Undergraduate Honors Thesis

The Effect of Temperature Volatility on Corn Futures Price Levels and Volatility

Cole Davis

Economics

University of Colorado Boulder

Advised by Dr. Adam McCloskey[1]

Defense Committee: Dr. Stephen Becker[2], Dr. Brian Cadena[3], Dr. Adam McCloskey[4]

April 18, 2023

[1]Assistant Professor in Economics at the University of Colorado Boulder
[2]Associate Professor in Applied Mathematics at the University of Colorado Boulder
[3]Associate Professor in Economics at the University of Colorado Boulder
[4]See footnote 1.

Abstract

In this paper, I investigate how temperature volatility affects corn futures price levels and corn futures price volatility. I develop a novel temperature volatility statistic by fitting a GARCH(1,1) model to aggregated demeaned weather-station level average daily temperature data from the lead corn-producing states in 2020 as reported by the USDA. I then regress futures price levels and volatility on temperature volatility and a vector of controls. I find no statistical evidence that temperature volatility affects corn futures returns, and no statistical evidence that temperature volatility affects corn futures returns volatility.

1 Introduction

Since its creation, the corn futures markets has proved to be a valuable tool for hedging against risk in the cash market for corn and has become a valuable asset for investors looking to diversify their portfolios. According to the Chicago Mercantile Exchange Group (CME Group), the leading American commodity futures exchange, hedgers include groups such as farmers, merchandisers, grain elevators, food processors, feed manufacturers, grain exporters, and grain importers. Speculators provide the market with liquidity, allowing both buyers and sellers to enter the market at an efficient price.

In this paper, I investigate the effect temperature volatility has on corn futures contract price levels and volatility. Increasing volatility in the corn futures market reduces the viability of corn futures as a risk hedging instrument. Furthermore, economic theory predicts that an increase in volatility should result in an increase in expected returns. Hence, futures price levels and volatility should be positively correlated. To perform my analysis, I construct a new type of temperature volatility statistic by fitting a generalized auto-regressive conditionally heteroskedastic (GARCH) model [4, Engle & Bollerslev, 1986] to daily regional temperature data.

Daily temperature data comes from National Oceanic and Atmospheric Administration (NOAA) National Climate Data Center (NCDC). I first aggregate weather station-level data to a state level then I aggregate state-level temperature data to a regional level via a weighted average. The weights in the state-level aggregation are determined by the share of corn produced in 2020 as reported by the USDA. The regions of interest are made up of contiguous, top corn producing states in 2020. Regional-level temperature data is then demeaned and a GARCH(1,1) model is fit to the demeaned data to construct GARCH filter estimates. I use the log returns of corn futures closing prices reported by investing.com to construct my log-returns and returns volatility variables. I test the effect temperature volatility has on corn futures price levels using a simple linear regression of corn futures returns on temperature volatility with precipitation and time to contract maturity controls. I also run a separate regression of corn futures returns volatility on temperature volatility with precipitation and time to contract maturity controls to test the effect temperature volatility has on corn futures price volatility.

The auto-regressive conditionally heteroskedastic (ARCH) model [3, Engle, 1982], which led to the creation of the GARCH model, was conceived to help economists empirically model time series that exhibit conditional non-stationary properties. More specifically, the ARCH model allows the conditional variance and covariance of a time series to vary over time, while maintaining an unconditional time invariant auto-covariance function. The ARCH and GARCH models are used throughout the financial literature and started making headway in temperature forecasting close to a decade after their implementation into financial econometrics [17, Tol, 1996].

To my knowledge, the existing economics literature does not make use of GARCH models to generate temperature volatility statistics, instead opting for quadratic temperature level variables [20, Urban, Roberts, Schlenker & Lobell, 2012], which don't measure temperature volatility, or agricultural degree-day (DD) 'variance'. My GARCH volatility improves upon DD variance by enabling me to estimate current temperature volatility with NOAA NCDC data and forego the need for what is often an arbitrary reference period. Moreover, the weighting scheme used in my average temperature variable construction allows my model to focus on areas that lead in corn production as opposed to using aggregated global or national data for which temperature variability of regions that do not significantly contribute to corn production are included.

Mu Provided evidence that natural gas futures price dynamics are affected by temperature volatility [10, 2007]. In this paper, we perform a semi-analogous analysis in the corn futures market. I expand on Mu's paper by using a temperature volatility statistic that is better at capturing temperature deviations from their expected level[5]. Moreover, my temperature volatility statistic uses temperature realizations, as opposed to temperature forecasts, allowing me to test how temperature deviations from their mean are affecting the futures market.

In an influential paper, Roll found statistical evidence that weather shocks affect orange juice futures returns [14, 1984], where weather shocks were measured as a difference in actual temperature and forecast temperature. In this paper, I opt for using realized demeaned temperature volatility as my 'weather shock' variable as this temperature deviation from its mean captures true, realized temperature abnormalities on a given day in my time period[6]. The geographical concentration of orange juice production is the primary reason Roll measured the effect of weather shocks on orange juice futures. I use a weighting scheme in the

[5]For a discussion on how my temperature volatility statistic may fall short of this, and a method for making a more robust temperature volatility statistic following a similar framework to the one laid out in section 3.2, see section 5.3.
[6]See footnote 5.

construction of my temperature volatility statistic to enable my analysis to be expanded to more general regions.

2 Literature Review

There is a rich literature discussing the different variables that effect commodity futures price dynamics and, more specifically, agriculture and corn futures price dynamics. Given the connection between corn futures prices and the cash market price of corn, I would be remiss not to mention popular work investigating potential factors affecting corn yields, as corn supply is a crucial factor in cash-market pricing and futures prices converge to cash market prices as contracts near maturity. Most of the literature assumes agricultural yields depend on temperature and precipitation (this dependence is often modeled as a quadratic dependence in both temperature and precipitation [20, Urban, Roberts, Schlenker, & Lobell, 2011] or a linear dependence on DDs [13, Ritchie & Nesmith, 1991]) and opts for simulating the effect of future climate change on agricultural yields. Perhaps the most popular example of this comes from a 2007 paper from Deschênes and Greenstone where they use an agricultural DD metric and simulated climate data to find evidence that agricultural profits increase with predicted climate change [1], and that predicted increases in temperature and precipitation will have virtually no effect on corn yields. Contradicting evidence is given by Schlenker and Roberts who claim that predicted global temperature increases will reduce future crop yields [16, 2008]. Further, Turvey finds evidence that county-level corn yield variability can be attributed to heat and rainfall events from June 1 to August 31 [19, 2001].

There is an extensive literature surrounding speculation in commodity futures markets. Speculation activity could feasibly be correlated with temperature volatility, which could lead to omitted variable bias in my primary regressions if not controlled for. The review from Haase, Zimmermann, and Zimmermann [7, 2016] evaluates 100 papers offering differing evidence on the impact speculation has on commodity futures. In their review, they find that the number of studies supporting and contradicting the criticized effects of speculation is about the same. In this paper we do not have the resources to account for speculation activity, so we assume that speculation does not impact corn futures price dynamics. At the very least, it's believable that any temperature-induced speculation activity is seasonal in the same manner that temperature is, so bias aggregated over a large enough time horizon would effectively be zero.

There is evidence suggesting that the implementation of biofuels has caused the global cash-market price of corn to increase [9, Mitchell, 2008]. Conversely, a 2012 paper from Trujillo-Barrera, Mallory, and Garcia [18] found evidence of US corn futures variability directly attributed to volatility in the crude oil market, and no evidence of transmission from the ethanol market to the corn market. In a 2013 paper, Gardebroek and Hernandez [5] found no evidence of cross-volatility effects from oil to US corn markets, and no evidence of cross-volatility effects from ethanol to US corn markets. In addition to this, fuel prices exhibit seasonality similar to temperature; however, fuel prices are determined on a global market and should be independent of demeaned daily temperature volatility in the American corn belt. Hence, I do not control for the implementation of any biofuel regulation in my analysis.

For measuring corn futures returns volatility and temperature volatility, a Generalized Auto-Regressive Conditionally Heteroskedastic (GARCH) (1,1) model is used [8, Hansen & Lunde, 2005]. The GARCH model [4, Engle & Bollerslev, 1986] is a generalization of Engle's Auto-Regressive Conditionally Heteroskedastic (ARCH) model [3, 1982]. The conditional variance of a time series passed through a GARCH filter can exhibit heteroskedasticity, but exhibits an unconditional auto-covariance function. There is an exhaustive list of financial papers that make use of (at least some variation of the) GARCH model that we will not list here. Use of the GARCH model in environmetrics started to become prevalent about a decade after its discovery [17, Tol, 1996], and I use it as my primary model to generate the temperature volatility statistic as it allows for temperature volatility to exhibit clustering. Moreover, the temperature volatility generated as a GARCH filter estimate allows me to forego the need to establish an arbitrary reference period, which is a commonplace technique in the construction of temperature shock variables.

The current standard for measuring temperature effects on agricultural yields is either a measure known as degree-days (DDs) [1, 13, Deschênes & Greenstone, 2007; Ritchie & Nesmith, 1991] or by use of including a quadratic temperature term [20, Urban, Roberts, Schlenker, & Lobell, 2011]. There are a variety of DD derivatives available for trade on the Chicago Mercantile Exchange with the primary DD being related to natural gas futures and calculated as the sum of heating-degree days (HDD) and cooling-degree days (CDD); that is,

$$DD_t = HDD_t + CDD_t$$

where t is a given day. Using degrees Fahrenheit, the formulas for HDD_t and CDD_t are given by

$$HDD_t = 65 - T_{\min}$$

$$CDD_t = T_{\max} - 65$$

where $T_{\min}$ is the minimum recorded temperature on day t and $T_{\max}$ is the maximum recorded temperature on day t. Thus, DD_t effectively measures the total degrees of heating or air conditioning being used on day t. The DD used for natural gas is not numerically identical to the DD used in agriculture; however, the agricultural DD calculation is structurally equivalent to the natural gas DD.

Mu provided evidence that natural gas futures price dynamics are affected by temperature volatility [10, 2007]. In the paper, he uses the following weather shock variable

$$W_t = \frac{1}{m}\sum_{i=1}^{m}(DD_{t+i} - DDNORM_{t+i})$$

where W_t is the weather shock on day t, m is the weather forecast horizon set to 7, DD_{t+i} is the forecast degree day on day $t+i$, and $DDNORM_{t+i}$ is the average degree days of the previous 30 years on day $t+i$.

Note that Mu's W_t variable is forward looking at temperature forecasts, and my temperature volatility uses temperature realizations. Even if W_t was modified to look back at temperature realizations, I still worry that extreme DD deviations in opposite directions occurring in the same week would cancel out and make the temperature shock variable small in magnitude, not capturing the effect of extreme deviations[7]. As I claimed above, Mu's W_t also relies on an average DD over an arbitrary reference period of 30 years despite there being little to no mathematical reasoning for the length of this reference period.

In an influential paper, Roll found statistical evidence that weather shocks affect orange juice futures returns [14, 1984], where weather shocks were measured as a difference in actual temperature and forecasted temperature. Temperature forecasts were reported 12, 24, and 36 hours in advance, and separate regressions were run for both AM maximum temperatures and PM minimum temperatures. In my paper, I measure temperature volatility which does not rely on temperature forecasts, allowing me to model how temperature is impacting the corn futures market as opposed to how weather reporters impact the corn futures market.

The use of forecast data also led to a messy interpretation of precipitation's effect on the orange juice futures market. Forecasts were given as a probability (measured in percent) rounded to the nearest factor of ten. This caused Roll to measure precipitation error as

$$\hat{\varepsilon}_t = A_t - \hat{a} - \hat{b}F_t$$

where $\hat{\varepsilon}_t$ is the precipitation error on day t, A_t is the actual type of precipitation on day t, $\hat{a}$ is a constant, $\hat{b}$ is a constant, and F_t is the precipitation forecast on day t as a percentage rounded to the nearest factor of ten. This is not an intuitive statistic because actual rainfall and predicted rain fall were weirdly binned. I only use actual rainfall measurements, allowing me to forego a binning process and making my precipitation variable easier to interpret.

[7] To his credit, Mu does note in his paper that he tried different weather shock constructions that accounted for this and found no change in his results. For a discussion on how my temperature volatility statistic may also fall victim to this, see section 5.3.

In Roll's paper, he explicitly states that orange juice futures are chosen because of the geographical concentration of orange production. In this paper, I am able to bypass the lack of a geographical concentration of corn production by weighting my temperature data by the share of corn produced by a state, so states with negligible corn production have a negligible affect on my temperature volatility statistic.

3 Data

In the following two subsections, I give an in-depth description of the data used and the variables constructed from the data, which I later use in my analysis.

3.1 Data Sources

To construct the temperature volatility statistic, my key independent variable, I use two different data sets: USDA corn production data and data from NOAA NCDC.

The first data set can be found using the Quick Stats tool on the USDA's National Agricultural Statistics Service website. I use a cross-sectional data set that reports the total bushels of corn produced by each U.S. state in 2020. I use this data to generate the share of corn produced by each state in 2020 and then use these shares as a weighting scheme to generate my temperature volatility statistic. The corn production data allows me to narrow my focus to states that lead in corn production as opposed to using national temperature averages. For example, in 2020 North Dakota produced approximately 1.75% of all the corn in the U.S. whereas Iowa produced approximately 16.18% of all the corn. Demeaned temperature data from North Dakota is included in the state-level aggregation for one of the regions analyzed in this study; however, it is weighted by approximately 0.0198 (not 0.0175 because weighting coefficients are scaled, so they sum to one), making data in this part of the region less significant than the data coming out of Iowa which is weighted by 0.1824. For specifics on what states are included in the regions analyzed in this paper, see section 4 and table 1.

For the second data set, I use NOAA's file transfer protocol to create data sets for each U.S. state, where each state file contains a variety of different daily measurements from each registered weather station in the state[8]. As mentioned above, I do not use the weather data for every state in the U.S., rather I focus on the data coming from states that led in corn production in 2020. The daily weather station data ranges from January 1, 1973 to December 31, 2020; however, weather stations are not guaranteed to have data available in this exact range. If this is the case, any available data in the range above is pulled. We use the max and

[8]A special thanks to Dr. Daniel Kaffine (Professor of Economics at the University of Colorado Boulder) who provided me with the MATLAB code that generated these data files.

min daily temperatures from the weather station data, measured in degrees Fahrenheit, to approximate daily average temperature by averaging the two values. This average daily temperature is then used to build our temperature volatility statistic. The weather station data files also contain a variable that measures total daily precipitation in inches which is used as a control variable in my primary regressions. After generating GARCH filter estimates, climate variables are restricted to trading day observations and exclude any leap day observations as well as any holidays where the corn futures market is closed. Specifically , the date range is limited to being between January 2, 1980 through December 31, 2020 because this is the range for which I was able to obtain corn futures end of day data.

To construct my corn futures price levels and volatility, my key dependent variables, I use daily closing prices from investing.com which records end of day market statistics as reported by the CME group. The closing price level is measured in U.S. dollars per contract. Throughout the course of the year, five different corn futures contracts reach maturity on the business day prior to the fifteenth of either March, May, July, September, or December which marks the last trading day of the contract month. Trading hours are 7:00 PM to 7:45 AM Central Time, Sunday through Friday and 8:30 AM to 1:20 PM Central Time, Monday through Friday with modified hours on specific holidays. My sample is the closing prices of the latter trading times excluding any leap day data and market closures. Further, the closing price data is of full-sized contracts which are worth 5,000 bushels of corn.

My analysis uses the closing price for the contract month that is closest to maturity. For example, if we are working with returns data from April 10, 2023, then the contract being observed is the May 2023 contract. Since I could only get corn futures price data for the corn futures contract closest to maturity, I have to counteract the Samuelson effect [15, Samuelson, 1965] which posits that a futures contract price volatility increases as the contract nears maturity. In my primary regressions, I control for the time to maturity with a collection of dummy variables which track how far out a contract is from reaching maturity. This is an admittedly crude way of controlling for the Samuelson effect, and a discussion on how this may impact my results can be found in section 5.2.

3.2 Data Processing

I start with a discussion on the construction of my temperature volatility statistic. I define T_0 as the set containing the enumeration of dates ranging from January 1, 1973 through December 31, 2020, excluding leap years.

I first aggregate weather-station level data up to a state level by simple averaging. That is, if there are N_{st}

weather stations in state s at time t, then

$$f_{st} = \frac{1}{N_{st}} \sum_{i=1}^{N_{st}} w_{ist} \tag{1}$$

where f_{st} is the average temperature in degrees Fahrenheit for state s at time t and w_{ist} is the weather-station level average temperature for weather station i in state s at time t.

Next, I want to aggregate the state level data to a regional level. In this paper, we analyze three different regions of corn production. Regions are comprised of contiguous, lead corn-producing states with similar climates. As mentioned in section 3.1, in aggregating the state-level average temperature to a regional-level average temperature, we weight the temperature data by the share of corn produced in state s such that the weights applied sum to one. The share of corn produced for each state in 2020 as reported by the USDA is listed in Table 1 below:

Table 1: Share of Corn Produced in 2020

Region	State	Share of Corn Produced (%)
	Iowa	16.1804787
	Illinois	15.023971
	Nebraska	12.6152885
Region 1	Minnesota	10.1648668
	Indiana	6.9571169
	Kansas	5.431618
	South Dakota	5.1086178
Region 2	Ohio	3.9988806
	Missouri	3.974645
	Wisconsin	3.5920478
	Michigan	2.1576098
	Kentucky	1.7733119
Region 3	North Dakota	1.7533281

Notes: Region 2 includes the states listed in region 1. Similarly, region 3 contains the states listed in region 1 and region 2. Cumulatively, region 1 accounts for approximately 53.98% of corn production in 2020, region 2 accounts for approximately 75.48% of corn production in 2020, and region 3 accounts for approximately 88.73% of corn production in 2020. Analysis using regions 2 and 3 act primarily as robustness checks for analysis done using temperature data from region 1 only.

I aggregate the data over the states in the region being used in my analysis via a weighted average; that is, if our analysis focuses on the states in region $\mathcal{R}$, then

$$F_t = \sum_{s \in \mathcal{R}} \lambda_s f_{st} \tag{2}$$

where F_t is the regional, average temperature at time t, λ_s is the weighting coefficient for state s such that $\sum_{s \in \mathcal{R}} \lambda_s = 1$, and f_{st} is the average temperature in state s at time t, as calculated in (1).

Next, I detrend and deseasonalize the time series $\{F_t\}_{t\in T_0}$ by regressing F_t on a linear trend and a collection of dummy variables where the j^{th} dummy variable takes value 1 if t is the j^{th} day of the year. Mathematically, this can be written as

$$F_t = \beta_0 + \beta_1 t + \sum_{j=1}^{364} \beta_{j+1} I_{day(t)=j} + \varepsilon_t \tag{3}$$

where F_t is the regional average daily temperature at time t, β_i is a regression coefficient for $i = 0, ..., 365$, $I_{day(t)=j}$ is an indicator function taking value 1 if t is the j^{th} day of the year, and ε_t is the error term. Note that the residuals of (3) are a demeaned time series.

Next, we fit a GARCH(1,1) model [4, Engle & Bollerslev, 1986] to our newly constructed time series from (3), assuming a Gaussian white noise error term[9], which has a skedastic function of the form

$$x_t^2 = a_0 + a_1\hat{\varepsilon}_{t-1}^2 + bx_{t-1}^2$$

where $a_1, b \geq 0$, $a_1 + b \leq 1$, and $\hat{\varepsilon}_t$ is the residual from (3). I then derive GARCH filter estimates, $\hat{x}_t$. Note again that this temperature volatility variable differs from the existing literature by giving heavier weight to states which produce more corn and by utilizing a GARCH filter to generate filter estimates.

I perform a near identical process to construct my precipitation control variable, except I do not fit a GARCH model to the data, leaving me with demeaned, regional precipitation p_t, measured in inches.

Note that seasonal adjustments are a necessary step because time series exhibiting seasonality are not stationary, which prohibits me from using a GARCH model and will lead to spurious regressions in my primary analysis [6, Granger & Newbold, 1974]. To see how seasonal data is non-stationary consider the expected average temperature in Iowa on January 1st of a given year. This almost certainly will not equal the expected average temperature in Iowa on August first of a given year. Thus, the expected value of the time series is not time invariant, implying that the process is non-stationary.

For my closing price data, I perform a log-differencing transformation to the data which is common throughout the financial literature. The log-differenced data, or the log-returns, give an approximation to the percentage change in daily closing prices of corn futures contracts. I use log-returns in estimating the effect of temperature volatility on corn futures price levels.

For constructing my price volatility data, I fit a GARCH(1,1) model [4, Engle & Bollerslev, 1986] to the log-returns of the corn futures closing price [8, Hansen & Lunde, 2005]. We assume a Gaussian white noise

[9] Use of a t distribution has little impact on the model coefficients.

error term[10] and a skedastic function of the form

$$\sigma_t^2 = a_0 + a_1 r_{t-1}^2 + b\sigma_{t-1}^2$$

where $a_1, b \geq 0$, $a_1 + b \leq 1$, $r_t = \log(P_t) - \log(P_{t-1})$, and P_t is the closing price at time t. Corn futures price volatility at time t is then the GARCH filter estimate, $\hat{\sigma}_t$.

3.3 Summary Statistics

Table 2 lists summary statistics of returns, returns volatility, temperature, temperature volatility, and precipitation. I don't list means for temperature and precipitation because they have been demeaned and, consequently, have empirically negligible means.

I perform both an Augmented Dickey-Fuller [2, Walter, Enders, & Granger, 1998] test and a Phillip-Perron unit root test [12, Phillip & Perron, 1988] on all the variables being used in my primary regressions which include log-returns, returns volatility, temperature volatility, and precipitation. For all variables, I am able to reject the null hypothesis that a unit root exists at the 1% significance level.

Note that the standard deviations of the temperature and precipitation statistics are similar across regions. This is evidence that the contiguous states contained in the different regions do in fact exhibit similar climates, or at the very least, weighting the climate data by the share of corn produced by each state is not allowing states with little corn production and substantially different climates to have a large effect on the climate variables.

The GARCH filter coefficient estimates for temperature provide statistically significant evidence that demeaned temperature exhibits volatility clustering. The maximum sum of a_1 and b across the regions is 0.961. This comes close to the boundary of 1; however, we have a sample size of 17,520, which means these estimates are quite precise and accurate. That is, we need not worry about integrating the temperature volatility statistic. This is further evidenced by the temperature volatility statistic rejecting the null hypotheses for the unit-root tests used.

On the other hand, the GARCH filter coefficient estimates for log-returns sum to 0.988, which, despite our sample size of 10,396 (we have data for December 28, 1979 and December 31, 1979 that get cut in the main analysis but are used in estimating model coefficients) and the rejection of the null hypotheses for the unit-root tests, is still reason to consider integrating our model due to this sum being so close to one. For a further discussion on this, see section 5.1

[10] Use of a t distribution has little impact on the model coefficients.

Table 2: Summary Statistics

	Returns		
Mean	$4.195{\times}10^{-5}$		
Standard Deviation	0.015		
	Returns Volatility		
a_0 ($\times 10^{-6}$)	1.924^{***}		
	(3.911×10^{-7})		
a_1	0.061^{***}		
	(0.004)		
b	0.927^{***}		
	(0.004)		
Mean	$2.428{\times}10^{-4}$		
Standard Deviation	$1.952{\times}10^{-4}$		
	Temperature_1	Temperature_2	Temperature_3
Standard Deviation	5.414	5.152	5.094
	Temp Volatility_1	Temp Volatility_2	Temp Volatility_3
a_0	2.455^{***}	1.916^{***}	1.729^{***}
	(0.061)	(0.054)	(0.049)
a_1	0.896^{***}	0.916^{***}	0.919^{***}
	(0.025)	(0.023)	(0.024)
b	0.045^{***}	0.039^{***}	0.042^{***}
	(0.005)	(0.005)	(0.005)
Mean	4.523	4.269	4.200
Standard Deviation	3.097	3.009	3.004
	Precipitation_1	Precipitation_2	Precipitation_3
Standard Deviation	0.068	0.084	0.162

Notes: The index on 'Temperature', 'Temp Volatility', and 'Precipitation' corresponds to the region being analyzed. For specifics on region construction, see Table 1 and section 4. Further, all of the climate variables listed are demeaned. In our regressions (detailed in section 4), we use the square-root of 'Temp Volatility', so the mean and standard deviation reported are for the square root of the GARCH filter estimate for temperature volatility. All values in parentheses correspond to standard errors of estimated coefficients.

4 Methodology

This paper focuses on the analysis of the following two regression models:

$$r_t = \beta_0 + \beta_1 \hat{x}_{t-m} + \beta_2 p_{t-m} + \beta_3 p_{t-m}^2 + \gamma_t \delta_3 + \mu_t \tag{4}$$

$$\hat{\sigma}_t = \delta_0 + \delta_1 \hat{x}_{t-m} + \delta_2 p_{t-m} + \delta_3 p_{t-m}^2 + \gamma_t \delta_3 + \nu_t \tag{5}$$

In equation (4), r_t is the log-returns of the closing price of the corn futures contracts at time t as discussed in section 3.2. In equation (5), $\hat{\sigma}_t$ is the volatility estimate for the closing price of corn futures contracts at time t as discussed in section 3.2. In both equation (4) and (5), $\hat{x}_{t-m}$ is the regional demeaned temperature volatility at time $t-m$ as discussed in section 3.2. In both equations (4) and (5), p_{t-m} is the regional demeaned precipitation level at time $t-m$. In both equations (4) and (5), γ_t is a row vector of dimension 90 containing the time to maturity controls, so the the j^{th} component of γ_t takes value 1 if and only if the corn futures contract being analyzed at time t is j days from reaching maturity. The error term at time t for regression (4) is μ_t, and ν_t is the error term at time t for regression (5). Additionally, t is restricted to only take on values from T, which we define as the set containing the enumeration of dates for which the corn futures market is open (see section 3.1).

Note that the climate variables in both regressions are lagged by an order of $m \geq 1$, so we can argue that the climate variables actually factor into the information set available to traders at time t. For the contracts we are analyzing, markets close at 1:20 PM (CT); however, temperature and precipitation measurements can not possibly be reported for time t by 1:20 PM (CT) as they are daily observations for time t. Hence, traders can not factor $\hat{x}_t$ and p_t into their information set until, at the earliest, time $t+1$. I experimented with values of $m \in \{1, 7, 30, 90, 365\}$ to simulate how climate conditions of the previous day, climate conditions a week ago, climate conditions a month ago, climate conditions three months ago, and climate conditions one year ago affect corn futures returns and returns volatility. I report regression results for $m = 7$ because the data is available to traders and the climate conditions from a week prior are likely very similar to current climate conditions. Moreover, my results are robust to all the values of m listed. To see the effect of varying m on the regression coefficient estimates, see the tables in section 8, appendix 1.

To account for the possibility of our returns volatility not being stationary, we also estimate (5) with integration order 1, giving us a model of the form

$$
\begin{aligned}
\Delta\sigma_t &= \Delta(\delta_0 + \delta_1\hat{x}_{t-m} + \delta_2 p_{t-m} + \delta_3 p_{t-m}^2 + \gamma_t\delta_3 + \nu_t) \\
\Delta\sigma_t &= \delta_0 + \delta_1\hat{x}_{t-m} + \delta_2 p_{t-m} + \delta_3 p_{t-m}^2 + \gamma_t\delta_3 + \nu_t \\
&\quad - (\delta_0 + \delta_1\hat{x}_{t-m-1} + \delta_2 p_{t-m-1} + \delta_3 p_{t-m-1}^2 + \gamma_{t-1}\delta_3 + \nu_{t-1}) \\
\Delta\sigma_t &= \delta_1(\Delta\hat{x}_{t-m}) + \delta_2(\Delta p_{t-m}) + \delta_3(\Delta p_{t-m}^2) + (\Delta\gamma_t)\delta_4 + residual_t
\end{aligned}
\tag{6}
$$

where $\Delta = I - B$, I is the identity operator, and B is the back-shift operator.

I will run regressions (4), (5), and (6) using three different regions (to see the share of corn produced by each state and the cumulative share of corn produced by the region, see table 1):

(i) I begin by only using data from the 4 leading corn-producing states in 2020 which are Iowa, Illinois, Nebraska, and Minnesota. Remark that these states are in a contiguous region and experience similar climates for large potions of the year.

(ii) As a first-round robustness check, I use data from the top 8 corn producing states in 2020 which are all the states mentioned in scenario (i), Indiana, Kansas, South Dakota, and Ohio. Remark again that these states are in a contiguous region and experience similar climates for large portions of the year.

(iii) As a final robustness check, I use data from the top 13 corn producing states which includes all the states mentioned in scenarios (i) and (ii), Missouri, Wisconsin, Michigan, Kentucky, and North Dakota. Again, these states are in a contiguous region and experience similar climates for large portions of the year.

The parameter of interest from (4) is β_1. If $\beta_1 > 0$, then the percentage change in daily closing prices at time t increases as temperatures deviate from their expected level at time $t - m$. If $\beta_1 < 0$, then the percentage change in daily closing prices at time t decreases as temperatures deviate from their expected level at time $t - m$. If $\beta_1 = 0$, then the percentage change in daily closing prices at time t is unaffected by temperature deviations from their expected level at time $t - m$.

The parameter of interest from (5) and (6) is δ_1. If $\delta_1 > 0$, then daily volatility of the percentage change of closing prices at time t increases as temperatures deviate from their expected level at time $t - m$. If $\delta_1 < 0$, then daily volatility of the percentage change of closing prices at time t decreases as temperatures deviate from their expected level at time $t - m$. If $\delta_1 = 0$, then daily volatility of percentage change of closing prices at time t is unaffected by temperature deviations from their expected level at time $t - m$. Specifically for regressions (5) and (6), if we see drastic difference in the level of significance and the value of our estimate for δ_1, then the results of (5) are likely spurious due to returns volatility being non-stationary.

Table 3: Results from Regression (4)

		Region 1		
Constant	0.000039	0.017874***	0.017856***	0.017875***
	(0.000282)	(0.005729)	(0.005726)	(0.005720)
$\hat{x}_{t-m}$	0.000001	-0.000010	-0.000012	-0.000013
	(0.000049)	(0.000049)	(0.000049)	(0.000049)
p_{t-m}			-0.002589	-0.002192
			(0.002272)	(0.003614)
p^2_{t-m}				-0.002479
				(0.015803)
		Region 2		
Constant	-0.000006	0.017832***	0.017804***	0.017838***
	(0.000278)	(0.005729)	(0.005727)	(0.005725)
$\hat{x}_{t-m}$	0.000012	0.000002	-0.000002	-0.000002
	(0.000051)	(0.000051)	(0.000051)	(0.000051)
p_{t-m}			-0.002415	-0.001922
			(0.001853)	(0.002884)
p^2_{t-m}				-0.002900
				(0.011765)
		Region 3		
Constant	-0.000047	0.017796***	0.017770***	0.017800***
	(0.000276)	(0.005729)	(0.005727)	(0.005723)
$\hat{x}_{t-m}$	0.000023	0.000013	0.000010	0.000009
	(0.000051)	(0.000051)	(0.000051)	(0.000051)
p_{t-m}			-0.001223	-0.000995
			(0.000952)	(0.001491)
p^2_{t-m}				-0.000657
				(0.002990)

5 Results and Discussion

5.1 Results

Regression results can be found in Tables 3, 4, and 5 for regressions (4), (5), and (6), respectively. Statistical significance of estimators is derived using a t-statistic with a Newey-West heteroskedastic autocorrelated consistent (HAC) standard error [11, 1987]. For all of the results tables, the first column gives the results of a regression model with no controls, the second column fives the results of a regression model including time to maturity controls (i.e. γ_t is a regressor), the third column gives the results of a regression model including time to maturity controls and a linear precipitation control (i.e., γ_t and p_t are regressors), and the fourth column gives the results of a regression model including time to maturity controls and a quadratic precipitation control (i.e., γ_t, p_t, and p_t^2 are regressors). The level of significance is denoted with asterisks, so *** (**, *) denote significance at the 1% (5%, 10%) level. All parenthetical values are Newey-West standard errors. As was mentioned in section 4, the lag coefficient on the climate variables for the regression results reported in the tables is 7.

Table 4: Results from Regression (5)

		Region 1		
Constant	0.016146***	0.016403***	0.016405***	0.016180***
	(0.000498)	(0.000596)	(0.000597)	(0.000625)
$\hat{x}_{t-m}$	-0.000322***	-0.000328***	-0.000328***	-0.000324***
	(0.000043)	(0.000043)	(0.000043)	(0.000043)
p_{t-m}			0.000181	-0.004463**
			(0.001064)	(0.002180)
p^2_{t-m}				0.029011***
				(0.009039)
		Region 2		
Constant	0.016081***	0.016335***	0.016335***	0.016091***
	(0.000499)	(0.000600)	(0.000601)	(0.000632)
$\hat{x}_{t-m}$	-0.000327***	-0.000334***	-0.000334***	-0.000329***
	(0.000046)	(0.000046)	(0.000046)	(0.000046)
p_{t-m}			-0.000014	-0.003558**
			(0.000925)	(0.001753)
p^2_{t-m}				0.020839***
				(0.006204)
		Region 3		
Constant	0.016087***	0.016321***	0.016320***	0.016089***
	(0.000497)	(0.000594)	(0.000596)	(0.000624)
$\hat{x}_{t-m}$	-0.000334***	-0.000341***	-0.000341***	-0.000337***
	(0.000046)	(0.000046)	(0.000046)	(0.000046)
p_{t-m}			-0.000018	-0.001811**
			(0.000465)	(0.000890)
p^2_{t-m}				0.005173***
				(0.001590)

Table 5: Results from Regression (6)

		Region 1		
$\hat{x}_{t-m}$	-0.000000	0.000001	0.000001	0.000001
	(0.000006)	(0.000006)	(0.000006)	(0.000006)
p_{t-m}			-0.000010	-0.000103
			(0.000117)	(0.000194)
p^2_{t-m}				0.000518
				(0.000702)
		Region 2		
$\hat{x}_{t-m}$	-0.000004	-0.000001	-0.000001	-0.000001
	(0.000007)	(0.000007)	(0.000007)	(0.000007)
p_{t-m}			0.000034	-0.000005
			(0.000095)	(0.000157)
p^2_{t-m}				0.000201
				(0.000525)
		Region 3		
$\hat{x}_{t-m}$	-0.000003	-0.000000	-0.000000	-0.000000
	(0.000008)	(0.000008)	(0.000008)	(0.000007)
p_{t-m}			0.000009	-0.000010
			(0.000049)	(0.000081)
p^2_{t-m}				0.000048
				(0.000131)

From table 3, we see that we only have statistical significance in our constant term; hence, we have no statistically significant evidence that temperature volatility affects corn futures returns. We discuss why this may be the case in section 5.2.

From tables 4 and 5, we see that the direction of the coefficients in both regressions are mostly the same with the exception of the coefficient on $\hat{x}_{t-m}$ in the region 1 regression. Magnitudes of the coefficients in the two regressions, however, differ substantially. Moreover, in our most robust regression model, we get at least 5% significance for all regression coefficients in each region. On the other hand, no regression coefficients are statistically significant after differencing, giving reason to believe that we encountered type one errors in our tests for unit roots in returns volatility.

5.2 Discussion

From an econometric prospective, I thought it possible that we may be failing to reject the null hypothesis that the temperature volatility coefficient is zero due to my non-parsimonious model. My method for controlling for the Samuelson Effect [15, Samuelson, 1965] is admittedly crude as it involves including 90 regressors in my model. This lack of parsimony could lead to model over-specification which makes my coefficient estimators inefficient, and could result in type two errors; however, we observe practically no change in the standard error for the temperature volatility coefficient as more regressors are included. This is likely because temperature volatility is, empirically, weakly correlated with the time to maturity of a futures contract.

The low hanging fruit explanation for why I find no statistical evidence of temperature volatility affecting corn futures returns and returns volatility is that temperature volatility has no effect on corn futures prices. In section 2, I discussed the use of quadratic temperature terms and DD variables to determine crop production. Despite corn futures prices being closely tied to corn production, it's possible that daily temperature deviations from their mean are too insignificant to have the same effect as looking at just temperature. In table 2, the largest mean and standard deviation of the temperature volatility is 4.523 and 3.097, respectively. Moreover, the largest temperature volatility observed is approximately 27.2; that is, the most I ever see temperature deviate from it's mean is by 27.2 degrees Fahrenheit. While this deviation could have a very notable effect on crop production if sustained for a long enough period, traders may not observe enough sustained deviations of this magnitude to make actions in the futures market that are statistically detectable using my models. In figure 1, I plot the temperature volatility statistic constructed in each region with the horizontal axis measuring the days from January 1, 1973 to help with illustrating this point.

Many of the studies cited in section 2 used climate variables that were forward looking; that is, the climate

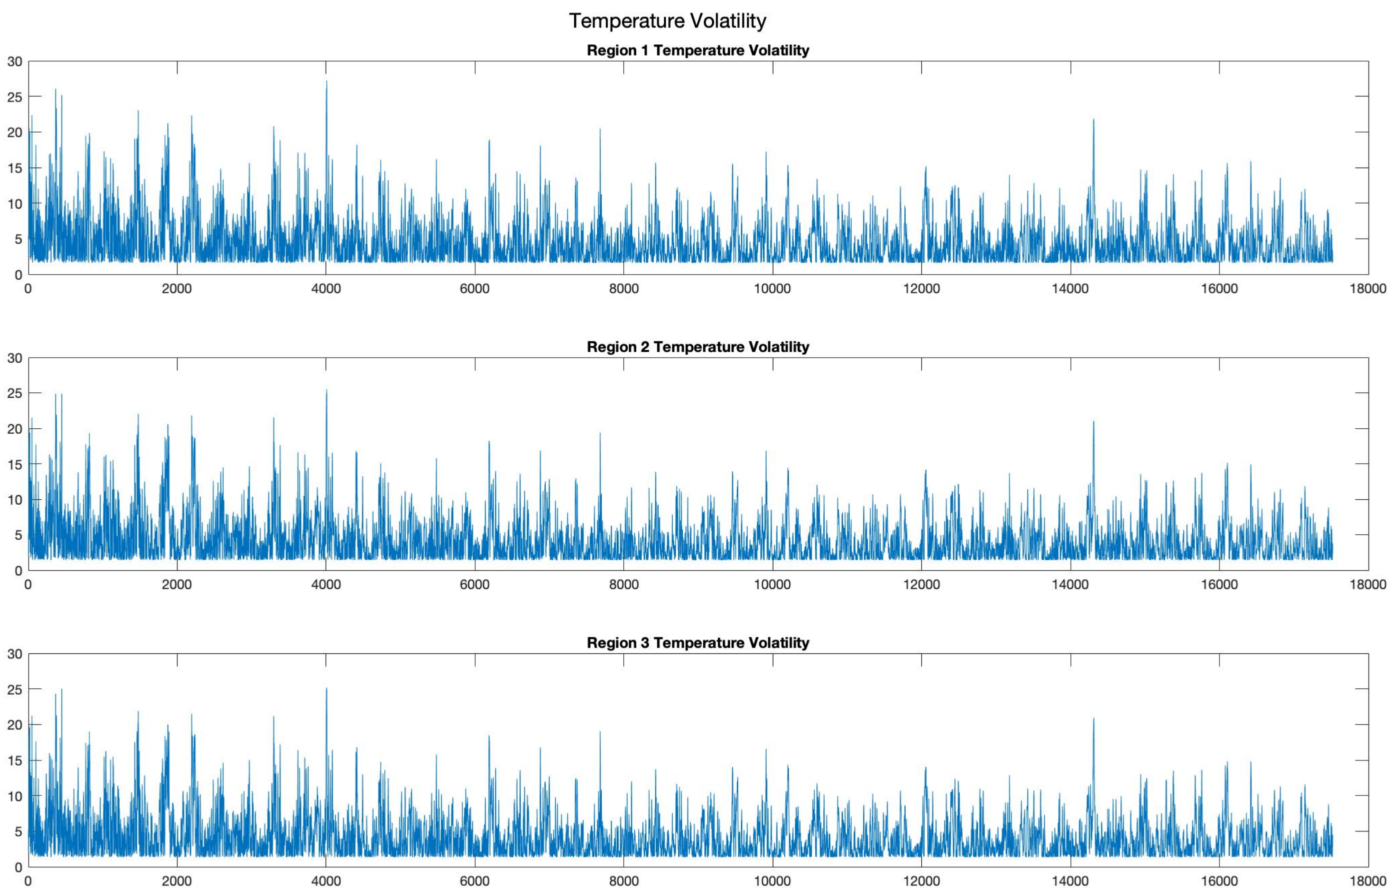

Figure 1: Temperature Volatility

variables used as regressors incorporated temperature and precipitation forecasts. The focus of this study is to see how climate realizations affect the corn futures market; however, it's possible that many traders are forward looking, and consequently, give much more weight to future forecasts as opposed to past observations when making decisions in the futures market. This could potentially lead to observed temperature deviations from their mean having a null effect on corn futures returns and returns volatility. The lack of statistical significance in the quadratic precipitation coefficients is possibly further evidence of this.

5.3 Limitations and Suggested Improvements

The Roll paper previously mentioned focused on orange juice futures because the vast majority of American orange production in America occurs around Orlando, Florida [14, 1984]. This allowed Roll to focus on climate conditions in a single geographical area. I used a brute force method to aggregate temperature and precipitation to a regional level by averaging granular weather station observations up to a regional level. There is almost certainly a cleaner way to go about data aggregation at this scale that could give a better measure of the temperature and precipitation experienced by the regions used in this study. My method inevitably resulted in a loss of information through repeated aggregating, which could have skewed my results.

I acknowledge that the weighting scheme I employed to aggregate state level climate data up to a regional level was fairly arbitrary. I chose 2020 because it was the last year I collected data for, and the lead corn producing states don't vary wildly from year to year. Future studies looking to enhance my climate variables could redefine the weights used every year, so the climate data in state s at time t is weighted by the normalized share of corn production in state s during year $year(t)$. This has the potential to redefine the states that make up each region on a yearly basis, which may change how results are interpreted, but should also make the temperature volatility and precipitation more representative of the lead corn producing regions throughout the time horizon.

Another potential oversight of my temperature volatility statistic is the order in which I aggregated, demeaned, and fit a GARCH model to the temperature data. The squared temperature volatility statistic used in my regressions has the following expanded form:

$$a_1 \left[\sum_{s \in \mathcal{R}} \left(\frac{\lambda_s}{N_{s(t-m-1)}} \sum_{i=1}^{N_{s(t-m-1)}} w_{is(t-m-1)} \right) - \beta_0 - \beta_1(t-m-1) - \sum_{j=1}^{364} \beta_{j+1} I_{\{day(t-m-1)=j\}} \right]^2 + a_0 + bx_{t-m-1}^2 \tag{7}$$

All variables in this equation are defined in sections 3.2 and 4. Note that in aggregating weather-station level data up to a regional level, we are potentially allowing for deviations of temperature from their expected level in opposite directions to cancel out. For example, if the temperature in Iowa at weather station i is above the expected level and the temperature in Iowa at weather station j ($i \neq j$) is below the expected level at the same magnitude for which the temperature at weather station i is higher than expected, then after averaging, the temperature will be at its expected level. To construct a temperature volatility that is robust to these deviations, I could have demeaned the weather-station level temperature data and fit a GARCH(1,1) model to the demeaned data. I then could aggregate the square-root of weather-station level temperature volatility up to a regional level using the same weighting scheme as before. This would give a temperature volatility, with the following expanded form

$$\begin{aligned} &a_1' \sum_{s\in\mathcal{R}} \frac{\lambda_s}{N_{s(t-m)}} \sum_{i=1}^{N_{s(t-m)}} \left(w_{is(t-m-1)} - \beta_0' - \beta_1'(t-m-1) \sum_{j=1}^{364} \beta_{j+1}' I_{\{day(t-m-1)=j\}} \right)^2 \\ &\quad + a_0' + b' y_{is(t-m-1)}^2 \end{aligned} \tag{8}$$

where y_{t-m-1} is the weather-station level skedastic function with coefficients a_0', a_1', and b'. The β coefficients in the demeaning process used to construct (8) are numerically different than the β's in (7), but all other variables are the same.

Furthermore, I could not find geographical coordinates for weather station data. This resulted in me using data from every weather station reporting temperature and precipitation in my states of interest throughout the time period in which I collected climate data. If geographical coordinates could be found, then the weather stations in more urban areas where farming does not occur could be removed from the study. Instead of using a uniform weight in the weather station data aggregation process, a weight could be devised that accounts for how much area we claim experiences the climate reported by the weather station.

In section 5.2, I argued that the inclusion of the time to maturity controls did not affect the inference of the study. That being said, there are simple ways of overcoming the need for these controls. The simpler method, which was not possible in this study due to limitations in corn futures data accessibility, is to use the contract next closest to maturity as opposed to using the contract closest to maturity. The trading volume of the contract next closest to maturity is comparable to the trading volume of the contract closest to maturity (prior to the contract reaching its maturity horizon). Moreover, there's little chance traders working with the contract next closest to maturity are weighting climate forecasts over past climate observations because climate conditions are generally not reliable past a week, and the next-week forecasts will not give a good gauge on climate conditions for the next growing season. Even null results in a study using closing prices of the futures contract next closest to maturity would at least narrow down the possible reasons for the null results.

6 Conclusion

In this study, I tested if temperature volatility had an effect on corn futures price levels and corn futures price volatility. I found no statistical evidence that temperature volatility affects corn futures returns and no statistical evidence that temperature volatility affects corn futures returns volatility. A better method for climate variable aggregation and an improved corn futures price data set could lead to different results. Future work may include expanding this study to different agricultural futures contracts and implementing different climate observation metrics.

7 Acknowledgements

Special thanks to my advisor, Dr. Adam McCloskey, who provided guidance and support throughout the writing of this thesis. Thank you also to my defense committee for taking the time to read and give feedback on my thesis. Special thanks also to Dr. Warren Hrung, J.P. Rothenberg, and the American Bankers Association for supporting me throughout the writing process and providing feedback. All errors remain my own.

8 Appendix: Varying the Time Lag of the Regressions

All of the tables below are regressions containing time to maturity and quadratic precipitation controls.

Table 6: Results from Regression (4) for Varying m

	$m = 1$	$m = 30$	$m = 90$	$m = 365$
		Region 1		
Constant	0.017554***	0.018095***	0.017879***	0.017811***
	(0.005742)	(0.005755)	(0.005722)	(0.005733)
$\hat{x}_{t-m}$	0.000049	0.000047	0.000017	0.000053
	(0.000046)	(0.000048)	(0.000053)	(0.000048)
p_{t-m}	-0.002557	0.001400	0.004704	0.004311
	(0.003600)	(0.004159)	(0.003448)	(0.003923)
p^2_{t-m}	0.012890	-0.060929**	-0.005218	-0.030237*
	(0.016245)	(0.025755)	(0.016170)	(0.018294)
		Region 2		
Constant	0.017551***	0.018142***	0.017796***	0.017750***
	(0.005742)	(0.005748)	(0.005718)	(0.005734)
$\hat{x}_{t-m}$	0.000064	0.000056	0.000010	0.000056
	(0.000047)	(0.000050)	(0.000055)	(0.000050)
p_{t-m}	-0.001221	0.002525	0.002820	0.003217
	(0.002927)	(0.003169)	(0.002809)	(0.003068)
p^2_{t-m}	0.006412	-0.044150***	0.003808	-0.015422
	(0.012509)	(0.017059)	(0.013341)	(0.012645)
		Region 3		
Constant	0.017545***	0.018098***	0.017863***	0.017774***
	(0.005746)	(0.005747)	(0.005718)	(0.005733)
$\hat{x}_{t-m}$	0.000065	0.000054	0.000011	0.000059
	(0.000047)	(0.000050)	(0.000055)	(0.000049)
p_{t-m}	-0.000644	0.001080	0.001835	0.001775
	(0.001494)	(0.001656)	(0.001435)	(0.001604)
p^2_{t-m}	0.001551	-0.011518**	-0.000352	-0.004677
	(0.003099)	(0.004485)	(0.003205)	(0.003336)

Table 7: Results from Regression (5) for Varying m

	$m = 1$	$m = 30$	$m = 90$	$m = 365$
		Region 1		
Constant	0.016135***	0.015981***	0.015649***	0.016056***
	(0.000558)	(0.000595)	(0.000621)	(0.000554)
$\hat{x}_{t-m}$	-0.000323***	-0.000292***	-0.000181***	-0.000312***
	(0.000041)	(0.000046)	(0.000058)	(0.000046)
p_{t-m}	-0.003478*	-0.005204**	-0.002636	-0.004471**
	(0.001916)	(0.002266)	(0.002115)	(0.002264)
p^2_{t-m}	0.023092***	0.028792***	0.018065**	0.031630***
	(0.007906)	(0.009748)	(0.007199)	(0.011361)
		Region 2		
Constant	0.016071***	0.015947***	0.015566***	0.016061***
	(0.000559)	(0.000587)	(0.000610)	(0.000557)
$\hat{x}_{t-m}$	-0.000328***	-0.000292***	-0.000177***	-0.000317***
	(0.000045)	(0.000049)	(0.000061)	(0.000048)
p_{t-m}	-0.002923*	-0.003842**	-0.002627	-0.002911*
	(0.001593)	(0.001834)	(0.001718)	(0.001769)
p^2_{t-m}	0.016104***	0.017679**	0.017116***	0.019258**
	(0.005793)	(0.006926)	(0.005650)	(0.007576)
		Region 3		
Constant	0.016037***	0.015960***	0.015580***	0.016027***
	(0.000556)	(0.000584)	(0.000614)	(0.000555)
$\hat{x}_{t-m}$	-0.000335***	-0.000299***	-0.000178***	-0.000321***
	(0.000045)	(0.000049)	(0.000062)	(0.000048)
p_{t-m}	-0.001486*	-0.001996**	-0.001236	-0.001572*
	(0.000809)	(0.000941)	(0.000899)	(0.000912)
p^2_{t-m}	0.004059***	0.004581**	0.003834***	0.005003**
	(0.001441)	(0.001786)	(0.001418)	(0.001947)

Table 8: Results from Regression (6) for Varying m

	$m = 1$	$m = 30$	$m = 90$	$m = 365$
		Region 1		
$\hat{x}_{t-m}$	-0.000001	-0.000008	0.000005	-0.000003
	(0.000005)	(0.000005)	(0.000005)	(0.000005)
p_{t-m}	-0.000241	0.000435	-0.000103	0.000072
	(0.000165)	(0.000269)	(0.000155)	(0.000296)
p^2_{t-m}	0.000535	-0.001969	-0.000116	0.000312
	(0.000718)	(0.001696)	(0.000625)	(0.001122)
		Region 2		
$\hat{x}_{t-m}$	0.000000	-0.000006	0.000004	-0.000005
	(0.000005)	(0.000005)	(0.000005)	(0.000005)
p_{t-m}	-0.000266*	0.000346*	-0.000104	0.000109
	(0.000145)	(0.000200)	(0.000125)	(0.000232)
p^2_{t-m}	0.000622	-0.001469	0.000133	-0.000081
	(0.000597)	(0.001095)	(0.000469)	(0.000748)
		Region 3		
$\hat{x}_{t-m}$	0.000001	-0.000006	0.000004	-0.000005
	(0.000005)	(0.000005)	(0.000005)	(0.000005)
p_{t-m}	-0.000130*	0.000185*	-0.000050	0.000064
	(0.000074)	(0.000104)	(0.000065)	(0.000122)
p^2_{t-m}	0.000161	-0.000385	0.000011	-0.000046
	(0.000146)	(0.000284)	(0.000118)	(0.000198)

References

[1] Olivier Deschênes and Michael Greenstone. The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *American economic review*, 97(1):354–385, 2007.

[2] Walter Enders and Clive William John Granger. Unit-root tests and asymmetric adjustment with an example using the term structure of interest rates. *Journal of Business & Economic Statistics*, 16(3):304–311, 1998.

[3] Robert F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.

[4] Robert F Engle and Tim Bollerslev. Modelling the persistence of conditional variances. *Econometric reviews*, 5(1):1–50, 1986.

[5] Cornelis Gardebroek and Manuel A Hernandez. Do energy prices stimulate food price volatility? examining volatility transmission between us oil, ethanol and corn markets. *Energy economics*, 40:119–129, 2013.

[6] Clive WJ Granger and Paul Newbold. Spurious regressions in econometrics. *Journal of econometrics*, 2(2):111–120, 1974.

[7] Marco Haase, Yvonne Seiler Zimmermann, and Heinz Zimmermann. The impact of speculation on commodity futures markets–a review of the findings of 100 empirical studies. *Journal of Commodity Markets*, 3(1):1–15, 2016.

[8] Peter R Hansen and Asger Lunde. A forecast comparison of volatility models: does anything beat a garch (1, 1)? *Journal of applied econometrics*, 20(7):873–889, 2005.

[9] Donald Mitchell. A note on rising food prices. *World bank policy research working paper*, (4682), 2008.

[10] Xiaoyi Mu. Weather, storage, and natural gas price dynamics: Fundamentals and volatility. *Energy Economics*, 29(1):46–63, 2007.

[11] Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.

[12] Peter C. B. Phillips and Pierre Perron. Testing for a unit root in time series regression. *Biometrika*, 75(2):335–346, 1988.

[13] Joe T Ritchie and D Scott Nesmith. Temperature and crop development. *Modeling plant and soil systems*, 31:5–29, 1991.

[14] Richard Roll. Orange juice and weather. *The American Economic Review*, 74(5):861–880, 1984.

[15] Paul A Samuelson. Proof that properly anticipated prices fluctuate randomly. In *The world scientific handbook of futures markets*, pages 25–38. World Scientific, 2016.

[16] Wolfram Schlenker and Michael J Roberts. Estimating the impact of climate change on crop yields: The importance of nonlinear temperature effects. Technical report, National Bureau of Economic Research, 2008.

[17] Richard SJ Tol. Autoregressive conditional heteroscedasticity in daily temperature measurements. *Environmetrics*, 7(1):67–75, 1996.

[18] Andrés Trujillo-Barrera, Mindy Mallory, and Philip Garcia. Volatility spillovers in us crude oil, ethanol, and corn futures markets. *Journal of Agricultural and Resource Economics*, pages 247–262, 2012.

[19] Calum G Turvey. Weather derivatives for specific event risks in agriculture. *Applied Economic Perspectives and Policy*, 23(2):333–351, 2001.

[20] Daniel Urban, Michael J Roberts, Wolfram Schlenker, and David B Lobell. Projected temperature changes indicate significant increase in interannual variability of us maize yields: A letter. *Climatic change*, 112:525–533, 2012.