Testing for Market Rationality: Lessons from the FCOJ Market

Jacob Boudoukh,[a] Matthew Richardson,[b] YuQing Shen[c] and Robert F. Whitelaw[b]*

April 22, 2002

*[a]Stern School of Business, New York University, IDC and NBER; [b]Stern School of Business, New York University and NBER; and [c]J.P.Morgan. We would like to thank Bill Silber and seminar participants at the Federal Reserve Board of Governors, Washington D.C., Princeton University, and New York University for their valuable comments. Contact: Prof. R. Whitelaw, NYU, Stern School of Business, 44 W. 4th St., Suite 9-190, New York, NY 10012, (212) 998-0338, rwhitela@stern.nyu.edu.

Testing for Market Rationality: Lessons from the FCOJ Market

Abstract

This paper questions the existing stylized fact that frozen concentrated orange juice (FCOJ) futures returns are not related to fundamentals, in particular, temperature. We show that when theory clearly identifies the fundamental, i.e., at temperatures close to or below freezing, there is a close link between FCOJ prices and that fundamental. Using a simple theoretical nonlinear model of the relation between FCOJ returns and temperature, we can explain approximately 50% of the return variation. Moreover, these R^2s can increase to almost 65% if the model is extended to include market information about temperature forecasts. This is important because while only 4.5% of the days in winter coincide with freezing temperatures, over two-thirds of the entire winter return variability occurs on these days. We also show that other fundamental information, such as USDA production forecasts, generate significant return variation that is consistent with theoretical predictions. These results cast doubts on other claims about irrationality and the lack of response to information in markets for much more complicated assets, most notably equities.

1 Introduction

There has been a flurry of research that investigates whether asset returns reflect their fundamental values.[1] The general conclusion from this research is that the variability of these returns is much greater than that implied by the assets' fundamentals, i.e., information about the asset's cash flows and expected returns. For example, regressions of asset returns on *ex post* information (relevant for pricing the assets) reveals little explanatory power.

There are three possible explanations for *estimated* excess volatility: (i) market prices are not completely rational (i.e., they do not reflect fundamentals), (ii) fundamental information is incorrectly measured or not measured at all (e.g., omitted variables), or (iii) the underlying model of asset price formation is false (e.g., the wrong functional form).[2] It has become popular as of late to focus on point (i) as the most reasonable explanation. In response, market efficiency proponents, such as Fama (1998), have questioned the way behavioral-based financial economists have measured fundamentals, such as discount rates (i.e., by ignoring risk factors) and cash flows (i.e., economic versus accounting earnings). Furthermore, the research designs of some of these previous studies have been put into question.[3]

Perhaps the strongest, and one of the forerunners, of the excess volatility literature is Roll's (1984) seminal paper on frozen concentrate orange juice (FCOJ) futures prices and weather. His paper is unique and clever in that the fundamental information appears to be identifiable and exogenous. A critical aspect of this market is that the relevant oranges for FCOJ are produced in a relatively small region of Florida. Thus, as Roll (1984) reasonably argues, the weather in this region should be a primary determinant of supply and thus of futures prices. While he does report a statistically significant relation, he finds little explanatory power in the relation between FCOJ

[1] For example, see Roley and Troll (1983), Urich and Wachtell (1984), and Balduzzi, Elton and Green (1997) in fixed income markets; Shiller (1981), Campbell and Shiller (1984), Black (1986), Roll (1988), Cutler, Poterba and Summers (1989), and Mitchell and Mulherin (1994) in equity markets; Frankel and Meese (1987), Ito and Roley (1987), and Ederington and Lee (1993) in foreign exchange markets; and French, Leftwich and Uhrig (1989), Fortenberry and Sumner (1993), and Bauer and Orazem (1994) in commodity markets.

[2] A fourth explanation could be that the data itself is mismeasured. For example, asset prices reflect a bid-ask spread, which, if large enough, could lead to erroneous conclusions about the asset return's volatility.

[3] For example, consider the debate generated by Kleidon (1986) and Marsh and Merton (1986) in regard to Shiller's (1981) excess volatility study, or, more recently, Fama's (1998) criticism of so-called anomalies in the corporate event study literature.

futures prices and weather. In his conclusion, Roll (1984, p.879) states that

> ...weather surprises explain only a small fraction of the observed variability in futures prices. The importance of weather is confirmed by the fact that it is the most frequent topic of stories concerning oranges in the financial press and by the ancillary fact that other topics are associated with even less price variability than is weather... There is a large amount of inexplicable price volatility.

Roll's study, and for that matter the majority of the aforementioned excess volatility studies, are performed in an environment with little economic structure imposed. Instead, researchers impose strong econometric structure, usually linearity, with respect to the relation between the asset returns and their fundamentals. Is this a reasonable approach? For many applications, the answer is most likely no. While a number of economic theories imply a linear relation between returns and aggregate factors, there is no reason to believe in general that assets respond linearly to information. For example, bond prices are nonlinear in discount rates, stock prices respond nonlinearly to earnings surprises, and spreads are nonlinearly related to trading volume. Linearity is, broadly speaking, an uncommon phenomenon. Consequently, this paper investigates nonlinear (and nonparametric) models of asset returns and their implications for fundamental valuation.

We take as our example Roll's (1984) study of FCOJ and weather, partly because it serves as a poster child for market irrationality and partly because the fundamental information is observable. For example, in a recent survey of behavioral finance, Hirshleifer (2001, p.1560) writes in a section covering mispricing effects that "little of stock price or orange futures price variability has been explained empirically by relevant public news."

Several important facts emerge from our analysis. First, the majority of the volatility of FCOJ futures returns occurs either around freezing temperatures or on days in which the United States Department of Agriculture (USDA) provides forecasts of orange production in Orlando. While these days represent less than 2% of the total days in our sample, they represent 40% of the total variance. This fact alone suggests that the market does not ignore fundamental information about production (either through weather's effect or the USDA forecast). Second, when we are confident that we can identify the fundamental information from theory (i.e., around freezing temperatures), we show that the weather has a highly nonlinear, and substantial effect on futures prices. For example, R^2s jump from 5% in the linear case to close to 50% in the nonlinear case, even ignoring information about weather forecasts. In fact, whenever we are sure we know the fundamental information, much of the price variability can be explained by these fundamentals. In our opinion, this empirical fact provides strong evidence that markets take into account this information for pricing. Third, though there is a remaining baseline volatility that our temperature-based analysis cannot explain (e.g., volatility in the spring and summer months), we provide anecdotal evidence that other supply factors (e.g., Brazilian production of FCOJ) and demand factors (e.g., shifts in demand) are important. Our analysis serves to illustrate the more general point that excess volatility estimates may have more to do with the researcher's ability to capture fundamentals than some view on market irrationality.

The paper is organized as follows. In Section 2, we describe the data and the basic stylized facts about weather, production and FCOJ futures prices. This section provides strong circumstantial evidence that these three variables are linked. In Section 3, we provide the theory for how FCOJ prices should be related to the weather. We then go on to test the theory, and, in general, confirm the theoretical relation empirically. In particular, there is close relation between FCOJ prices and temperature when the theory suggests there should be one. Section 4 explores the fact that there is remaining volatility in FCOJ prices not explained by the weather. In particular, we explore the effect of other factors on FCOJ prices. Section 5 concludes.

2 Data

Futures contracts in frozen concentrated orange juice (FCOJ) have traded on the New York Cotton Exchange since September of 1967. At any given time, there are usually nine to eleven contracts outstanding with expiration schedules every second month (January, March, May, etc.) and with at least two January months listed at all times. The contract is for 15,000 pounds of frozen concentrated orange juice, which represents about 2,400 ninety-pound field boxes of oranges, with specific requirements for color, favor, and defects. Due to these requirements, 95% of the total U.S. processed orange production takes place in central Florida in and around the Orlando area.[4]

[4]California is also a large producer of oranges; however, oranges produced in that region are not suitable for FCOJ. Internationally, Brazil is, along with the U.S., the largest producer of oranges used for FCOJ. While Brazil primarily exports these oranges to countries other than the U.S., the majority of Florida production is consumed domestically. Though Brazilian oranges are subject to significant tariffs, they can provide a close substitute to Florida oranges.

There are two types of oranges produced in the Orlando area with their main distinction being the harvest period, namely early and mid-season (EM) varietals and Valencia oranges, which get harvested from November through March (though primarily in January) and April though June, respectively.

This geographic concentration is highly unusual for agricultural commodities, and allows for a unique opportunity to study the interaction between asset returns (i.e., FCOJ futures prices) and an important exogenous variable (i.e., Orlando weather). *Ex ante* there are strong reasons why weather should be an important variable for FCOJ prices and why Roll's (1984) result is such a puzzle. In particular, freezes can have devastating effects on orange production. Attaway (1997) cites several freezes over the past 170 years that have had a major impact on fruit production. A full list of the impact freezes is provided in Table 1. For example, in the freeze of February 7-9, 1835, temperatures dropped to as low as 11 degrees in northern Florida, destroying almost all existing orange trees. In fact, prior to this date, a number of more northern states (such as South Carolina and Georgia) produced oranges. The freeze of 1835 essentially forever ended the desire to produce oranges in those states. In the decade from 1894 to 1905, Florida was hit by ten freezes, one of which (February 8, 1895 when the temperature reached 17 degrees) led to almost all the orange crop being destroyed. In 1893-94, 5.06 million boxes of oranges were produced; by 1895-96, production had fallen to 0.147 million boxes (Attaway (1997)). The industry did not fully recovery until 1909-10. Of interest to this paper, Table 1 shows that a number of freezes occurred during our sample period, that is, post 1967.

We collected data from September 1967 to December 1998 on several series: (i) the minimum temperature in Orlando, (ii) a limited amount of minimum temperature forecasts from the National Weather Service (primarily for winter months and for only some of the years within our sample), (iii) USDA production forecasts, and (iv) FCOJ futures prices, volume and open interest on every contract. We also collected other relevant data series such as Brazilian orange production and FCOJ exports to the U.S., FCOJ pack and movement information, and news events associated with FCOJ. We describe the majority of the data in detail as it becomes relevant throughout the paper.

The issue of Brazil's production and how it affects the FCOJ market is discussed in Section 4.2.

2.1 Orlando Weather

The most important fundamental for FCOJ prices is the presence of a production shock to a given year's orange crop. The largest factor in this regard is weather, and, in particular, winter freezes. As mentioned above, winter freezes can severely impact orange production, thus raising the prices of oranges in the commodity market.

Figure 1 graphs the minimum temperature in Orlando, Florida from 1967 through the end of 1998. Over 80% of the winter years include days (and, in many cases, multiple days) with temperatures close to or below freezing during periods relevant for orange juice production. Specifically, within our sample period of thirty-one winters, only six winters can be strictly described as completely freeze-free winters. Twenty-five winters experienced one or more nights when the minimum temperature dipped to 32 degrees or below. Approximately 50% of the winters have two to three nights of freeze, and the probability of four or more freeze nights is approximately 25%. Thus, there is clearly a considerable amount of return-relevant information during the winter season.

However, it is important to point out that not every freeze has material consequences for the orange crop. As mentioned above, Attaway (1997) provides a detailed description of the history of impact freezes that happened in Florida over the past 170 years (see Table 1). While Table 1 documents the actual dates of these freezes, Table 2 describes the critical conditions for freezes to have a substantial effect on production. These conditions are basically the temperature itself and the duration of this temperature level.[5] For example, if the temperature falls below 28 degrees for six hours or more, then, without intervention, there will be extensive damage done to the oranges, and hence production of FCOJ. With respect to our particular sample period, according to Attaway (1997), there are eleven relevant freeze seasons, some of which are multiple, from 1967 to 1998.[6] The identification of these freezes will form the basis for our comparison between freeze and non-freeze years.

[5] Of course, the magnitude of a freeze alone is not sufficient to explain the impact on prices. Also important are coincidental effects on the determinants of the supply of Florida FCOJ substitutes such as imported FCOJ from Brazil. Examples of these effects are whether there was a drought in the Sao Paolo area of Brazil, tariff rates, and the costs of transportation.

[6] In a separate study by Hebert (1993), similar freezes are also identified using different freeze identification methodology.

2.2 Production

The documentation of severe freezes in Table 1 and the theory of freeze impacts in Table 2 suggest that production of FCOJ should be impacted by the weather. Every October, the United States Department of Agriculture (USDA) provides forecasts of the upcoming season's orange production. This forecast is conditioned on a non-freeze season, that is, the USDA does not take into account the possibility of a freeze when forming its forecast.

Table 3 documents the percentage difference between the October USDA forecast and final production, which is generally reported in August of the following year, for all the years in our 30-year sample. Freeze years, as classified by Attaway (1997), are noted, and we also report the average across the 11 freeze years and 19 non-freeze years. Several observations are of interest. First, the USDA's forecasts in non-freeze years are nearly unbiased, with only a 1.3% difference on average between the forecast and the realization. However, there is substantial variation in forecast accuracy, with underestimates and overestimates as large as 6.2% and 8.1%, respectively. Second, and most important, there is a considerable reduction in orange production in freeze years, with 12.7% less production than forecast. Again, there is substantial variation, with a maximum decline of 30.5%. If the demand for oranges is downward sloping and there is no perfect substitute for Florida oranges (at the prevailing price), then the spot price of oranges should rise in response to a decline in supply. Given the economic laws of supply and demand, therefore, futures prices on FCOJ should move in the same direction.

2.3 Futures Prices

In order to investigate the effect of freezes on prices, we obtained daily closing prices of FCOJ futures from September 1967 to August 1998. On each trading day 3 prices of the three near maturity contracts were collected from Bridge/CRB and Datastream. We also collected volume and open interest data. There are a number of ways to generate a continuous return series from the three price series. In general, we attempt to get the most accurate proxy for price changes in spot FCOJ prices. We need to take into consideration two important characteristics of futures' prices that may affect the accuracy of calculated returns - liquidity and limit days. First, it is well known that near maturity contracts are the most liquid, a fact that we verified with the volume and open interest data. However, liquidity diminishes rapidly as a contract comes close to expiration, particularly in the expiration month. These factors were considered in constructing the continuous return series. A second issue is that price moves on futures contracts are limited by the exchange. The magnitude of these price limits and the contracts to which they applied changed over the course of our sample, but the general impact of price limits is to prevent prices from fully incorporating information on days with important information releases. We adjust for this phenomenon in the standard way by aggregating returns over consecutive days with price moves that hit either the lower or upper limit. For example, when a severe freeze occurs, the price may hit the upper limit for 4 consecutive days. In this case, we aggregate returns over these 4 days plus the next day, which is a no-limit day, and use this as the 1-day return. The precise algorithm we use to construct the spliced data series is described in Appendix 1.

Given this continuous daily return series, Table 3 provides the cumulative return over the winter season (December through February) for each year and the average of these returns for freeze and non-freeze years. The average return is almost 19% higher in freeze than in non-freeze years, namely 12.7% versus -6.1%, although there is substantial variation in returns on an annual basis. This suggests that, on average, futures prices are indeed higher in years in which production has been impacted negatively.

It is also worthwhile documenting some additional stylized facts that provide substantial circumstantial evidence of an important relation between FCOJ futures prices and temperature. Table 4, Panel A reports the volatility of FCOJ returns in different seasons. The variability (variance) of FCOJ futures returns in winter months is four times that of returns in spring and summer months and three times higher than in the fall season. This should not be surprising as weather, in particular freezes, has its only impact in these winter months. We also divided the winter season into two separate periods: pre-freeze and post-freeze. The pre-freeze period includes days up to an including the first freeze of the season in a freeze season and all the days in a non-freeze season; the post-freeze period includes the days subsequent to a freeze in a freeze season. Note that the post-freeze period can include a second freeze, and, in fact, two seasons had multiple freezes (see Table 1). Interestingly, the variability in winter months is greater pre-freeze than post-freeze in spite of the fact that freeze frequencies are similar in the two periods. Of course, this phenomenon may be due to the fact that the second freeze in a season happened to be less severe in our sample. Alternatively, the second freeze in a season may be less important because damage has already been done to the orange crop. Due to uncertainty surrounding the impact of multiple freezes, the remainder of the paper focuses on the period up to and including the first freeze of the season, i.e., the pre-freeze period.

The fact that winter months produce greater variation in futures prices does not in itself suggest a relation between temperature and FCOJ prices. As shown in Tables 1 and 2, the main impact on production comes from a particular event, namely a freeze. Thus, it is the change in the likelihood of a freeze that should move FCOJ futures prices. Table 4, Panel B presents the volatility of FCOJ futures prices in winter months conditional on various contemporaneous minimum temperature realizations, namely below 36 degrees, 36-40 degrees, 41-45 degrees and above 45 degrees. As can be seen from the table, at or around freezing temperatures (i.e., 35 and below), the daily volatility of FCOJ returns is 11.81% compared to between 1.79% and 2.16% for the other temperature buckets. That is, there is over 50 times greater variability (variance) in FCOJ futures returns near freezing temperatures than when temperatures are warmer. This result strongly suggests that freezes have a substantial impact of FCOJ futures pricing.

While there are many factors that can affect the FCOJ futures price, Tables 3 and 4 show that (i) much of the variability in FCOJ prices occurs around freezes, and (ii) the movement in FCOJ prices seems to be in the right direction as characterized by the laws of supply and demand. In the next section, we provide the theory and empirical results that demonstrate a strong relation between FCOJ futures prices and temperature. We also show why the literature has mistakenly concluded no relation exists.

3 FCOJ Prices and Temperature: Theory and Results

Because storage costs of oranges and, in particular, FCOJ are high, only a fraction of the juice gets carried forward from year-to-year. Therefore, long-run shifts in either demand or supply, are expected to have little impact on FCOJ prices. What are the major factors affecting FCOJ prices? The agricultural literature has identified three major determinants of prices. First, as mentioned above, short-run shocks to the production of oranges shift the supply curve and move prices. Second, due to there being worldwide production of oranges suitable for FCOJ, the ability of distributors to shift away from Florida oranges to other sources also affects prices. In particular, Brazil is the largest exporter of oranges worldwide. Given tariffs in the U.S. on these oranges, as well as transportation costs, the prices of Brazilian oranges can have a major impact on the available supply in the U.S. The U.S. also imports orange juice from Mexico and the Caribbean, and these countries, in general, face lower tariffs. However, their production is substantially lower than that of Brazil. Third, there could be short-run demand shifts either to or away from FCOJ to other citrus products, such as grapefruit juice, lemonade or freshly squeezed orange juice, or to other fruit juices such as apple juice.

If these are the major factors affecting the demand and supply of FCOJ, and, therefore, the underlying price, the question remains whether these factors get rationally incorporated into the day-to-day movement in the prices of financial assets, i.e., futures prices on FCOJ. As described in the introduction, there is a growing literature that argues asset prices have excess volatility due to irrational behavior on the part of economic agents. While some of this excess volatility could be due to typical market-microstructure biases, the argument is that the volatility is too large to justify this possibility. In fact, Roll's (1984) work on orange juice and weather makes the twofold observation that (i) there is considerable volatility in the FCOJ futures market, and (ii) this volatility cannot be explained by the obvious fundamental factors.

In this section, we focus on what economic theory would tell us about the relation between FCOJ futures price changes and supply shocks due to temperature (freezes), while we address other potential factors in Section 4. As a starting point, the stylized facts of Section 2 are quite informative. First, over 80% of the years have freezing weather in Orlando, Florida (Figure 1). Moreover, 40% of these years are considered freeze years (Table 1); that is, years in which the temperature was low enough for a long enough time to be relevant for orange production. Second, adjusting for expectations about production, there is a 14% difference in production levels for freeze versus non-freeze years (Table 3). Consistent with agricultural theory, this result confirms that freezes have substantial impacts on U.S.-based production. Third, with respect to financial markets, the distribution of futures returns differs in freeze versus non-freeze years. Specifically, returns are on average 19% higher in freeze years (Table 3). In contrast to the finance literature's view of FCOJ futures returns and weather, this is consistent with financial markets relating asset prices to the fundamentals, i.e., futures prices are much higher in years with negative shocks to production. Fourth, the majority of the well-documented variation in futures returns occurs in the temperature region one would expect, namely around freezing temperatures (Table 4B). For example, while only 5.4% of the days during winter are in this region, 64% of all of the variation (variance) in futures returns occurs on these days.

Given these four facts, it seems surprising then that one of the major results in the excess volatility literature is the lack of a relation between FCOJ futures and temperature. In a rational market the FCOJ futures return over a given period should reflect changes in expectations about the level of short-run production, i.e., the supply shocks over that period. With respect to temperature, these shocks correspond to changes in the market's perceived likelihood and severity of a freeze.

Perhaps the finance literature has found little or no relation because the change in likelihood of a freeze is a difficult variable for the econometrician to measure. Clearly, there is no theoretical reason to believe that the relation between FCOJ futures returns and temperature is linear. In fact, there is considerable reason to believe this is not the case. Specifically, consider the following points:

- The change in likelihood of a freeze clearly depends very differently on a number of factors. First, and foremost, a freeze is going to be seasonal as fall, spring and summer months are unlikely to produce freezing weather in Florida. Second, a temperature surprise, or change in temperature, will only matter around freezing temperatures. That is, a 10-degree surprise at 70 degrees has very different implications for the likelihood and severity of a freeze than it does at a temperature of 35 degrees. Even during winter, at most temperature levels, theory tells us that there is no relation between FCOJ prices and temperature. The relation exists only to the extent we learn new information about the likelihood and severity of a freeze.

- The distribution of orange production should differ depending on whether a freeze has occurred or not. For example, if a freeze has a substantial impact on orange production, then another freeze may be expected to have less impact to the extent the damage is already done. Of course, this depends on the relative severity of the freezes, the timing of the harvests, etc. In addition, conditional on a freeze, the uncertainty regarding its effects on production can be quite high until the damage is surveyed.

- The severity of a freeze depends on two weather factors: (i) the level of temperature, and (ii) the duration of a freeze. Table 2 provides a brief description of these relevant factors. In particular, the severity of a freeze is not linear in temperature, as 30-31 degree temperatures have relatively mild effects compared to the 27-29 degree range. In the latter instance, there is severe damage done to both the leaves and fruit, which dramatically affects production.

In the next subsection, we try and relate these theories more closely to the data by explicitly incorporating some of these ideas into an empirical analysis of futures prices and temperature.

3.1 A First Look

While it is difficult for the econometrician to measure changes in the market's perceived likelihood and severity of a freeze, there is one instance in which this is straightforward to do, that is, when a freeze actually occurs. In this case, the likelihood is probability one and the severity of the freeze maps to the temperature level. Even in this simplified case, the temperature level is just one of two pieces of information. The other piece is the length of time at each temperature level, which we discuss in the next subsection. Of course, market participants have access to greater information such as verbal assessments of the probable freeze damage.

If one assumes that the conditional probability of a freeze is constant, then the realization of a freeze and its corresponding temperature level will measure the supply shock (albeit functionally). There are strong reasons to question the assumption that the probability of a freeze is constant. For example, at the very least, the probability will vary with the time of the year. However, more important, there is information provided to the marketplace about the likelihood and severity of freezes throughout the year. In particular, the US National Weather Service provides 12, 24 and 72 hour forecasts that the market should incorporate in its assessment. These issues are addressed in the next subsection in some detail.

The functional form of the return-temperature relation is unknown, but is certainly nonlinear. We assume that such a relation exists, and attempt to characterize the function empirically. We investigate four different specifications for this relation.

The first model is linear:

$$R_t = \alpha + \beta W_t + \epsilon_t,$$

where R_t is the close-to-close return and W_t is the contemporaneous realized minimum temperature. Thus, the return and the minimum temperature are aligned in time. We consider this the benchmark model. The clear problem with the linear model is that it imposes the same relation between returns and temperature at temperatures both above and below freezing.

In contrast, an alternative specification, the piecewise linear model, estimates a number of linear relations,

$$R_t = \alpha_n + \beta_n W_t + \epsilon_t,$$

each for a different temperature bucket $W_t \in (W^n_{min},\ W^n_{max})$. Each return-temperature pair is placed in a temperature bucket and is used in estimation of the relevant coefficients α_n and β_n. In theory, this model provides complete flexibility. Given enough data the econometrician can specify a large number of small buckets, effectively estimating the functional form of the relationship piece-by-piece. Note that we do not impose the restriction that the lines connect at the bucket borders.

The most general model estimates an unspecified functional form linking returns to temperature:

$$R_t = f(W_t) + \epsilon_t.$$

In this paper, we estimate this functional form using the kernel regression methodology described in Ullah (1988) and others. The kernel estimation provides a standard nonparametric regression of returns on realized temperatures based on nonparametric density estimation. The econometrician chooses the kernel function and the bandwidth or degree of smoothing. The results are insensitive to the choice of kernel function and we use a normal density. The bandwidth is chosen via cross-validation. Specifically, we estimate the fitted returns at each temperature observation in the data using all the observations except that specific point and choose the bandwidth that minimizes the resulting mean squared error.

Finally, the fourth model imposes priors on the functional form. This nonlinear model is in the spirit of Roll (1984), i.e., using the relevant theory in a parsimonious way. The specification,

$$R_t = \alpha + \beta_1 Max(0,\ W^* - W_t) + \beta_2[(Max(0,\ W^* - W_t))]^2 + \epsilon_t,$$

is a second order polynomial in the regressor $Max(0,\ W^* - W_t)$, where W^* represents the critical temperature threshold. In our analysis, we use 32 degrees as this threshold, although the results are insensitive to any reasonable choice. To the extent FCOJ prices are not affected by non-freezing temperatures, the regressor is zero. Then, as the temperature gets below a certain level, the regressor rises quickly as the temperature falls. Thus, this model incorporates the nonlinearity implied by the theory of freeze impacts described in Table 2.

Table 5 presents the results for these four different specifications. The results are in stark contrast to the existing literature's view of the relation between FCOJ futures prices and temperature (e.g., Hirshleifer (2001)). First, while temperature is statistically significant in the linear model, it is clear that the evidence derives almost entirely from low temperatures. The piecewise linear models show that weather only matters at low temperatures. Of course, this is precisely what the theory would predict. That is, in terms of the weather, it is only information about a freeze, in this case, the realization of one, that matters.

Second, the R^2s jump from 1.5% for the linear model to as high as 33.3% and 33.6% for the nonlinear model and kernel nonparametric models, respectively. Thus, across the entire pre-freeze winter period, over one-third of the variation can now be explained. To get a finer partition of the explanatory power of weather, Figure 2 breaks up the R^2s into different buckets. This is important because the theory suggests that weather should only matter at low temperatures. The results confirm this theoretical prediction, with the more general models capturing close to 50% of the return variation. In fact, outside of the low temperature ranges, there is almost no relation between weather and FCOJ futures returns. Note that the linear model produces negative R^2s for the three higher temperature ranges. A negative R^2 is possible because we estimate the model on all the data and then calculate the R^2 for specific groups of observations sorted by temperature. In theory, the other globally estimated models (the nonlinear and nonparametric models) could also produce negative R^2s in some regions, but the actual R^2s are positive but close to zero.

These results show that researchers need to be cautious in how they interpret the relation between asset prices and fundamentals. In this case, the theory predicts a strong relation only under certain conditions, i.e., freezing temperatures. In fact, when those conditions are met, the R^2s are very high. These R^2s are all the more remarkable given the simple nature of the models, which ignore both market forecasts and specific information about the freeze, such as its duration. The existing literature concludes that FCOJ futures returns are unrelated to news about the weather because it implicitly focuses on the 99% of observations that by theory predict no relation. The aggregation of these into one simple linear model across all observations in all the months drowns out the strong theoretical relation between FCOJ prices and temperature.

Third, the theory predicts that the relation between FCOJ futures returns and the temperature should be nonlinear. As the temperature drops further below freezing, the impact on production (duration of the freeze aside) should be more severe. The regression results of Table 5 strongly support the nonlinear nature of the theoretical relation. The coefficient on the nonlinear term in the polynomial regression is highly significant with a t-statistic of approximately 7. Figure 3 graphs the fitted relations of all four models, and both the kernel regression and nonlinear specification show a highly convex relation between FCOJ returns and freezing temperature levels. The piecewise linear model captures the convexity between freezing and above freezing temperature but fails to capture the convexity within the low temperature region. Further subdividing the low temperature range may permit a better fit, i.e., less misspecification, but estimation error would also increase.

3.2 Another Look

The results of Section 3.1 are in stark contrast to the existing literature. That is, when theory expects a strong relation between FCOJ futures returns and weather, the data strongly supports it. One should note, however, that the R^2s are on the order of 50%. This suggests there is substantial variation left to be explained. Of course, the empirical models in Section 3.1 ignore potentially important information and assume that the actual temperature describes the entire shift in the likelihood of there being a freeze and its magnitude. While the true likelihood will be difficult for the researcher to uncover, it is possible to gather more evidence on the true functional form. In particular, beyond the actual realization of a freeze, what other information might be helpful in describing the shift in the freeze likelihood?

3.2.1 Forecasts

The most obvious missing factor is information regarding the market's forecast of a freeze. For example, if the market knew with probability one that there would be a freeze the next day, then the actual realization would have no information. That is, all the price movement would occur the day before. Thus, the model of Section 3.1, which ignores forecasts, is a lower bound on the true model's ability to explain the relation between FCOJ futures returns and temperature.

There are two primary ways forecasts can impact FCOJ futures returns. First, as described in the extreme example above, forecasts measure the current expectation, though not the entire distribution, of the future temperature and thus contain information about a freeze. Thus, FCOJ futures returns might move prior to the realization of a freeze because the market's likelihood of a freeze has already incorporated the forecast. Second, the actual realization of the temperature might be a surprise relative to the forecast, which also signals a shift in the likelihood (and magnitude) of a freeze. For example, consider the case in which FCOJ returns are negative when we get positive surprises in the weather, that is, a freeze was forecasted but one did not materialize. This is a case that we do not even try to explain in Section 3.1. Alternatively, there are cases when the freeze was either worse or better than expected relative to the forecast. This too in theory would cause a shift in the distribution of the freeze magnitude and thus FCOJ returns.

Forecasts are only relevant to the extent that the National Weather Service has some forecasting power. Roll (1984) provides a detailed examination of the NWS's forecasting ability, and, in general, finds they produce unbiased forecasts with R^2s on the order of 50%. We duplicate Roll's tests for a limited 10-year sample of National Weather Service (NWS) temperature forecasts of the next-day minimum temperature for the Orlando area and also look at a time series approach to temperature forecasting. We estimate two basic types of regressions, one including lagged minimum temperatures and forecasts of the form

$$W_t = \alpha + \beta_1 W_{t-1} + \beta_2 W_{t-2} + \gamma F_{t-1}(W_t) + \epsilon_t$$

and one including lagged temperatures and monthly dummy variables of the form

$$W_t = \alpha_i D_{it} + \beta_1 W_{t-1} + \beta_2 W_{t-2} + \epsilon_t,$$

where $F_{t-1}(W_t)$ is the one-day ahead forecast and D_{it} are monthly dummies. We only consider two temperature lags because further lags add little or nothing to the explanatory power of the regressions. Also, we do not include monthly dummies in the forecast regressions because again they do not significantly improve the fit. The estimation results are reported in Table 6. It is important to note that the pure time series regressions are run only on winter season data and that the forecast regressions are for November through March. Variability is much higher in the winter months and forecastability is much lower.

There are several interesting results from the regressions. First, the NWS forecast is not conditionally unbiased since the coefficient is significantly different from 1, but this may simply be due to rounding in high temperature regions as will be apparent shortly. However, the forecast is quite powerful with an R^2 of 56%. Second, lagged temperatures appear to provide additional forecasting power, increasing the R^2 to 73%, although again NWS rounding may be an issue. Minimum temperatures are persistent, and a fall (rise) in minimum temperatures over the previous 2 days suggest a continuing downward (upward) trend as indicated by the negative coefficient on the second lag, although this coefficient is marginally significant at best. Third, when forecast data is unavailable, monthly temperature dummies are extremely significant, but persistence and trend continuation are still evident. The R^2 is a meaningful 41%. These latter results are consistent with the time series evidence in Campbell and Diebold (2001), who use a much larger sample to forecast temperatures in 10 U.S. cities (although not Orlando).

Unfortunately, the overall ability of the models to forecast temperature is not that relevant. What matters is their ability to forecast temperatures in the freezing range, since that is the only temperature level that affects production and thus FCOJ pricing. That is, the models may be very efficient at predicting 50-degree nights in Florida, but what market participants care about is their ability to predict 25-degree nights. On this issue the pure time series model is essentially useless. It never predicts freezes due to the strong reversion to a monthly mean that is well above freezing, even in January. Similarly, the time series augmentation of the NWS forecast does not appear to help in predicting freezes. Consequently we focus only on the NWS forecast. Figure 4 graphs forecasts against temperature realizations for the observations in our limited sample of forecasts. The tendency of the NWS to round to an even 5 degrees at high temperature forecasts is clear. However, the important areas are the three quadrants (i) when a freeze is forecast but one did not occur, (ii) when a freeze is forecast and one occurred, and (iii) when a freeze occurred but it was not forecast. Ideally, the observations should lie on the 45 degree line if the NWS always forecasts the temperature perfectly.

While our 10-year sample size is small, some observations are in order. First, in terms of temperature level, there were three severe freezes (i.e., below 25 degrees), but only one severe freeze was predicted. Second, while a number of freezes were predicted which did not take place, only one freeze materialized that was not predicted. Third, there are several observations in the relevant quadrants in which the NWS was far away in terms of temperature. However, the magnitude of this effect really depends on how far their forecast was from freezing. In one instance, the NWS predicted mid 40's and the realization was 27 degrees, but it is not clear that a prediction of mid 30's and the same realization would be any better. In general, the limited data in Figure 4 shows some, though clearly imperfect, ability to predict freeze-level temperatures.

The NWS's forecasting ability aside, the remainder of this section explores how well forecast information explains FCOJ return variability. We use the following methodology to investigate this issue. First, we generate a set of pricing errors from the kernel regression pricing model of Section 2.1. That is, we take the FCOJ futures return minus the fitted value. We then look at whether forecasts can help improve the model, either through forecasting future freezes or through temperature surprises, by regressing the pricing errors on the forecast information. If forecasts are not helpful, then the R^2s and coefficients should be close to zero. If forecasts are helpful, then we could in principle incorporate them in a more complete model of FCOJ returns. However, our limited sample of forecast data prevents us from estimating a multi-dimensional nonlinear model with any accuracy.

Table 7 reports results from regressing the pricing errors on information about forecasts through two different sources: (i) forecasts today of freezing temperatures tomorrow, and (ii) realizations of freezing temperatures tomorrow.[7] We also break up the sample into periods in which either a freeze was forecasted or a freeze occurred. That is, we consider both univariate and multivariate versions of the regression,

$$\epsilon_t = \alpha + \beta_1 Max[0, 32 - F_t(W_{t+1})] + \beta_2 Max[0, 32 - W_{t+1}] + \nu_t,$$

where ϵ_t is the pricing error from the nonparametric regression, $Max[0, 32 - F_t(W_{t+1})]$ is tomorrow's forecast of a freeze, and $Max[0, 32 - W_{t+1}]$ is the realization of a freeze. The table also reports R^2s calculated within buckets, sorted by the forecast or the future realized temperature.

The overall conclusion from the table is that forecasts have significant but not overwhelming explanatory power for returns, i.e., the relation between FCOJ futures returns and temperature is even stronger than implied by Section 3.1. In the regression of returns on contemporaneous forecasts of tomorrow's minimum temperature, the coefficient is significant and has the correct sign (i.e., forecasts of lower temperatures generate higher positive returns), and the R^2 is 17% in the only relevant bucket (i.e., when the forecast is low). When used on its own, the future temperature has limited explanatory power (e.g., 4% in the relevant temperature range for the regression that uses the full set of winter pre-freeze data) but the correct sign. When both the forecast and future

[7] The latter piece of information can be broken up into the forecast and the unexpected component of the temperature. We include this latter variable as a noisy signal of the forecast because we have data for the entire sample.

temperature are combined, the R^2 jumps to 24% in the relevant forecast bucket, although the coefficient on the future temperature reverses sign. This sign reversal is puzzling, but it may be attributable to a combination of multicollinearity between the two variables and the relatively small sample size in the relevant region.

As an alternative, Table 8 investigates whether temperature forecast surprises help explain the kernel regression model's pricing errors, and thus potentially improve the model's fit. Because in theory the temperature surprise has a different effect depending on whether a freeze was forecast and/or realized, we consider error analysis regressions of the following sort:

$$\epsilon_t = \alpha + \beta_1 I_{1,t} Z_t + \beta_2 I_{2,t} Z_t + \beta_3 I_{3,t} Z_t + \beta_4 I_{4,t} Z_t + \nu_t$$

where

$$\begin{aligned} Z_t &= W_t - F_{t-1}(W_t) \\ I_{1,t} &= 1 \text{ for } F_{t-1}(W_t) \leq 32^o \text{ and } W_t > 32^o,\ 0 \text{ otherwise} \\ I_{2,t} &= 1 \text{ for } F_{t-1}(W_t) \leq 32^o \text{ and } W_t \leq 32^o,\ 0 \text{ otherwise} \\ I_{3,t} &= 1 \text{ for } F_{t-1}(W_t) > 32^o \text{ and } W_t \leq 32^o,\ 0 \text{ otherwise} \\ I_{4,t} &= I_{1,t} + I_{2,t} + I_{3,t}. \end{aligned}$$

The R^2s are calculated within buckets, sorted by the forecast or the temperature. Forecast surprises seem to matter under two circumstances–when a freeze was forecast and realized, and when a freeze occurred but was not forecast. The former case generates a negative and significant coefficient (β_2), as predicted by the theory, and is responsible for the 24% R^2 in the low forecast bucket. When the freeze was more severe than predicted, returns are higher. The latter case generates a significant coefficient of the wrong sign and is responsible for the 27% R^2 in the 41 to 45 degree forecast bucket. When the observations are sorted by temperature, the explanatory power of these two regressors is combined in the low temperature bucket, with a resulting R^2 of 31%. Although this estimation is carried out on a limited sample due to the unavailability of extensive forecast data, one could think of extrapolating the results to the full sample. Specifically, the contemporaneous temperature alone explains approximately 50% of the variation of returns in the low temperature region (see Figure 2), and the temperature surprise explains approximately one third of the remaining variation. The combined explanatory power is therefore approximately 65%. In other words, the contemporaneous minimum temperature plus the temperature surprise (relative to the previous day's forecast) can explain almost two thirds of the return variation in by far the most volatile subset of days, namely when temperatures are low.

3.2.2 What's Missing?

Realized temperatures alone explain a significant fraction of variation in FCOJ futures returns, and forecasts contribute additional explanatory power. Nevertheless, there is still some unexplained variation in returns in winter months relative to the spring and summer. There are three possible explanations for this phenomenon.

First, the realized temperature minimum temperature may not be a perfect proxy for the severity of the freeze. Table 2 indicates that freeze duration is also an important variable. In an attempt to address this issue, we analyzed hourly Orlando temperatures, as provided by the National Weather Service, during freeze episodes for the 1983-1998 period. Using these data, we constructed a number of measures of duration based on the critical temperature levels in Table 2. Perhaps unsurprisingly, our duration measures were highly correlated with the minimum temperatures. As a result, given the small sample size, it was impossible to detect a significant effect.

Second, the severity of the freeze may not be a perfect proxy for the resulting supply shock for a number of reasons. For example, the timing of the freeze is important, especially with respect to the EM varietals, which are harvested primarily in January. In our sample, the earliest freeze is in late December and the latest is in late February (see Table 1). Clearly, even controlling for severity, these freezes will not have the same impact on production. There is also the possibility of time dependence. Freezes that occur in years following a freeze year or sequence of freeze years may have different effects on production. In addition, there is a degree of nonstationarity induced by the fact the Florida orange production has gradually been moving southward over the sample period. Orlando temperatures are slowly becoming less relevant, and freezes are generally becoming less of an issue. Tables 1 and 3 provide casual evidence of these combined effects. There is far from a one-to-one correspondence between the minimum temperature reported in Table 1 and the change in the production forecast reported in Table 3. Moreover, looking more closely at the monthly production forecast updates leads to the same conclusion. Ideally, we would relate daily returns to the market's daily unexpected production shock, but this latter variable is unavailable. However, we do look more closely at the monthly production forecast data in the next section.

Finally, even knowing the shock to the market's production forecast is not sufficient. The impact on returns of a given shock to production will depend on a variety of other factors. For example, the shape of the demand curve will influence the price effect of a fixed production shock, given different initial levels of anticipated supply. The availability and price of substitutes (e.g., imports from Brazil, Mexico and the Caribbean) will also influence the price impact. This availability is itself determined by myriad other factors, including the weather in other orange growing regions and tariffs, which changed considerably over the sample period. We take a look at some of these factors in the next section.

Given the complexities outlined above, the explanatory power we document using a simple and naive model is even more remarkable.

4 Other Components of FCOJ Return Variation

The results of Sections 2 and 3 demonstrate three important results: (i) there is large variation in FCOJ futures returns consistent with theory, i.e., around freezing temperatures, (ii) the relation between FCOJ futures returns and temperature conforms to theory, being highly nonlinear and convex at temperatures around 32 degrees and below, and (iii) other information, such as forecasts, have an important (though difficult to measure) impact on FCOJ futures returns.

We can use these results to better understand the sources of return variation of FCOJ futures returns. Consider Table 4A, which breaks up the variation of FCOJ returns by season. Approximately 50% of all the variation comes in winter months, 20% in fall months, and the remaining 30% is shared equally between spring and summer months. To the extent we can explain almost 35% of the variation in the winter months with the simple model (and more using forecast data), much of the winter variation reduces to the baseline FCOJ futures return variation of the other months. Moreover, freezing temperature observations represent only 4.3% of winter months, but two-thirds of winter return variability. In fact, they represent only 1% of all the observations yet capture one-third of all return variability! That is, the puzzle is not whether FCOJ futures returns and weather are related - they clearly are - but why is there variation in the other months of the year?

We leave to other research issues related to the microstructure of this market and to the stochastic convenience yield of futures, although these factors may certainly be important. Instead, we focus on FCOJ-specific factors, namely (i) news about Florida FCOJ production, (ii) news about Brazil's FCOJ production, and (iii) demand for FCOJ. We provide some rudimentary evidence that these factors may also play an important role in FCOJ futures pricing.

To supplement our analysis, we searched all Wall Street Journal (WSJ) articles about FCOJ futures from January 1, 1984 to November 11, 1998. During this period of 3686 days, there were a total of 384 identified articles in the WSJ about FCOJ. These articles are then divided into one of six categories related to the focal point of the article, namely news about (i) weather, (ii) production, (iii) Brazil, (iv) demand, (v) technical factors, and (vi) miscellaneous. For example, the news article about Brazil might relate to information about export prices, shipment delays, tariff rates, and Brazilian weather such as droughts. In contrast, demand stories relate to changes in the movement of FCOJ, behavior patterns of consumers, competing products, and retail sales. The results are reported in Tables 9 and 10.

Table 9 shows that on days associated with WSJ news, the volatility is substantially higher than on days associated with no WSJ news. In particular, the variability (variance) is 3.85 times larger on news days. Furthermore, the direction of the FCOJ price movement coincides with the news content, e.g., 0.97% versus -1.23% on news about low versus high temperatures, respectively.[8] Most interesting is that of the 384 WSJ stories, less than one-third were about weather, and almost one-half were either about production, Brazil or demand. This is strong anecdotal evidence that other information seems relevant for the FCOJ market.

Table 10 provides the distribution of news stories across months. Unsurprisingly, the weather-related articles generally occurred in winter months. However, stories about demand factors and Brazilian FCOJ production occur throughout the year, which suggests that important information gets released during spring and summer months. Finally, while the WSJ writes some stories about production in winter and early spring months, the majority of stories are written in October. As discussed below, October is an important month as it represents the first official government forecast

[8]It should be noted that there is a self-selection problem at work here. On days in which FCOJ futures prices moved, WSJ reporters might look for a story to fit the futures price movement. Alternatively, on days with more interesting stories, the WSJ might ignore news about FCOJ futures. In either case, the reader should be cautious interpreting these results beyond the anecdotal nature of the evidence.

of orange production for the forthcoming season.

4.1 Production Data

Table 4A shows that the variability of FCOJ futures returns in the fall season is about 40% greater than the spring and summer seasons. Why? Is there something special about fall? It clearly cannot be the temperature level as no freeze has actually occurred in the fall season (September through November).[9] However, as mentioned above, an important event occurs every October with respect to orange production–the USDA releases its first production forecasts of the season. In fact, the USDA provides a prediction of year-end total orange production on a monthly basis from October to July. Usually, in August, the USDA issues a report stating the actual annual production for the previous crop year. At the beginning of each month, the USDA sends out surveys to growers to report their expected production as of the first of the month. Then, the production forecasts are complied from the survey data and released in the second week of the month at 2pm after the FCOJ futures market has been closed. Since October represents the first release of the upcoming orange production forecast, it seems reasonable to assume that it may have special importance; however, other monthly forecast releases may also be relevant.

We collected the monthly USDA production forecast, which is the USDA forecast of the year-end orange production of Florida, for the period 1967-1996. Consequently, we have 30 announcement days in each month, which represents about 5% of the total trading days within the month. In order for these forecasts to affect returns, in a rational setting, they must contain useful information about future production. One way to assess their information content is analyze the implied temporal resolution of uncertainty. Table 11, Panel A documents the cross-sectional volatility of USDA production forecasts errors, where the forecast errors are calculated as the percentage difference between the current forecast and final production level. Volatilities are calculated for each month across the 30 years in the sample. Two observations are in order. First, there is significant production uncertainty in October in both freeze and non-freeze years, confirming the evidence from Table 3. Thus, whether a freeze takes place is not the only relevant factor for production and thus pricing. Second, the resolution of uncertainty occurs slowly throughout the year. Little or no relevant information comes out in November or December, but after that there is a noticeable,

[9]This is not quite correct as a mild freeze did occur within our sample on November 25, 1970.

albeit somewhat uneven, monthly decline in uncertainty. Not surprisingly, a great deal of production information is revealed during January and February in freeze years. In non-freeze years, the smaller amount of uncertainty is resolved later in the season.

These facts together give support for the existence of important fundamental information other than temperature. One possibility is that this information is actually revealed to the market on USDA announcement dates. However, it is also plausible that the USDA production forecasts lag market information. For example, the State of Florida releases crop damage estimates after significant freezes, which may precede and subsume the information in the subsequent USDA announcement. Table 11, Panel B, examines return volatility on USDA announcement dates. We report the daily volatility of returns on these days and the daily volatility on the remaining non-announcement days. Freeze days are excluded from the sample since they have already been studied extensively. The results are quite striking and consistent with the theory. The volatility of FCOJ futures returns on the October USDA announcement day is 9.15%, which is in the same range as the volatility around freezing temperatures. Incredibly, this one day represents over 65% of October return variance, and 40% of the total variability in the fall season. While the results for other months are less impressive, announcement days still coincide with up to 16% of the return variance within the month despite the fact that they only account for 5% of the days. The only exceptions to this pattern are November and December, when it is clear from Panel A that the USDA forecasts contain little or no information.

Consistent with theory, volatility is high on the days with the most news about production. The direction of these FCOJ futures return movements are also consistent with theory. For example, Bauer and Orazem (1994) look at the relation between production surprises and FCOJ futures returns, and find a strong negative correlation.[10] Moreover, the above results are consistent with the WSJ news story analysis documented in Tables 9 and 10. Both the volatility and direction of FCOJ futures returns are consistent with the WSJ writing stories about FCOJ in October. As mentioned in an earlier footnote, while there is a sample selection issue, the WSJ analysis, and volatility results here, are consistent with a relation between asset prices and fundamentals.

[10] Using our sample, we also confirm Bauer and Orazem's (1994) findings. It is also possible to improve their model's fit by taking into account nonlinearities similar in spirit to Section 3.1. For brevity, the results are not reported in the tables.

4.2 Brazil

Table 3 documents that the average return for the winter season is related to weather as one might expect. While this result represents only an average finding, the data itself is fairly noisy. For example, both the winter season of 1983 and 1984 experienced freezes. In the 1983/1984 winter, prices rose 30.1% due to the freeze, while the following winter prices fell 3.3% in spite of a freeze. How do we explain such a chain of events?

Market imperfections aside (e.g., tariffs, import quotas, and transactions costs), the U.S. price of oranges relevant for FCOJ should reflect the world price of these oranges. This world price in turn depends on the interaction of the worldwide demand and supply for oranges. Thus, in response to the question posed above, a downward shift in U.S. supply does not necessarily map one-to-one with an increase in FCOJ prices. Moreover, volatility of FCOJ prices in general will depend on information about aggregate supply and demand beyond the U.S.

To understand this more fully, consider the following facts about FCOJ: (i) the U.S. is a net importer of FCOJ, (ii) the worldwide supply curve of oranges in any given year is fairly inelastic (i.e., the production levels are fixed, irrespective of price), and (iii) the U.S. demand curve has a higher slope than worldwide demand because, at the very least, it is a subset of that demand.

Consider also the following headline from a Miami Herald article, "Florida - The King of Citrus No More", in October 1984 that described the end of Florida's reign as the largest producer worldwide of oranges. The article discussed Brazil's succession as a result of severe freezes in Florida that began in early 1981. The origin of the FCOJ production in Brazil dates to the 1960's. The motivation for the development was the big freeze that destroyed a great part of the American groves, namely the 13 million adult orange trees that died in 1962 (see Table 1). Since the U.S. industry did not have enough raw materials to supply both their domestic market and the worldwide market, Brazil filled the gap by accelerating the development of their orange production and processing plants.

Interestingly, similar to the U.S. market and its reliance on the Orlando area, the majority of oranges are produced in one area of Brazil, namely Sao Paolo. However, unlike Orlando's predicament with respect to weather, it is droughts rather than freezes that affect Brazilian production.[11]

[11] Since Brazilian groves are frost-free, summer droughts are usually the weather risk. In general, the state of Sao Paulo is especially suitable for the production of high quality oranges. It is characterized by an absence of frost,

The drought season in Brazil starts in July and ends in November. Total Brazilian supply is hence known by the winter season in the U.S., but may vary year-to-year. Dry weather may retard or damage buds that bear next season's harvest. Therefore, Brazil's news on production usually starts at the beginning of March.

Relevant for our discussion, note that there is not enough U.S. production to meet U.S. domestic demand. Therefore, imports of FCOJ are needed. Though Brazil and the U.S. account for over 85% of worldwide production of orange juice, the U.S. does import from other countries, most notably and recently Mexico due to its favorable import duties relative to Brazil since the passage of NAFTA, and the Caribbean countries, which export duty-free to the U.S. under CBERA. Nevertheless, Brazil has historically been the primary exporter of FCOJ to the U.S. Figure 5A graphs total U.S. FCOJ production (i.e., pack), U.S. supply (i.e., pack plus inventories), and Brazil FCOJ imports over the period 1977 through 1996. As can be seen from the figure, imports generally rose threefold during the early to mid 1980's while somewhat subsiding during the 1990's. These imports coincide with an overall production loss or slowing in the U.S. that is very transparent starting in early 1981, and only fully recovering in the 1990's. At a glance, this suggests that Brazil FCOJ production is used as a substitute for U.S. FCOJ production in fairly significant amounts.

Even though there are clearly many factors that affect Brazil's exports to the U.S., Figure 5B graphs the change in the U.S.'s imports of Brazilian FCOJ divided into freeze and non-freeze years. The data is annual and covers the available period from 1977 to 1996. The results are quite startling. In particular, in almost very year there was a freeze, Brazil's exports to the U.S. increased from the previous year. Similarly, in non-freeze years, Brazil tended to reduce imports. The correlation is strongly significant. The two exceptions are in freeze years in which the U.S. FCOJ pack were not strongly affected.

While this analysis is casual at best, it does illustrate two important facts: (i) Brazilian imports are significant in magnitude, and (ii) these imports are used as substitutes for domestic production, especially during freeze years. To the extent that there is information about Brazil's FCOJ production, such as transport costs, tariffs, Sao Paolo weather, and, more directly, Brazil's FCOJ production and prices, this should get incorporated into FCOJ futures prices. Thus, this factor

coupled with dry winter and heavy summer precipitation. This annual moisture pattern tends to result in a single blooming period, but does subject FCOJ production to droughts.

represents another important component, not related to the temperature level in Orlando. As supporting evidence, the WSJ news story analysis in Table 9 shows that approximately 15% of all FCOJ stories are related to Brazil, and that on these days the variability of FCOJ returns is more than three times that on non-news days.

4.3 Demand

On a daily basis, Americans drink over 79 million six-ounce glasses of orange juice, that is, over half of all fruit juice sold in the U.S. There are a variety of choices, however, ranging from fresh orange juice to grapefruit juice. Not surprisingly, there is a substantial advertising industry aimed at FCOJ consumers. Because of the size of this industry, movement reports are closely tracked by market participants. Significant changes in movement numbers may suggest changes in the consumer behavior and can hence influence the short-term behavior of prices.

We have collected annual data from the Movement and Pack report from the New York Cotton Exchange. The period covers 1978-1998. This annual report contains information on the movement of FCOJ in the U.S., which many consider to be a measure of demand. It covers the movement within the U.S. of FCOJ, which includes both the FCOJ pack (i.e., U.S. production), imports and inventories. Since the data is annual, we divided the movement data into two categories: increasing demand and decreasing demand seasons. An increasing (decreasing) demand season is simply one with movement higher (lower) than the previous season.

Conditional on the different trends in demand, we compare the cumulative returns on FCOJ over spring and summer months. We avoid the fall and winter months to avoid contamination with information already incorporated in returns due to supply shocks (i.e., production news and weather). In terms of evidence, when demand is increasing (decreasing), the FCOJ prices on average increase (decrease) during spring and summer. In particular, the difference in returns for increasing and decreasing demand is around 15% at the end of a season. In addition, since the demand and FCOJ price change move together in the right direction, the price change is more likely to be a shift of the demand curve, rather than a movement along the demand curve. This suggests on a heuristic level that demand factors play a role in FCOJ futures price valuation. As supporting evidence, Table 9 documents that approximately 7% of the WSJ news stories are demand-related. Moreover, the variability of FCOJ returns on these days is more than three times that of non-news days. While this evidence is not proof of a fundamental relation between FCOJ prices and demand, it is consistent with that notion.

5 Concluding Remarks

This paper reverses the existing stylized fact that FCOJ futures returns are not related to fundamentals, particularly temperature. In fact, we present the following empirical facts which are strongly supported by theory:

- Though covering only 4.5% of the days in winter, two-thirds of all winter return variability in FCOJ futures coincide with freezing temperatures.

- On these days, using a simple, but theoretically appropriate nonlinear model of the relation between FCOJ returns and weather, we can explain almost 50% of the variation. Moreover, these R^2s can increase to almost 65% if the model is extended to include market information about forecasts.

- Other factors also play an important role in explaining FCOJ price variability. For example, we can document that over 40% of the FCOJ futures return variation occurs on 1.5% of the days in the fall season. This day is easy to identify; it is the first USDA forecast of orange production for the upcoming season, released in mid-October.

The bottom line from this paper is that when we seem to have an understanding of the fundamental information, both the FCOJ futures return variability and direction seems to coincide with theory. Thus, this research has produced a strong link between FCOJ futures returns and their underlying fundamentals.

Why is there volatility in FCOJ futures returns on days other than around low temperatures or USDA production news? The pricing of FCOJ futures is, like any asset, extremely complicated. The futures convenience yield, market microstructure, long-run news about weather, market news about orange tree moisture, news events about Brazil, and short-run demand shifts, among many other factors, can all have important effects. Cursory evidence suggests this is the case. For example, on news days about Brazilian production, which tend to be in summer, there is substantially higher volatility in the FCOJ futures market. What is important is that when we can clearly identify the fundamental, there is a close link between prices and that fundamental.

What this result suggests about other assets and fundamentals, such as stock prices and earnings announcements, is open for debate. However, our view is that it is unlikely that we can answer this question for equity prices when, in the apparently simple world of FCOJ futures and weather, the relation between prices and fundamentals is complex, nonlinear, multi-dimensional and state and path dependent.

Appendix 1: Return Calculation - Discussion of Limit Moves

In general we want to focus on the most liquid contracts, which are usually the ones closest to maturity except when close to expiration. To that end we develop an algorithm that results in a spliced return series which utilizes the most liquid contracts. Specifically, we make a distinction between two types of periods: (i) trading days in an expiration month prior to expiration (we denote this period EXP), and (ii) the rest of the time, i.e., trading days in months in which no contracts expire and trading days in an expiration month and past the expiration date of the contract (we denote this period NoEXP). In period NoEXP we want to focus on the two near maturity contracts C1 and C2:

- If neither contract hits its price limit, the return is the average return on C1 and C2.
- If C2 hits the limit and C1 does not, we calculate the return from C1 only.
- If C1 hits the limit and C2 does not, we average the returns on C1 and C2.
- If both contracts limit, we average the returns on C1 and C2 return and aggregate forward until either is off-limit (and record the minimum temperature for the trading "day" as the minimum temperature for the entire limiting period).

In period EXP we want to avoid using C1 due to the illiquidity that develops in futures contracts close to expiration. Our rule above is then applied to C2 and C3 in a similar manner.[12]

[12]During the sample period a number of rules relevant for the calculation of returns were changed. First, with respect to contract expiration dates, prior to November 1994 the expiration date was 10 business days prior to the end of every other month. In the following period expiration was 14 business days prior to the end of the month. As for the size of the price limits, initially it was 3 cents, and then it moved up to 5 cents in January 1979. From May 1986 there is no price limit on C1. From January 1993 there is no limit on the move in C2 during the expiration month of C1.

6 References

Attaway, J., 1997, *A History of Florida Citrus Freezes*, Florida Science Source, Inc.

Attaway, J., 1999, *Hurricanes and Florida Agriculture*, Florida Science Source, Inc.

Balduzzi, P., E. Elton and C. Green, 2001, Economic News and the Bond Prices: Evidence from the U.S. Treasury Market, forthcoming *Journal of Financial and Quantitative Analysis*.

Bauer, R. and P. Orazem, 1994, The rationality and price effects of U.S. Department of Agriculture forecasts of oranges, *Journal of Finance*, 49, 681-695.

Black, F., 1986, Noise, *Journal of Finance*, 41, 529-543.

Campbell, S. and F. Diebold, 2001, Weather forecasting for weather derivatives, PIER Working Paper 01-131, University of Pennsylvania.

Campbell, J. and R. Shiller, 1987, Cointegration and tests of present value models, *Journal of Political Economy*, 95, 1062-1088.

Cutler, J. Poterba and L. Summers, 1989, What moves stock prices?, *Journal of Portfolio Management*, 15, 4-12.

Ederington, L. and H. Lee, 1993, How markets process information: News releases and volatility, *Journal of Finance*, 48, 1161-1191.

Fama, E., 1998, Market efficiency, long-term returns and behavioral finance, *Journal of Financial Economics*, 49, 283-306.

Frankel, J. and R. Meese, 1987, Are exchange rates excessively variable, in S. Fischer, ed., *NBER Macroeconomics Annual 1987*, Cambridge: MIT press, 1117-152.

French, K., R. Leftwich, and W. Uhrig, 1989, The effects of scheduled announcements on futures markets, CRSP working paper 273, University of Chicago.

Fortenberry, R. and D. Sumner, 1993, The effects of USDA reports in futures and options markets, *Journal of Futures Markets*, 13, 157-173.

Hebert, Paul, 1993, Freezes and hurricanes in Florida, *National Weather Digest*, vol. 18, no. 2.

Hirshleifer, D., 2001, Investor psychology and asset pricing, *Journal of Finance*, 56, 1533-1598.

Kleidon, A., 1986, Variance bounds tests and stock price valuation models, *Journal of Political Economy*, 94, 953-1001.

Ito, T. and V. Roley, 1987, News from the U.S. and Japan: Which moves the Yen/Dollar exchange rate? *Journal of Monetary Economics*, 19, 255-277.

Marsh, T. and R. Merton, 1986, Dividend variability and variance bounds tests for the rationality of stock market prices, *American Economic Review*, 76, 483-498.

Mitchell, M. and J. Mulherin, 1994, The impact of public information in the stock market, *Journal of Finance*, 49, 923-950.

Poterba, J. and L. Summers, 1988, Mean reversion in stock prices, *Journal of Financial Economics*, 22, 27-59.

Roley, V. and R. Troll, 1983, The impact of new economic information on the volatility of short-term interest rates, *Economic Review*, Federal Reserve Bank of Kansas City, 3-15.

Roll, R. 1984, Orange juice and weather, *American Economic Review*, 74, 861-880.

Roll, 1988, R^2, *Journal of Finance*, 44, 1-17.

Shiller, R., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review*, 71, 421-498.

Ullah, A., 1988, Nonparametric estimation of econometric fundamentals, *Canadian Journal of Economics*, 21, 625-658.

Urich, T. and P. Wachtell, 1984, The effects of inflation and money supply announcements on interest rates, *Journal of Finance*, 39, 1177-1188.

Date	Low Temperature (Orlando area)	Damage
February 7-9, 1835	11^o	Catastrophic
1857	NA	Moderate
December 1, 1876	NA	Moderate/Severe
January 9-13, 1886	19^o	Severe
December 29-30, 1894	18^o	Severe
February 8-9, 1895	17^o	Catastrophic
January 28-29, 1897	25^o	Moderate
January 2-3, 1898	23^o	Severe
February 13-14, 1899	20^o	Catastrophic
January 1900	27^o	Moderate
February 17-18, 1900	26^o	Moderate
December 21, 1901	26^o	Moderate/Severe
January 14, 1902	26^o	Severe
January 26-27, 1905	21^o	Severe
February 2-6, 1917	22^o	Severe
December 12-13, 1934	22^o	Severe
January 27-29, 1940	22^o	Severe
February 5-6& 11, 1947	24^o	Moderate/Severe
December 12-13, 1957	21^o	Moderate/Severe
January 9-10, 1958	26^0	Severe
February 4-5, 14, 17-19, 21, 1958	$24^o - 26^o$	Severe
December 12-13, 1962	19^0	Severe/Catastrophic
January 31, 1966	23^o	Moderate
January 8-11, 1970	24^o	Moderate/Severe
January 20-22, 1971	22^o	Severe
January 19, 1977	20^o	Severe
January 12-14, 1981	18^o	Severe
January 11-12, 1982	23^o	Severe
December 24-25, 1983	24^o	Severe
January 20-22, 1985	19^o	Severe*
December 26-27, 1985	26^o	Moderate
January 28, 1986	27^o	Moderate
February 24, 1989	29^o	Moderate
December 23-24, 1989	25^o	Severe
January 8-9, 1996	29^o	Moderate
February 5-6, 1996	26^o	Moderate

Table 1: History of Freezes

The table lists the important freezes from 1835 to 1998 as documented by Attaway (1997). An important freeze is one in which the temperature in the Orlando area (i) dropped substantially below freezing, and (ii) remained at this level for a significant period of time. The description of the freezes, varying from moderate to catastrophic, is subjective, based on a description of the economic impact of the freezes provided in Attaway (1997). Note that the * refers to the fact that, although each individual freeze in the 1981-85 period was severe, the combined effect of four freezes in five years was generally considered catastrophic. Our sample period is 1967-98 and therefore covers the freezes below the horizontal line.

Critical Temperatures for Oranges		
	Winter (Dec to mid-Feb)	Spring (mid-Feb to Mar)
Bloom	NA	29-30
Fruit	26	29
Leaves	24	27
Twigs	22	24
Branches	20	22

Severity of Freeze (Damage vs. Time)		
Temperature	Duration	Consequence
28	6 hours or more	Extensive Fruit Damage
26	4 hours or more	Extensive Fruit Damage
24	2 hours or more	Extensive Fruit Damage
20	4 hours or more	Extensive Tree Damage

Freeze Recovery Periods		
Freeze Type	Minimum Temperature	No. of Years to Recovery
Light	32-29	0
Moderate	28-25	0
Hard	24-21	1-3
Severe	Less than 20	3-7

Table 2: Temperatures and Orange Production Damage

Source: Citrus Associates of New York Cotton Exchange, Inc, 1994

Year	Freeze	Prod. %Change	Cum. Ret.(%)
67		4.7	-12.6
68		4.6	31.1
69	Yes	-3.7	-8.4
70	Yes	-15.3	16.2
71		5.0	-14.6
72		-2.5	-6.2
73		3.6	-16.1
74		-0.4	-19.5
75		5.3	-2.1
76	Yes	-10.6	52.7
77		2.3	-0.7
78		-1.8	-7.8
79		3.3	-12.2
80	Yes	-15.1	44.5
81	Yes	-24.2	-0.6
82		-2.4	-12.6
83	Yes	-30.5	30.1
84	Yes	-12.7	-3.3
85	Yes	-9.7	-27.7
86		-7.2	0.2
87		6.2	-4.1
88	Yes	-3.6	-14.9
89	Yes	-15.2	55.0
90		-8.1	1.6
91		2.8	-16.1
92		0.3	-34.3
93		1.4	7.1
94		4.8	-9.8
95	Yes	0.6	-3.5
96		2.8	-23.6
Frz		-12.7	12.7
Non-Frz		1.3	-6.1

Table 3: Freezes, Production and Returns

Year by year percentage differences between actual orange production and the October USDA forecast at the beginning of the season, and the cumulative FCOJ futures return over the corresponding winter season (December to February). The bottom two rows show the averages of these numbers over freeze years and non-freeze years. Seasons are classified as freeze/non-freeze based on Attaway (1997).

Panel A: Seasonal Returns			
Season	Nobs	Mean	SD
All	7543	0.033	2.105
Winter	1772	0.009	3.160
Winter Pre-Freeze	1475	0.015	3.305
Winter Post-Freeze	297	-0.068	2.315
Fall	1878	0.068	1.847
Spring	1939	0.059	1.555
Summer	1954	-0.002	1.539
Panel B: Temperature-Sorted Returns			
From-To	Nobs	Mean	SD
0-35	77	2.721	11.815
36-40	134	-0.311	2.163
41-45	198	-0.273	2.081
46-99	1066	-0.086	1.790

Table 4: FCOJ Futures Returns

Means and standard deviations of daily FCOJ futures returns in percent for different seasons and temperature ranges. Continuous futures daily return series are constructed off of the three closest maturity contracts in a manner described in Appendix 1. The data period is September 1967 to August 1998. Temperatures are minimum night temperature for the Orlando region, contemporaneous with the close-to-close futures return. Winter season is defined as December, January and February, Spring is March, April, May, Summer is June, July, August and Fall is September, October, November. The pre-freeze period includes days up to and including the first freeze of the season, if a freeze occurs.

	α	β	β_1	β_2	R^2
Linear	2.187	-0.042	–	–	0.015
	(0.474)	(0.009)			
Piecewise					0.281
$0-35$	58.113	-1.741			
	(8.409)	(0.262)			
$36-40$	7.714	-0.211			
	(5.419)	(0.142)			
$41-45$	6.620	-0.159			
	(4.590)	(0.106)			
$46-99$	-0.558	0.008			
	(0.499)	(0.009)			
Nonlinear	-0.135		0.095	0.202	0.333
	(0.071)		(0.339)	(0.031)	
Kernel					0.336

Table 5: Model Estimates and R^2s

Data period is for winter pre-freeze period September 1967 to August 1998. Continuous futures return series are constructed off of the three closest maturity contracts in a manner described in Appendix 1. Returns are in percent per day. Temperatures are minimum night temperature for the Orlando region, contemporaneous with the close-to-close futures return. The linear model is

$$R_t = \alpha + \beta W_t + \epsilon_t$$

The piecewise linear model is a set of four linear regressions within temperature buckets. The nonlinear model is

$$R_t = \alpha + \beta_1 Max[0,\ W^* - W_t] + \beta_2 (Max[0,\ W^* - W_t])^2 + \epsilon_t$$

with the threshold being $W^* = 32^o F$. Kernel provides nonparametric estimates

$$R_t = f(W_t) + \epsilon_t$$

as described in the paper.

	Const.	W_{t-1}	W_{t-2}	$F_{t-1}(W_t)$	R^2
W_t	13.96			0.74	0.56
	(1.22)			(0.02)	853
W_t	18.11	0.70	-0.04		0.46
	(1.49)	(0.03)	(0.03)		853
W_t	0.10	0.46	-0.03	0.56	0.73
	(1.23)	(0.03)	(0.02)	(0.02)	853

	W_{t-1}	W_{t-2}	Dec.	Jan.	Feb.	R^2
W_t	0.67	-0.07	21.05	19.77	20.51	0.41
	(0.02)	(0.02)	(1.03)	(1.05)	(1.08)	1877

Table 6: Temperature Forecasting

Estimates of temperature forecasting models. The first set of models includes some combination of NWS temperature forecasts and lagged temperatures and is estimated over a 10-year sample period with data from November through March. The most general regression is of the form

$$W_t = \alpha + \beta_1 W_{t-1} + \beta_2 W_{t-2} + \gamma F_{t-1}(W_t) + \epsilon_t.$$

The second set of models is estimated on the full sample period, winter only, and uses temperatures and monthly dummy variables:

$$W_t = \alpha_i D_{it} + \beta_1 W_{t-1} + \beta_2 W_{t-2} + \epsilon_t,$$

Standard errors are in parentheses under the coefficient estimates, and the number of observations in each regression appear under the regression R^2s.

α	β_1	β_2	Bucket	R^2_{0-35}	R^2_{36-40}	R^2_{41-45}	R^2_{45-99}	R^2_{0-99}
-0.09	0.24	–	$F_t(W_{t+1})$	0.17	0.00	0.00	0.00	0.01
(0.11)	(0.15)			26	38	70	309	443
-0.07	–	0.02	W_{t+1}	0.00	0.00	0.00	0.00	0.00
(0.11)		(0.11)		21	30	60	332	443
-0.09	0.52	-0.26	$F_t(W_{t+1})$	0.24	0.00	0.03	0.00	0.01
(0.11)	(0.23)	(0.17)		26	38	70	309	443
			W_{t+1}	0.02	0.00	0.05	0.01	0.01
				21	30	60	332	443
0.01	–	0.12	W_{t+1}	0.04	0.00	0.00	0.00	0.00
(0.06)		(0.07)		79	132	197	1056	1464

Table 7: Error Analysis - Forecasts

Error analysis regressions, where the errors, ϵ_t, are from a winter pre-freeze kernel regression of returns R_t on contemporaneous temperatures W_t. Error analysis regressions involve the multivariate regression

$$\epsilon_t = \alpha + \beta_1 Max[0, 32 - F_t(W_{t+1})] + \beta_2 Max[0, 32 - W_{t+1})] + \nu_t$$

and univariate versions of the above. The R^2 is calculated within buckets, sorted by the forecast or the temperature, as denoted in column "Bucket". Data is for days where forecasts are available, except the final regression, which is for all winter pre-freeze days less the days where ϵ_t is a freeze day (hence 1464 days: 1475 winter pre-freeze observations less 11 freeze days in our sample). Standard errors are in parentheses under the regression estimates, and the number of observations in each bucket appear under the bucket R^2s.

α	β_1	β_2	β_3	β_4	Bucket	R^2_{0-35}	R^2_{36-40}	R^2_{41-45}	R^2_{45-99}	R^2_{0-99}
-0.07	0.00	-1.58	0.48	-	$F_{t-1}(W_t)$	0.24	-0.01	0.27	0.00	0.11
(0.11)	(0.06)	(0.27)	(0.11)			26	37	70	310	443
					W_t	0.31	0.00	0.00	0.00	0.11
						20	30	59	334	443
-0.07	-	-	-	0.05	$F_{t-1}(W_t)$	-0.02	0.00	0.05	0.00	0.00
(0.11)				(0.06)		26	37	70	310	443
					W_t	0.01	0.00	-0.01	0.00	0.00
						20	30	59	334	443

Table 8: Error Analysis - Temperature Surprises

Error analysis regressions, where the errors, ϵ_t, are from a winter pre-freeze kernel regression of returns R_t on contemporaneous temperatures W_t. Error analysis regressions involve combinations of the multivariate regression

$$\epsilon_t = \alpha + \beta_1 I_{1,t} Z_t + \beta_2 I_{2,t} Z_t + \beta_3 I_{3,t} Z_t + \beta_4 I_{4,t} Z_t + \nu_t$$

where

$$\begin{aligned} Z_t &= W_t - F_{t-1}(W_t) \\ I_{1,t} &= 1 \text{ for } F_{t-1}(W_t) \leq 32^o \text{ and } W_t > 32^o \\ I_{2,t} &= 1 \text{ for } F_{t-1}(W_t) \leq 32^o \text{ and } W_t \leq 32^o \\ I_{3,t} &= 1 \text{ for } F_{t-1}(W_t) > 32^o \text{ and } W_t \leq 32^o \\ I_{4,t} &= I_{1,t} + I_{2,t} + I_{3,t}. \end{aligned}$$

The R^2 is calculated within a bucket, sorted by the forecast or the temperature, as denoted in column "Bucket". Data is for days where forecasts are available, for a total of 445 days. Standard errors in parenthesis under regression estimates, and number of observations in each bucket appear under bucket R^2s.

	Sample Size	Mean (%)	Vol (%)
News Day	384	0.015	3.06
No News Day	3302	-0.008	1.55
Weather	121	0.320	3.30
Production	101	-0.034	3.85
Brazil	58	0.554	2.95
Demand	21	0.065	2.90
Technical	50	0.307	2.80
Miscellaneous	33	-0.042	1.27
Weather (low)	73	0.966	3.112
(high)	48	-1.230	2.637
Production (low)	46	0.928	3.534
(high)	55	-0.869	2.437

Table 9: FCOJ Futures Returns and WSJ Events

The table reports the mean and volatility of FCOJ futures returns on all the days the Wall Street Journal (WSJ) reported some story about FCOJ futures. In particular, the sample covers all WSJ articles from January 1, 1984 to November 11, 1998 in which FCOJ futures contracts were discussed. This sample represents 3,686 trading days. Based on the tenor of the article, the stories were broken down into one of six categories related to a particular fundamental factor: (i) weather, (ii) production, (iii) Brazil, (iv) demand, (v) technical, or (vi) miscellaneous.

	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sept
Weather	4	12	41	45	21	7	0	0	1	0	2	0
Production	31	1	2	12	18	12	11	3	3	1	6	6
Brazil	6	3	3	9	3	9	3	6	9	4	3	2
Demand	0	2	3	0	0	3	1	4	1	3	3	1
Technical	3	7	3	3	3	5	8	3	6	7	2	4
Misc.	4	1	0	3	3	0	3	0	3	5	6	7

Table 10: Number of News Stories by Month

Distribution of news stories by month. News stories related to FCOJ as reported in the Wall Street Journal from 1984-1998. The stories are broken down into one of six categories: (i) weather, (ii) production, (iii) Brazil, (iv) demand, (v) technical and (vi) miscellaneous. The table reports the timing of the news story by month.

Panel A: Production Uncertainty										
	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul
Frz	12.9	12.9	13.6	11.9	6.1	3.5	3.7	2.4	1.8	0.4
Non-Frz	4.2	4.1	4.3	3.8	3.8	2.8	2.4	1.8	1.2	0.5
Panel B: Announcement Day Volatility										
	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul
Ann.	9.149	1.522	2.494	2.451	3.035	2.203	2.336	2.765	2.107	1.675
Non-Ann.	1.463	1.509	2.529	2.330	1.650	1.650	1.371	1.451	1.454	1.595
% Ann. Days	4.6	5.0	4.9	4.9	5.3	4.6	4.7	4.6	4.6	4.7
% Var.	65.3	5.1	4.8	5.4	16.0	7.9	12.6	15.0	9.2	5.2

Table 11: Production Uncertainty

Panel A reports month by month volatility of the USDA production forecasts. Forecast errors are calculated as the percentage difference between the forecast and the actual production relative to the initial (October) forecast:

$$100\left(\frac{P_{i,t} - P_{i,Final}}{P_{i,Oct}}\right)$$

where $P_{i,t}$ is year i month $t =$Oct, Nov, ... , Jul. The table reports the cross-sectional (across years) standard deviation of these forecast errors. Panel B reports return volatility on the announcement day for each month, the volatility for others days in the month which are not freeze days, the percentage of days which are announcement days, and the percentage of non-freeze day variance occuring on announcement days.

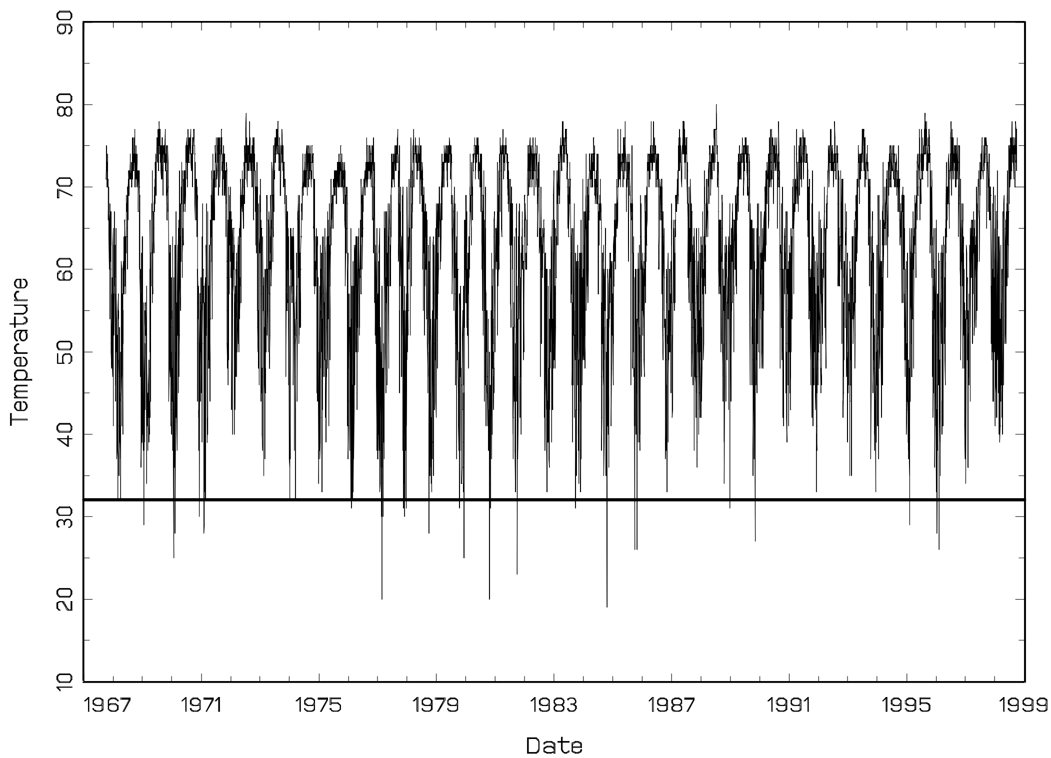

Figure 1: Minimum Temperatures

Daily minimum temperatures over the sample period, September 1967 to August 1998. The Horizontal line represents 32^o.

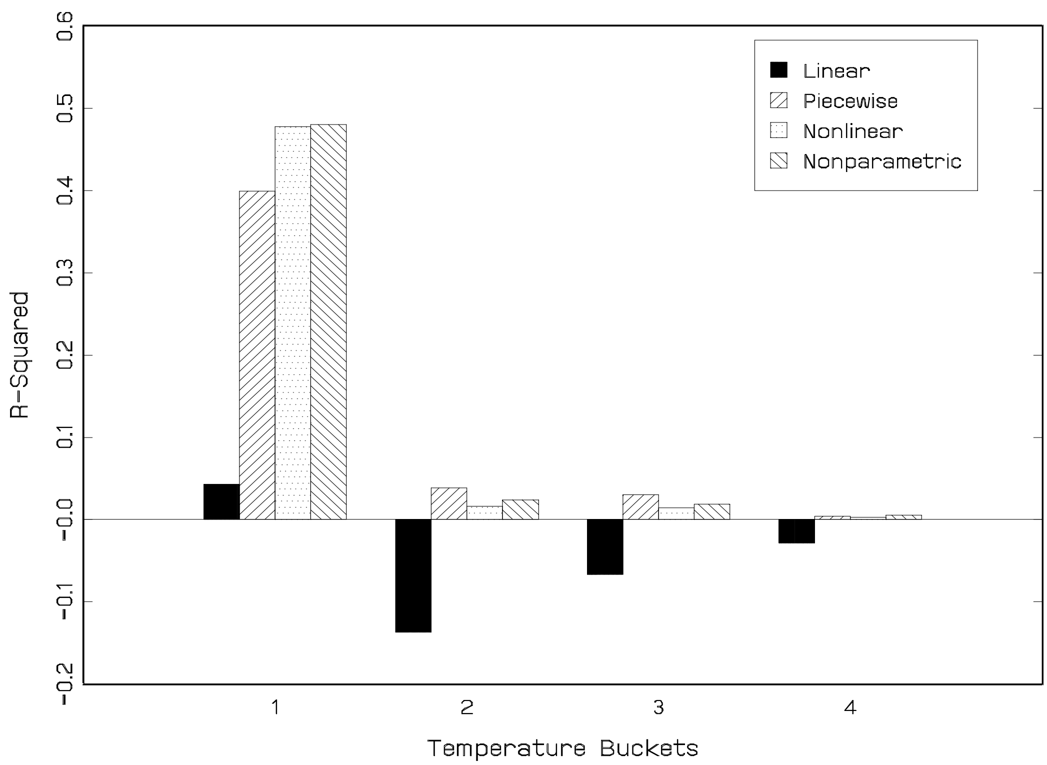

Figure 2: R^2s by Temperature Bucket

R^2s by temperature buckets for the four models described in Section 3.1. The temperature buckets are ≤ 35, 36-40, 41-45 and ≥ 46. The R^2s refer to models performed in winter months, pre-freeze.

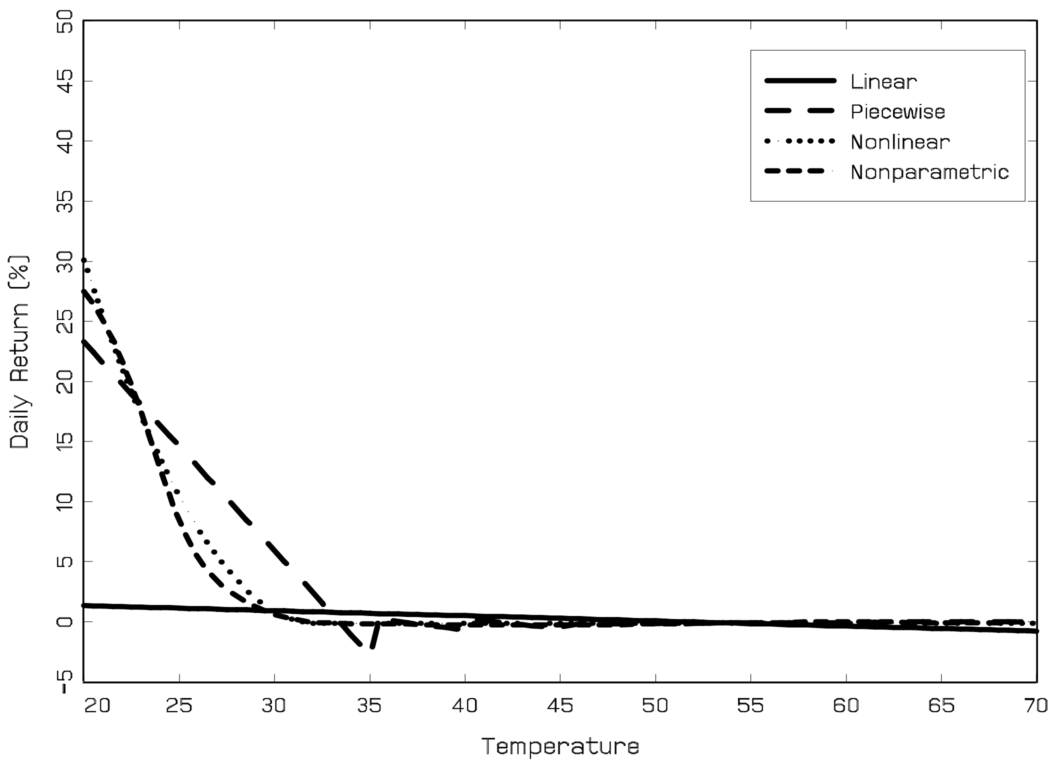

Figure 3: Fitted Returns

Fitted daily returns versus contemporaneous realized temperature for the four models described in Section 3.1. The analysis was performed in winter months, pre-freeze.

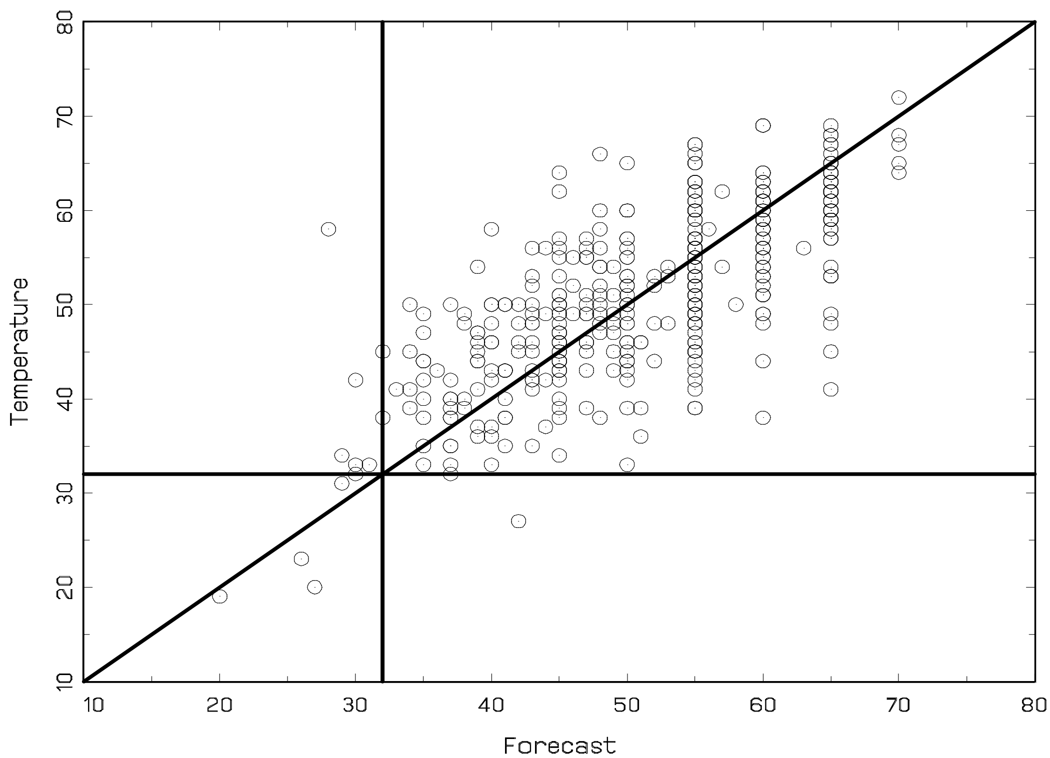

Figure 4: Temperatures and Forecasts

A scatter plot of minimum temperature forecasts from the National Weather Service (NWS) versus the corresponding minimum temperature the following day. The horizontal and vertical line represents 32^o.

Figure 5A and 5B:
Figure 5A graphs total imports of Brazilian FCOJ compared to U.S. production. Figure 5B graphs the change in U.S. imports of Brazilian FCOJ in freeze and non-freeze years. The period covers 1977 through 1996.