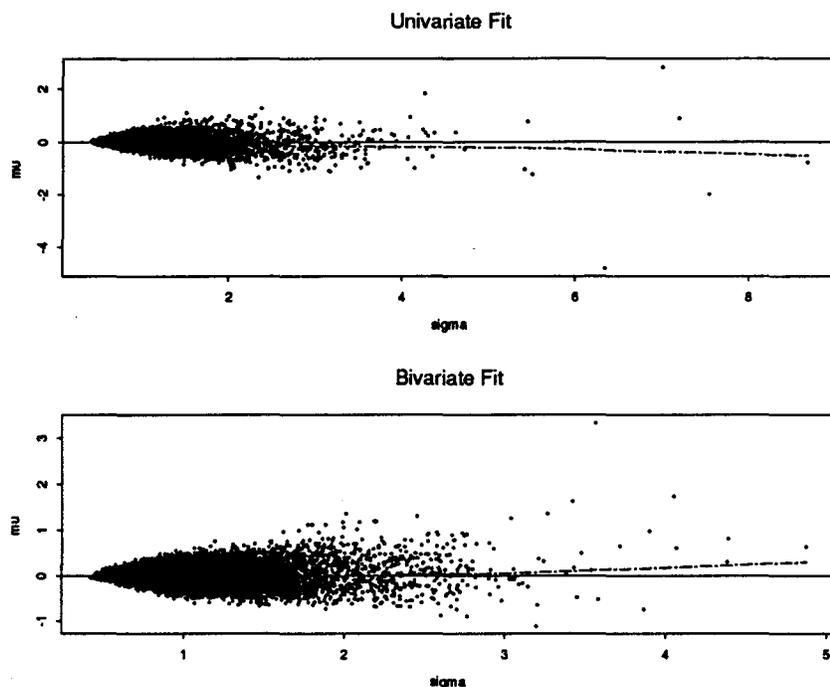


Figure 13: Conditional Mean vs. Standard Deviation

rate of substitution (Hansen and Jagannathan, 1989), and not directly to their own internal second moment structure. Under special assumptions (Merton, 1973), there will be a direct link between the risk premium and the conditional variance. Backus and Gregory (1988) and Tauchen and Hussey (1990) study the reduced form relationships between risk premium and conditional variance that emerge from more general asset pricing models. They find that under the familiar CRR (power) utility function, the direction of the relationship between the risk premium and conditional variance can go either way, as it is sensitive to assumptions regarding the stochastic properties of the consumption endowment. At the same time, Tauchen and Hussey (1990) find that the relationship is monotone and increasing when the law of motion for the consumption endowment is calibrated in a realistic manner from actual time series consumption data.

This discussion makes clear that the characteristics of the relationship between the risk premium and conditional variance have not been fully determined, either theoretically or empirically. More evidence is clearly warranted.

Figure 13 summarizes the available evidence from our estimated conditional densities. The figure consists of two scatterplots of the pairs (conditional mean of price change, conditional standard deviation) computed from the fitted conditional densities, each evaluated at

every sample point. The top panel of Figure 13 presents the scatterplot for the conditional moments derived from the univariate fit while the bottom panel presents the scatterplot for the moments from the bivariate fit. The dashed curve is a smoothed estimate of the regression function obtained using Cleveland's (1979) locally-weighted, robust, regression procedure. For the univariate fit, the curve shows a slight downward slope. However, the curve is not monotone decreasing and appears to be influenced by several points in the four to six sigma range. The negative slope is similar to the findings of Pagan and Hong (1989) and Nelson (1989a). In contrast to the results from the univariate fit, the bivariate fit shows a monotone increasing relationship between the conditional standard deviation and the mean. This finding is consistent with the finding of French, Schwert, and Stambaugh (1987) regarding the relationship between predictable volatility and the conditional mean.

In order to assess the extent to which these fitted risk premium functions are influenced by outlying points, we identified all of the points with very large sigma values. Most of these points are associated with the crash of October 1987. When we divide the data into three equal subperiods, the monotone increasing risk premium result for the bivariate fit holds up in each of the subperiods. In the middle, "quiet" period of the data, the difference between the univariate and bivariate fits is most apparent.

All told, our findings suggest that after conditioning on lagged volume, there is a positive relationship between the risk premium and the conditional variance of the return. The fact that the positive relationship holds up over each of the three subperiods indicates that it is robust finding and is not the result of a few extreme observations.

Two caveats are in order regarding this discussion of the risk premium. First, the mean price change is not expressed as a return in excess of the return on a risk-free asset, which is correct theoretically and also nets out inflation. Although using excess returns will not change the shape of the risk premium function, it may alter the magnitude of the premium. Second, we should note that our measure of return is the nominal, daily, percentage, capital gain on the Standard and Poor's Composite Index. Thus, it excludes the dividend component of the total return. The data required to make these adjustments on a daily basis are unavailable over the long time period of our sample. Still, one can plausibly argue that, on a daily basis, the capital gain component dominates other components. Nelson (1989a)

presents some empirical evidence on the extent to which the capital gain overwhelmingly dominates. Thus we believe that our findings from Figure 13 are robust with respect to these adjustments, if they could be made.

5 Summary and Conclusion

The main objective has been to investigate the characteristics of price and volume movements on the stock market. Motivating this effort were the recent events on the stock market together with a desire to provide a comprehensive set of empirical regularities that economic models of financial trading will ultimately need to confront. We organized the effort around the tasks of estimating and interpreting the conditional one-step-ahead density of joint price change and volume process. For a stationary process, the one-step density is a time invariant population statistic that subsumes all probabilistic information about the process. In particular, issues concerning predictability, volatility, and other conditional moment relationships can be addressed by examining the conditional density. Indeed, such issues seem more naturally thought of in terms of features of the population conditional density, and not in terms of the signs and magnitudes of specific parameters.

The raw S&P price change and NYSE aggregate volume data display systematic calendar and trend effects in both mean and variance, and thus are not stationary. Prior to estimation, we undertook a fairly extensive effort to remove these systematic effects. This effort resulted in series on adjusted logarithmic price changes and adjusted log volume which appear to be reasonably modeled as jointly stationary. All subsequent statements concerning the price changes and volume pertain to these adjusted series.

The SNP estimation technique entails fitting a series expansion to the bivariate condition density. The leading term of the expansion is a VAR model with an ARCH-like error process; higher order terms accommodate departures from that model. There is substantial evidence that the higher order terms are needed to capture all of the complex structure of the data. These complexities include, among other things, the complicated structure of the bivariate conditional variance function, the thick-tailed error density characteristic of financial price movements data, the nonlinear interactions between volume and prices, and the temporal dependence of the volume series.

Examination of the fitted conditional density reveals four major findings regarding the interactions between stock prices and volume:

(1) Contemporaneous Volume-Volatility Correlation

The daily trading volume is positively and nonlinearly related to the magnitude of the daily price change. This association is a characteristic of both the unconditional distribution of price changes and volume and the conditional distribution given past price changes and volume constant.

The finding of an unconditional volume-volatility relationship is consistent with many other studies (see Tauchen and Pitts, 1983; Karpoff, 1987), though it was obtained with a rather different data set. We use a very long time series on changes in a market-wide index and overall volume, while other studies almost exclusively examine price changes and volume for individual financial assets.

The finding of a conditional volume-volatility relationship is more interesting. It means that the volume-volatility association is still observable after taking account of non-normalities, stochastic volatility, and other forms of conditional heterogeneity. The finding extends other recent work on volume and volatility. Using daily individual security data, 1981-1983, Lamoureux and Lastrapes (1989) find a positive conditional volume-volatility relationship in models with Gaussian errors and GARCH-type volatility specifications. Using monthly measures, 1885-1987, Schwert (1989) finds a positive relationship in linear Koyck distributed lag regression of estimated volatility on current and lagged volume growth.

(2) Large Price Movements Associated with Higher Subsequent Volume

Price changes lead to volume movements. The effect is fairly symmetric, with large price declines having nearly the same impact on subsequent volume as large price increases.

(3) Attenuated Leverage Effect after Conditioning on Lagged Volume

For bivariate price-volume estimations, the conditional variance of the price change is essentially a symmetric function of the lagged price change. This finding stands in contrast to those of Nelson (1989a) and Pagan and Schwert (1989). Using univariate price data, they report evidence for an asymmetric conditional variance function, with price decreases being

associated with higher levels of subsequent volatility than price increases. This asymmetry is called the leverage effect. Using price data alone, we also find a leverage effect.

The reason we do not observe a leverage effect in the bivariate estimation, but we do in the univariate estimation, is that introducing volume has a robustifying effect. In essence, bringing volume into the analysis downweights a few extreme observations that appear to be mainly responsible for the findings of a leverage effect in univariate data.

(4) Positive Conditional Risk/Return Relation after Conditioning on Lagged Volume

For bivariate price-volume estimation, there is evidence for a positive association between the conditional mean and the conditional variance of daily stock returns. The finding is useful in view of the fact that equilibrium asset pricing theory is silent on the manner in which the conditional first two moments of the market return co-vary in response to shocks to the economy. As we discussed above, in some special models the conditional mean and variance are positively related, which is consistent with the intuitive notion that stocks should command a higher return in periods of high volatility. In general, however, the direction of the relationship is indeterminate, as it is sensitive to the specification of the dynamics of the consumption endowment (Backus and Gregory, 1988; Tauchen and Hussey, 1990).

The finding of a positive conditional mean-variance relationship is also interesting in view of other empirical work on this issue. As we note above, some studies using univariate price data find a negative relationship between the conditional mean and variance (Pagan and Hong, 1989; Nelson, 1989a). On the other hand, French, Schwert, and Stambaugh (1987) find evidence for a positive relationship between the risk premium and predictable volatility. Using conditional moments from our univariate estimation, we find a negative relationship. But we also find evidence that lagged volume contains some additional predictive power (over and above that embodied in past prices) for both the mean and the variance of the price change. With volume incorporated into the analysis, we find a positive relationship between the conditional mean and variance.

In closing, we note that there are models that can account for various subsets of the known characteristics of financial prices and volume, but no single model seems capable of explaining all of them jointly. For instance, familiar representative agent asset pricing models

can produce persistent volatility and leptokurtic price change densities, but are silent on the contemporaneous relationship between price and volume as well as price/volume dynamics. For the effect of volume on the risk/return relation, representative agent models suggest that volume should be incorporated into the information sets but fail to explain why volume might affect volatility. On the other hand, models based on random mixing (Tauchen and Pitts, 1983; Harris, 1986) can exhibit persistent volatility and accommodate the observed contemporaneous and dynamic price/volume relationship, but have no direct bearing on the attenuation of leverage effects or the risk/return relation. Furthermore, as we noted in the introduction, these mixing models are closer to being statistical models than economic models. An interesting theoretical challenge is to develop a complete equilibrium model that can jointly account for all of the above-discussed characteristics.

6 References

- Admati, A. and P. Pfleiderer (1988), "A Theory of Intraday Patterns: Volume and Price Variability," *The Review of Financial Studies* 1, 3–40.
- Admati, A. and P. Pfleiderer (1989), "Divide and Conquer: A Theory of Intraday and Day-of-the-Week Mean Effects," *The Review of Financial Studies* 2, 189–223.
- Ariel, R. (1988), "Evidence on Intra-month Seasonality in Stock Returns," in E. Dimson (ed.), *Stock Market Anomalies*. Cambridge: Cambridge University Press.
- Backus, D. K. and A. W. Gregory (1988), "Theoretical Relations Between Risk Premiums and Conditional Variances," manuscript, Federal Reserve Bank of Minneapolis.
- Black, F. (1976), "Studies of Stock Market Volatility Changes," *Proceedings of the American Statistical Association, Business and Economics Section*, 177–181.
- Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics* 31, 307–327.
- Bollerslev, T. (1987), "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *The Review of Economics and Statistics* 64, 542–547.

- Bollerslev, T., R. Engle, and J. Wooldridge (1988), "A Capital Asset Pricing Model with Time Varying Covariances," *Journal of Political Economy* 96, 116–131.
- Braun P., D. Nelson, and A. Sunier, (1990), "Good News, Bad News, Volatility and Betas," manuscript, Graduate School of Business, University of Chicago.
- Christie, A. (1982), "The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects," *Journal of Financial Economics* 10, 407–432.
- Cleveland, W. S. (1979), "Robust Locally Weighted Regression and Smoothing Scatterplots," *Journal of the American Statistical Association* 74, 829–836.
- Constantinides, G. (1984), "Optimal Stock Trading with Personal Taxes," *Journal of Financial Economics*, 65–90.
- Davidian, M. and R. J. Carroll (1987), "Variance Function Estimation," *Journal of the American Statistical Association* 82, 1079–1091.
- Diebold, Francis X. (1987), *Empirical Modeling of Exchange Rates*. Berlin: Springer-Verlag.
- Duffie, Darrell and Kenneth J. Singleton (1989), "Simulated Moments Estimation of Markov Models of Asset Prices," manuscript, Stanford University.
- Engle, R. F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica* 50, 987–1007.
- Engle, R. and T. Bollerslev (1986), "Modeling the Persistence of Conditional Variance," *Econometric Reviews* 5, 1–50.
- Epps, T. and M. Epps (1976), "The Stochastic Dependence of Security Price Changes and Transaction Volumes, Implications for the Mixture of Distributions Hypothesis," *Econometrica* 44, 305–321.
- French, K. (1980), "Stock Returns and the Weekend Effect," *Journal of Financial Economics* 8, 55–69.

- French, K. and R. Roll (1986), "Stock Return Variances: The Arrival of Information and the Reaction of Traders," *Journal of Financial Economics* 17, 5–26.
- French, K., G. W. Schwert, and R. Stambaugh (1987), "Expected Stock Returns and Volatility," *Journal of Financial Economics* 19, 3–30.
- Foster, D. F. and S. Viswanathan (1990), "Variations on Volumes, Variances, and Trading Costs," manuscript, Fuqua School of Business, Duke University.
- Gallant, A. R. (1981), "On the Bias in Flexible Functional Forms and an Essentially Unbiased Form: The Fourier Flexible Form," *Journal of Econometrics* 15, 211–244.
- Gallant, A. Ronald, Lars Peter Hansen, and George Tauchen (1990), "Using Conditional Moments of Asset Payoffs to Infer the Volatility of Intertemporal Marginal Rates of Substitution," *Journal of Econometrics* 45, 141–180.
- Gallant, A. Ronald, David A. Hsieh, and George E. Tauchen (1989), "On Fitting a Recalcitrant Series: The Pound/Dollar Exchange Rate, 1974–83," in William A. Barnett, James Powell, George E. Tauchen, eds. *Nonparametric and Semiparametric Methods in Econometrics and Statistics, Proceedings of the Fifth International Symposium in Economic Theory and Econometrics*. Cambridge: Cambridge University Press, forthcoming.
- Gallant, A. Ronald and Douglas W. Nychka (1987) "Seminonparametric Maximum Likelihood Estimation" *Econometrica* 55, 363–390.
- Gallant, A. Ronald and George Tauchen (1989), "Seminonparametric Estimation of Conditionally Constrained Heterogeneous Processes: Asset Pricing Applications," *Econometrica* 57, 1091–1120.
- Gallant, A. Ronald and George Tauchen (1990), "A Nonparametric Approach to Nonlinear Time Series Analysis: Estimation and Simulation," manuscript, North Carolina State University.

- Glosten L., R. Jaganannathan, and D. Runkel (1989), "Relationship Between the Expected Value Volatility of the Nominal Excess Return on Stocks," Banking Research Paper #166, Northwestern University.
- Harris, Lawrence (1986), "Cross-Security Tests of the Mixture of Distributions Hypothesis," *Journal of Financial and Quantitative Analysis* 21, 39–46.
- Harvey, Campbell (1989), "Is the Expected Compensation for Market Volatility Constant Through Time?" manuscript, Fuqua School of Business, Duke University.
- Harvey, Campbell and Robert E. Whaley (1989), "The Impact of Discrete Cash Dividends on the Valuation of S&P 100 Index Options," manuscript, Fuqua School of Business, Duke University.
- Hansen, Lars Peter and Ravi Jagannathan (1989), "Restrictions on Intertemporal Marginal Rates of Substitution Implied by Asset Returns," manuscript, Northwestern University.
- Huffman, Gregory W. (1987), "A Dynamic Equilibrium Model of Asset Prices and Transactions Volume," *Journal of Political Economy* 95, 138–159.
- Huffman, Gregory W. (1988), "A Representative Agent Model of Asset Pricing and Transaction Volume," manuscript, Federal Reserve Bank of Minneapolis.
- Karpoff, J. (1987), "The Relation between Price Changes and Trading Volume: A Survey," *Journal of Financial and Quantitative Analysis* 22, 109–126.
- Keim, D. (1983), "Size-Related Anomalies and Stock Return Seasonality: Further Statistical Evidence," *Journal of Financial Economics* 12, 13–32.
- Ketterer, Joan A. and Albert Marcet (1989), "Introduction of Derivative Securities: A General Equilibrium Approach," manuscript, Carnegie-Mellon University.
- Lamoureux, C. and W. Lastrapes (1989), "Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects," manuscript, Louisiana State University.
- Lo, A. and A. MacKinlay (1988), "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *The Review of Financial Studies* 1, 41–66.

- McCulloch, R. and P. Rossi (1989), "Predictive, Posterior and Utility-based Approaches to Testing the Arbitrage Pricing Theory," CRSP manuscript #269, University of Chicago.
- Merton, R. C. (1973), "A Intertemporal Capital Asset Pricing Model," *Econometrica* 41, 867-888.
- Mulherin, J. H. and M. Gerety (1988), "Trading Volume on the NYSE During the Twentieth Century: A Daily and Hourly Analysis," manuscript, Securities and Exchange Commission.
- Nelson, D. (1989a), "Conditional Heteroskedasticity in Asset Returns: A New Approach," manuscript, University of Chicago.
- Nelson, D. (1989b), "Modeling Stock Market Volatility Changes," manuscript, University of Chicago.
- Newey, W. K. (1985), "Maximum Likelihood Specification Testing and Conditional Moment Tests," *Econometrica* 53, 1047-1071.
- Officer, R. R. (1973), "The Variability of the Market Factor of New York Stock Exchange," *Journal of Business* 46, 434-453.
- Pagan, A. and Y. Hong (1989), "Nonparametric Estimation and the Risk Premium," in William A. Barnett, James Powell, George E. Tauchen, eds. *Nonparametric and Semiparametric Methods in Econometrics and Statistics, Proceedings of the Fifth International Symposium in Economic Theory and Econometrics*. Cambridge: Cambridge University Press, forthcoming.
- Pagan, A. and G. W. Schwert (1989), "Alternative Models for Conditional Stock Volatility," manuscript, University of Rochester.
- Potscher, B. M. (1989), "Model Selection under Nonstationarity: Autoregressive Models and Stochastic Linear Regression Models," *The Annals of Statistics* 17, 1257-1274.
- Robinson, P. M. (1983), "Nonparametric Estimators for Time Series," *Journal of Time Series Analysis* 4, 185-207.

- Rozeff, M. and W. Kinney (1976), "Capital Market Seasonality: The Case of Stock Returns," *Journal of Financial Economics* 3, 379–402.
- Schwarz, G. (1978), "Estimating the Dimension of a Model," *Annals of Statistics* 6, 461–464.
- Schwert, G. William (1989), "Why Does Stock Volatility Change Over Time?" *Journal of Finance* 44, 1115–1154.
- Schwert, G. William (1990a), "Indexes of Stock Prices from 1802 to 1987," *Journal of Business* 63, 399–426.
- Schwert, G. William (1990b), "Stock Volatility and the Crash of '87," *Review of Financial Studies* 3, 77–102.
- Standard and Poor's (1988), *Security Price Index Record*. New York: Standard and Poor's Corporation.
- Tauchen, George and Mark Pitts (1983), "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica* 51, 485–505.
- Tauchen, George and Robert Hussey (1990), "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," forthcoming in *Econometrica*.
- Tauchen, George (1985), "Diagnostic Testing and Evaluation of Maximum Likelihood Models," *Journal of Econometrics* 30, 415–443.