

Information, Trading, and Volatility: Evidence from Weather-Sensitive Markets

JEFF FLEMING, CHRIS KIRBY, and BARBARA OSTDIEK*

ABSTRACT

We find that trading- versus nontrading-period variance ratios in weather-sensitive markets are lower than those in the equity market and higher than those in the currency market. The variance ratios are also substantially lower during periods of the year when prices are most sensitive to the weather. Moreover, the comovement of returns and volatilities for related commodities is stronger during the weather-sensitive season, largely due to stronger comovement during nontrading periods. These results are consistent with a strong link between prices and public information flow and cannot be explained by pricing errors or changes in trading activity.

A LARGE SEGMENT OF THE FINANCE LITERATURE investigates the link between information and prices. Theory suggests that prices are a function of public information and order flow (see, e.g., Grossman and Stiglitz (1980), Kyle (1985), and Glosten and Milgrom (1985)). Order flow is driven by both public and private information as well as investor shocks, which may be either rational (e.g., noninformation-based liquidity trades) or irrational (e.g., trades based on noise as described by Black (1986)). Prices can deviate from fundamental value due to market microstructure, liquidity, and hedging effects, and pricing errors can arise from noise trading and the systematic underreaction or overreaction to information.

Much of the empirical literature attempts to discriminate among public information, private information, market frictions, and pricing errors as sources of return volatility. Some researchers argue that observed volatility is much higher than can be explained by fundamentals, by market microstructure or

*Fleming and Ostdiek are at the Jones Graduate School, Rice University; Kirby is at the John E. Walker Department of Economics, Clemson University. Much of this research was completed while Ostdiek was visiting at the London Business School and Kirby was at the University of Texas at Dallas. We thank Gordon Bennett, Steve Koenig, and Dave Ostdiek for providing information on the agricultural commodity markets, Robert Gaudette for providing information on the natural gas markets, Craig Solberg for providing information on the weather-sensitive seasons in these markets, and Matthew Richardson and Robert Whitelaw for sharing their information and understanding of the orange juice market. We thank Tom George, Joel Hasbrouck, George Oldfield, Sergei Sarkissian, Tom Smith, Rob Stambaugh, and an anonymous referee for providing useful comments on a previous draft. We also thank seminar participants at City University, Clemson University, London Business School, Rice University, Southern Methodist University, University of Texas at Dallas, University of Warwick, and the 2004 Western Finance Association meetings (especially Duane Seppi, the discussant) for providing many useful comments. Hubert Marcus, Elise McCutchen, and Parel Patel provided invaluable research assistance.

liquidity effects, or by measurement error. Some argue further that this “excess volatility” must be caused by irrational trading (see, e.g., Shiller (2003) for a review of this evidence). Other researchers, however, point out several empirical challenges to work in this area: (1) information generation is endogenous, which makes it likely that information-based price changes and pricing errors occur at the same time; (2) information flow is unobservable, which presents identification and measurement problems; (3) an agent’s motivation for trading is difficult to ascertain; and (4) multiple dimensions of a market’s structure impact the nature of price discovery and the extent of pricing errors.

One approach to disentangling the sources of volatility is to analyze trading-versus nontrading-period variances across different market structures. It is well known, for example, that trading-period returns in stock markets display higher variance than overnight or weekend returns (see, e.g., Fama (1965), Oldfield and Rogalski (1980), French and Roll (1986), Amihud and Mendelson (1987), and Stoll and Whaley (1990)). Why this occurs is still a matter of debate. Stoll and Whaley (1990) attribute it to greater public information flow during the trading period since this period largely overlaps with the normal business day.¹ French and Roll (1986) argue that the higher trading-period variance is due to greater private information flow because traders are more likely to become informed and act on their information during business (trading) hours.² Still another hypothesis is that the trading process itself generates volatility beyond that attributable to information flow. It is difficult to distinguish among these hypotheses, however, because both trading and information flow are concentrated in one part of the calendar day.

Harvey and Huang (1991) provide additional insights from the currency markets wherein both trading and information flow occur around the clock. They find that the variance during U.S. trading hours is much closer to the variance during nontrading hours than it is for stocks. Moreover, the variance for currencies quoted against the dollar is highest during U.S. trading hours, even though trading activity is greater during London trading hours, and the variance increases around U.S. macroeconomic announcements. These findings suggest that the higher U.S. trading period volatility is related to public information flow. Again, however, we cannot rule out the effect of trading because trading and information flow occur contemporaneously.

Weather-sensitive agricultural and energy markets provide a setting that allows us to further discriminate among the sources of volatility. A key component

¹ Consistent with this view, Jones, Kaul, and Lipson (1994) find that volatility is higher even on days when exchanges (and businesses) are open but no trades occur than when exchanges are closed. George and Hwang (2001) find that the lower overnight variance cannot be explained by the absence of trading and conclude that it must be due to a reduction in public information flow.

² Consistent with this view, Ito and Lin (1992) find lower volatility on the Tokyo Stock Exchange (TSE) when the market is closed for the lunch break and lower volatility around the noon hour in New York when the market is open. Barclay, Litzenberger, and Warner (1990) find higher weekend volatility on the TSE when the exchange is open on Saturdays. Barclay and Hendershott (2003) find that even with after-hours trading most private information price discovery for Nasdaq stocks occurs during the trading day because liquidity is higher.

of the information flow in these markets (weather conditions) evolves 24 hours a day, yet the trading day is short with little after-hours trading. Thus, we have a setting in which a large component of the public information flow occurs outside of trading hours. Under the extreme assumption that weather conditions are the only relevant information in these markets and that this information evolves randomly over the 24-hour day, the trading- and nontrading-period variances should be equal per unit of time. In this case, a higher variance during the trading period would be evidence of excess volatility due to market frictions or pricing errors. More generally, though, if the information flow also includes private information and other public information that is not randomly distributed throughout the day, predictions about the relative variances in the trading and nontrading periods become ambiguous.

This ambiguity can be at least partially resolved by recognizing that the importance of weather conditions in some agricultural and energy markets varies in a seasonal pattern throughout the year. Thus, even if the trading-period variance is normally higher than the nontrading-period variance due to factors unrelated to weather conditions, the difference between the two variances should narrow during the weather-sensitive season. Moreover, if the private information flow also increases during the weather-sensitive season, this would increase the trading-period variance and work against the public information effect. As a result, finding that the difference between the variances narrows during the weather-sensitive season would provide strong indirect evidence of a link between public information flow and volatility. This research design also provides a degree of robustness to many of the concerns that have been raised about variance ratio tests because any biases that affect the tests must vary systematically with the seasons to influence our results.³

Some previous studies attempt to directly model the relation between prices and information in weather-sensitive markets using weather data. Roll (1984) is perhaps the most famous example. He finds that temperature has little power in explaining futures prices for frozen-concentrated orange juice (FCOJ). Some researchers point to this finding and argue that the lack of a relation between prices and fundamentals in such a simple market makes it unlikely that a relation exists in more complex markets (see, e.g., Shleifer (2000), Hirshleifer (2001), and Daniel, Hirshleifer, and Teoh (2002)). Boudoukh, Richardson, Shen, and Whitelaw (BRSW) (2005), however, present opposing evidence. Using a nonlinear model of the price-temperature relation, they find that FCOJ futures prices do react to temperature changes as predicted. Their model has an explanatory power of 50% versus just 5% obtained with a linear model. The contrast between these results raises concerns about our ability to develop appropriate models in more complex markets and highlights the appeal of our methodology. By simply comparing the seasonal variance ratios, we can assess the contribution of public information flow to volatility without imposing much econometric structure.

³ George and Hwang (2001) highlight the potential impact of pricing errors (from a variety of sources) on variance ratio tests. We explicitly consider the impact of pricing errors in our analysis.

Our empirical results provide strong support for our hypotheses. The unconditional variance ratios for the weather-sensitive markets are generally lower than the variance ratio for the U.S. equity market and higher than the variance ratio for the currency market. This is consistent with the view that information flow is more evenly distributed around the clock in weather-sensitive markets than in the equity market but less than in the currency market. We also find that the differences between the variance ratios in weather-sensitive markets and the equity market are greater on weekends and during the weather-sensitive seasons. The variance ratios for the grains and natural gas, for example, are 30% lower than the variance ratio for the equity market on average, but 60% lower on weekends and 75% lower on weekends during the weather-sensitive seasons. This is consistent with the view that the higher nontrading-period volatility during these periods is driven by public information about the weather.

We find that the seasonal variance ratios in weather-sensitive markets are substantially lower during the weather-sensitive seasons than during the rest of the year. For example, for the grains and natural gas, the weekend variance ratios are 60–70% lower on season than off season. These results cannot be explained by seasonal variation in either trading activity or return autocorrelations (i.e., pricing errors). We also find that the comovement of corn, wheat, and soybean returns and absolute returns is stronger during the growing season and that the increased comovement is most dramatic during the nontrading period: The nontrading-period return (absolute return) correlations are 20% higher (35% higher) on season than off season. These results provide further evidence of strong seasonal information flow and suggest that public information flow during the nontrading period is an important source of volatility.

The remainder of the paper is organized as follows. Section I develops our hypotheses regarding the variance ratios in weather-sensitive markets and explains our test methodology. Section II provides details on the markets, our season definitions, and the data. Section III tests our hypotheses. Section IV examines several robustness tests, and Section V concludes.

I. Methodology

A. Information in Weather-Sensitive Markets

We define weather-sensitive markets as agricultural and energy markets in which weather conditions in a concentrated geographic area are a key determinant of supply and/or demand, and therefore price. Consider the corn market. Corn yields depend on rainfall and temperature levels at various stages of the growing season (see, e.g., Hu and Buyanovsky (2003)) and crops can be damaged by flooding, hail storms, and freezes. Since about 20% of the world's corn is grown in the Midwestern U.S. (Food & Agriculture Organization of the United Nations (FAO) Statistical Database (2004), and United States Department of Agriculture (USDA) (2004)), weather conditions in this region can have a material affect on world supply, and hence price.

News about weather conditions is therefore an important public information flow in weather-sensitive markets. In general, this news consists of two components, weather forecasts and forecast errors. We expect that on average the arrival rate for both of these components is nearly constant across the calendar day. The primary source of weather forecasts in the U.S. is the National Weather Service (NWS). Their information forms the basis for forecasts made by radio and television stations and private forecasting services (National Weather Service (2002)). According to the NWS, forecast updates were issued every 12 hours prior to 1994, every 3 hours from 1994 to 1998, and every hour starting in 1998. Many commodity traders also subscribe to private services that frequently update their forecasts.

Weather forecast errors also occur around the clock. Since weather conditions evolve 24 hours a day and even intraday temperature and precipitation forecasts are not perfect,⁴ unexpected weather developments that affect prices can occur at any time. Moreover, as BRSW (2005) observe, even weather conditions that conform with point forecasts can move prices. For example, if a freeze is forecast, but there is some probability that the freeze will not occur, the realization of the freeze is likely to move prices. Finally, since weather conditions are persistent (see, e.g., Wilks (1995)), weather forecast errors can also affect future forecasts.

We also expect that the importance of daily weather news varies systematically across the calendar year. In the case of corn, for example, weather conditions are most important during the growing season. Because changing weather conditions during this period can affect supply, forecast revisions and forecast errors during the growing season have greater potential price impact than similar events during the nongrowing period. This intuition is consistent with the seasonality in daily return volatility documented for futures contracts on seasonal commodities (see Anderson (1985)).

B. Testable Hypotheses

We frame our hypotheses regarding the trading- to nontrading-period variances in weather-sensitive markets in terms of variance ratios. Let σ_{oc}^2 and σ_{co}^2 denote the variances of open-to-close and close-to-open returns, respectively. We consider the total variance ratio (TVR), which is simply $\sigma_{oc}^2/\sigma_{co}^2$, and the variance rate ratio (VRR), which divides each variance by the length of its trading or nontrading period. If these lengths are constant over the sample period, the TVR and VRR differ only by a multiplicative constant. Therefore, we rely on the TVR for most of our analysis and consider the VRR only when the distinction is relevant.

Our hypotheses regarding the variance ratios in weather-sensitive markets arise from three propositions: (1) overnight and weekend weather events are an important source of public information in weather-sensitive markets;

⁴ NWS forecast verification statistics are available on the NWS website, <http://www.nws.noaa.gov/mdl/verif/>.

(2) the importance of this information changes in a predictable seasonal pattern over the calendar year; and (3) there is a positive relation between information flow and variance. These propositions lead to several predictions about how the variance ratios in weather-sensitive markets compare to those in the stock and currency markets and how they vary with the seasons.

HYPOTHESIS 1: The variance ratios in weather-sensitive markets are lower than the variance ratio in the stock market and higher than the variance ratio in the currency market.

Hypothesis 1 reflects the view that the structure of information flow and trading hours for weather-sensitive markets lies between the structures for the stock and currency markets. With respect to the stock market, both trading and information flow are concentrated in the business day. As French and Roll (1986) hypothesize, both public and private information is more likely to be generated during normal business hours when markets are open. Although stock-specific announcements are often made during nontrading hours, these should have a small impact on stock index portfolio values.⁵ In contrast, a large component of the public information flow in weather-sensitive markets occurs around the clock, independent of trading hours. To the extent that information flow generates volatility, we should observe lower variance ratios in weather-sensitive markets than in the stock market. With respect to the currency market, on the other hand, both trading and information flow occur around the clock. To the extent that private information and nonweather-related public information affect prices during the trading day, we should observe higher variance ratios in weather-sensitive markets than in the currency market.

HYPOTHESIS 2: The variance ratios in weather-sensitive markets are lower during the weather-sensitive season than during the rest of the year.

Hypothesis 2 reflects the view that if public information about the weather is important to price discovery, the effects should be more pronounced during periods of the year in which commodity prices are most sensitive to the weather. Although both trading- and nontrading-period variances should be higher during the weather-sensitive season, to the extent that private information is incorporated during the trading period, the proportional increase in information flow should be greater for the nontrading period, producing a lower variance ratio. Seasonality that is consistent with Hypothesis 2 would also strengthen the support for Hypothesis 1. On any given day, there may be many nonweather-related news events, such as crop reports, changes in demand forecasts, and government policy announcements that affect prices in weather-sensitive markets. If this news is more likely to be released during the trading day, this would work against Hypothesis 1. However, during the weather-sensitive season, weather news should be the dominant information

⁵ Macroeconomic announcements are usually made in the morning before the stock market opens (see Bauwens, Ben Omrane, and Giot (2005)). However, this information flow can impact both stock prices and commodity prices.

flow. Therefore, we expect that Hypothesis 1 should hold, in particular, during the weather-sensitive season.

Hypothesis 1A: *The variance ratios in weather-sensitive markets are lower and more (less) comparable to those in the currency (stock) market for weekend returns than for weekday returns.*

Hypothesis 2A: *The seasonal difference in variance ratios for weather-sensitive markets is greater for weekend returns than for weekday returns.*

Hypotheses 1A and 2A reflect the view that the length of the nontrading period in weather-sensitive markets is important because it affects the nature of the information flow. Obviously, markets are closed for a longer period on weekends than overnight. This affects both the likelihood that a weather event occurs (e.g., temperatures can fall dramatically over a weekend) and the potential cumulative impact of multiple weather events (e.g., three straight days of rain). In contrast, information events in the stock market (and, to some extent, in the currency market) are more likely to occur during the week than on weekends. Therefore, the variance ratios in weather-sensitive markets should be much lower than the variance ratio in the stock market on weekends and closer to the variance ratio in the currency market. Moreover, the distinction between weekday and weekend weather information flow should be more important during the weather-sensitive season. To investigate these issues, we compute two sets of variance ratios, one based on the weekday nontrading period, using the variances of open-to-close returns on Tuesday through Friday and overnight returns on Monday through Thursday, and another based on the weekend nontrading period, using the variances of open-to-close returns on Monday and weekend returns from Friday close to Monday open.⁶

Hypothesis 3: *The comovement of returns and absolute returns in markets sensitive to the same weather information is stronger during the weather-sensitive season than during the rest of the year, and the seasonal difference is greater in the nontrading period than in the trading period.*

Hypothesis 3 reflects the view that if weather information is more important during certain periods of the year, commodities that depend on the same weather information should comove more closely during these periods. Moreover, this effect should be more pronounced during the nontrading period, when (common) information about the weather is the most important information, as opposed to during the trading period, when many other idiosyncratic events can affect prices.

⁶ Note that our weekend and weekday variance ratios each have the same number of observations in the numerator and the denominator, thus avoiding the bias discussed by Jones, Kaul, and Lipson (1994).

C. Test Methodology

We test our hypotheses using a bootstrap approach to evaluate statistical significance. This approach avoids the size distortions associated with asymptotic variance ratio tests in the presence of strong persistence in conditional variances (see Andersen, Bollerslev, and Das (2001)) and also allows us to easily construct joint hypothesis tests, even though the length of the data series varies across markets. Specifically, we employ the stationary bootstrap, a block bootstrap procedure in which the length of each block is chosen at random from the geometric distribution (see Politis and Romano (1994)).

To illustrate, consider Hypotheses 1 and 1A. If our variance ratio estimates do not satisfy all of the inequality restrictions implied by one of these hypotheses, then we need to determine the sampling distribution of our variance ratio estimators to assess whether the violations are statistically significant. We do this by generating bootstrap draws from the empirical distribution of $z_t = (z_{co,t}, z_{oc,t})'$, the vector of standardized returns obtained by dividing the demeaned close-to-open and open-to-close returns by their sample standard deviations. First, we construct a resample, $Z^* = (z_1^*, z_2^*, \dots, z_T^*)$, using the stationary bootstrap. The resample is such that, in general, if $z_t^* = z_t$, then $z_{t+1}^* = z_{t+1}$ with probability p and z_{t+1}^* is drawn randomly from $Z = (z_1, z_2, \dots, z_T)$ with probability $1 - p$. This delivers an expected block length of $\bar{L} = 1/(1 - p)$. Second, we multiply each element of Z^* by the sample standard deviation of the same element of Z . Third, we compute the sample variance ratios for this bootstrapped data set and record the results. By performing a large number of bootstrap replications, we can approximate the sampling distribution of the variance ratio estimators.⁷

The actual procedure is more involved because the starting date for the data differs across markets. To account for this structure, we divide the sample into subperiods based on data availability. For example, we might have one subperiod in which data are available for five markets, a second subperiod in which data are available for six markets, and a third in which data are available for seven markets. We then build the bootstrap data set in stages. Each stage consists of implementing the first two steps of our bootstrap procedure for a given subperiod. Using this approach, all of the standardized returns for a given date are drawn as a vector, which preserves the correlation structure of the data, and avoids the problem of missing values in the data set.

Suppose we want to test Hypothesis 1. There are two statistics of interest for each weather-sensitive market, the difference between the TVRs for the stock market and the weather-sensitive market, and the difference between the TVRs for the weather-sensitive market and the currency market. We reject the hypothesis if the bootstrap confidence intervals indicate that either statistic is negative and statistically significant. It would also be useful to have a

⁷ We use a similar approach to test Hypotheses 2, 2A, and 3. The main difference is that we standardize the returns using the sample standard deviations for the weekday and weekend on- and off-season returns. In testing Hypothesis 3, we also account for the seasonality in the sample correlations.

summary measure of the evidence against the null, particularly for evaluating Hypotheses 2 and 3, since seasonal comparisons are at the core of our information story. However, the multivariate statistics normally used to test multiple inequality restrictions are not well suited to this purpose because they are functions of the asymptotic covariance matrix of the vector of parameter estimators, and it is not clear how to estimate this matrix when the length of the data series differs across markets.

Instead, we employ the approach developed by Fleming, Kirby, and Ostdiek (FKO) (2006) for testing multiple inequality restrictions on variance ratios. Multiple inequality restrictions are more challenging to test than equality restrictions because they do not specify a unique parameter configuration. FKO (2006) employ the usual strategy of using the least favorable configuration (LFC) with respect to the alternative hypothesis to construct critical values (see, e.g., Perlman (1969), Wolak (1987), and Wolak (1989)). In particular, they construct critical values using a procedure that reduces the multivariate testing problem to a univariate problem by focusing on the variance ratio statistic that provides the strongest evidence against the null. Their approach can easily be extended to test the restrictions implied by Hypotheses 2 and 3.

For example, if there are a total of n weather-sensitive markets, we can express the joint null and alternative hypotheses implied by Hypothesis 2 as

$$H_0: \Delta_i \leq 0 \forall i = 1, 2, \dots, n; \quad H_A: \Delta_i > 0 \text{ for some } i \in (1, 2, \dots, n), \quad (1)$$

where Δ_i denotes the difference between the population on- and off-season variance ratios for market i . In this case, our test of H_0 is based on the statistic

$$\hat{\theta} = \max_{1 \leq i \leq n} \sqrt{T_i}(\hat{\Delta}_i / \hat{\gamma}_i), \quad (2)$$

where $\hat{\Delta}_i$ is an estimator of Δ_i based on the sample variances and $\hat{\gamma}_i$ is a consistent estimator of the standard deviation of $\sqrt{T_i}(\hat{\Delta}_i - \Delta_i)$. It is easy to show that $\sqrt{T_i}(\hat{\Delta}_i - \Delta_i)$ is asymptotically normal.⁸ Thus, $\hat{\theta}$ is simply the largest of the t -statistics obtained by setting $\Delta_i = 0$ for each $i = 1, 2, \dots, n$. Using a test based on the maximum t -statistic has two advantages. First, it does not require an estimate of the covariance matrix of $(\hat{\Delta}_1, \dots, \hat{\Delta}_n)$. Second, it should have good power against alternatives for which the number of inequality restrictions violated is small.

We implement the test by using the stationary bootstrap to approximate $P(\hat{\theta} \leq c)$, the probability that $\hat{\theta}$ is less than or equal to c . First, we generate a bootstrap data set using the procedure described earlier and calculate

$$\hat{\theta}^{*(1)} = \max_{1 \leq i \leq n} \sqrt{T_i}(\hat{\Delta}_i^* - \hat{\Delta}_i) / \hat{\gamma}_i^*, \quad (3)$$

⁸ The sample variances of the off- and on-season returns are method of moments estimators whose asymptotic normality follows from standard results. Therefore, the asymptotic normality of $\hat{\Delta}_i$ follows immediately by application of the Delta Method (see Campbell, Lo, and MacKinlay (1997), p. 540).

where $\hat{\Delta}_i^*$ and $\hat{\gamma}_i^*$ are the values of $\hat{\Delta}_i$ and $\hat{\gamma}_i$ obtained using the bootstrap data. Next, after replicating this process M times to obtain $\hat{\theta}^{*(1)}, \dots, \hat{\theta}^{*(M)}$, we approximate $P(\hat{\theta} \leq c)$ by

$$\hat{P}(\hat{\theta} \leq c) = \frac{1}{M} \sum_{m=1}^M I(\hat{\theta}^{*(m)} - \hat{\theta} \leq c), \quad (4)$$

where $I(\cdot)$ denotes the indicator function. This corresponds to bootstrapping the distribution of the test statistic under the LFC, which occurs when all the inequalities are binding (Wolak (1987)).⁹ The critical value for testing H_0 at significance level ν is given by $\inf\{c : \hat{P}(\hat{\theta} \leq c) \geq 1 - \nu\}$.

D. Robustness Analysis

Our hypothesis tests regarding seasonal differences assume that we have correctly defined the weather-sensitive season in each market. We evaluate the robustness of our results to these definitions using kernel regression. Specifically, we estimate the variance ratios conditional on the day of year and then we evaluate the patterns in the conditional estimates.

To implement the kernel regression estimator, we assume that the squared demeaned close-to-open and open-to-close returns can be expressed as

$$e_{\text{co},t}^2 = \sigma_{\text{co}}^2(x_t) + u_{\text{co},t}, \quad (5)$$

$$e_{\text{oc},t}^2 = \sigma_{\text{oc}}^2(x_t) + u_{\text{oc},t}, \quad (6)$$

where x_t is a day-of-year indicator variable that takes on values from 1 to 365.¹⁰ We obtain the kernel estimator of the unknown regression function by taking a weighted average of the observations on the response variable with the weights determined by the distance of the observations on the explanatory variables from the point at which the function value is desired. For instance, the kernel estimate of the value of $\sigma_{\text{co}}^2(x_t)$ at $x_t = x$ is given by

$$\hat{\sigma}_{\text{co}}^2(x) = \frac{\sum_{t=1}^T K_h(x - x_t) e_{\text{co},t}^2}{\sum_{t=1}^T K_h(x - x_t)}, \quad (7)$$

where

$$K_h(x - x_t) = \frac{1}{h\sqrt{2\pi}} \exp\left(-\frac{(x - x_t)^2}{2h^2}\right) \quad (8)$$

is the Gaussian kernel with bandwidth h .

⁹ The LFC implies that $\sqrt{T_i} \hat{\Delta}_i$ is asymptotically distributed as $N(0, \gamma_i^2) \forall i = 1, 2, \dots, n$. Since the bootstrap distribution of $\sqrt{T_i}(\hat{\Delta}_i^* - \hat{\Delta}_i)$ is centered at zero, it follows that the bootstrap delivers critical values for the LFC.

¹⁰ We handle leap years by assigning February 29 the same indicator as February 28.

Kernel regression is also useful for investigating alternative explanations for our empirical results. For example, Stoll and Whaley (1990) and George and Hwang (2001) point out that pricing errors can influence variance ratio tests if they are correlated with information. Pricing errors can be generated by bid–ask bounce and other market microstructure effects or by systematic under- or overreaction to information.

We first consider broadly the level of trading activity as a source of bias. Suppose the kernel estimates of the variance ratios display a seasonal pattern that is consistent with the hypothesized seasonality in the information flows. To assess the likelihood that trading-related pricing errors could explain this pattern, we use kernel regression to estimate the expected level of trading activity by day of year and then we compare the pattern in these estimates to the pattern in the kernel estimates of the variance ratios. Even if the patterns are different, it is possible that volatility is linked to trading and that the strength of the linkage is seasonal. To investigate this possibility, we examine kernel estimates of the correlation between daily trading activity and contemporaneous open-to-close absolute demeaned returns.

To assess the impact of pricing errors more directly, we consider the autocorrelation structure of returns in a multiperiod variance ratio framework. Bid–ask bounce and pricing errors that are subsequently corrected (i.e., temporary errors) induce serial correlation in returns. Our interest centers on whether seasonal variation in the pricing errors can explain our results. We address this issue by comparing the variances of m - and one-period returns. Since the m -period variance is m times the one-period variance if returns are serially uncorrelated, variation in the multiperiod variance ratio can reflect the potential impact of pricing errors.

II. Data

A. Contract Selection

We select the commodity futures contracts for our analysis based on two factors, trading activity and price sensitivity to weather conditions in a concentrated geographic area. Contracts on agricultural and energy commodities are the most likely to satisfy the weather-sensitivity condition. The most active agricultural futures contracts are soybeans, corn, live cattle, and wheat. Over the past 5 years, the next-most active contracts, coffee and world sugar, traded less than half the dollar volume of any of these contracts (Commodity Research Bureau (2004)). Thus, we select the soybean, corn, and wheat contracts, and we exclude live cattle because it is not likely to be weather sensitive. The most active energy futures contracts are crude oil, natural gas, heating oil, and unleaded gasoline. Of these, we select the natural gas contract because, as we explain below, its prices are likely to be the most sensitive to weather conditions.

We expect futures prices for the grains to be weather sensitive because crop yields depend on weather conditions during the growing season.¹¹ Corn and

¹¹ Stevens (1991) documents a persistent weather effect on grain futures prices during the growing season.

soybean production are both concentrated. The U.S. produces 40% of the world's corn and 50% of the world's soybeans, with the next largest shares being 18% (China) for corn and 20% (Brazil) for soybeans.¹² Therefore, weather conditions in the primary U.S. growing states (the grain belt) are likely to be a key determinant of corn and soybean prices. The production of wheat, on the other hand, is less concentrated. The U.S. produces just 10% of the world's wheat and China, Russia, and India are also major producers. Production in these countries is largely for local consumption, however, so the U.S. accounts for a large share of the wheat available on the world market (over 33%). We therefore expect wheat prices to have some sensitivity to weather conditions in the U.S. grain belt, but less sensitivity than corn and soybean prices.

We expect natural gas futures prices to be weather sensitive because demand for natural gas is sensitive to temperatures in the winter and winter supply is constrained. The U.S. typically consumes twice as much natural gas in the winter as in the summer (due to space heating), and 70% of this increase is due to consumption in the Northeast and the Upper Midwest (see *Natural Gas Monthly*, various issues, published by the Energy Information Administration (EIA)). Thus, colder temperatures in these states can lead to large increases in demand. At the same time, the supply of natural gas in the winter is essentially fixed. U.S. production is relatively constant throughout the year, and the seasonality in consumption is satisfied by storing gas in the summer and withdrawing it in the winter (see International Energy Agency (IEA) (2002), Fig. 97). On average, these storage withdrawals supply 20% of consumption in the winter and as much as 50% on peak demand days (EIA (1995)). Overseas imports cannot provide much additional supply (EIA (2001a)). Therefore, natural gas prices can spike during peak periods in the winter in order to balance supply and demand.

We exclude heating oil, gasoline, and crude oil futures because these contracts are much less weather sensitive. While heating oil consumption is seasonal (EIA (2001b)), supply is less constrained in the winter than natural gas because higher heating oil prices stimulate additional supply. Refiners can increase production and imports can increase from overseas. Both of these sources take 2 to 3 weeks to reach consumers (EIA (2001b)), but their availability means that a short-term supply shortage that increases the spot price has a more muted effect on futures prices. Consistent with these fundamentals, the volatility of heating oil futures is much less seasonal during our sample period (low of 27% in May, high of 38% in January, excluding the Gulf War in 1990 and 1991) than the volatility of natural gas futures (36% in April, 77% in December). We expect even less weather sensitivity for gasoline and crude oil futures. Gasoline consumption is seasonal (i.e., summer driving months) but does not vary with daily weather conditions in any particular geographic area. Crude oil prices are not typically weather sensitive because over 90% of U.S. oil consumption is for

¹² All of the crop production and export statistics cited in this paragraph are from the FAO Statistical Database (2004) for the period from 1980 to 2004.

transportation (gasoline, jet fuel, and diesel) and industrial uses (EIA (2004)), which are not sensitive to the weather.¹³

In addition to the grain and natural gas contracts, we include futures contracts on FCOJ, S&P 500, and Japanese yen. We include FCOJ due to its prominence in prior research and its sensitivity to freezing temperatures around Orlando, Florida (BRSW (2005)). However, this contract is far less active than the other contracts we consider. The average dollar volume (Table I) is less than a tenth that of wheat, the next least active contract. In addition, the seasonality of FCOJ may be less well defined because orange trees are sensitive to general weather conditions throughout the year and to weather events other than freezes (e.g., hurricanes). We include the financial contracts (S&P 500 and Japanese yen) to serve as benchmarks for our analysis of the weather-sensitive markets. We select the yen because it is an actively traded currency and Japanese business hours do not overlap with trading hours for the futures contract, increasing the importance of overnight information flow.

B. Sample Construction

The data set consists of daily observations for open and close prices, number of trades, and trading volume. We obtain the prices and number of trades from Tick Data and trading volume from Datastream. Table I reports the sample period for each contract. The sample periods for the commodity contracts start with the first day of data available from Tick Data. We choose the start date for the financial contracts to match the first date for the grains. All of the sample periods end on December 31, 2004. We eliminate days with missing price observations for any of the contracts. The sample size ranges from 2,855 observations for natural gas to 5,370 observations for the grains and the financials.

We compute the daily open-to-close and close-to-open returns for each market using, in most cases, the nearby futures contract. However, we switch to the second nearby contract on the 20th day of the month prior to expiration for the commodities and a week prior to expiration for the financials to avoid expiration effects on prices and on the level of trading activity. In addition, we exclude the September contract for corn and wheat and the August and September contracts for soybeans because these contracts straddle the harvest cycle and are therefore much less actively traded (see Smith (2005)). Some of the contracts have exchange-imposed price limits (see Table I). In constructing the return series for these contracts, we retain the return for an interval in which a limit move occurs but we exclude the return over the subsequent interval unless the

¹³ Business press articles occasionally mention a link between crude oil prices and weather conditions in the Northeast (e.g., see "Crude oil up on high demand, cold weather, fears of tight supplies," *Washington Post*, March 4, 2005). This occurs when heating oil supplies are low and cold weather could increase demand for crude oil to produce heating oil (see Mussa (2000)). However, these conditions are not prevalent and, when they do occur, the impact on crude oil futures prices is much less than the impact on the spot price.

Table I
Data Description

This table provides descriptive data on the futures contracts used in the analysis. All of the price data are from the Tick Data futures transactions files. The sample includes all observations in our sample period except for days following a holiday or a missing date in the Tick Data data set, which are eliminated across all contracts. The market size data are the average daily dollar volume and year-end dollar open interest for the last year reported in the 2004 Commodity Research Bureau Yearbook CD. The market hours are the trading hours in effect at the end of the sample period. The mean trading hours are the average number of trading hours across the days in our sample. Mean trades per day are reported for the specific contract maturity used to generate each return time series ("Active") and across all contract maturities ("All"). The last column reports whether the contract is subject to daily price limits.

Contract	Sample Period	Sample Size	Exchange	Market Size (millions)		Market Hours		Mean Trading Hrs	Mean Trades/Day		Price Limits
				\$ Volume	\$ Open Int.	Open	Close		Active	All	
Corn	19820701–20041231	5,370	CBOT	935	5,175	93000	131500	CT	454	997	Yes
Soybean	19820701–20041231	5,370	CBOT	2,764	9,004	93000	131500	CT	818	1,707	Yes
Wheat	19820701–20041231	5,370	CBOT	521	1,979	93000	131500	CT	379	791	Yes
FCOJ	19870706–20041231	4,165	NYBOT	25	255	100000	133000	ET	387	234	Yes
Natural Gas	19930104–20041231	2,855	NYMEX	4,676	22,681	100000	143000	ET	144	641	Yes ^a
JY	19820701–20041231	5,370	CME	2,817	13,986	72000	140000	CT	5.22	1,334	Yes ^a
SP500	19820701–20041231	5,370	CME	57,677	200,863	83000	151500	CT	6.47	1,304	No
									2,795	3,302	Yes ^b

^aThe data in our sample are not affected by the price limits.

^bThere were no limits prior to 1988. After 1988, the potentially most binding constraints are lifted each day 45 minutes before the market close.

Table II
Weather-Sensitive Seasons

This table reports the weather-sensitive period for each contract and the criteria used to define these seasons.

Contract	On-Season Period		Identification Criteria
	Start	End	
Corn	May 15	October 31	Season begins when 50–75% of the crop is likely planted and ends when 75% of the crop is likely harvested.
Soybean	June 1	October 15	Season begins when 50–75% of the crop is likely planted and ends when 75% of the crop is likely harvested.
Wheat	March 15	August 31	Season begins when winter wheat is likely sprouting and ends when 75% of the spring wheat crop is likely harvested.
FCOJ	December 1	February 28	Months in which freezing temperatures around Orlando, Florida are most likely.
Natural Gas	December 1	February 28	Months in which the market is in a state of storage withdrawal and average temperatures are below freezing in the states that account for the majority of withdrawal.

return is a reversal. This removes zero-return trading days following a limit open as well as overnight continuations following a limit close.¹⁴

C. Season Definitions

Table II reports our definition of the weather-sensitive season for each contract. In general, the on season for the grains is the period of the year in which the crops are in the ground, the on season for natural gas is the period of the year in which the average temperature is below freezing in the states whose demand accounts for the majority of natural gas storage withdrawal, and the on season for FCOJ is the period of the year in which damaging freezes are most likely to occur around Orlando, Florida. The on season starting and ending dates reported in the table are based on our assessment of the potential for weather conditions to impact supply or demand in each market. The Appendix provides additional details.

III. Variance Ratio Hypothesis Tests

A. Unconditional Comparisons

Our first set of hypotheses (1 and 1A) relates to the unconditional variance ratios. Table III reports the close-to-close and trading- and nontrading-period

¹⁴ We provide more details regarding price limits and trading halts and our treatment of the affected price observations in the Supplemental Appendix, which is available at www.ruf.rice.edu/~jflaming/pub/tnt-app.pdf.

Table III
Variance Ratio Estimates

This table reports the trading- to nontrading-period total variance ratios (TVR) for each contract and the differences between the TVR for each weather-sensitive contract and the TVRs for the S&P 500 (excluding the 1987 crash) and the yen. The weekend results (Panel A) are based on returns from Friday close to Monday close and the weekday results (Panel B) are based on returns from Monday close to Friday close. The table also reports the number of observations and volatility (annualized by $\sqrt{252}$) for the close-to-close (CC), nontrading (CO), and trading (OC) periods. The confidence intervals (95% CI) reported for the TVRs and the differences are based on the stationary bootstrap procedure described in the text with 10,000 replications. The sample period starts on July 1, 1982 for all contracts except FCOJ (July 6, 1987) and natural gas (January 4, 1993) and ends on December 31, 2004. The crash period excluded for the S&P 500 is October 19–23, 1987.

Contract	Num. Obs.			Ann. Vol. (%)			Variance Ratio			TVR vs. S&P 500 (Ex. '87 Crash)			TVR vs. JY	
	CO	OC		CC	CO	OC	TVR	(95% CI)	Diff.	(95% CI)	Diff.	(95% CI)		
Panel A: Weekends														
Corn	982	972		23.62	16.22	16.65	1.05	(0.82, 1.37)	3.03	(1.55, 4.63)	0.65	(0.39, 0.96)		
Soybean	983	977		24.83	17.41	17.92	1.06	(0.80, 1.44)	3.01	(1.56, 4.66)	0.65	(0.36, 1.03)		
Wheat	985	983		23.47	13.05	19.60	2.26	(1.76, 2.91)	1.83	(0.29, 3.52)	1.85	(1.33, 2.50)		
FCOJ	752	755		31.42	17.51	27.18	2.41	(1.10, 4.84)	1.65	(-1.24, 3.62)	2.00	(0.69, 4.43)		
Natural Gas	515	515		63.58	37.17	46.68	1.58	(1.11, 2.19)	2.48	(0.79, 4.15)	1.17	(0.69, 1.78)		
JY	985	985		11.23	9.38	5.99	0.41	(0.32, 0.52)	—	—	—	—		
S&P 500	985	985		24.58	8.67	20.99	5.86	(2.86, 11.31)	—	—	—	—		
S&P 500 – Ex. '87 Crash	984	984		18.15	8.05	16.21	4.06	(2.65, 5.71)	—	—	—	—		
Panel B: Weekdays														
Corn	4,334	4,367		18.35	10.40	15.69	2.28	(1.97, 2.63)	1.87	(1.00, 2.87)	1.53	(1.22, 1.88)		
Soybean	4,343	4,374		19.62	10.63	16.98	2.55	(2.15, 3.07)	1.61	(0.63, 2.62)	1.81	(1.40, 2.32)		
Wheat	4,374	4,384		20.93	9.19	19.38	4.45	(3.84, 5.17)	-0.29	(-1.40, 0.83)	3.70	(3.09, 4.42)		
FCOJ	3,386	3,403		29.72	15.74	24.90	2.50	(1.68, 3.80)	1.66	(0.06, 2.92)	1.76	(0.95, 3.02)		
Natural Gas	2,340	2,340		49.58	26.58	42.64	2.57	(2.10, 3.20)	1.59	(0.59, 2.63)	1.83	(1.34, 2.45)		
JY	4,385	4,385		11.28	8.41	7.24	0.74	(0.64, 0.85)	—	—	—	—		
S&P 500	4,385	4,385		17.73	10.30	16.47	2.56	(1.66, 4.34)	—	—	—	—		
S&P 500 – Ex. '87 Crash	4,381	4,381		17.07	7.70	15.71	4.16	(3.41, 5.13)	—	—	—	—		

sample volatilities for each market, as well as the TVR estimates, for weekends (Panel A) and weekdays (Panel B). In each case, natural gas has the highest volatility, followed by FCOJ. The grains, which all have similar volatilities, are somewhat more volatile than the S&P 500 (excluding the 1987 crash) and the yen has the lowest volatility. The close-to-close volatilities are higher on weekends than on weekdays for all contracts except the yen. Turning to the TVR estimates, the estimates for weekend returns range from 0.41 for the yen to 5.86 for the S&P 500. The TVR estimates for weekday returns range from 0.74 for the yen to 4.45 for wheat.

We evaluate the precision of the TVR estimates using the stationary bootstrap with the expected block length $\bar{L}$ set equal to 40. This choice of $\bar{L}$ is based on results reported by FKO (2006). FKO use simulations to investigate the small sample properties of the bootstrap inequality tests and find (not surprisingly) that the test size is sensitive to the degree of volatility persistence implied by the data generating process. In particular, they find that $\bar{L}$ must be at least 40 for the test to have good size when the conditional return variances are highly persistent. Their results also suggest that there is little disadvantage to using a large $\bar{L}$ even if the returns display constant volatility: In this case, the power of the test with $\bar{L} = 40$ is only slightly lower than that obtained with $\bar{L} = 1$.

Table III reports the 95% confidence intervals (CIs) based on 10,000 bootstrap replications. If the CI does not include one, then we conclude that the trading- and nontrading-period variances are significantly different. The weekend results indicate that the trading-period variance is significantly higher than the nontrading-period variance for wheat, FCOJ, natural gas, and the S&P 500. For the yen, the trading-period variance is significantly lower. The weekday results indicate that the trading-period variance is significantly higher than the nontrading-period variance for all contracts except the yen. Again, for the yen, the trading-period variance is significantly lower. The results for the yen are not surprising given that the Japanese business day occurs during the (U.S.) nontrading period.

Table III also reports the differences between the TVR estimates for the weather-sensitive contracts and those for the S&P 500 (excluding the 1987 crash) and the yen. Hypothesis 1 says that all of the differences reported in the table should be positive. Consistent with the hypothesis, all of the weekend point estimates are positive and, except for the difference between FCOJ and the S&P 500, none of the CIs include zero. Similarly, all of the weekday point estimates are positive and significantly different from zero, except for the difference between wheat and the S&P 500, which is insignificantly negative. Since this difference would appear even less significant if we were to treat it as the minimum of five univariate statistics, we can forgo implementing a multivariate test and conclude that the data provide no evidence against Hypothesis 1.

The results in Table III also provide support for Hypothesis 1A. The TVR estimates for the weather-sensitive contracts are lower on weekends than on weekdays and the differences between the estimates for the S&P 500 and the weather-sensitive contracts are generally much greater on weekends than on weekdays. This occurs because the weather-sensitive contracts, unlike the

S&P 500, have much higher volatility on weekends than overnight, consistent with the view that weather information flow evolves randomly around the clock while news about stocks is concentrated during the business day. On the other hand, the differences between the TVRs for the weather-sensitive contracts and for the yen tend to be smaller on weekends than on weekdays. This is expected because weather-sensitive markets and currency markets both have important information flow during the weekend nontrading period.

As an additional test, we evaluate whether Hypothesis 1 holds for variance rate ratios. The weekend VRR estimates (not reported in the table) are 4.2 for the yen, 39.7 for the noncrash S&P 500, and between 19.0 and 42.6 for the weather-sensitive markets. The weekday VRR estimates are 2.0 for the yen, 10.8 for the noncrash S&P 500, and between 9.5 and 23.7 for the weather-sensitive markets. Thus, the weekend estimates are consistent with Hypothesis 1, but the weekday estimates are not. Although this might be interpreted as evidence against the hypothesis, it may also reflect the problem of comparing VRRs across markets with trading days of different lengths. The VRR reflects the average variance rate per unit of time and the trading day is much shorter for the weather-sensitive contracts than for stocks (see Table I). To the extent that trading is necessary to incorporate private information flow, a given amount of information would have to be incorporated at a faster rate in these markets than in the stock market.

B. Seasonal Comparisons

Although the unconditional results are generally consistent with our information story, seasonal comparisons in the weather-sensitive markets (Hypotheses 2 and 2A) are at the core of the story. If public information about weather conditions is more important during the on season, the proportional increase in information flow should be greater during nontrading periods than during trading periods. Therefore, if information drives volatility, we should find lower variance ratios during the on season.

Table IV reports the results of the seasonal comparisons. Consider first the on- and off-season sample volatilities. In all cases, the close-to-close volatility is higher on season than off season, which is consistent with the premise that weather information flow is more important during the on season. Consider next the on- and off-season TVR estimates. Both the weekend and weekday estimates are lower on season for four of the five contracts.¹⁵ Lower estimates on season indicate that the on-season increase in volatility is proportionally greater for the nontrading period, as predicted by Hypothesis 2. Only the estimates for FCOJ are inconsistent with the hypothesis.

The 95% CIs for the seasonal differences indicate that the TVR is significantly lower on season than off season in most cases. The estimates for FCOJ and the weekday estimate for wheat are not significantly different from zero. However, the upper confidence limits for these estimates suggest that the

¹⁵ Note that the inference from these difference statistics is independent of the distinction between TVRs and VRRs because the two variance ratios differ only by a multiplicative constant.

Table IV
Seasonal Variance Ratio Estimates

This table reports the on- and off-season total variance ratios (TVR) for the weather-sensitive contracts and the difference between the on- and off-season estimates ("Seasonal Diff."). The weekend results (Panel A) are based on returns from Friday close to Monday close and the weekday results (Panel B) are based on returns from Monday close to Friday close. The table also reports the number of observations and volatility (annualized by $\sqrt{252}$) for the close-to-close (CC), nontrading (CO), and trading (OC) periods. The confidence intervals (95% CI) reported for the TVRs and the seasonal differences are based on the stationary bootstrap procedure described in the text with 10,000 replications. The sample period starts on July 1, 1982 for all contracts except FCOJ (July 6, 1987) and natural gas (January 4, 1993) and ends on December 31, 2004. The on-season period is May 15 to October 31 for corn, June 1 to October 15 for soybeans, March 15 to August 31 for wheat, and December 1 to February 28 for FCOJ and natural gas.

Contract	Seas.	Obs.		Ann. Volatility (%)			TVR	(95% CI)	Seasonal Diff.
		CO	OC	CC	CO	OC			(95% CI)
Panel A: Weekends									
Corn	On	469	460	28.18	20.81	18.79	0.82	(0.57, 1.18)	1.20
	Off	513	512	18.32	10.19	14.45	2.01	(1.44, 2.85)	(0.51, 2.05)
Soybean	On	388	382	31.36	24.18	20.48	0.72	(0.48, 1.08)	1.45
	Off	595	595	19.42	10.88	16.03	2.17	(1.57, 2.99)	(0.75, 2.28)
Wheat	On	469	467	26.40	16.40	20.78	1.61	(1.15, 2.23)	2.62
	Off	516	516	20.43	8.99	18.47	4.22	(3.08, 5.77)	(1.30, 4.19)
FCOJ	On	160	161	37.88	16.06	32.31	4.05	(0.93, 12.32)	-1.99
	Off	592	594	29.31	17.82	25.58	2.06	(0.91, 4.29)	(-10.08, 1.82)
Natural Gas	On	109	109	92.11	60.57	60.32	0.99	(0.56, 1.72)	1.32
	Off	406	406	53.24	27.70	42.10	2.31	(1.72, 3.11)	(0.37, 2.21)
Panel B: Weekdays									
Corn	On	2,037	2,064	21.11	12.54	17.59	1.97	(1.65, 2.37)	0.96
	Off	2,297	2,303	15.44	8.04	13.75	2.93	(2.47, 3.48)	(0.35, 1.60)
Soybean	On	1,639	1,666	23.01	14.22	18.84	1.76	(1.40, 2.23)	2.44
	Off	2,704	2,708	17.19	7.66	15.69	4.20	(3.50, 5.09)	(1.60, 3.38)
Wheat	On	2,037	2,047	22.28	10.02	20.28	4.09	(3.45, 4.93)	0.80
	Off	2,337	2,337	19.67	8.38	18.54	4.89	(4.13, 5.81)	(-0.30, 1.90)
FCOJ	On	798	806	29.76	14.69	26.53	3.26	(1.51, 6.55)	-0.96
	Off	2,588	2,597	29.69	16.05	24.37	2.31	(1.47, 3.65)	(-4.32, 1.33)
Natural Gas	On	559	559	63.82	37.21	51.22	1.89	(1.43, 2.54)	1.31
	Off	1,781	1,781	44.14	22.08	39.56	3.21	(2.73, 3.81)	(0.51, 2.06)

p -value generated by a formal test of Hypothesis 2 might be reasonably low. We investigate this further using 10,000 bootstrap replications to approximate the p -value implied by the multivariate test described in Section I.C.¹⁶ The maximum t -statistic for the weekend TVRs is 1.75, which has a p -value of 0.3431, and the maximum t -statistic for weekday estimates is 0.96, which has a p -value of 0.7040. Thus, we conclude that the data provide no reliable evidence against Hypothesis 2.

¹⁶ We estimate γ_i by treating the moment conditions for the off- and on-season TVR estimators as an exactly identified generalized method of moments (GMM) system and using the Newey and West (1987) estimator of the GMM covariance matrix. Following FKO (2005), we set the lag truncation for the estimator equal to $1.4\bar{L}$.

Despite this conclusion, the FCOJ results stand out as the least consistent with Hypothesis 2. Several features of the FCOJ market might explain these results. First, the most important weather event is dichotomous, freeze or no freeze. Since there are no damaging freezes in many years, we may face a peso problem. Second, temperature impacts FCOJ prices nonlinearly: Prices only move dramatically when the expectation of a freeze is high (see BRSW (2005)). Moreover, once a damaging freeze occurs, it essentially fixes supply and ends the weather sensitivity of prices until the next growing season.¹⁷ Finally, the nature of weather news on weekends is different for FCOJ than for the other contracts. For example, a freeze is more likely to occur over the weekend than overnight, but the impact of a freeze on supply (and hence price) is the same regardless of when it occurs. In contrast, a rainstorm is more likely to occur on weekends and there is greater potential for cumulative rainfall that has a large impact on supply. This could explain why the weekend and weekday volatilities for FCOJ are comparable, while the volatilities for the other contracts are sharply higher on weekends.

Turning to Hypothesis 2A, the seasonal differences in the TVR estimates in Table IV tend to be greater for weekends than weekdays, as predicted. The raw differences for corn, wheat, and natural gas are all greater on weekends and, in percentage terms, the differences are greater for all contracts except FCOJ. For the grains, the weekend estimates are 60–70% smaller on season than off season, while the weekday estimates are 20–60% smaller on season. For natural gas, the weekend estimate is 60% smaller on season than off season, while the weekday estimate is 40% smaller on season.

Finally, note that the results of our seasonal analysis reinforce our earlier conclusions regarding Hypothesis 1. Specifically, all of the on-season TVRs for the weather-sensitive contracts (Table IV) lie between the TVRs for the yen and the S&P 500 (Table III). This was not the case for the unconditional TVRs. However, we expect to find stronger evidence in favor of Hypothesis 1 during the on season because weather news represents a greater share of the total information flow during this period of the year. This evidence is consistent with public information about the weather being a key determinant of volatility in weather-sensitive markets.

C. Seasonal Comovements

Our final hypothesis relates to seasonal changes in the comovement of markets sensitive to the same weather conditions. All of the grains, for example, are grown in a similar geographic area during a similar growing season. Therefore, if weather information generates price movement, the comovement of returns and absolute returns across the grains should be stronger during the growing season and the seasonal increase should be larger for the nontrading period.

Table V reports the seasonal correlation estimates for the trading and nontrading periods, the seasonal difference in the correlation estimates, the

¹⁷ For this reason, BRSW (2005) exclude all returns following the first damaging freeze of a season.

Table V
Seasonal Comovement Analysis Across the Grains

This table reports the correlation of returns (Panel A) and absolute returns (Panel B) across the grain complex. We report the sample correlations for both the nontrading (CO) and trading (OC) periods, both on and off season, the difference between the on- and off-season estimates ("Seasonal Difference"), and the difference between the nontrading- and trading-period estimates ("CO - OC Difference"). We use both weekday and weekend returns to obtain these estimates and the on-season period for each pair of markets is defined by the earliest on-season date and latest on-season date for the pair. The confidence intervals (95% CI) reported for the differences are based on the stationary bootstrap procedure described in the text with 10,000 replications. The sample period starts on July 1, 1982 for all contracts except FCOJ (July 6, 1987) and natural gas (January 4, 1993) and ends on December 31, 2004. The on-season period is May 15 to October 31 for corn, June 1 to October 15 for soybeans, March 15 to August 31 for wheat, and December 1 to February 28 for FCOJ and natural gas.

			Obs.		Seas.	CO	OC	Seasonal Difference (95% CI)	OC	Seasonal Difference (95% CI)	CO – OC	CO – OC Difference (95% CI)
			CO	OC								
Panel A: Return Correlations												
Corn	Soybean	On	2506	2524	0.8109	0.1314	0.6336	0.0408	0.1774	(0.1364, 0.2160)		
		Off	2810	2815	0.6796	(0.0739, 0.1916)	0.5927	(−0.0075, 0.0886)	0.0868	(0.0357, 0.1342)		
Corn	Wheat	On	3407	3429	0.6783	0.1013	0.5528	0.0495	0.1255	(0.0808, 0.1697)		
		Off	1909	1910	0.5770	(0.0320, 0.1736)	0.5033	(−0.0061, 0.1063)	0.0737	(0.0059, 0.1404)		
Soybean	Wheat	On	3141	3164	0.6492	0.1253	0.4573	0.0528	0.1919	(0.1409, 0.2409)		
		Off	2185	2187	0.5239	(0.0491, 0.2032)	0.4044	(−0.0040, 0.1094)	0.1195	(0.0511, 0.1857)		
Panel B: Absolute Return Correlations												
Corn	Soybean	On	2506	2524	0.7498	0.1488	0.4447	0.0334	0.3052	(0.2468, 0.3628)		
		Off	2810	2815	0.6011	(0.0721, 0.2299)	0.4113	(−0.0367, 0.1033)	0.1898	(0.1200, 0.2568)		
Corn	Wheat	On	3407	3429	0.5977	0.1310	0.3636	0.0575	0.2340	(0.1634, 0.3006)		
		Off	1909	1910	0.4667	(0.0260, 0.2473)	0.3061	(−0.0246, 0.1415)	0.1606	(0.0537, 0.2611)		
Soybean	Wheat	On	3141	3164	0.5601	0.1815	0.2387	0.0278	0.3214	(0.2397, 0.4006)		
		Off	2185	2187	0.3786	(0.0601, 0.3064)	0.2109	(−0.0495, 0.1057)	0.1677	(0.0579, 0.2721)		

within-season difference in the correlation estimates, and 95% CIs for the differences based on 10,000 bootstrap replications.¹⁸ The return correlations (Panel A) are consistent with Hypothesis 3. For each pair of markets, the correlation of nontrading period returns is higher on season than off season and the differences are all statistically significant. The correlations of trading period returns are also all higher on season but none of the differences is significant. Finally, for each pair of markets, the difference between the correlation estimates for the nontrading and trading periods is greater on season and the on-season differences are all statistically significant.

The absolute return correlations (Panel B) are also consistent with Hypothesis 3. Again, for each pair of markets, the estimate for the nontrading period is significantly higher on season than off season, the estimate for the trading period is (insignificantly) higher on season than off season, and the difference between the estimates for the nontrading and trading periods is greater and highly significant on season. These findings are consistent with the view that public information about the weather is an important source of nontrading-period volatility in the grain markets.

IV. Robustness Tests

In this section, we use kernel regression to investigate the robustness of our information story as an explanation for our results. First, we estimate the TVRs conditional on the day of year to assess the robustness of our season definitions and to establish the degree of seasonal variation in the estimates. Second, we assess whether seasonal patterns in trading activity and the correlation between trading activity and volatility are consistent with the seasonal pattern in the TVRs. Third, we estimate conditional multiperiod variance ratios to evaluate the possible impact of pricing errors. Finally, we consider the choice of sample period, the impact of scheduled public information releases, and the impact of our treatment of price limits.

A. Conditional TVRs

We begin by using kernel regression to estimate the conditional TVRs (CTVRs) for each day of the year. For all of the kernel regressions in our robustness tests, we set the bandwidth (h) equal to 30. We determine this value based on visual inspection of the results, that is, we start with a large bandwidth and decrease the value until it appears to produce too much local variation in the CTVR estimates.¹⁹ To put this choice in perspective, recall that h acts as

¹⁸ We combine the weekday and weekend returns in computing the correlation estimates. The on-season period for each pair of markets is defined by the first on-season date and the last on-season date for the pair. However, the results are robust to using only the common on-season period for each pair.

¹⁹ We do not use cross validation, the most common approach to bandwidth selection, because it has two important drawbacks in our application. First, it is essentially a least-squares procedure, so it is not robust to outliers. This is problematic because we are applying the kernel estimator to squared demeaned returns. Second, cross validation is a univariate procedure. Thus, we could obtain a different value of h for each contract, raising concerns about the validity of cross-contract comparisons.

the standard deviation in the Gaussian kernel. Hence, about 95% of the kernel weight is placed on observations within ± 60 days of the day of the year for which the CTVR is desired.

Figure 1 plots the CTVR estimates for each market. The solid horizontal line in each plot shows the average estimate for the full year. For the weather-sensitive contracts, the shaded area shows the defined on-season period and the dashed horizontal line in the shaded (non-shaded) area shows the average estimate for the on (off) season. In general, the patterns in the CTVR estimates for the weather-sensitive contracts correspond closely to our defined seasons. For corn and soybeans, the estimates decrease as the growing season begins, reach a low in July, and increase again as the crops approach harvest. Wheat exhibits a similar but less dramatic pattern. Corn and soybeans also exhibit a secondary cycle, with decreasing estimates from December into February. Consistent with our information story, this pattern may be driven by the impact of weather conditions in Argentina and Brazil. These countries produce over 30% of the world's soybeans and nearly 10% of the world's corn (FAO Statistical Database (2004)) and the most weather-sensitive period for the South American crops is December to February (see Solberg (1999)).²⁰

The CTVR estimates for natural gas also correspond to our defined seasons, reaching a peak in early summer and falling as winter approaches. The plot also shows an interesting plateau in August. This may reflect demand for electricity generation during summer heat waves (see Sturm (1997)), as greater sensitivity to information regarding daytime temperatures would tend to produce higher variance ratios. On average, the CTVR estimates for FCOJ are also lower during the on season and higher during the off season. However, the estimates reach a low in October, well ahead of our defined weather-sensitive season. This may be associated with the USDA's first production forecast for the growing season, which is released during nontrading hours in the second week of October (see BRSW (2005)). If we eliminate USDA announcement dates, the plot for FCOJ conforms much more closely with our defined seasons. In contrast to the plots for the weather-sensitive markets, the CTVR estimates for the yen and the S&P 500 stay within a relatively narrow range throughout the year. The only clear pattern is that the CTVRs for the S&P 500 tend to be lower during the last quarter of the year.

B. Conditional Trading Activity Analysis

Although the seasonality in the CTVR estimates is consistent with the seasonal impact of weather information flow, perhaps there are other explanations. Suppose, for example, that trading generates pricing errors, and hence volatility that is unrelated to information flow. If this is the case, then the observed seasonality in the CTVR estimates may be somehow related to seasonal changes in trading activity. We investigate this possibility by extending our kernel regression analysis to include measures of trading activity.

²⁰ In contrast, South America produces just 3% of the world's wheat (FAO Statistical Database (2004)).




Figure 1. Kernel estimation of total variance ratios by day of year. These graphs display the total variance ratios (TVR) based on kernel estimation of the mean trading- and nontrading-period demeaned squared returns for each day of the calendar year. The shaded areas in the graphs for the weather-sensitive contracts represent the defined on-season periods. The thin horizontal line in each graph represents the average TVR across the full year and the dashed horizontal lines represent the average TVRs in the defined on- and off-season periods. The sample period starts on July 1, 1982 for all contracts except FCOJ (July 6, 1987) and natural gas (January 4, 1993) and ends on December 31, 2004.

Figure 2 shows the conditional expected volume (left-hand column) and expected number of trades (right-hand column) for each market by day of year. For corn and natural gas, both measures of trading activity are greater, on average, during the off season. For the other contracts, trading activity tends to be greater on season. In each case, however, the seasonal differences are fairly small. Moreover, the seasonality in trading activity does not mimic the




Figure 2. Kernel estimation of trading activity by day of year. These graphs display the conditional expected volume (left-hand column) and number of trades (right-hand column), as well as the correlation between each trading activity variable and absolute open-to-close returns, based on kernel estimation for each day of the calendar year. The solid line represents the total volume (100s of contracts) or total number of trades (100s of contracts) and the dashed line represents the correlation between contemporaneous open-to-close returns and the trading activity variable. The vertical dashed lines in each graph represent the days of the year on which the kernel estimates of the variance ratio cross the annual average (Figure 1). The sample period starts on July 1, 1982 for all contracts except FCOJ (July 6, 1987) and natural gas (January 4, 1993) and ends on December 31, 2004.

seasonality in the CTVR estimates (Figure 1). Trading activity for soybeans, for example, reaches a peak at the beginning of the on season and falls below the annual average just after mid-season. In contrast, the CTVR estimates fall into July and remain below the annual average through the end of the season.

Nonetheless, it is still possible that trading activity is linked to volatility and that the strength of the linkage is seasonal. To investigate this, we examine whether there is seasonal variation in the correlation between daily trading activity and contemporaneous open-to-close absolute demeaned returns. The dashed curves in Figure 2 plot the correlation estimates by day of year. In general, the estimates are relatively flat throughout the year, and there is no evidence that they are systematically different during the on season. Therefore, it seems unlikely that trading activity is the cause of the seasonal differences observed in the CTVRs.

C. Conditional Pricing Error Analysis

We directly assess whether pricing errors can explain the seasonal variation in the CTVR estimates by examining the multiperiod variance ratios for both trading- and nontrading-period returns. If the returns for a given interval (trading- or nontrading-period) are serially uncorrelated, then the m -period variance should equal m times the one-period variance. However, if pricing errors induce negative serial correlation (i.e., a pattern of overreaction and subsequent correction), the ratio of the m -period variance to the one-period variance will be less than m ; similarly, if pricing errors induce positive serial correlation (i.e., a pattern of underreaction and subsequent correction), the multiperiod variance ratio will be greater than m .

Figure 3 plots the kernel regression estimates of the 2- and 20-period variance ratios for the trading and nontrading periods by day of year. The 2-period estimates are close to 2, consistent with no serial correlation, but the estimates for the nontrading period tend to be greater than those for the trading period. The 20-period estimates for the trading period are close to 20 for soybeans, wheat, and natural gas, and tend to be greater (less) than 20 for corn (FCOJ). The estimates for the nontrading period tend to be greater than 20, on average, for all of the contracts except soybeans, and are substantially higher during some periods of the year for soybeans and FCOJ. Note, however, that the 20-period estimates are much less precise than the 2-period estimates because of the reduction in the number of observations.²¹

In order for pricing errors to explain our results, they must induce a bias that decreases the trading-period variance and/or increases the nontrading-period variance during the on season. If this is the case, we should find that, on season, the m -period variance ratios are greater than m for the trading period, less than m for the nontrading period, or at least that the variance ratio for the trading period increases relative to that for the nontrading period. Figure 3 shows some evidence of this for wheat. However, for corn, soybeans, and natural gas, the multiperiod variance ratios for the nontrading period increase relative to those for the trading period on season. The evidence regarding FCOJ is mixed. Therefore, it seems unlikely that our results can be attributed to the impact of pricing errors.

²¹ Although we use overlapping returns to improve the efficiency of the multiperiod variance ratio estimators, the efficiency gains are relatively small. See Richardson and Smith (1991).




Figure 3. Kernel estimation of multiperiod variance ratios by day of year. These graphs display the conditional 2-period (left-hand column) and 20-period (right-hand column) variance ratios for close-to-open (solid line) and open-to-close (dashed line) returns based on kernel estimation for each day of the calendar year. The vertical dashed lines in each graph represent the days of the year on which the kernel estimates of the total variance ratio (TVR) cross the annual average (Figure 1). The sample period starts on July 1, 1982 for all contracts except FCOJ (July 6, 1987) and natural gas (January 4, 1993) and ends on December 31, 2004.

D. Other Robustness Tests

We consider three additional robustness issues: (1) changing market fundamentals; (2) scheduled public information releases; and (3) our treatment of price limits. To assess whether our results are robust to these issues, we

investigate their impact on the CTVR estimates. Below we summarize the main findings. We report the complete results in the Supplemental Appendix.

Market fundamentals have changed over our sample period, including changes in the U.S. share of world agricultural production and changes in the supply and demand of natural gas. In addition, the NWS increased the frequency of its forecast update cycle in 1994 and again in 1998. To evaluate the impact of these changes on our results, we divide the sample period in half and examine the CTVRs for each subperiod. The CTVRs for corn, soybeans, and natural gas are stable across the two subperiods. The CTVRs for wheat are greater and exhibit less seasonality in the second subperiod. This may be due to the declining importance of U.S. production on the world market: The U.S. share of wheat exports fell from 36% in the 1980s to 23% since 1997 (FAO Statistical Database (2004)). The CTVRs for FCOJ are different across the two subperiods, but there is no clear pattern to the differences. The CTVR estimates for the yen (S&P 500 excluding 1987 crash) are about 25% (10%) less on average in the second subperiod. None of the differences, however, is sufficient to alter our basic conclusions.

Public information releases in the agricultural and energy markets include monthly government crop reports and weekly industry reports on storage levels for natural gas. Since these reports are an easily identifiable source of public information that is unrelated to current weather conditions, we examine how excluding announcement interval returns affects our results. In general, we find that excluding these returns has little effect on our results. The most notable difference is that the CTVR estimates for FCOJ show less seasonality in October, which likely reflects the importance of nontrading-period news about orange production estimates entering the growing season.

Price limits for the agricultural contracts were hit on a number of days in our sample. Section II.B describes our procedure for handling returns on and subsequent to limit move days. To assess whether our results are sensitive to the choice of procedure, we estimate the CTVRs using two alternative procedures: (1) ignoring price limits entirely; and (2) using the approach commonly used for full-day returns. We find that the CTVRs are not sensitive to the choice of procedure.

V. Conclusions

In stock and currency markets, it is difficult to distinguish between volatility caused by public and private information flow and volatility caused by irrational phenomena such as overreaction and excess trading because, in most cases, information flow and trading occur contemporaneously. Weather-sensitive commodity markets provide an opportunity to shed light on this issue. The information environment in these markets differs from that in financial markets in two key respects. First, a major component of the public information flow (i.e., news about weather conditions) evolves randomly over the 24-hour day and is independent of whether the market is open, the level of trading activity, and the timing of normal business hours. Second, the importance of this

component of the public information flow varies in a predictable seasonal pattern across the calendar year. Based on these differences, we generate hypotheses about how the trading- to nontrading-period variance ratios in weather-sensitive markets compare to those in the stock and currency markets and how the ratios vary across seasons.

We test our hypotheses using futures data for corn, soybeans, wheat, natural gas, and FCOJ. In general, the results are consistent with our information hypotheses. The variance ratios in these markets are lower than the variance ratio in the stock market and higher than the variance ratio in the currency market, the variance ratios tend to be much lower during the weather-sensitive season than during the rest of the year, and the comovements of returns (and absolute returns) across the grain complex are stronger during the weather-sensitive season, with a greater increase in comovements during the nontrading period. We also find that the seasonality in the variance ratios cannot be explained by seasonal changes in trading activity or by pricing errors. Therefore, we conclude that a large component of the volatility in weather-sensitive markets is driven by public information flow regarding weather conditions.

Appendix: Defining the Weather-Sensitive Seasons

A. Grains

We define the seasons for the grains according to the planting and harvesting times in states that account for the majority of the production of each crop as reported in USDA (1997). Specifically, for corn and soybeans, the beginning of the season corresponds to the date when approximately 50% to 75% of the U.S. crop acreage is likely to be planted and the end of the season corresponds to the date when approximately 75% of those acres are likely to have been harvested. The corn season is five and a half months long, spanning the summer. The soybean season is four weeks shorter than the corn season, with soybeans planted two weeks later and harvested two weeks earlier than corn.

Defining the wheat season is more complicated because there are two growing seasons in the U.S., winter and spring. Winter wheat accounts for about two-thirds of total U.S. wheat acreage and is planted in the fall, usually by early- to mid-October. Depending on conditions, this wheat might sprout in the fall and be grazed before going dormant for the winter. In the spring, this wheat sprouts sometime after mid-March and is harvested in June and July. Yields are sensitive to weather conditions (particularly moisture) during the entire growing period, but they are much less sensitive during the dormant months. For spring wheat, 50% to 75% of the crop is typically planted by mid-May and 75% is harvested by the end of August. To capture the weather-sensitive periods for both wheat crops, we define the season as starting in mid-March when the winter wheat sprouts and running through the end of August when the spring wheat is harvested.²²

²² Several varieties of wheat are accepted as deliverable grades against the CBOT wheat futures contract including No. 2 Hard Red Winter Wheat, No. 2 Dark North Spring Wheat, and No. 2 Northern Spring Wheat.

B. FCOJ

The supply of oranges for FCOJ is sensitive to weather-related growing conditions throughout the year. Since about 95% of the oranges produced in the U.S. for FCOJ are grown around Orlando, Florida, supply is particularly sensitive to freezing temperatures in this area. However, defining the weather-sensitive season is complicated by at least three factors. First, while winter freezes pose the greatest threat to FCOJ production in the U.S., the orange trees, buds, and blossoms are sensitive to general weather conditions throughout the year. This is in contrast to the more concentrated period of weather sensitivity during the shorter growing seasons for the grains. Second, demand fluctuations may be more important for orange juice, a more discretionary item, than for staples such as grains. Third, the impact of (relatively weather-insensitive) Brazilian production may cause specific weather conditions to have less effect on supply and, hence, prices.²³

Nonetheless, since the potential impact of freezes is the source of weather sensitivity most commonly studied in the literature, we define the weather-sensitive season for FCOJ to encompass the months when damaging freezes around Orlando are most likely to occur. Attaway (1997) documents the important freezes in Florida from 1835 to 1998. All of these freezes occurred in December, January, or February. Hence, we define these three months as the weather-sensitive season for FCOJ.

C. Natural Gas

Our definition of the weather-sensitive season for natural gas is based on two considerations. The first is identifying the months in which natural gas is withdrawn from storage and therefore supply is constrained.²⁴ Using data reported in the EIA's *Natural Gas Monthly* (various issues), we determine that natural gas withdrawals normally occur from November through March. The second consideration is identifying when there is a high probability of extreme temperatures that can cause unexpectedly high withdrawals. Using data reported by the National Climate Data Center in *Surface Data, Monthly*, we identify the months with average temperatures below freezing for the states whose demand accounts for the majority of natural gas storage withdrawals. We define these months, December through February, as our weather-sensitive season for natural gas.

REFERENCES

- Amihud, Yakov, and Haim Mendelson, 1987, Trading mechanisms and stock returns: An empirical investigation, *Journal of Finance* 42, 533–553.
- Andersen, Torben G., Tim Bollerslev, and Ashish Das, 2001, Variance-ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns, *Journal of Finance* 56, 305–327.

²³ See BRSW (2005) for a full discussion of the FCOJ market.

²⁴ Susmel and Thompson (1997) model monthly natural gas price and storage data using a regime switching model and find seasonality in price and variance related to storage levels.

- Anderson, Ronald W., 1985, Some determinants of the volatility of futures prices, *Journal of Futures Markets* 5, 331–348.
- Attaway, John A., 1997, *A History of Florida Citrus Freezes* (Florida Science Source, Longboat Key, FL).
- Barclay, Michael J., and Terrance Hendershott, 2003, Price discovery and trading after hours, *Review of Financial Studies* 16, 1041–1073.
- Barclay, Michael J., Robert H. Litzenberger, and Jerold B. Warner, 1990, Private information, trading volume, and stock return variances, *Review of Financial Studies* 3, 233–253.
- Bauwens, Luc, Walid Ben Omrane, and Pierre Giot, 2005, News announcements, market activity and volatility in the euro/dollar foreign exchange market, *Journal of International Money and Finance* 24, 1108–1125.
- Black, Fisher, 1986, Noise, *Journal of Finance* 41, 529–543.
- Boudoukh, Jacob, Matthew Richardson, YuQing Shen, and Robert F. Whitelaw, 2005, Do asset prices reflect fundamentals? Freshly squeezed evidence from the OJ market, *Journal of Financial Economics* (forthcoming).
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).
- Commodity Research Bureau, 2004, *CRB Yearbook CD* (Commodity Research Bureau, Chicago, IL).
- Daniel, Kent, David Hirshleifer, and Siew Hong Teoh, 2002, Investor psychology in capital markets: Evidence and policy implications, *Journal of Monetary Economics* 49, 139–209.
- Energy Information Administration, 1995, The value of underground storage in today's natural gas industry (U.S. Department of Energy, Washington DC).
- Energy Information Administration, 2001a, Impact of interruptible natural gas service on Northeast heating oil demand (U.S. Department of Energy, Washington DC).
- Energy Information Administration, 2001b, U.S. natural gas markets: Mid-term prospects for natural gas supply (U.S. Department of Energy, Washington DC).
- Energy Information Administration, 2004, *Annual Energy Review 2003* (U.S. Department of Energy, Washington DC).
- Fama, Eugene F., 1965, The behavior of stock-market prices, *Journal of Business* 38, 34–105.
- FAO Statistical Database, 2004, (Food and Agriculture Organization of the United Nations, New York) (<http://www.fao.org/>).
- Fleming, Jeff, Chris Kirby, and Barbara Ostdiek, 2006, Bootstrap tests of multiple inequality restrictions on variance ratios, *Economic Letters* 91, 343–348.
- French, Kenneth R., and Richard Roll, 1986, Stock return variances: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 5–26.
- George, Thomas J., and Chuan-Yang Hwang, 2001, Information flow and pricing errors: A unified approach to estimation and testing, *Review of Financial Studies* 14, 979–1020.
- Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.
- Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.
- Harvey, Campbell R., and Roger D. Huang, 1991, Volatility in the foreign currency futures market, *Review of Financial Studies* 4, 543–569.
- Hirshleifer, David, 2001, Investor psychology and asset pricing, *Journal of Finance* 56, 1533–1597.
- Hu, Qi, and Gregory Buyanovsky, 2003, Climate effects on corn yield in Missouri, *Journal of Applied Meteorology* 42, 1626–1635.
- International Energy Agency, 2002, *Flexibility in Natural Gas Supply and Demand* (OECD, Paris).
- Ito, Takatoshi, and Wen-Ling Lin, 1992, Lunch break and intraday volatility of stock returns: An hourly data analysis of Tokyo and New York stock markets, *Economics Letters* 39, 85–90.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson, 1994, Information, trading, and volatility, *Journal of Financial Economics* 36, 127–154.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.

- Mussa, Michael, 2000, The impact of higher oil prices in the global economy, Working paper, International Monetary Fund.
- National Weather Service, 2002, *NOAA Backgrounder: National Weather Service* (<http://www.publicaffairs.noaa.gov/grounders/pdf/nws.pdf>).
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Oldfield, George S., and Richard J. Rogalski, 1980, A theory of common stock returns over trading and non-trading periods, *Journal of Finance* 35, 729–751.
- Perlman, Michael D., 1969, One-sided testing problems in multivariate analysis, *Annals of Mathematical Statistics* 40, 549–567.
- Politis, Dimitris N., and Joseph P. Romano, 1994, The stationary bootstrap, *Journal of the American Statistical Association* 89, 1303–1313.
- Richardson, Matthew, and Tom Smith, 1991, Tests of financial models in the presence of overlapping observations, *Review of Financial Studies* 4, 227–254.
- Roll, Richard, 1984, Orange juice and weather, *American Economic Review* 74, 861–880.
- Shiller, Robert J., 2003, From efficient markets theory to behavioral finance, *Journal of Economic Perspectives* 17, 83–104.
- Shleifer, Andrei, 2000, *Clarendon Lectures: Inefficient Markets* (Oxford University Press, Oxford).
- Smith, Aaron, 2005, Partially overlapping time series: A new model for volatility dynamics in commodity futures, *Journal of Applied Econometrics* 20, 405–422.
- Solberg, Craig, 1999, Weather markets: Mother nature's role in commodity trading, *CRB Trader* 8:4 (<http://www.crbtrader.com/trader/>).
- Stevens, Stanley C., 1991, Evidence for a weather persistence effect on the corn, wheat, and soybean growing season price dynamics, *Journal of Futures Markets* 11, 81–88.
- Stoll, Hans R., and Robert E. Whaley, 1990, Stock market structure and volatility, *Review of Financial Studies* 3, 37–71.
- Sturm, Fletcher J., 1997, *Trading Natural Gas* (PennWell, Tulsa, OK).
- Susmel, Raul, and Andrew Thompson, 1997, Volatility, storage and convenience: Evidence from natural gas markets, *Journal of Futures Markets* 17, 17–43.
- United States Department of Agriculture, 1997, *USDA Agricultural Handbook Number 628: Usual Planting and Harvesting Dates for U.S. Field Crops*.
- United States Department of Agriculture, 2004, *USDA-NASS Agricultural Statistics Data Base* (<http://www.usda.gov/nass/>).
- Wilks, Daniel S., 1995, *Statistical Methods in the Atmospheric Sciences* (Academic Press, San Diego, CA).
- Wolak, Frank A., 1987, An exact test for multiple inequality and equality constraints in the linear regression model, *Journal of the American Statistical Association* 82, 782–793.
- Wolak, Frank A., 1989, Testing inequality constraints in linear econometric models, *Journal of Econometrics* 41, 205–235.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.