ANALYZING THE IMPACT OF RECENT DEVELOPMENTS IN AGRICULTURAL COMMODITY MARKETS, AND USING MACHINE LEARNING TOOLS TO IMPROVE FORECASTING TECHNIQUES

by

RAGHAV GOYAL

(Under the Direction of Michael Adjemian)

ABSTRACT

USDA reports are highly informative and valued by market participants. Markets quickly absorb their information, and exhibit volatility spikes following report announcements that dissipate quickly (Adjemian and Irwin - 2018). A government shutdown in January 2019 prevented the USDA from publishing information about the situation and outlook for major U.S. agricultural commodities. I show that, as a result, Chicago Mercantile Exchange Board of Trade markets for corn and soybeans experienced heightened market uncertainty, elevating the cost of managing risk using options. Having established the importance of government information, I decompose ending stocks forecast errors into errors of the other supply and demand components using a Machine Learning (ML) algorithm. The results show that export and production misses are the major contributors to ending stocks projection errors. Lastly, using past yield shocks as instruments to endogenous futures prices, I estimate a supply elasticity for corn at 0.26 percent. Recently, due to the COVID-19 pandemic, demand for ethanol plummeted – resulting in an inward shift in the demand curve for corn. Consequently, my analysis shows that corn producers suffered losses worth $5.4 billion, 27% lower compared to the actual payments USDA made under Coronavirus Food Assistance Program (CFAP). Moreover, all my estimates suggest that the USDA was able to fully compensate producers for any losses they faced due to demand reduction.

INDEX WORDS: WASDE, Corn, Soybean, Wheat, COVID-19, Forecast Error, Implied Volatility, Options, Shutdown

ANALYZING THE IMPACT OF RECENT DEVELOPMENTS IN AGRICULTURAL COMMODITY MARKETS, AND USING MACHINE LEARNING TOOLS TO IMPROVE FORECASTING TECHNIQUES

by

RAGHAV GOYAL

Bachelor of Technology in Biotechnology, Panjab University, 2014

Post Graduate Diploma in Rural Management, Institute of Rural Management Anand, India, 2017

A Dissertation Submitted to the Graduate Faculty of The University of Georgia in Partial Fulfillment of the Requirements for the Degree

DOCTOR OF PHILOSOPHY

ATHENS, GEORGIA

2022

ANALYZING THE IMPACT OF RECENT DEVELOPMENTS IN AGRICULTURAL COMMODITY MARKETS, AND USING MACHINE LEARNING TOOLS TO IMPROVE FORECASTING TECHNIQUES

by

RAGHAV GOYAL

Major Professor: Michael Adjemian
Committee: Berna Karali
William Secor

Electronic Version Approved:

Ron Walcott
Vice Provost for Graduate Education and Dean of the Graduate School
The University of Georgia
August 2022

DEDICATION

I would like to dedicate this work to my family and friends.

ACKNOWLEDGEMENTS

I would like to acknowledge and thank my advisor, Dr. Michael Adjemian, whose guidance helped me throughout the past four years and made this work possible. I would also like to thank my committee members for taking out time to read my thesis and provide valuable feedback.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Page

CHAPTER 1

INTRODUCTION

Agricultural markets are characterized by volatile supply due to weather patterns, crop diseases, and pests, and therefore volatile prices. Given that regularity, to assist market participants the U.S. Department of Agriculture (USDA) has for decades provided the public with price, demand, and supply forecasts for major agricultural commodities over each marketing year. These reports play a key role in setting expectations and guiding private production, consumption, and inventory decisions along the supply chain. Markets react to their news. Isengildina-Massa et al. (2008) find that WASDE reports significantly impact corn and soybean futures markets. Adjemian (2012) demonstrates that futures prices rapidly incorporate government information following the publication of these reports. Irwin, Gerlow, and Liu (1994) and Isengildina-Massa et al. (2006) establish how market participants consider USDA forecasts as benchmarks while making supply chain decisions. Isengildina-Massa et al. (2021) show that USDA's January report clusters, which include final forecasts for several commodities' marketing year, have a significant impact on nearby futures prices.

During an appropriations lapse in 2013, USDA suspended its October World Agricultural Supply and Demand Estimates (WASDE) report. Adjemian et al. (2018) conclude that corn and soybean markets did not experience normally expected changes in prices and reductions in implied volatility (IV) around the release of the USDA information. As a result, market participants had less information with which to plan or conduct their operations, and likely faced increased hedging costs. In this first chapter, I offer an obvious extension of that work and study the government shutdown in January 2019, which likewise forced USDA to suspend its publication of various reports, including its normally-scheduled WASDE, Crop Production, and Grain Storage releases. I explore its consequences for corn and soybean markets, and quantify how the price of managing risk using options increased due to the additional commodity market uncertainty caused by the missing reports. These findings make a policy-relevant case for the continued release of public information to align market expectations.

The second chapter deals with decomposing the source of errors in commodity demand & supply forecasts. Ending stocks are among the most important balance sheet forecasts because they provide an estimate of the end-of-marketing-year inventory of a particular commodity, summarizing supply and demand outlook. Anticipated ending stocks levels guide production and use plans, influence carrying costs, impact prices paid to farmers, and influence import and export decisions. Overestimates (underestimates) of ending stocks could lower (higher) producer prices and distort storage and production signals. Understanding the source of errors in ending stocks projections will help organizations like USDA understand and perhaps minimize their misses if they are able to make improvements to forecasting techniques. In the chapter, I decompose ending stocks forecast errors into forecast misses for the component demand and supply-side variables. I use machine learning models to handle multicollinearity and non-linearities in the data. My results underscore the importance of production and foreign trade activities to ending stocks misses.

In the third chapter, I intend to estimate the price elasticity of supply for agricultural commodities. Large, recent negative shocks such as (1) the trade war effect on soybean exports, and (2) COVID-19's destruction of the demand for biofuels—which most directly affected corn—may reveal new portions of commodity supply curves. Understanding and quantifying these elasticities is important to policy formation and analysis, especially since they were used to generate remunerative payments to affected producers. I exploit USDA's twice-annual plantings forecasts (in March and again in June) alongside expected price changes to estimate these elasticities. Fortuitously, the COVID-19 shock occurred between the periods during which the data for these 2020 projections was collected. With these elasticities, I forecast planting behavior and farmer returns in the absence of the COVID shock, and compare the realized damage it did to producers to the payments the government made.

CHAPTER 2

THE 2019 GOVERNMENT SHUTDOWN INCREASED UNCERTAINITY IN MAJOR AGRICULTURAL COMMODITY MARKETS[1]

[1] Goyal, R., and M. K. Adjemian. 2021. *Food Policy*. 102:102064.
Reprinted here with permission of the publisher

Abstract:

In January 2019, a government shutdown prevented the U.S. Department of Agriculture from publishing information about the situation and outlook for major U.S. agricultural commodities. We show that, as a result, Chicago Mercantile Exchange Board of Trade markets for corn and soybeans experienced heightened market uncertainty, elevating the cost of managing risk using options. We use historical options data to estimate that, on the first day of trading following the normally scheduled USDA publication time, the additional commodity market uncertainty caused by the government shutdown increased the price of managing risk using ATM corn and soybean options by additional 2.95% and 1.66%, respectively for corn and soybeans, using an approach that assumes a normal January report impact. Using a different counterfactual approach–assuming that the observed, abnormally large implied volatility reduction following the February 2019 publication would have been experienced in January, we find that the increase in risk management costs due to missing information was actually about 11.5% higher for corn and 4.4% higher for soybeans.

2.1. Introduction

For decades, USDA has published information about the market situation for a wide range of domestic commodities and has estimated the elements of supply and demand for major agricultural commodities both in the United States and around the world. USDA publishes these reports free of charge in order to align market expectations, resolve uncertainty, and improve efficiency and economic activity in these markets. The most prominent examples of USDA reports include the World Agricultural Supply and Demand Estimates (WASDE), Grain Stocks, Crop Progress, Feed Outlook, Oil Crops Outlook, Crop Production, Acreage and Prospective Plantings (APP).

Market participants like farmers, elevators, commercial grain firms, and end-users rely on agricultural commodity futures and options markets to manage their business risks. Options strategies can limit potential losses, or provide additional income, depending on the side a trader decides to take – writer or purchaser. Information provided by USDA reports is closely followed by market participants and observers to make informed supply management and speculative decisions: realigning expectations about commodity market conditions, re-allocating resources, and maximizing their profitability. A growing empirical literature shows that these reports provide the market with important information, impacting prices and implied volatilities (IVs) at announcement time (see, e.g., Sumner and Mueller, 1989; McNew and Espinosa, 1994; Isengildina-Massa et al., 2008a and 2008b; Adjemian, 2012; Adjemian and Irwin, 2018; Ying, Chen and Dorfman, 2019).

Indeed, recent work has shown that missing USDA reports may elevate uncertainty levels in agricultural markets. During an appropriations lapse in 2013, USDA suspended its October World Agricultural Supply and Demand Estimates (WASDE) report. Adjemian et al. (2018) concluded that corn and soybean markets did not experience normally expected changes in prices and reductions in implied volatility (IV) around the release of the USDA information. As a result, market participants had less information about planning or conducting their operations, and likely faced increased hedging costs. A natural extension of that work is to study the government shutdown in January 2019, which likewise forced USDA to suspend its publication of various reports, including its normally-scheduled WASDE, Crop Production, and Grain Storage releases. Once the 2019 shutdown ended, USDA published that information on February 8th, 2019, about a month later than in no shutdown years (USDA, 2019). The shutdown is, therefore, a natural experiment that provides exogenous variation to study the impact of missing information on agricultural commodity markets. We exploit this natural experiment and study its consequences in this article.

Figures 2.1 and 2.2 show how daily changes in the logarithm of IV (which, if they are small, approximate percent changes) for corn and soybeans, respectively, relate to USDA report publication days. In both charts, USDA publications tend to reduce IV levels, which we interpret as a decrease in market uncertainty about each commodity. Without accounting for any other factors (as we do later on in our regression framework), average WASDE report days correspond to a 3.4% daily reduction for corn, and a 3% reduction for soybeans. Average January report days (which include a WASDE, Crop Production, and Grain Stocks publication) occur alongside 3.6% and 2.7% IV decreases for corn and soybeans, respectively, while February WASDEs produce similar values of 2.8% and 1.4%. The figures also show that the February 2019 report—the first issued after the government shutdown that year resolved—coincides with a historically large reduction in IV.

Figure 2.1: Corn: Daily change in log IV

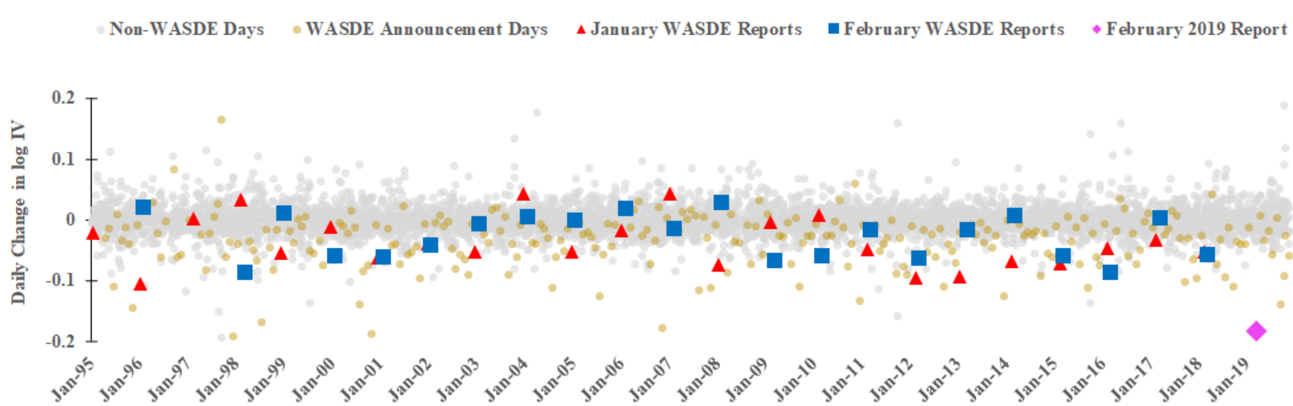

Figure 2.2: Soybeans: Daily change in log IV

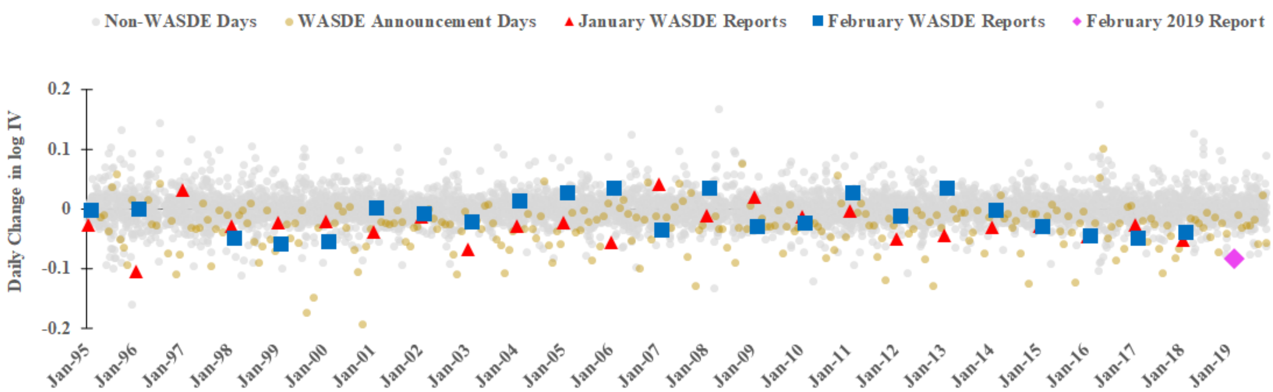

Figure 2.3 plots the average pre-2019 path of log IV for CBOT corn in a 5-day window around the release of a January report against the actual log IV observed around (1) the originally-scheduled release date during the 2019 shutdown, and (2) the February 2019 report (the first that USDA published after the shutdown). The dotted line in figure 2.3 marks the normal, scheduled USDA WASDE and/or Grain Stock report release day. In general, USDA January announcements produce a sharp decline in CBOT corn implied volatility. Similarly, a sharp decline is also observed following the release of the February 2019 report. However, in January 2019 no such decline is observed, indicating that USDA's failure to publish a report led to a higher level of uncertainty in the market than would have otherwise been observed. CBOT soybeans also exhibit similar patterns (see figure A2.1 in Appendix). In this paper, we estimate the degree to which missing government information increased uncertainty in major domestic commodity markets and its impact on market participants.

Several important factors differentiate the 2013 and 2019 government shutdowns. The 2013 shutdown prevented the publication of an October report during a relatively tranquil crop year. But the 2019 shutdown curtailed the release of a January report, and USDA publications during that time are among the most important produced by the Department (e.g., Adjemian et al. 2018). They provide the initial “final” production estimates for the forecasting cycle in the form of an annual summary and include one of the USDA’s quarterly reports of inventories, highly anticipated by traders since it signals crop disappearance and on- and off-farm stocks levels at the State level— and January is the first stocks release to following the completion of both crop’s harvest. It is the only report day of the year that includes the simultaneous publication of both WASDE and Grain Stocks reports; generally, these are separated by about two weeks. Given the federal debt burden, it is likely that public budgetary pressures will increase in the future, so understanding how markets are affected by missing government information is an important area of study for both participants and policymakers.

Figure 2.3: CBOT CORN: Indexed Logged IV Levels

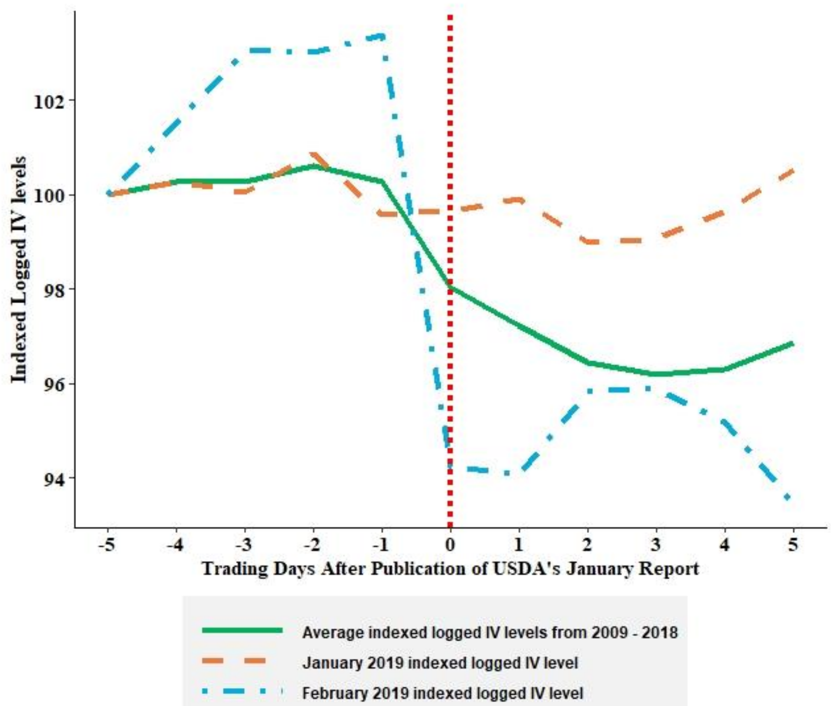

Notes: The red dotted vertical line represents the trading when historic January USDA reports were released, the day when missing January 2019 USDA reports were scheduled for release, and the day when February 2019 USDA reports were published. The index log IV value = 100 on day T = -5, or five trading days before relevant release day.

Adjemian et al. (2018) suggest that the price of hedging is higher in the absence of the WASDE report. We quantify the additional risk management cost for at-the-money (ATM) options caused by missing public news.[2] We also contribute to the literature by providing a full portrait of the impact of missing government information across moneyness levels. To construct the counterfactual for the estimated changes in IV if the January 2019 reports had been published on time, we use two measures: (1) the average change in IV around the release of pre-2019 January reports, and (2) the actual observed change in IV following the release of the next published report (i.e., in February 2019). We use the option *vega*, a measure of how much the price of an option changes conditional on a 1% change in implied volatility, to estimate the likely changes in options prices that would have occurred if the missing reports were published as originally scheduled. Given that the report was not released for several weeks, we later extrapolate those counterfactuals across the shutdown period to estimate the additional cost it posed to managing commodity risk using options.

Risk averse commodity market participants with short positions (i.e., those with interest in selling the underlying commodity) can buy put options on futures contracts to hedge against the risk that prices fall; the option's price is referred to as a *premium*. A put holder has the right to exercise the option–meaning that s/he can short a futures contract to the option writer at the strike

[2] An option is said to be at-the-money if its strike price is (nearly) equivalent to the price of the underlying security. These options only have time value and are usually the most highly traded. Corn ATM options accounted for 94.5% of the nearby option trade from 1/11/2019 – 1/25/2019. During the same period, Soybeans ATM options made up 91.1% of all nearby options traded in that market.

price. This becomes intrinsically valuable as the futures price falls below the strike price (plus the price of the option), since a profit can then be earned by taking a long at the lower price: selling high and buying low. If futures prices rise, or at least do not fall below the specified strike price, a put option would not become profitable. Since they provide the perspective of a short hedger, we use put options to evaluate the additional cost of managing risk due to publication curtailment of the January 2019 WASDE report.[3] We control for the general market uncertainty caused by the government shutdown (and other phenomena in the world, including the ongoing trade war) using the Chicago Board Options Exchange Volatility Index (VIX), as well as standard trend, seasonality, and day-of-the-week effects.

We estimate that, on the first day of trading following the normally scheduled USDA publication time, the additional commodity market uncertainty caused by the government shutdown increased the price of managing risk using ATM corn and soybean puts by an additional 2.95% (95% CI: 2.938% - 2.97%) and 1.66% (95% CI:1.64% - 1.677%), respectively for corn and soybeans, using an approach that assumes a normal January report impact. Using a different counterfactual approach–assuming that the IV reduction following the February 2019 publication would have been experienced in January, we find that the increase in risk management costs due to the missing information was actually about 11.5% (95% CI: 10.97% - 11.8%) higher for corn and 4.43% (95% CI: 4.28% - 4.58%) higher for soybeans. We find significant but lower increases for deep out-of-the-money or in-the-money options for these commodities.

[3] We generate nearly identical results when we calculate these cost increases using ATM calls.

2.2. Background

Every month USDA publishes a series of comprehensive reports defining the current situation and projections for supply and demand fundamentals in agricultural commodity markets. The World Agricultural Supply and Demand Estimates (WASDE) report provides a detailed balance sheet of variables for the agricultural commodity markets. WASDE is prepared by combining insights from several USDA agencies represented in the Interagency Commodity Estimates Committees (ICECs), chaired by the USDA World Agricultural Outlook Board (WAOB). ICECs relies on Foreign Agricultural Service (FAS) reports for overseas commodity market developments, the National Agricultural Statistics Service (NASS) for agricultural commodity and livestock projections, the Economic Research Service (ERS) for regional assessments, the Agricultural Marketing Service (AMS), and the Farm Service Agency (FSA) for market information and impact of domestic policies. ICEC committees generate consensus forecasts based on all these data sources. Since it contains highly market-sensitive information, WASDE is prepared under secure lock-up conditions and is released to all users simultaneously each month (Vogel and Bange, 1999).

Uncertainty about these fundamentals naturally decreases as the harvest comes in, and USDA updates its projections. While the May WASDE report contains the first projections for the new marketing year, the first survey-based corn and soybean production forecasts are included in the August report (Vogel and Bange, 1999). January reports contain the final annual projections each year. Between February and April, USDA makes only minor changes to its reports. NASS also issues Grain Stock reports, providing an extensive analysis of the on- and off-farm stock positions for major crops. Grain Stocks reports are released four times a year (January, March, June, September). Moreover, the January Grain Stocks report and the WASDE report are released concurrently (Adjemian et al., 2018).

Many researchers have found USDA reports affect market expectations of prices and uncertainty. Sumner and Mueller (1989) show that market prices change following the release of USDA harvest forecasts. McNew and Espinosa (1994) report reductions in market uncertainty following report release. Isengildina-Massa et al. (2008a, 2008b) show that corn and soybean markets react to announcement days, reflected in increased futures return variance and reduced options IV. Adjemian (2012) shows that agricultural commodity futures markets rapidly incorporate new information, confirming the persistence of impact across contract positions. Following report release, futures markets experience heightened intraday trading volumes and volatility dissipating in a few minutes (Adjemian and Irwin, 2018).

Over the last decade, private market advisory alternatives have grown, and so has their projections' accuracy. Xiao et al. (2014) find that the accuracy of private forecasts is highest for wheat and lowest for soybeans. Recent literature suggests that, even in the presence of private forecasts, USDA reports are informative (Ying, Chen, and Dorfman, 2019; Karali et al., 2019). Karali et al. (2019) observed significant futures commodity price movements in response to market surprises, even for the reports most likely affected by the availability of private sources of information. As freely-available products of the federal government, USDA commodity reports are publicly-provided goods, and they have a strong track record of forecast reliability. Private forecasts, on the contrary, are not freely accessible. Even in the presence of private sources of information, many USDA reports remain highly valued and trusted by end users (Karali et al., 2019, Isengildina-Massa et al., 2020).

2.3. Data

Several popular techniques are used to represent backward-looking, historical market volatility (e.g., GARCH). When making plans, market participants also take into consideration market expectations. Implied volatility, a forward-looking measure, represents the volatility expected over the remaining life of the related futures contract. IV is also a basic component of option premia. We draw historical ATM options IV data for individual CBOT corn and soybean contracts from Bloomberg, from 1995 to 2019. Corn contracts deliver in March, May, July, September, and December, while soybean futures do so in January, March, May, July, August, September, and November. To form a continuous series of close-to-close changes in implied volatility for each commodity, we rollover from the nearby to the next deferred contract fifteen days prior to the contract expiration month. Bloomberg Listed Implied Volatility Engine (LIVE) IV values are calculated in a similar way to the inversion of the canonical Black-Scholes formula. It is important to note that researchers have identified limitations associated with the Black-Scholes formula, including the assumptions of lognormality of underlying asset prices, constant volatility over time, and zero transaction costs. Beckers (1980) and Wiggins (1987) suggest that these assumptions result in higher Black-Scholes prices for out-of-money options and lower prices for in-the-money and at-the-money options. However, as suggested in Whaley (1986) and Barone-Adesi and Whaley (1987), the difference between the prices derived from the Black-Scholes and an alternative option valuation model is less than 1% of the actual price. Recently, Jankova (2018) suggests that irrespective of the shortcomings of Black-Scholes method, it is still regarded as a universal pricing model. Moreover, several papers use IV as a variable in their regression models (see, e.g., Adjemian et al., 2018, and Davidson et al., 2001).

We match the change in implied volatility with the historical *WASDE* and *Grain Stocks* report announcements; the Cornell University Library maintains these data. For both commodities, figures A2.2 and A2.3 in the appendix compare kernel density plots based on the IVs observed on USDA announcement days relative to other days. Unfortunately, Bloomberg does not offer historical ATM options implied volatility data are not available by *moneyness*, or the ratio of the strike price of the option in question to the ATM strike levels. Therefore, to study the impact of missing information by moneyness levels, we extract Bloomberg data on 30-day implied volatility by moneyness for corn and soybeans from 2005 to 2019 (those data are not available from 1995-2004, the time periods of analysis do not match exactly). The 30-Day IV refers to the forecasted daily implied volatility for the coming 30 days. In this way, it is similar to the way the VIX is constructed, since that measure represents the forecasted daily implied volatility associated with the S&P 500 index over the next thirty days. Table 2.1 provides summary statistics for each of the series used in our analysis. Comparing the daily change in IV for ATM options across the two data sets we used (which differ by time period and unit of analysis: contract level versus 30-day), the means, medians, and standard deviations are approximately the same for both commodities. Similarly, within each commodity, the statistics in the table for the daily change in IV are quite similar across moneyness levels.

USDA published a total of 398 WASDE and Grain Stocks reports (298 WASDE, 100 Grain Stocks reports) on 373 announcement days from January 1995 - February 2019. On 25 of those announcement days, WASDE and Grain Stocks reports were released concurrently. To control for general market uncertainty and investor sentiment, we collect CBOE Volatility Index (VIX) data. We draw options vegas from DataStream, a database maintained by Thomson Reuters

Table 2.1: CBOT Corn and Soybeans: Change in Daily IV

	Time	Moneyness	Statistic	Mean	Median	Std. Dev.	Skewness
Corn	1995-2019	At the Money	ΔIV	-0.0006	-0.0004	0.028	0.62
			$\|\Delta IV\|$	0.0173	0.0113	0.022	6.11
	2005-2019	90% Moneyness	ΔIV	$-2.8e^{-5}$	-0.0023	0.036	-2.6
			$\|\Delta IV\|$	0.021	0.013	0.029	5.47
		95% Moneyness	ΔIV	$-3.4e^{-5}$	0.0007	0.032	-1.5
			$\|\Delta IV\|$	0.02	0.014	0.025	5.8
		At the Money	ΔIV	$-2.35e^{-5}$	0.0004	0.031	-1.06
			$\|\Delta IV\|$	0.019	0.0137	0.024	7.27
		105% Moneyness	ΔIV	$-8.1e^{-5}$	0.001	0.033	-1.58
			$\|\Delta IV\|$	0.019	0.013	0.026	7.4
		110% Moneyness	ΔIV	$-1.5e^{-5}$	0.002	0.034	-2.25
			$\|\Delta IV\|$	0.02	0.013	0.027	6.18
Soybean	1995-2019	At the Money	ΔIV	-0.0004	-0.0004	0.026	-0.066
			$\|\Delta IV\|$	0.018	0.0125	0.02	4.22
	2005-2019	90% Moneyness	ΔIV	$-7.8e^{-5}$	0.002	0.031	-1.9
			$\|\Delta IV\|$	0.02	0.013	0.024	4.6
		95% Moneyness	ΔIV	-0.0001	0.0003	0.027	-0.19
			$\|\Delta IV\|$	0.018	0.0135	0.02	3.65
		At the money	ΔIV	-0.0001	-0.0002	0.025	0.13
			$\|\Delta IV\|$	0.018	0.014	0.017	2.88
		105% Moneyness	ΔIV	-0.0001	0.0007	0.03	1.04
			$\|\Delta IV\|$	0.018	0.013	0.024	11.46
		110% Moneyness	ΔIV	$-7.9e^{-5}$	0.0013	0.031	-0.17
			$\|\Delta IV\|$	0.018	0.013	0.025	8.57

Note: Δ IV is computed as the daily difference between the logged IVs. There are a total of 6133 and 6115 observations for corn and soybeans respectively from 1995 to 2019. From 2005 to 2019, there are 3494 observations for corn and 3441 observations for soybeans.

2.4. Methodology

2.4.1. Regression Models

In this paper, we estimate how implied volatility for major agricultural futures contracts changed following USDA report announcements and calculate the daily change in implied volatility as:

$$\Delta \mathrm{IV}_{i,t} = \log(IV_{i,t}) - \log(\mathrm{IV}_{i,t-1}) \quad (1)$$

Where the commodity (corn or soybeans) is indexed by i, and the trading day is indexed by t. The logarithmic transformation helps to normalize the data, and differencing ensures stationarity. Our basic regression model is specified as:

$$\begin{aligned}\Delta \mathrm{IV}_{i,t} = \beta_0 + \beta_1 \Delta \mathrm{IV}_{i,t-1} + \beta_2 D_{WASDE} + \beta_3 D_{Grain\ Stocks} \\ + \beta_4 D_{Pre-2019\ Jan\ WASDE} + \beta_5 D_{Pre-2019\ Feb\ WASDE} \\ + \beta_6 D_{Jan\ 19\ WASDE} + \beta_7 D_{Feb\ 19\ WASDE} + \beta_8 \Delta VIX_t + \beta_M M \\ + \epsilon_{i,t}\end{aligned} \quad (2)$$

where VIX refers to CBOE's Volatility Index, IV is the implied volatility for commodity *i* on day *t*. *M* is a vector of other regressors including month and day-of-week dummies and time trends. We specifically include dummies for WASDE and Grain stock report release days (D_{WASDE} & $D_{Grain\ Stocks}$), for all trading days that include WASDE, and Grain stock reports released in January and February ($D_{Pre-2019\ Jan\ WASDE}$ & $D_{Pre-2019\ Feb\ WASDE}$), for the missing January 2019 reports ($D_{Jan19\ WASDE}$) and for the reports released in the following month ($D_{Feb19\ WASDE}$).

Because the residuals from (2) are characterized by auto-correlation and heteroskedasticity (as discussed in detail in the next section), we also estimate GARCH models. In doing so, we attempt to capture the volatility clustering present in the IV series. Figure A2.4 plots the (discrete) first differenced logged daily IV series for nearby Chicago Mercantile Exchange (CME) corn contracts, from 1995-2019; it is clearly characterized by volatility clusters, interspersed with relatively tranquil periods. Our GARCH (1,1) model with ARMA (1,1) specification is as follows:

$$\Delta\text{IV}_{i,t} = \beta_0 + \beta_1\Delta\text{IV}_{i,t-1} + \beta_2 D_{WASDE} + \beta_3 D_{Grain\ Stocks}$$
$$+ \beta_4 D_{Pre-2019\ Jan\ WASDE} + \beta_5 D_{Pre-2019\ Feb\ WASDE}$$
$$+ \beta_6 D_{Jan\ 19\ WASDE} + \beta_7 D_{Feb\ 19\ WASDE} + \beta_8 \Delta VIX_t + \beta_M M$$
$$+ \epsilon_{i,t} + \epsilon_{i,t-1} \qquad (3)$$
$$\epsilon_{i,t} = \sigma_{i,t} Z_{i,t}$$
$$\sigma_{i,t}^2 = \omega + \alpha\epsilon_{i,t-1}^2 + \beta\sigma_{i,t-1}^2 + \gamma_R R$$
$$Z_{i,t} \sim N(0,1)$$

Where, ω, α and β are GARCH parameters. R is the vector of the same regressors that are added to the mean equation. $Z_{i,t}$ is a random variable with a standard normal distribution. In order to estimate the change in IV that would have occurred had the January 2019 reports been published on time, we construct two counterfactual measures:

- the normal change exhibited by the IV series in response to pre-2019 January reports = $\beta_2 + \beta_3 + \beta_4$
- The actual observed change in IV following the release of the next published report (i.e., February 2019), compared to the impact of a pre-2019 February release = $\beta_7 - \beta_5$

Using the variance covariance matrix of the parameters estimated from the regression framework, we apply the delta method to construct standard errors for the two counterfactuals. We also report Hodrick (1992) standard errors for regression models (1) and (2), which impose the null hypothesis of no serial correlation but do not impose conditional homoskedasticity

.

2.4.2. Counterfactual-Implied Options Price

To compute the increase in put option risk management costs, we calculate the counterfactual implied put options price, i.e., $P_{t,c}$, where *t* refers to the day scheduled for the publication of the January 2019 WASDE report, and *c* refers to a counterfactual (either 1 or 2), using the following approach:

$$P_{t,c} = P_{t,b} + \Delta\mathrm{IV}_{t,c} * v \tag{4}$$

$$\Delta\mathrm{IV}_{t,c} = IV_{t-1,b} * (1 + c) * \left(1 + \frac{IV_{t,b} - IV_{t-1,b}}{IV_{t-1,b}}\right) - IV_{t-1,b} \tag{5}$$

$P_{t,b}$ is the observed put options price on the (missing) report day, and the subscript *b* refers to an observed value. $IV_{t-1,b}$ is the IV observed a day prior. The observed options vega is represented by *v*, and $\Delta IV_{t,c}$is the difference between the counterfactual IV (adjusted by the actual observed change in IV from *t-1* to *t*.) and the previous days IV.

2.4.3. Moneyness-level Analysis using 2005-2019 30-Day Data

To investigate the impact of the missing reports across moneyness levels, for each commodity, we construct a 5x1 IV vector of daily changes in the 30-day IV, for 90%, 95%, 100%, 105%, and 110% moneyness, as our dependent variable and estimate a dynamic conditional correlation (DCC) GARCH model (Engle, 2002) to generate the marginal effects of the missing report. We take this approach because for each IV vector the underlying commodity is the same. Since each series is influenced by the same supply demand dynamics, like weather shocks and the general level of uncertainty in the markets, there is a high correlation among each component in the vector. A DCC GARCH model is flexible to those relationships while accounting for serial correlation and conditional heteroskedasticity.

2.5. Results and Discussion

2.5.1. Regression Models

Table 2.2 summarizes the results of the CBOT corn market IV models we estimate to measure the impact of missing government data. Model (1) presents the OLS regression results with no month and day-of-week effects (i.e., we do not include any dummies for month or weekdays). After controlling for other factors, the model associates historical (1995-2018) January WASDE and Grain-stock reports (reported in the table as counterfactual 1) with a highly statistically significant decline in daily IV of about 3.7%. Historic February WASDE reports have no significant effect on IV, while the release of the February 2019 WASDE report carried a substantial IV decline, after controlling for other factors. Under counterfactual 2, we estimate the impact of missing information is even larger: about13.8%. Model(2)controls for additional month and day-of-week effects, but produces no significant change in any of the coefficients; indicating that the results are robust to these effects.

However, both models (1) and (2) suffer from auto-correlation – errors are serially correlated, as observed from the LM and LBP test statistics in the table. LM tests for the null hypothesis of no serial correlation. LBP tests for the null hypothesis of independence by testing overall randomness. We, therefore, estimate a standard GARCH (1,1) (Bollerslev, 1986) an exponential GARCH (1,1) model (Nelson, 1991) with ARMA (1,1) process, respectively referred to as models (3) and (4) in the table. Both make only minor changes to the counterfactual values, and the estimates remain highly statistically significant. We use Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) values to compare different models.[4]

Figure 2.4: Corn: Counterfactual 1 and 2 Bands

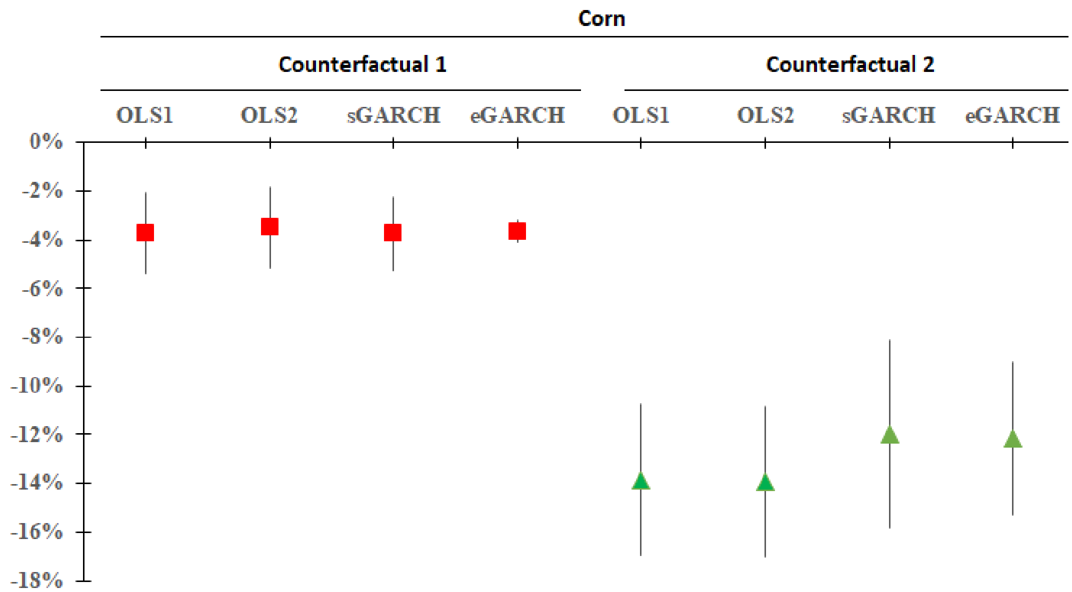

Note: CF1 refers to counterfactual 1 and CF2 refers to counterfactual 2

Model (4) exhibits the lowest AIC and BIC values, so it is preferred according to model selection criteria. As a result, we report headline point estimates and confidence intervals from this model, and refer to it as preferred, although we note that at the 95% confidence level, Models (1) – (4) produce statistically indistinguishable results for each counterfactual, as shown graphically in figure 2.4. The similarity of our findings for each counterfactual across model specifications is a sign of the robustness of our results.

[4] AIC measures the relative distance between the fitted likelihood function and the true (unknown) likelihood function, while BIC estimates the posterior probability of the estimated model being true; lower AIC and BIC values are preferred.

Table 2.2: Factors that Affect the Daily Change in Log Implied Volatility for CBOT Corn, 1995-2019

		Regression Models			
		OLS		sGARCH [1]	eGARCH [1]
		(1)	(2)	(3)	(4)
(A)	Lagged change in Daily Log IV	0.0878 ** (0.0373) [0.0028]	−0.0868** (0.0372) [0.0027]	0.9607*** (0.0057)	0.9535*** (0.0085)
(B)	Generic Grain Stock Announcement	−0.0231*** (0.0040) [0.0046]	−0.0245*** (0.0041) [0.0047]	−0.0228*** (0.0040)	−0.0208*** (0.0012)
(C)	Generic WASDE Announcement	−0.0347*** (0.0029) [0.0097]	−0.0343*** (0.0029) [0.0096]	−-0.0297*** (0.0025)	−0.0252*** (0.0016)
(D)	Pre-2019 January WASDE/Grain Stock	0.0207** (0.0098) [0.0026]	0.0237** (0.0099) [0.0026]	0.0150** (0.0089)	0.0095*** (0.0009)
(E)	Pre-2019 February WASDE	0.0065 (0.0080) [0.0021]	0.008 (0.0081) [0.0021]	0.0005 (0.0087)	-0.002 (0.0078)
(F)	January 2019 WASDE	0.003 (0.003) [0.0001]	0.0027* (0.0019) [0.00003]	0.0064*** (0.0021)	0.0048*** (0.0011)
(G)	February 2019 WASDE/Grain Stock Announcement	−0.132*** (0.0085) [<0.0000]	−0.131*** (0.0085) [<0.0000]	−0.119*** (0.0132)	−0.123*** (0.0094)
(H)	Daily change in VIX	0.0454*** (0.0133) [0.0008]	0.0345** (0.0133) [0.0008]	0.045** (0.0172)	0.0492*** (0.0072)
	Week Dummies? [2]	No	Yes	No	No
	Month Dummies? [2]	No	Yes	No	No
		Impact Of Missing WASDE			
	Counterfactual 1 [3]	−0.0371*** (0.0085)	−0.0349*** (0.0085)	−0.03755*** (0.0078)	−0.03653*** (0.0023)
	Counterfactual 2 [4]	−0.138*** (0.0158)	−0.139*** (0.0157)	−0.119*** (0.0196)	−0.121*** (0.016)
		Test Statistics			
	AIC	-27040	-27097	-4.69	-4.71
	BIC	-26960	-26915	-4.67	-4.69
	LM Test [5]	59***	99***	0.63	0.57
	LBP Test (Lag 9)	44***	55***	7.3	6.6

Note: Robust standard errors are reported in parentheses. Counterfactual 1 = (B) + (C) + (D), and Counterfactual 2 = (G) - (E). Standard errors for both are calculated using the Delta Method. Squared Brackets report Hodrick (1992) standard errors. These standard errors do not change the significance/ statistical importance of counterfactuals.

Table 2.3: Factors that Affect the Daily Change in Log Implied Volatility for CBOT Soybeans, 1995-2019

		Regression Models			
		OLS		sGARCH [1]	eGARCH [1]
		(1)	(2)	(3)	(4)
(A)	Lagged change in Daily Log IV	−0.0617 (0.0387) [0.0027]	−0.0598 (0.0386) [0.0027]	0.7363*** (0.0935)	0.7153*** (0.0734)
(B)	Generic Grain Stock Announcement	−0.0293*** (0.0043) [0.0044]	−0.0312*** (0.0043) [0.0044]	−0.0306*** (0.0047)	−0.0303*** (0.0015)
(C)	Generic WASDE Announcement	−0.0316*** (0.0026) [0.0082]	−0.0315*** (0.0025) [0.0081]	−0.0298*** (0.0025)	−0.0261*** (0.0013)
(D)	Pre-2019 January WASDE/Grain Stock	0.0323*** (0.0079) [0.0019]	0.0371*** (0.0080) [0.0019]	0.0300*** (0.0078)	0.0268*** (0.0012)
(E)	Pre-2019 February WASDE	0.0164** (0.0070) [0.0018]	0.0166** (0.0071) [0.0018]	0.0123** (0.0059)	0.0101 (0.0062)
(F)	January 2019 WASDE	0.0098 (0.0113) [0.0003]	0.0108 (0.0106) [0.0003]	0.0103 (0.0402)	0.0103 (0.0093)
(G)	February 2019 WASDE/Grain Stock Announcement	−0.0381*** (0.0079) [<0.0000]	−0.0375*** (0.0078) [<0.0000]	−0.0389*** (0.0084)	−0.0453*** (0.0068)
(H)	Daily change in VIX	0.0264** (0.0117) [0.0007]	0.0189 (0.0117) [0.0007]	0.0398*** (0.0109)	0.0343*** (0.0074)
	Week Dummies? [2]	No	Yes	No	No
	Month Dummies? [2]	No	Yes	No	No
		Impact Of Missing WASDE			
	Counterfactual 1 [3]	−0.0286*** (0.0062)	−0.0257*** (0.0062)	−0.0305*** (0.0059)	−0.0295*** (0.0028)
	Counterfactual 2 [4]	−0.0545*** (0.0140)	−0.0541*** (0.0139)	−0.0512*** (0.0130)	−0.0554*** (0.0126)
		Test Statistics			
	AIC	-27467	-27497	-4.72	-4.75
	BIC	-27386	-27316	-4.7	-4.72
	LM Test [5]	62***	97***	1.1	1.2
	LBP Test	37***	45***	9	11

Note: Robust standard errors are reported in parentheses. Counterfactual 1 = (B) + (C) + (D), and Counterfactual 2 = (G) - (E). Standard errors for both are calculated using the Delta Method. Squared Brackets report Hodrick (1992) standard errors. These standard errors do not change the significance/ statistical importance of counterfactuals.

Figure 2.5: Soybeans: Counterfactual 1 and 2 Bands

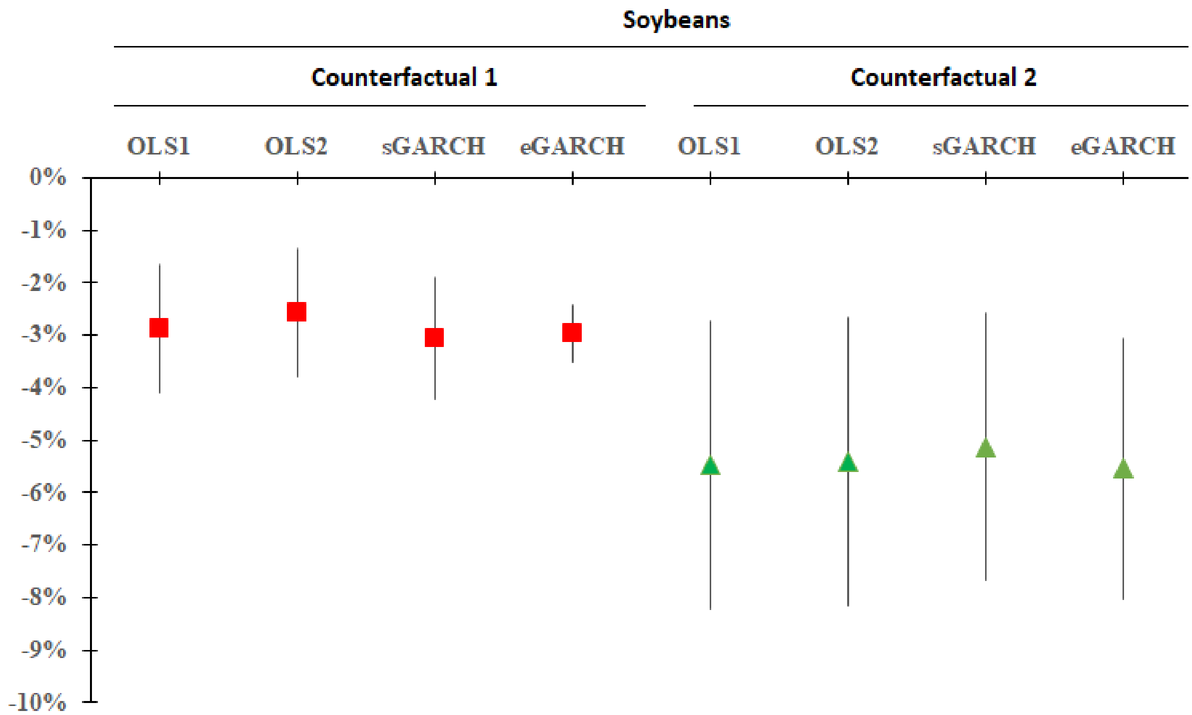

Note: CF1 refers to counterfactual 1 and CF2 refers to counterfactual 2

Table 2.3 provides the same results for CBOT soybeans. To correct for serial correlation in OLS regression models (1) and (2), we estimate standard GARCH (1,1) and exponential GARCH (1,1) models with ARMA (3,1) processes. Across model specifications, we estimate the counterfactual 1 mean impact on IV in the range of negative 2.6% - 3.1%, and counterfactual 2 as a daily decline in IV of between 5.1% - 5.5%. As in the case of corn, our eGARCH model is preferred according to both AIC and BIC. Once again, however, these results are statistically indistinguishable at the 95% confidence level, as shown in figure 2.5.

The tables and figures above compare the two counterfactual values we estimate. Recall that the first approach reconstructs what likely would have been observed had the market behaved as it normally does on a January release day. The second counterfactual measures the difference between the observed February reaction by the markets against what would normally be expected in February. Because the latter involves an actual market response, it could be interpreted as more realistic. However, although it is stronger at the mean for both commodities, and statistically stronger in the case of corn, we hesitate to make that case because the February 2019 report contained information that had always previously been published in January. As a result, we generate both counterfactuals in order to bound the market effect of the shutdown-generated disruption of the normal flow of government news about commodities. Figures 1.4 and 1.5 depict our estimates of the 95% confidence intervals around both counterfactuals, for corn and soybeans, respectively. Both approaches are robust across four different model specifications, but counterfactual 2 is clearly larger (in an absolute sense) for corn. Confidence bounds overlap for counterfactuals 1 and 2 for soybeans, indicating that they are not distinguishable, statistically

Our results are consistent with the literature's findings. For example, Adjemian et al. (2018), using data from 1995-2015, indicate that January WASDE reports are associated with a 5.75% (95% CI: -10.84%, -0.65%) reduction in uncertainty for corn and a 2.2% (95% CI: -7.93%, 3.48%) reduction for soybeans. Our findings are statistically indistinguishable from theirs, which is expected given that our study periods mostly overlap. Likewise, Isengildina-Massa et al. (2008a) study commodity markets over 1985-2002 and show that January WASDE report sessions coincide with a 0.5 percentage point and a 0.86 percentage point reduction in IV for corn and soybeans, respectively. Our findings are in the same direction, even though our study period doesn't share much overlap.

2.5.2. Counterfactual-Implied Options Price

Following the results in tables 2.2 and 2.3, we approximate the counterfactual price of ATM corn and soybean put options due to publication curtailment of the January 2019 WASDE report. We estimate that, on the first day of trading following the normally scheduled USDA publication time, the additional commodity market uncertainty caused by the government shutdown increased the price of managing risk using ATM corn and soybean puts by an additional 2.95% per contract per bushel[5] (95% CI: 2.938% - 2.97%) and 1.66% per contract per bushel (95% CI:1.64% - 1.677%), respectively for corn and soybeans, using an approach that assumes a normal January report impact. Using a second counterfactual approach–assuming that the IV reduction following the February 2019 publication would have been experienced in January, we find that the increase in risk management costs due to missing information was actually about 11.54% per contract per bushel (95% CI: 10.97% - 11.80%) higher for corn and 4.43% per contract per bushel (95% CI: 4.28% - 4.58%) higher for soybeans.

2.5.3. Moneyness-Level Analysis using 2005-2019 30-Day IV Data

Options can be characterized by their moneyness levels, or the ratio of the strike price of the option in question to the ATM strike. An ATM option only has time value, because its strike price is (nearly) equivalent to the price of the underlying security. In the money (ITM) options possess both intrinsic value (or monetary value), since they can be exercised profitably immediately, and time value, since they have the possibility of becoming even more profitable

[5] For example, our mean estimate for corn counterfactual 2 is -12.14%. Substituting this and an observed v=0.511, IV_{t-1}= 15.22%, $IV_{t,b}$= 15.28%, and $P_{t,b}$ = 8.875 cents per bushel into equations 4 and 5 generates the counterfactual options price of 7.96 cents per bushel. That compared to the observed price of 8.875 cents per bushel resulted in an 11.5% increase in ATM put options risk management cost.

with the passage of time. Out-of-the money (OTM) options have no intrinsic value but do possess time value. The trading volume of an option depends, to some extent, on the moneyness levels. ATM options are usually the most highly traded. Figure A2.5 in the appendix depicts trading volume data for the corn and soybeans options we studied, segregated by moneyness level, from January 11^{th}, 2019 (the day when the missing reports were scheduled for publication) to January 25^{th}, 2019. For both the commodities, ATM options represent nearly all market trading interest during the period of observation, highlighting the usefulness of our approach.

Implied volatility surfaces in commodity markets are generally skewed, and the impact of news may be affected by option moneyness. That is, deep out-of-the-money (OTM) or in-the-money (ITM) options, meaning those with strike prices associated with a very low or high probability of being executed, may be affected differently by news than ATM options. To explore this, we use the Bloomberg 30-day IV data at moneyness levels of 90%, 95%, 100%, 105%, and 110%, for each commodity. The 30-Day IV refers to the forecasted daily implied volatility for the coming 30 days. It is similar to the way the VIX is constructed, since that measure represents the forecasted daily implied volatility associated with the S&P 500 index over the next thirty days. Due to data availability constraints, it is different from the IV values used for analysis in the previous section (There we use contract level IV).

We estimate a multivariate dynamic conditional correlation (DCC) GARCH model - Engle 2002 - (DCC(1,1) eGARCH(1,1) ARMA (1,1)).[6] Tables A2.1 and A2.2 in the appendix provide

[6] We conduct a formal test for the relevance of DCC GARCH against a constant conditional correlation (CCC) GARCH model. For both the commodities, we reject the null hypothesis at 1% significance level, confirming the relevance of DCC GARCH model for our analysis.

the DCC GARCH outputs for the two commodities. The counterfactual 1 estimates for corn and soybeans are statistically significant, at 5% significance level, for all moneyness levels. For corn, counterfactual 2 is statistically significant at 10% significance level. For soybeans only ATM options have a statistically significant value (at 10% level) for counterfactual 2.

To measure how the impact of the government shutdown and missing January 2019 reports affected options markets across moneyness levels, we use a cubic spline to interpolate the counterfactual values for the intervals between the counterfactual values we estimate at observed moneyness levels using a DCC GARCH model, shown as red dots in figures 2.6 and 2.7. The yellow line joining those dots are imputations estimated using a cubic spline method. We caveat this exercise by noting that splines pose risks, including over-fitting and deterioration of predictive ability in the tails of the distribution (Ruppert, 2002; Gauthier et al. 2020).

Figure 2.6 provides the point estimate results and the 95 % confidence bands for counterfactual 1 for both corn and soybean across moneyness levels. ATM corn contracts experience the largest impact of government news (-7.4%, 95% CI: 4.5% - 10.3%), with smaller impacts estimated at other levels of option moneyness. Further, confidence intervals are narrower for deep OTM or ITM options and wider for other moneyness levels. Soybeans offer a smaller variation in marginal effects across different moneyness levels, but near-the-money options still display the maximum impact. Similar to corn, confidence intervals for soybeans are narrower towards both the ends of moneyness levels. Overall, for both the commodities, marginal effects closely follow a convex function. ATM options are generally among the most heavily traded[7], so they are highly sensitive to market uncertainty changes. Although our ATM results in figure 2.6 and 2.7 are slightly smaller in absolute value than our findings in section 2.5.1, their similarity in magnitude and direction–despite being estimated using a different data set an over a different time period–offers a robustness check on our findings.

Figure 2.6: IMPACT OF MISSING JANUARY 2019 REPORTS: Marginal Effects across Moneyness Levels (Based on Counterfactual 1)

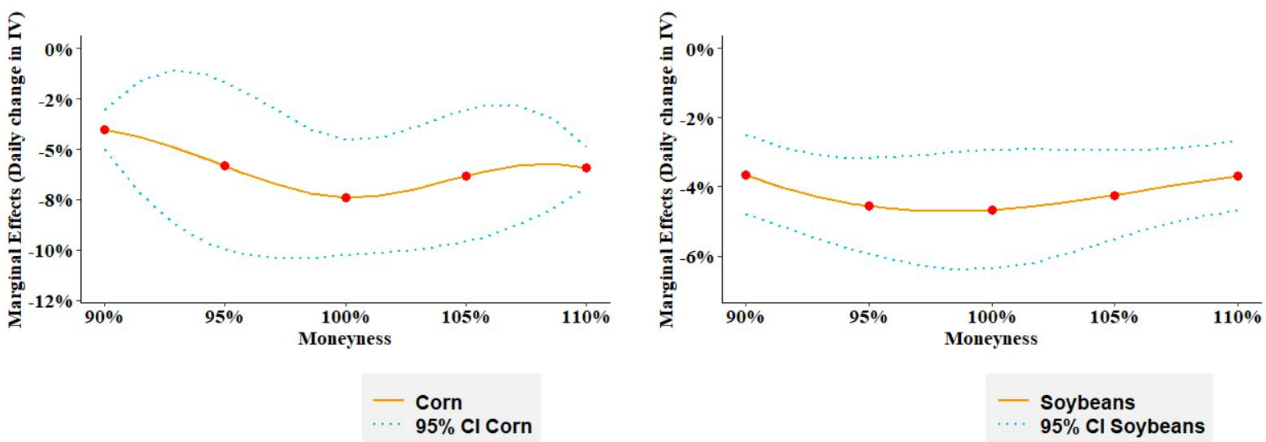

Figure 2.7 provides the same results for counterfactual 2 for both the commodities across moneyness levels. Once again, near-the-money corn and soybean contracts experience the largest impact of government news, and their similarity in to our ATM results in section 2.5.1, once again offers us a helpful robustness check on our analysis.

[7] As a result, we believe we capture the bulk of the observable risk management costs by focusing our analysis on the effect on ATM options. We are unable to observe trades that were not initiated due to increased costs

Based on the marginal effects from figure 2.6 and 2.7 and the actual IVs observed on each of these days across the moneyness level, we calculate the counterfactual IVs that would have been observed on January 11^{th} 2019, if the missing reports had been published. Figures 2.8 and 2.9 report our mean level results for corn and soybeans, respectively. We project that, had January 2019 WASDE and Grain Stocks reports been published on time, each market would have experienced notable IV declines across moneyness levels.

Figure 2.7: IMPACT OF MISSING JANUARY 2019 REPORTS: Marginal Effects across Moneyness Levels (Based on Counterfactual 2)

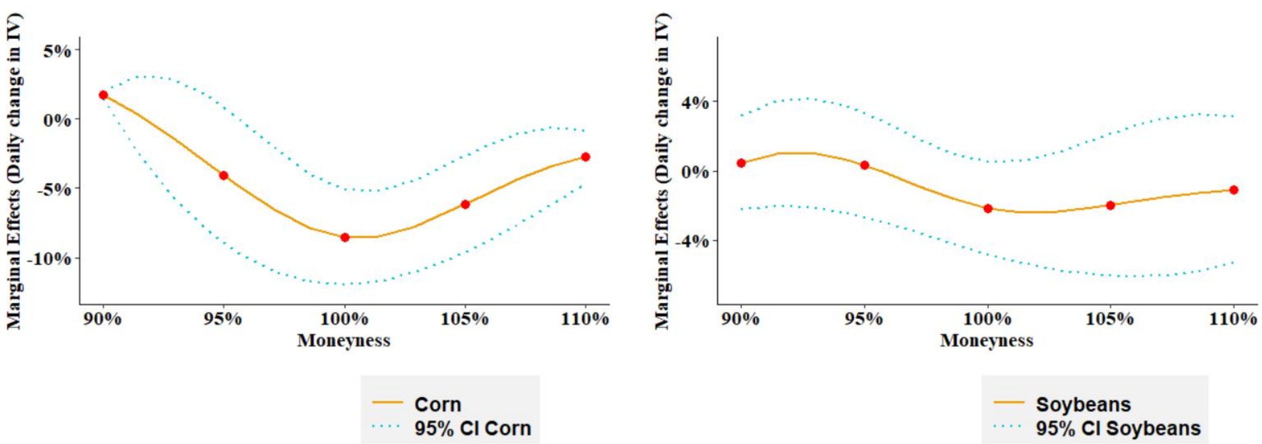

2.6.Conclusion and Policy Implications

USDA reports are highly informative and valued by market participants. Markets quickly absorb their information, and exhibit volatility spikes following report announcements that dissipate quickly (Adjemian and Irwin - 2018). A government shutdown in January 2019 prevented the USDA from publishing information about the situation and outlook for major U.S. agricultural commodities. As an exogenous shock, the shutdown presents a natural experiment to study how markets behave in the absence of important government information about commodity market conditions.

According to our results, corn and soybean markets in January 2019 did not experience reductions in uncertainty normally observed around the release of important government crop news. Uncertainty is an important factor in determining options values; more uncertainty raises their premia, all else equal. The economics literature documents a strong link between uncertainty and reduced economic vitality, raising concerning implications about our findings–were governments to curtail access to important supply and demand information for sustained periods. Hassett and Metcalf (2011) find that increases in tax uncertainty adversely impact capital formation. Nodari (2014) relates uncertainty over financial regulation to increases in the cost of borrowing. Baker et al., (2016), using a new measure of economic policy uncertainty, show that high policy uncertainty leads to lower investment rates and higher unemployment growth.

Using volume data from the options markets, we us a conservative approach to estimate that, on the first day of trading following the normally scheduled USDA publication time, the additional commodity market uncertainty caused by the government shutdown increased the cost of managing risk with options by over $117,000 for corn and almost $37,000 for soybeans. Extrapolating this total over the shutdown period implies corn and soybean market costs of over $1.4 million and more than $938,000, respectively. Under our more aggressive counterfactual approach, we estimate the shutdown raised one-day risk management costs by about $425,000 for corn and $95,000 for soybeans, and more than $5.1 million and $2.4 million for the two commodities if we extend the period to the full nineteen trading days until the shutdown resolved (and the next report was published). Note that because we use observed volume figures to generate these estimates, we are unable to observe options trades that did not take place possibly due to higher costs caused by the shutdown.

ATM options experience the largest effects of missing information, at least in percentage terms. These options are usually the most highly traded. Hence, our results offer important insights for traders and market participants, as well as government decision makers who face budget constraints regarding the collection and provision of information about supply and demand fundamentals for agricultural commodities. Indeed, were private entities able to seamlessly replace the information provided by public sources during the shutdown, we would likely not have found statistically and economically significant effects associated with the first government report following the resolution of the government's budget impasse. Researchers have investigated the substitutability between public and private sources of information about agricultural commodities, and generally found that public reports offer information that is—up to now—more accurate that the alternative shared by private firms. Egelkraut et al. (2003) compare USDA forecasts to Conrad Leslie and Sparks Companies' private forecasts for corn and soybean production (which were made ahead of government announcements). They find USDA reports are more accurate for corn, and soybeans in October and November. Theorizing that private companies may have benefited over time from advances in computing power and big data, Karali et al. (2019) study whether USDA reports remain informative by modeling the market surprise, which they define as the difference between USDA forecasts and private expectations; the authors show that the "news" component of USDA reports has not decreased over time. They also show that the price impact of public reports has increased since January 2007, indicating that markets have grown more sensitive to government news. Isengildina-Massa et al. (2020) compare USDA production and acreage forecasts (for corn, soybeans and wheat) to those produced over time by a series of private forecasters such as Conrad Leslie, Informa Economics, Dow Jones Newswire survey, and the Bloomberg Survey. The authors observe USDA errors to be smaller, and find increases in relative accuracy over time for USDA reports.

Figure 2.8: Impact of Missing 2019 Reports on CBOT Corn IV: January 11, 2019

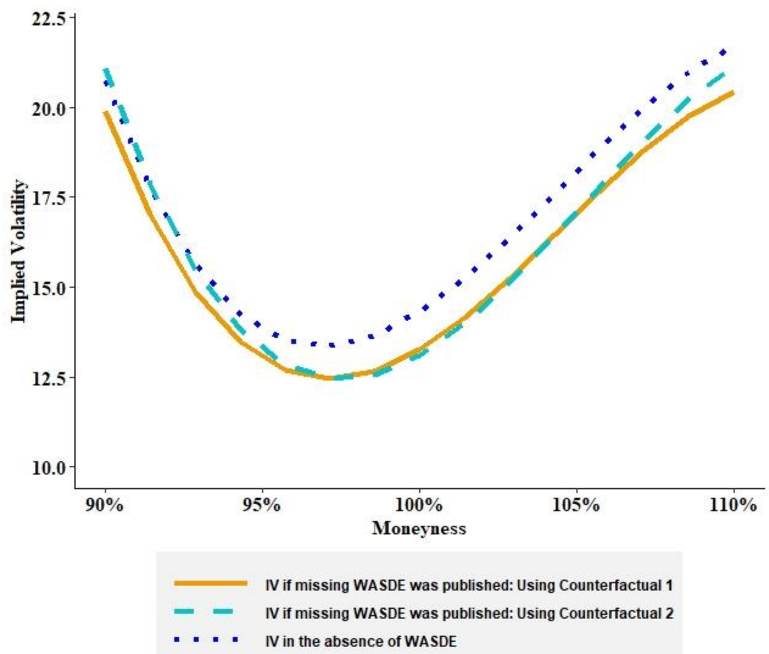

One limitation of our approach is that the length of the shutdown period cannot be varied to determine how much more costly it would have been to deprive the market of government statistics for a longer period of time. However, it is likely that future budget pressures will make shutdown events more common and offer researchers the ability to consider this aspect of the absence of important, publicly-provided goods.

Figure 2.9: Impact of Missing 2019 Reports on CBOT Soybeans IV: January 11, 2019

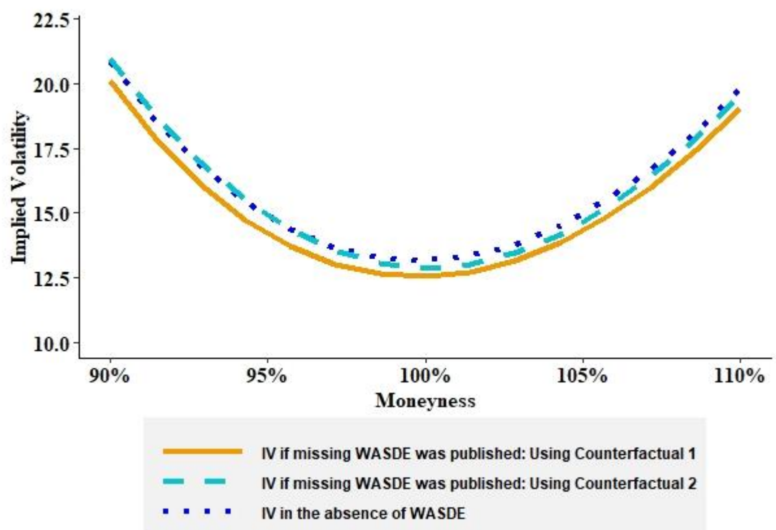

2.7. APPENDIX

Figure A2.1: CBOT SOYBEANS: Indexed Logged IV Levels.

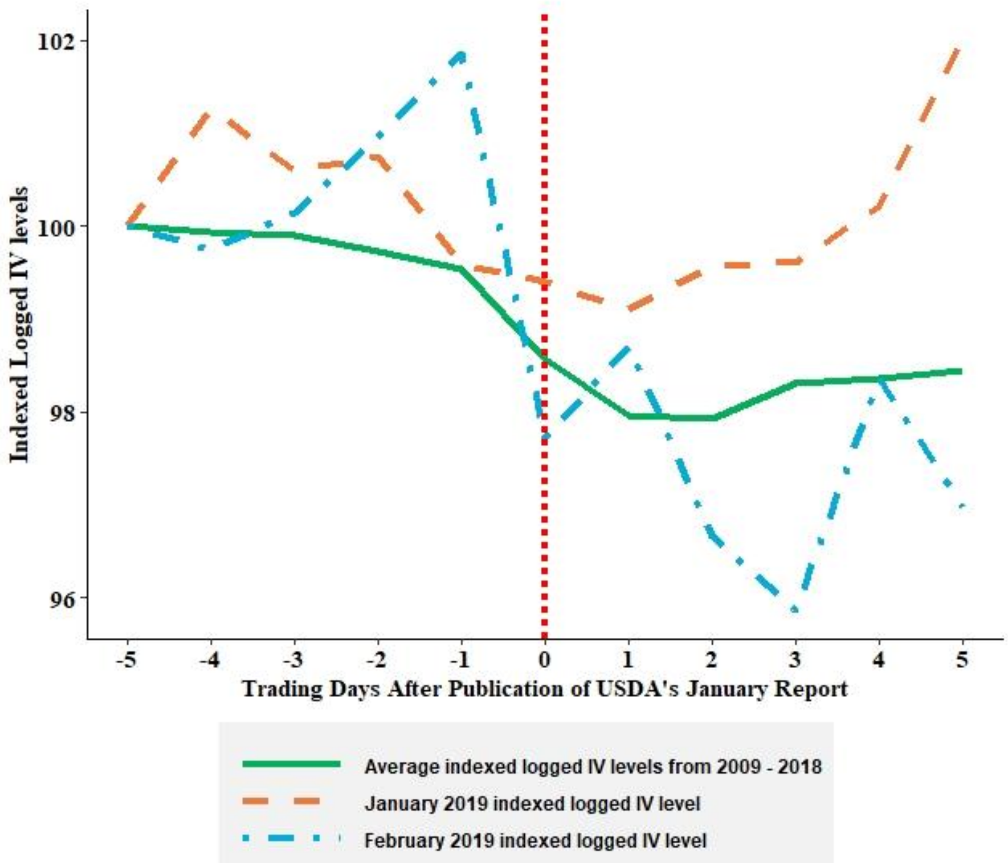

The lines are the time series of the indexed logged CBOT soybean IV levels in a 5-day window around the normal release of a January report averaged from 2009 - 2018. The red dotted line marks the day when the historic January USDA reports were released, the day when missing January 2019 USDA reports were scheduled for release and the day when February 2019 USDA reports were released. The index log IV value = 100 on day T = -5, or five trading days before normally-scheduled January report release.

Figure A2.2: Corn: Kernel Density Plot

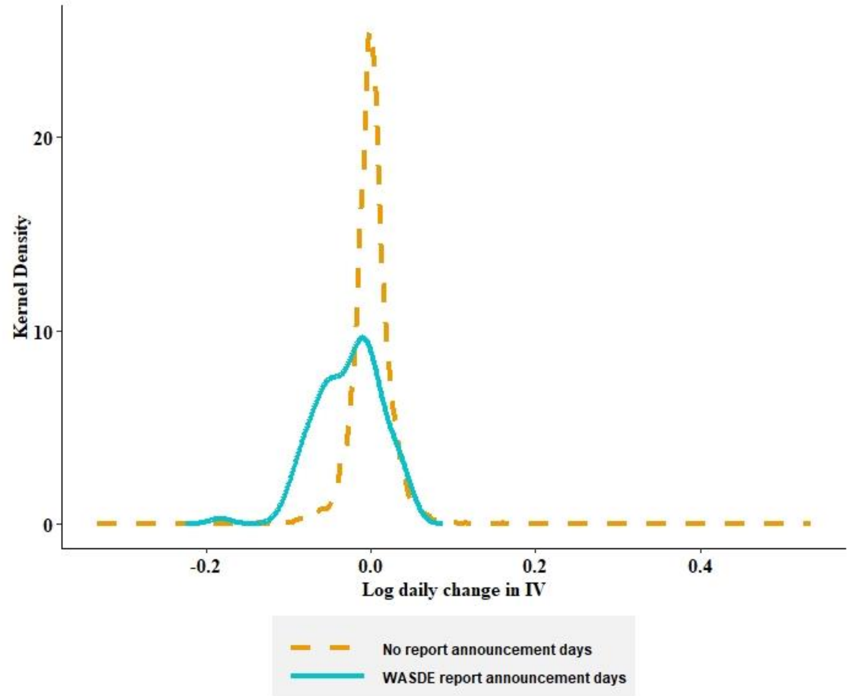

Figure A2.3: Soybeans: Kernel Density Plot

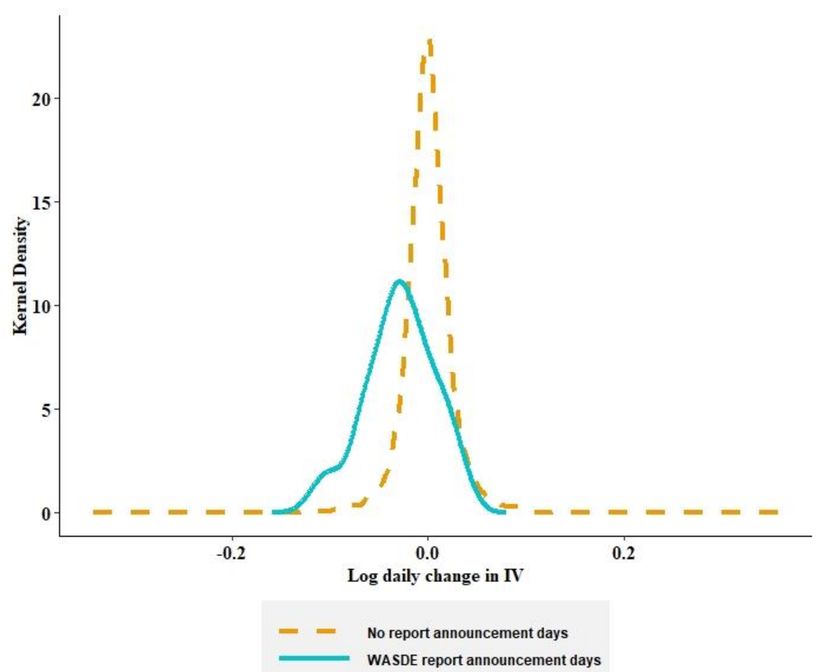

Figure A2.4: CME CORN (1995 - 2019))

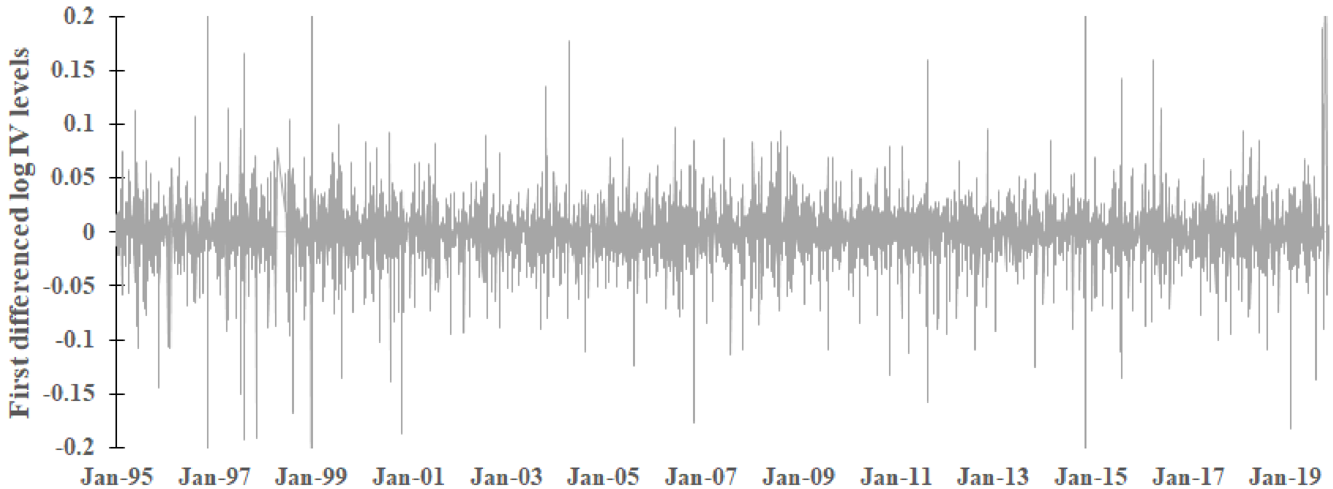

Figure A2.5: Options Trading Volume from 1/11/2019 – 1/25/2019

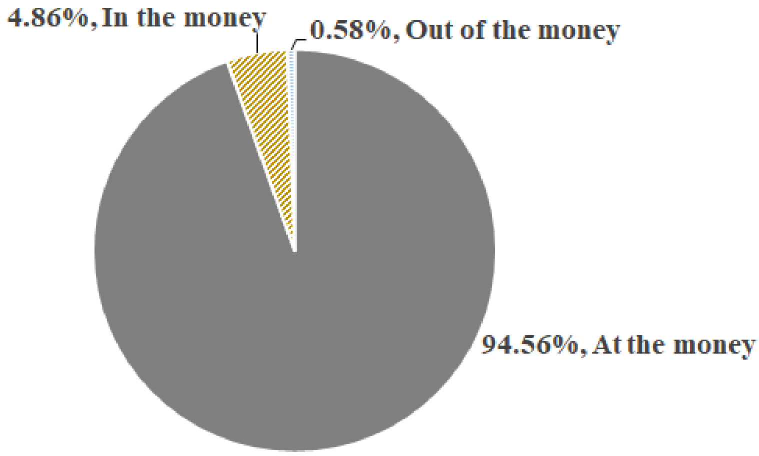

PANEL A: CORN

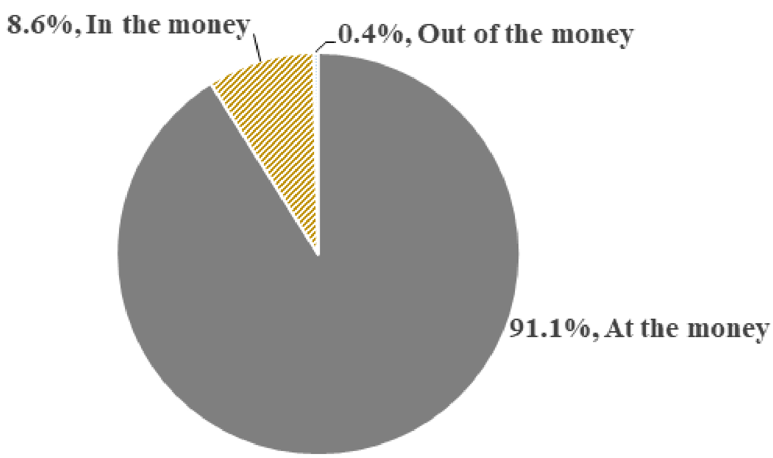

PANEL B: SOYBEANS

Table A2.1: CBOT Corn DCC GARCH Results

		DCC GARCH Model				
		90%IV	95%IV	100%IV	105%IV	110%IV
		(1)	(2)	(3)	(4)	(5)
(A)	Lagged change in Daily Log IV	0.7701*** (0.0143)	0.7294*** (0.0303)	0.7763*** (0.0140)	0.7809*** (0.0753)	0.7363*** (0.0157)
(B)	Generic Grain Stock Announcement	−0.0078** (0.0031)	−0.0224* (0.0115)	−0.0307*** (0.0085)	−0.0470*** (0.0111)	−0.0387*** (0.0047)
(C)	Generic WASDE Announcement	−0.0199*** (0.0011)	−0.0308*** (0.0032)	−0.0354*** (0.0034)	−0.0329*** (0.0034)	−0.0255*** (0.0017)
(D)	Pre-2019 January WASDE/Grain Stock	-0.0127** (0.0052)	-0.0051 (0.0237)	-0.0079 (0.0158)	0.0168 (0.0196)	0.0052*** (0.00076)
(E)	Pre-2019 February WASDE	0.0119*** (0.0004)	0.0065 (0.0076)	0.0079 (0.0076)	0.0147** (0.0065)	0.0122*** (0.0018)
(F)	January 2019 WASDE	0.0134*** (0.0016)	0.0043 (0.0075)	0.0085 (0.0077)	0.0140*** (0.0043)	0.0193*** (0.0029)
(G)	February 2019 WASDE/Grain Stock Announcement	0.0291*** (0.0012)	−0.0342 (0.022)	−0.0772*** (0.0120)	−0.0466*** (0.0138)	−0.0154* (0.0091)
(H)	Daily change in VIX	0.0455*** (0.0034)	0.0686*** (0.0109)	0.0649*** (0.0163)	0.0348** (0.0152)	0.0389*** (0.0103)
		Impact Of Missing WASDE				
	Counterfactual 1 [1]	-0.0403*** (0.0049)	-0.0583** (0.0213)	−0.0741*** (0.0146)	−0.0631*** (0.0167)	−0.0589*** (0.0051)
	Counterfactual 2 [2]	0.0171*** (0.0013)	−0.0401* (0.0251)	−0.0852*** (0.0175)	−0.0613*** (0.0177)	-0.0275** (0.0096)
		Test Statistics				
	AIC			-26.5		
	BIC			-26.4		

Note: Standard errors are reported in parentheses. Counterfactual 1 = (B) + (C) + (D), and Counterfactual 2 = (G) - (E). Standard errors for both are calculated using the Delta Method.

Table A2.2: CBOT Soybeans DCC GARCH Results

		DCC GARCH Model				
		90%IV	95%IV	100%IV	105%IV	110%IV
		(1)	(2)	(3)	(4)	(5)
(A)	Lagged change in Daily Log IV	0.8084*** (0.0124)	0.7787*** (0.1713)	0.9119*** (0.0081)	0.8801*** (0.0064)	0.7619*** (0.0384)
(B)	Generic Grain Stock Announcement	−0.0225*** (0.0046)	−0.0289*** (0.0049)	−0.0331*** (0.0054)	−0.0261*** (0.0008)	−0.0172*** (0.0053)
(C)	Generic WASDE Announcement	−0.0111*** (0.0017)	−0.0200*** (0.0035)	−0.0251*** (0.0023)	−0.0238*** (0.0027)	−0.0149*** (0.0025)
(D)	Pre-2019 January WASDE/Grain Stock	-0.0029 (0.0097)	0.0033 (0.0096)	0.0116 (0.0104)	0.0074 (0.0071)	-0.0048 (0.0076)
(E)	Pre-2019 February WASDE	0.0051 (0.0059)	0.0075 (0.0071)	0.0123** (0.0053)	0.0119 (0.0076)	0.0069 (0.0069)
(F)	January 2019 WASDE	0.0090 (0.0105)	0.0120 (0.0161)	0.0017 (0.0132)	-0.0018 (0.0126)	0.0088 (0.0097)
(G)	February 2019 WASDE/Grain Stock Announcement	0.0097 (0.0099)	0.0108 (0.0104)	−0.0093 (0.0106)	−0.0079 (0.0163)	−0.0042 (0.0185)
(H)	Daily change in VIX	0.0596*** (0.0098)	0.0608*** (0.0125)	0.0427*** (0.0068)	0.0295*** (0.0082)	0.0183*** (0.0061)
				Impact Of Missing WASDE		
	Counterfactual 1 [1]	-0.0365*** (0.0059)	-0.0457*** (0.0071)	−0.0466*** (0.0088)	−0.0424*** (0.0065)	−0.0369*** (0.0052)
	Counterfactual 2 [2]	0.0046 (0.0139)	0.0034 (0.0154)	−0.022* (0.0136)	-0.0198 (0.0208)	-0.0111 (0.0216)
				Test Statistics		
	AIC			-29.4		
	BIC			-29.3		

Note: Standard errors are reported in parentheses. Counterfactual 1 = (B) + (C) + (D), and

Counterfactual 2 = (G) - (E). Standard errors for both are calculated using the Delta Method.

CHAPTER 3

DECOMPOSING USDA ENDING STOCKS FORECAST ERRORS[8]

[8] Goyal, R., M.K. Adjemian, J. Glauber, and S. Meyer. 2022. Forthcoming in the *Journal of Agricultural and Resource Economics*.
Reprinted here with permission of the publisher

Abstract:

The U.S. Department of Agriculture (USDA) publishes monthly Ending Stocks projections, providing an estimate of the end-of-marketing-year inventory of a particular commodity, which effectively summarizes its supply and demand outlook. By comparing USDA's projections of balance sheet variables against their realized values from marketing years 1992/3 to 2019/20, we decompose ending stocks forecast errors into errors of the other supply and demand components. We apply a decision-tree-based ensemble Machine Learning (ML) algorithm, the Extreme Gradient Boost Tree (EGBT), that uses a gradient boosting framework and is robust to multicollinearity. Our results indicate that export and production misses are the key contributors to ending stocks projection errors. Because foreign imports of U.S. products are likely tied to foreign production deficits, we likewise investigate how U.S. export errors are linked to USDA's foreign production and export forecast misses, country-by-country, and show that misses on production and export levels in China, Mexico, Brazil, and European Union cost USDA the most. Overall, our results make a strong case that better information about production expectations, both domestically and worldwide, will contribute to more efficient agricultural balance sheet forecasts.

3.1. Introduction

The U.S. Department of Agriculture (USDA) publishes monthly supply and demand forecasts in a balance sheet format for key domestic agricultural commodities; balance sheet elements include beginning stocks, production, domestic consumption, exports, imports, and ending stocks. Recent agricultural economics literature shows that these reports provide the market with important information, align expectations, resolve uncertainty, and enhance economic activity (see, e.g., Isengildina-Massa et al, 2008; Adjemian, 2012; Ying, Chen, & Dorfman, 2019). Following announcements, USDA reports reduce implied volatilities and elevate realized volatilities in futures markets, indicating that market participants react to new commodity information collected and published by the government.

Ending Stocks projections represent the expected end-of-marketing-year inventory of a particular commodity, summarizing its supply and demand outlook. Because they signify the residual after all demand values are differenced from the supply side; ending stocks *balance* the balance sheet. Estimates of ending stocks signal the level of scarcity in the market, guide production and use plans, influence carrying costs, impact prices paid to farmers, and affect import and export decisions. Good & Irwin (2016) propose a curvilinear relationship between farm prices and the forecasted ending stocks-to-use ratio, a projection of the amount of inventories that will be carried to the next marketing year as a percentage of the total amount consumed in the current marketing year.

The task of forecasters is complicated by tail events and structural changes to markets; during our period of observation, these include changes to ethanol policy; the rise of China as an export market; late plantings, droughts, bumper harvests; the U.S. trade war with China; and the COVID-19 pandemic (which significantly reduced the derived demand for corn in the form of ethanol). Forecast errors that events like these produce can negatively impact the welfare of the market participants. For example, overestimates (underestimates) of ending stocks could lead to lower (higher) producer prices and distort storage and production signals, ultimately generating inefficient intertemporal resource allocation decisions. Federal government budgets, payments, and expenditures based on these estimates will also be inefficient (Xiao et al., 2017). Given that, several studies evaluate the accuracy and efficiency of USDA forecasts. Isengildina-Massa et al. (2006) find evidence of forecast smoothing—positive correlations in forecast revisions possibly resulting from strategic behavior on the part of forecasters. Isengildina-Massa et al. (2013) uncover multiple sources of average absolute forecast errors stemming from macroeconomic and behavioral sources consistent with the patterns of optimism and pessimism. No and Salassi (2009) observe time-based deterioration in USDA estimates, and Isengildina-Massa et al. (2020) draw attention to persistent under-estimation in USDA's net cash income forecasts.

In this article, we decompose ending stocks errors into the errors of other supply and demand components in order to study how different balance sheet elements contribute to ending stocks projection errors. No previous study has conducted an in-depth analysis or developed a comprehensive model to answer this question. One candidate approach to find the contribution of each balance sheet element is to use a regression framework and observe what proportion of the variation in the dependent variable (ending stocks) is explained by each (balance sheet) regressor. Common choices include applying LMG (Lindeman, Merenda, and Gold, 1980) and proportional marginal variance decomposition (PMVD) (Feldman, 2005) methods to find the relative importance of regressors in a model. LMG reports the average contributions of sequentially-added regressors to the R-squared of an ordinary least squares model. PMVD builds upon LMG and ensures that an independent variable's regression coefficient with a true coefficient of zero asymptotically approaches zero. However, by construction, ending stocks are a linear combination of other balance sheet variables under consideration, so any regression of ending stocks on other balance sheet values suffers from multicollinearity—contaminating the LMG and PMVD measures. To resolve this problem, we propose a decision-tree-based ensemble Machine Learning (ML) algorithm, Extreme Gradient Boost Tree (EGBT), that uses a gradient boosting framework and is robust to multicollinearity. In this paper, we conduct two Monte-Carlo simulation exercises to demonstrate the superior performance of EGBTs in the presence of multicollinearities.

This study makes important contributions to the existing literature. First, our empirical approach proposes a machine learning model that indicates how misses in each balance sheet element contributes to ending stocks errors, and is flexible to capture non-linear relationships in the data. Second, extant work generally explores balance sheet errors for one or two commodities. We decompose USDA’s ending stocks projection errors for four major U.S. commodities: corn, cotton, soybeans, and wheat.

According to our results, for all four commodities U.S. export and production forecast errors are the primary contributors to ending stocks misses. Predictably, production misses tend to cluster pre-harvest, while export misses are generally highest post-harvest. Given the importance of exports projections in explaining ending stocks forecast errors, in another step we first study the contribution of USDA’s foreign imports forecast misses, country-by-country, to U.S. export forecast errors, and find that—across commodities—import forecast misses by China, European Union, Brazil, Mexico stand out. Because foreign imports are likely tied to foreign production deficits, we likewise investigate how U.S. export errors are linked to USDA’s foreign production forecast misses, country-by-country, and show that misses on production levels in China, Brazil, European Union, and Southeast Asia cost USDA the most. That is, our findings indicate production forecast misses—both domestic and international—to be the main culprit. We interpret these results to mean that better information about global production expectations would contribute to more accurate agricultural balance sheet forecasts.

By using ML to decompose agricultural commodity projection errors our work avoids multicollinearity problems associated with traditional regression analysis and identifies likely sources of forecast misses, indicating to government agencies like USDA where they should focus to improve the informational value of their reports. Ultimately, more accurate forecasts will provide better information to traders and market participants, improving the efficiency of the agricultural supply chain.

3.2. Background

Botto et al. (2006) study USDA's ending stocks forecasts for corn and soybeans from 1980/81 – 2003/04, and use linear regression to show that almost every balance sheet element contributes to ending stocks forecast errors. The authors find evidence of overestimation in the ending stocks estimates for soybeans, and a high degree of bias in the early months of a given marketing year; they do not observe any significant bias in ending stocks projections for corn. Isengildina-Massa, Karali, and Irwin (2013) use correlation analysis and determine that production forecast misses are the primary source of error in ending stocks estimates for corn, soybeans, and wheat from 1987/88 – 2009/10. The authors identify macroeconomic factors (like producer price and exchange rate indices) and indications of behavioral sources of error (such as lagged forecast errors and percent change in forecasts from the previous year's values) in the ending stocks misses for all three commodities.

MacDonald and Ash (2016) identify bias in soybean export forecasts as the likely source of errors in ending stocks forecasts. The authors find incidences of smoothing (Isengildina-Massa et al., 2006) in all balance sheet projections for soybeans. They also observe a bias in foreign trade forecasts consistent with the bias in ending stocks and export projections. MacDonald et al. (2017) analyze soybean balance sheet elements using mean absolute percentage error and Diebold-Mariano tests. They show soybean production errors influence USDA's export forecasts; they also link price forecast errors and ending stock errors. From 2004-2015, the authors note a downward bias in export projections, implying that forecast errors in foreign balance sheet elements are associated with errors in U.S. exports. Xiao et al. (2017) decompose ending stocks forecast errors (of corn, wheat, and soybeans from 1985/86 through 2014/15) into idiosyncratic errors and unpredictable shocks using a Bayesian Markov Chain Monte Carlo (MCMC) method. The authors find inefficiency in projections for each commodity—soybeans worst of all. They detect bias in soybeans forecast revisions but not in corn and wheat revisions. Corroborating the previous literature, they also find evidence of USDA's tendency to overestimate soybean ending stocks.

3.3. Methodology

Most of the research studies in the previous section use linear regression methods to check for bias and make inferences about the importance of different variables (based on their statistical significance). This approach has several limitations, including that (1) regression coefficients capture the average effect of a regressor on the outcome (given exogeneity), but not how much it contributes to the explained sum of squares, and (2) by construction, ending stocks are a linear combination of other balance sheet variables. Regression of ending stocks on other balance sheet values suffers from multicollinearity, contaminating the regression estimates. Moreover, these studies do not include all the balance sheet elements in their analysis. Given our objective of determining the contribution of misses in each balance sheet element to ending stocks forecast errors, OLS is unsuitable because such a regression suffers from perfect multicollinearity; in that case the covariate matrix $X'X$ in $\widehat{\beta_{OLS}} = (X'X)^{-1}X'Y$ is not invertible, and so OLS coefficients and the contribution of each element to the explained sum of squares is not estimable. Common methods to address multicollinearity include partial least squares, and principal component analysis. However, each of these models limits explanatory power, preventing us from identifying the relative importance of the component balance sheet factors in ending stocks forecast errors.

Another candidate choice is a machine learning (ML) technique termed gradient boosting decision trees (GBTs) that provides predictions for the dependent variable from multiple "trees", trained in succession. ML models offer flexibility in extracting associations from the data. Moreover, since these methods select parameters via a grid search, they depend less on the researcher's priors. (Breiman, 2001; Liaw & Wiener, 2002; Amin et al., 2021; Badruddoza et al., 2021). Decision trees establish a tree-like structure, designed to capture the relevant information in the independent variables, subject to constraints. At every branch of the tree, the model splits the data into two parts. GBTs connect several such trees to provide robust estimates. Recent studies have explored the applicability of GBTs for economic applications. Hossain et al. (2019) study the effectiveness of different variables in predicting calorie-based food securing in rural Bangladesh. The authors show that random forests and GBT perform equivalently to non-machine learning methods. Amin et al. (2021) use extreme gradient boosting trees (EGBTs) to predict the modified Retail Food Environment Index using the data from U.S. census tracts; they use Shapley values to extract features and use EGBTs, random forests, and LASSO to predict access to healthful food retailers, with 72% out-of-sample prediction accuracy. (EGBTs, compared to GBTs, leverage computational abilities of computer hardware to speed up the calculations.) Badruddoza et al. (2021) study the long-term dynamics of yogurt, eggs, organic whole milk, and other fluid milk, using vector autoregression to identify the determinants of price premiums for the above products. They also use EGBTs to capture data-driven insights and find that sellers maintain the price of milk, but increase the premiums for yoghurt and egg as input costs increase. In a 2020 working paper, Woloszko notes the importance of GBTs for macroeconomic modeling, particularly in the presence of structural breaks, and favors the easy interpretability of these models. In a 2017 Bank of England working paper, Chakraborty and Joseph study feature importance values obtained from decision trees. Following this literature, we apply EGBTs in our analysis and report the relative importance of each regressor.

We also report Shapley additive explanations (SHAP) values (Lundberg and Lee, 2017). Shapley values quantify by how much does the ending stocks prediction error increases or decreases from the average prediction error due to the inclusion of a given predictor variable. In a 2019 working paper by the Bank of England, Joseph uses SHAP values to calculate contributions for an underlying random forest (a collection of decision trees) to model unemployment in the U.S. and U.K. In our application, we use SHAP values to report the relative contribution of misses on each balance sheet element to the ending stocks projection errors for each observation in the dataset. They measure the influence of the independent variables by comparing the model with and without that variable.

Because we find that exports play a crucial role in determining ending stocks errors, we study how USDA's country-by-country import projections contribute to the Department's export forecast misses. Given the influence of foreign production deficits on the demand for foreign imports, we also study how country-by-country production estimates link to U.S. export misses. Below, we describe in detail our methods for constructing forecast errors and then calculating the relative contributions of balance sheet components to ending stocks errors.[9] Lastly, we run a relative importance simulation exercise to gauge the performance of EGBT relative to linear regression, and report those results in the appendix.

3.3.1. Forecast Errors

We compute forecast errors as follows, depending on which side of the balance sheet the variable resides. If a positive shock to a given element increases the domestic commodity supply, as in the case of beginning stocks, production, and imports, we calculate its forecast errors as:

$$FE_{s,i,t} = Actual_{s,i=19,t} - Forecast_{s,i,t} \tag{1}$$

where ***i*** is the forecast index (which takes on values from 1 to 19, representing USDA's marketing year forecast horizon), ***s*** refers to balance sheet elements in question, and ***t*** refers to the forecast year.

Alternatively, if the positive shock contributes to the commodity's domestic disappearance, as in the case of exports and use, we calculate its forecast errors as:

$$FE_{d,i,t} = Forecast_{d,i,t} - Actual_{d,i=19,t} \tag{2}$$

[9] We use a similar approach to calculate the contribution of foreign import forecast (and production forecast) errors to USDA export forecast errors.

where the other subscripts are the same and $\boldsymbol{d}$ refers to the balance sheet elements in question. By construction, it follows that the forecast errors for all these elements sum to the ending stocks forecast error for all $\boldsymbol{s},\boldsymbol{i},\boldsymbol{t}$, with ending stocks forecast errors calculated according to equation (1).

3.3.2. Using Extreme Gradient Boosting Trees to Measure Relative Importance

Decision trees are a type of supervised machine learning tool that sequentially split data to generate predictions for a dependent variable.[10] Decision trees pick a regressor and split the dataset into two parts. Different split points are tested, and the one resulting in the lowest prediction loss (defined in equation 3) is selected. They repeat this procedure for all regressors in the dataset. Of all the independent variables, the one resulting in the lowest loss is selected for the first split point. This procedure repeats until it is no longer possible the split the data (Quinlan, 1993). However, these trees are often associated with overfitting (poor out-of-sample performance), high variance (making them highly susceptible to outliers), and bias (Leiva et al., 2019). Chen and Guestrin, 2016, propose "extreme boosting" decision trees. Boosting refers to obtaining a robust model by sequentially combining several models—decision trees in our case. The task of each sequential tree is to fit the model by reducing the errors made by the previous model.

Let $\hat{y}^{i,h}$ be the prediction for iteration h, observation i. The gradient boosting tree minimizes the following loss function, Q^h:

$$Q^h = \sum_{i=1}^{n}\left(y^i - (\hat{y}^{i,(h-1)} + g^h\left(x^i\right))\right)^2 + \Omega(g^h) \tag{3}$$

[10] Supervised machine learning trains the model using the dataset containing the independent and dependent variables. Unsupervised learning, on the other hand, has no data on the dependent variable.

where Ω is the regularization term, imposing a penalty to the loss function and preventing overfitting. The vector of independent variables for observation i is given by x_i. The loss function adds a decision tree g^h to further reduce the error between the actual value y^i and the prediction, $\hat{y}^i$.

The EGBT approach uses a gradient descent algorithm to minimize the loss function in equation (3). Gradient descent tweaks the parameters in this equation in the direction of the steepest descent of the loss function. The process is run through several iterations to arrive at the local minimum. The EGBT algorithm has a few parameters that deserve further discussion. Its "learning rate" determines the step size to approach the minimum of the models defined in equation (3), and prevents overfitting, while its "depth" equals the total number of splits to be made. In addition, the practitioner chooses the number of independent variables to select at each split point, and the number of sequential decision trees to fit for the EGBT model. Since these parameters control optimization, we select a grid of values for each of them. For each combination of these values, we fit an EGBT model. The combination resulting in the optimal resampling statistic (mean and standard deviation) is selected.

After fitting the EGBT model, we compute the relative contributions of each regressor to the dependent variable by following the approach specified in Woloszko (2020). We calculate the average number of splits using that regressor and the average "height" of that variable (the point in the tree where a variable is selected to make a split) in all the sequential trees to arrive at its importance measure. We conduct this exercise for all independent variables in our dataset.

3.3.3. Shapley Additive Explanations

The relative contributions obtained from EGBT models are recovered using the entire sample. To account for time-varying contributions of different balance sheet elements to ending stocks projection errors, we built the EGBT models for each five-year period from 1992/3 – 2018/19. However, it might be beneficial to note average contributions for each observation, which is not possible using EGBTs – since we cannot fit a tree on a single observation. We, therefore, use Shapley Additive Explanations (SHAP), proposed by Lundberg and Lee (2017) to help interpret the predictions we make in each five-year period EGBT model. Those authors define the explanation model (v) as:

$$v(z') = \phi_0 + \sum_{j=0}^{M} \phi_j z'_j \tag{4}$$

Where z is the simplified independent variable (also known as features), and $z' \in \{0,1\}^M$ If feature j is present, $z'_j = 1, else\ z'_j = 0$. ϕ_j is the feature attribution for feature j, and $\phi j \in \mathbb{R}$. ϕ is also known as SHAP values. M is the total number of features in the dataset.

To illustrate how we use SHAP values, consider the dataset for corn with 500 observations. The dependent variable is the absolute ending stocks forecast errors. The independent variables are projection errors for beginning stocks (*B.S.*), exports (*Exp*), imports (*Imp*), food, seed & industrial (*FSI*), feed & residual (*FR*), ethanol (*E*), and production (*P*). For observation 5, equation (4) simplifies to:

$$v(5) = \phi_{BS,5} + \phi_{Exp,5} + \phi_{Imp,5} + \phi_{FSI,5} + \phi_{FR,5} + \phi_{E,5} + \phi_{P,5} \tag{5}$$

Where *v (5)* is the SHAP value for the fifth observation. In effect, equation (5) decomposes observation 5 into the linear influence of each feature; we do this for all the observations in the data set.

Lundberg et al. (2018) use game theory to derive the feature attribution for each independent variable. The model assumes the values of the independent variables (of each observation) as the players and the value of the dependent variable as the payout. An equitable distribution of the payout among each independent variable is computed using Shapley decomposition. SHAP values exhibit specific desirable properties such as consistency, missingness, and local accuracy. The values are "consistent" in that if the importance of one feature increases in a new model, its SHAP value will never decrease. "Missingness" just signifies that any missing features get a zero attribution in equation (5), while "local accuracy" constrains that all feature attributes sum to the value of the dependent variable (Lundberg and Lee, 2017; 2018). Therefore, for robustness, we report SHAP values along with our EGBT results for each commodity.

Along with EGBT, as an additional robustness check, we also estimate Random Forests in this article.[11] For each of these models, we split our data randomly (following Bajari et al., 2015) into a training set (80%) and a test set (20%). We train our models on the training dataset and use that model to make predictions of ending stocks misses on the test dataset. We compare the test model predictions with the actual values of ending stocks forecast errors by calculating the standardized root mean squared error between the two. These results are reported in the appendix

[11] To conserve space, we only provide results for our EGBT model.

table A3.6. We find the performance of both models to be quite similar, with Random Forests performing slightly better compared to EGBT for some commodities, while EGBT does for others. Amin et al. (2021) note that Shapley values are computationally expensive, and hence they use EGBT for faster calculations. We follow them and report the Shapley values for the EGBT model only. Also, we report the parameter tuning figures for EGBT in the appendix table A3.7. We fit between 100 to 200 trees, and the change in parameter values does not impact the model output.

3.4. Data

We draw forecast data from USDA's monthly World Agricultural Demand and Supply Estimates (WASDE) data for corn, soybeans, wheat, and cotton, maintained by the Cornell University library system. These reports provide government projections of marketing-year beginning stocks, production, domestic consumption, exports, and ending stocks for several commodities, in both domestic and international markets. Each forecast cycle runs 19 months, beginning in May and ending in November of the following calendar year. We extract these forecasts for corn, soybeans, wheat, and cotton from the 1992/1993 to 2018/2019 marketing years. This procedure yields 510 observations for corn, soybeans, and wheat, and 434 observations for cotton (since those data are only available beginning in 1996/97).[12] Table 3.1 provides summary statistics for USDA forecast errors of these elements for the four commodities. As we describe in section 3.1, we split the elements into those that either increase or decrease commodity.

[12] Our timeframe includes three missing reports relative to the normal and expected schedule. Both are associated with U.S. government shutdowns, which curtailed normal publication of USDA forecasts (Goyal and Adjemian, 2021; Adjemian et al., 2017). Two missing reports in 2013, and one in 2019.

In table 3.1, units differ across commodities so we restrict our comparisons within each commodity's balance sheet. Export forecast errors are highest for cotton, while the feed & residual category are largest for corn and wheat. Average export projection errors are negative for all commodities, indicating the USDA's tendency to underestimate U.S. exports (Isengildina-Massa et al., 2020); ending stocks projections for corn, wheat, and cotton exhibit a similar pattern. In contrast, the Department tends to overestimate ending stocks for soybeans.

Table 3.1: Average forecast errors for USDA balance sheet elements

	Commodity Mean Forecast Errors			
Balance Sheet Elements	**Corn (Million Bu.)**	**Soybeans (Million Bu.)**	**Wheat (Million Bu.)**	**Cotton (Million 480-pound bales)**
Variables increasing commodity supply				
Beginning Stocks	7.48	-4.9	1.65	-0.05
	(2.66)	(1.23)	(0.45)	(0.01)
Production	-10.7	9.93	3.57	0.07
	(18.92)	(4.5)	(2.28)	(0.04)
Imports	3.97	1.53	2.38	0.004
	(0.7)	(0.39)	(0.53)	(0.002)
Variables decreasing commodity supply				
Exports	-6.04	-33.51	-3.85	-0.25
	(10.2)	(4.51)	(3.06)	(0.06)
Feed & Residual	33.47		11.22	
	(9.34)		(2.12)	
Food, Seed & Industrial*	-1.07	4.28	3.76	
	(2.08)	(1.29)	(16.47)	
Ethanol	-9.45			
	(4.99)			
Crushings		-20.48		
		(2.09)		
Domestic Use				0.09
				(0.02)
Ending Stocks	17.64	-43	17.79	-0.13
	(12.68)	(4.11)	(3.07)	(0.06)

Standard errors are reported in parentheses. Mean forecast errors for corn, soybeans, and wheat are in million bushels. Whereas for cotton, they are in million 480-pound bales.
** We provide seed & residual for soybeans and food & seed for wheat in the food, seed & industrial category.*
Source: Authors' calculations based on USDA data.

Tables A3.1 to A3.4 in the appendix provide summary statistics for USDA's projection errors for country-by-country imports and production. For corn and wheat, exports to ROW countries have highest (in absolute terms) forecast errors. China's import projection misses are highest for soybeans, while exports to Turkey dominate for cotton. On average, absolute corn production misses are highest for China, absolute production errors for Brazil dominate for soybeans, while ROW production errors are highest for cotton and wheat. As a robustness check, we also conduct our entire analysis using the logarithm of balance sheet forecast errors, to avoid the impact of changing production levels over the study period. Appendix table A3.5 reports summary statistics for these forecast errors. Although not presented here to conserve space, the results, are broadly similar to what we observe with normal forecast errors. Therefore, we continue with forecast errors computed using equations (1), and (2). Before proceeding with any analysis, for each rolling five-year period, we normalize the forecast errors to mean zero, and variance one.

3.5. Results

Below, we discuss the results for corn, soybeans, wheat, and cotton.

3.5.1. Corn

Figure 3.1 plots USDA's historical forecast errors corn's balance sheet elements. By construction, large production (calculated as actual > forecast) and export errors (calculated as forecast > actual) lead to high USDA ending stock projection misses (calculated as actual > forecast). In both cases, misses mean that more commodity is on-hand domestically, raising ending stocks. As an example in 2004/2005—especially early in the season—USDA missed high on both these elements, so its corn ending stocks forecasts also missed high. In 2012/2013, USDA missed high on exports but missed low enough on production (due to a historic drought) that its ending stocks errors were also low (especially early in the forecast cycle).

Figure 3.2 presents EGBT results for the relative importance of each balance sheet element forecast to corn's absolute ending stock errors over a rolling five-year period. The figure also shows the average absolute ending stocks forecast errors within each period in order to provide contextual comparison for errors over time. Our results show that production and export misses are significant contributors to the ending stocks misses. Exports contribute as much as 65% to ending stocks errors, while production contributes as much as 77%. Feed & Residual, on the other hand, contribute as high as 69%. Ethanol's importance increases beginning in the early 2010s, rising to 73% of ending stocks forecast errors, on average, during the last three five-year cycles.

Figure 3.3 provides the SHAP plot for corn. The vertical axis of the left panel labels the errors in descending order of their relative importance in determining USDA's ending stocks forecast errors for corn. SHAP scores in figure 3.3 attribute the highest importance to export projections, followed by feed & residual and production. These conclusions confirm our findings in figure 3.2 demonstrating the robustness of our results. To read the chart, note that each circle represents a single observation; its color indicates the value of the independent variable (on a scale of low to high), while its position signifies how much it influences ending stocks errors. For example, consider region A in the left panel. The purple color of circles in this region indicates that export errors are moderately-high to high (based on the scale given at the bottom of the plot), and their location indicates that a standard deviation increase in such errors raise ending stocks projection errors between 1.5-2 standard deviations. Averaging over all the SHAP values related to exports produces a value of 0.214 standard deviations, which is located next to the exports entry in the figure.

Since exports are the most crucial factor in determining USDA's ending stocks errors, we also show their marginal effect on the latter in the right panel of figure 3.3. The nonlinear relationship depicted in the figure indicates that export predictions misses make a decreasing marginal contribution to ending stocks errors as they grow larger (because other variables exert more influence).

Due to the importance of export projection errors in corn ending stocks misses, we next conduct a trade analysis to determine which export destinations contribute the most by running a similar analysis on USDA's country-level corn import forecast misses. Figure 3.4 presents EGBT relative importance results for our international trade analysis. Egypt (especially during the latter half of the sample period), Former Soviet Union, and ROW (rest of world) corn import errors are the most important.[13] Because foreign import projections are closely tied to foreign production forecasts, we conduct a follow-on analysis of the relative importance of the latter to USDA export forecast misses. Figure 3.5 shows the relative importance of the errors made in estimating production levels in these countries to U.S. corn export misses. In the figure, Argentina, South Africa, the European Union, and Mexico misses are among the most important, although Brazil and China also appear in certain time periods. USDA's December 2021 WASDE (USDA, 2021)

[13] ROW includes Colombia, which imports a high quantity of corn from the U.S. (FAS, 2019). Although not presented here to conserve space, in an auxiliary analysis SHAP values confirm our findings.

indicates that Argentina, Brazil, and South Africa are among the world's major corn exporters, while the European Union and Mexico are among its major importers. China is the world's second leading corn producer after the United States.

3.5.2. Soybeans

Figure A3.3 in the appendix plots the balance sheet element forecast errors for soybeans. Like corn, higher production and export forecasts (in absolute value) play an important role in soybean ending stocks errors. In general, USDA misses positive on production and negative on exports (underestimating both categories, which in our analysis leads to opposite signs). The same story applies to soybean crushings. In 2018/19, however, USDA missed high—very sharply—on exports, coinciding with the onset of the trade war between the United States and China (Adjemian et al., 2021). This is notable also because it broke with over a decade of practice of USDA missing high (underestimating) U.S. soybean exports. As with corn, USDA tends to balance large misses with similarly large misses on the opposite side of the balance sheet: production and export misses tend to work in opposite directions. For example, in 2016/17 USDA missed very high on production and very low on exports.

Appendix figure A3.4 presents our EGBT relative contribution results for soybean ending stocks errors over a rolling five-year period. In the figure, production (32% contribution, on average) and export misses (31%, on average) are the lead contributors. Crushings and seed residual errors each contribute 8-10% to ending stocks errors. Figure A3.5 in the appendix presents observation-level SHAP scores, for robustness. They confirm that exports and production are the most crucial balance sheet elements, followed by crushings. Just as in the case of corn, as USDA's export misses increase, they contribute less to ending stocks errors on the margin.

Since the early 2000s (coinciding with China's accession to the World Trade Organization –Agarwal and Wu, 2003), USDA tends to underestimate Chinese soybean imports (at least partially explaining its underestimate of U.S. soybean exports). However, in 2018/19—again, the onset of the trade war—USDA overestimated both China's soybean imports and U.S. soybean exports. Appendix figure A3.6 presents EGBT relative importance results for international trade and confirms the intuition that USDA's Chinese import misses are responsible for the bulk of its export misses, especially as those misses rise (in absolute value) during the early-to-mid 2000's. China's average influence in explaining U.S. ending stocks errors peaks during two periods: following China's World Trade Organization (WTO) accession (66%), and during the ongoing trade war (55%). USDA's forecast of Brazilian imports contributed notably during the first half of the sample period (generally before U.S. exports jumped in size), while the relative importance of its errors in Mexican imports increased since the 2010s. In addition, figure A3.7 (in the appendix) shows how the errors in the world production estimates relate to U.S. soybeans export misses. As opposed to figure A3.6, China's production errors play a less important role. This is intuitive, because while China is the world's foremost soybean importer (forecasted by USDA in December 2021 at 100 million metric tons in 2021/22), its production ranks far down the list (projected at 16.4 million metric tons in the same report) (USDA, 2021). Rather, USDA's domestic export misses are explained to a large degree by its production errors for the world's other large feed soybean producers: Argentina, Brazil, and ROW (which includes Paraguay) (FAO, 2017).

3.5.3. Wheat

Appendix figure A3.8 plots USDA forecast errors for wheat from 1992/93 through 2018/19. As in the cases of corn and soybeans, high production and export errors are associated with large ending stocks misses. The feed residual and food & seed categories also appear prominent at times. But, unlike corn and soybeans, beginning in the late 2000s USDA mostly overestimated, rather than underestimated, domestic wheat exports.

Figure A3.9 (refer to the appendix) explains how these balance sheet elements contribute to ending stocks misses. Production contributes as much as 80% to ending stocks errors, food & seed contribute as much as 60%, and exports' contribution peaks at 68% in the last five-year period. In the appendix figure A3.10, SHAP results confirm the important role played by production, exports, and food & seed elements misses to ending stocks errors.

Appendix figure A3.11 in the appendix presents EGBT results for how international trade contributes to U.S. wheat export errors. Imports by China (16%, on average), ROW (16%), North Africa (19%) and Brazil (9%) are the most influential. Likewise, appendix figure A3.12 relates USDA's country-level foreign production forecasts to its U.S. wheat export errors. Clearly, Brazil, and North Africa feature in both the production and import stories, but European Union's production misses play a more prominent role than do USDA misses on its forecasted imports. This is consistent with the fact that while the European Union is among the world's largest producers and exporters of wheat, its import level is fractionally small (USDA, 2021).

3.5.4. Cotton

Figure A3.13 in the appendix plots the balance sheet element projection errors for cotton. As with other commodities, USDA tends to underestimate both exports and production. During the 2008 recession, which substantially weakened apparel and clothing demand cotton export prices declined by 33% (BLS, 2011), possibly leading USDA to overestimate U.S. exports in that timeframe. Because of the way the balance sheet works, overestimating exports leads USDA to underpredict ending stocks.

Figure A3.14 in the appendix shows the contribution of the various balance sheet elements to absolute USDA ending stocks errors. Export projection misses are clearly dominant; on average, they contribute 57% to ending stocks forecast errors. Specifically, exports become much more important than use from the early-2000's and on, coinciding with China's accession to the WTO; China is the world's leading cotton importer (USDA, 2021). SHAP summary plots in the appendix figure A3.15 confirm the oversized role of forecast misses for exports and domestic use in explaining USDA's ending stocks errors. Unlike other commodities, the relative contribution of export errors maintains importance even as export levels rise.

Since 2011, following China's large buildup of cotton stocks (MacDonald et al., 2015) USDA has mostly underestimated the country's imports, but this hasn't always weighed down its U.S. export forecasts: domestic forecast errors were small in 2011/12 and 2012/13, even though the Department underestimated Chinese imports quite significantly. At the same time, USDA actually overestimated imports by both India and Mexico, contributing to, at times, aggregate U.S. export overestimates. Figure A3.16 in the appendix confirms that intuition: USDA import misses on the part of China, India, and Mexico play significant parts in U.S. export misses. Appendix figure A3.17 explores how U.S. export errors are linked to misses in USDA's country-level foreign production forecasts. European Union, Central Asia, India, and China's influence is notable, with Bangladesh growing in importance over the latter half of the period of observation.

3.5.5. Analysis by Forecast Step

We also compute the contribution of balance sheet elements to USDA's ending stock forecast errors, by forecast step, and present these results in figures A1.18 through A1.21 in the appendix. Doing so provides a portrait of how USDA misses contribute to ending stocks errors along the dimension of the average forecast cycle. For corn, production errors play a significant role in the beginning periods, and the importance of feed & residual errors emerge later in the cycle. For wheat (figure A3.20), food & seed, exports, and feed residuals have higher contributions during the forecast period's beginning, middle, and ending periods, respectively.

For soybeans, wheat, and cotton (figures A3.19 – A3.21), exports play a much important role throughout the forecast cycle. This is consistent with the fact that relative to corn, these three crops are far more export dependent. According to the latest WASDE (USDA, 2021), virtually all the cotton produced in the United States is exported. Over half of the soybeans and wheat produced this year are expected to be exported, as well. In contrast, the USDA expects that just 17% of the corn produced domestically will find its way to foreign consumers. It is therefore unsurprising that over the forecast cycle, exports matter less in explaining corn's balance sheet errors.

3.6. Conclusion and Discussion

Every month USDA publishes forecasts of production, exports, imports, domestic consumption, and ending stocks for major commodities, in both domestic and international markets. These projections help market participants form expectations and make decisions. Ending stocks summarize commodity fundamentals, and so are a natural focal point. Because forecast errors can distort decisions along the supply chain, researchers have focused on explaining them. We contribute to this literature by decomposing USDA's ending stocks forecast errors for corn, wheat, cotton, and soybeans into the influence of different balance sheet variables. As both sides of the balance sheet are equivalent, perfect multicollinearity makes ordinary regression analysis unsuitable for analysis. Common approaches (like principal component analysis, LMG, and least absolute shrinkage and selection operator) to address multicollinearity limit the interpretability of results on the part of the researcher, so we address this problem by using EGBTs.

Our EGBT results consistently highlight the importance of USDA's domestic production and export misses in explaining its ending stocks errors. Export errors tend to exhibit a declining marginal contribution to explanatory power; their importance in generating ending stocks misses falls as export levels get very large. Summarized SHAP values with correlation coefficients with the predictors (figure A3.22), confirm the central contribution of exports and production errors to ending stocks misses for corn, soybeans, wheat, and cotton. Because U.S. exports are clearly linked to foreign production and import demand, we also use EGBTs to explore their role. Our findings indicate that misses in these forecasts among large international purchasers and export competitors are key contributors to USDA's projection misses for domestic ending stocks. And for almost every commodity, USDA forecasts for China's imports and production stand out in their importance, especially after its accession to the WTO.

Our results offer a portrait (really, several portraits) of which balance sheet elements affect USDA ending stocks forecasts. Ultimately, improvements in production forecasts could lead to progress in the USDA's and the world's ability to anticipate the global supply situation (and therefore reduce errors balance-sheet-wide). Going forward, technological innovations in satellite sensing, mapping, and geographic information systems may help realize these gains, in addition to assessing agricultural conditions and mitigating risk (Benami et al., 2021).

Another source of USDA forecast errors could be incomplete adjustment to new information. For example, Isengildina-Massa et al. (2006) estimate statistically and economically significant smoothing in corn and soybean production forecasts, possibly due to strategic behavior on the part of forecasters. Obviously, correcting those tendencies would also improve the efficiency of USDA forecasts. However, it may instead be the case that information rigidities (see, e.g., Coibion and Gorodnichenko, 2015) prevent USDA forecasters from accessing the information they need to make fully efficient forecasts. In ongoing work (Goyal and Adjemian, 2021), we test for the presence of information rigidities among agricultural production and yield data. If they are present, investments in technologies like the ones we discuss above could reduce the noise or cost associated with gathering new data, and improve the forecasts themselves.

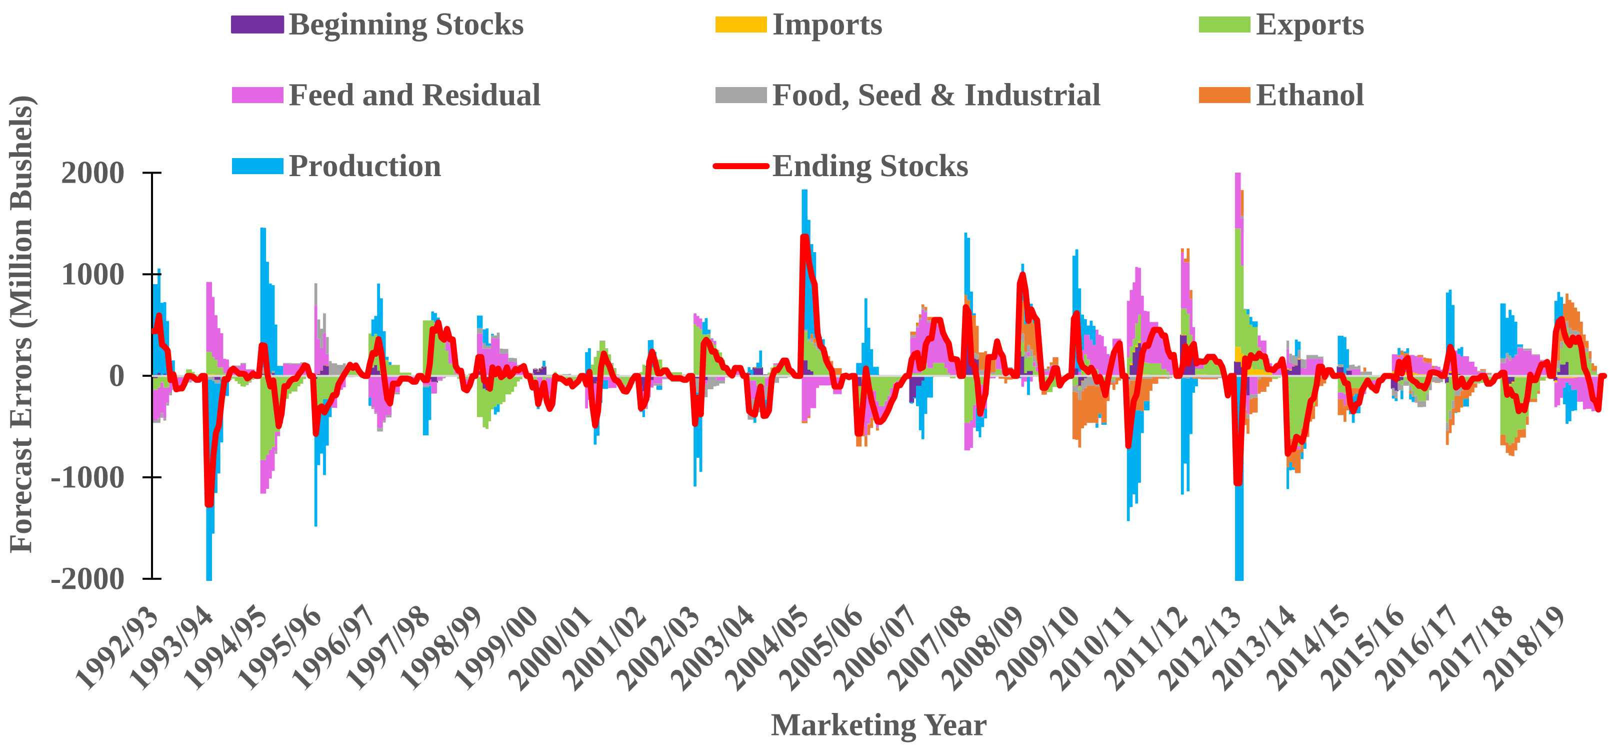

Figure 3.1: USDA Corn balance sheet element forecast errors, 1992-2019

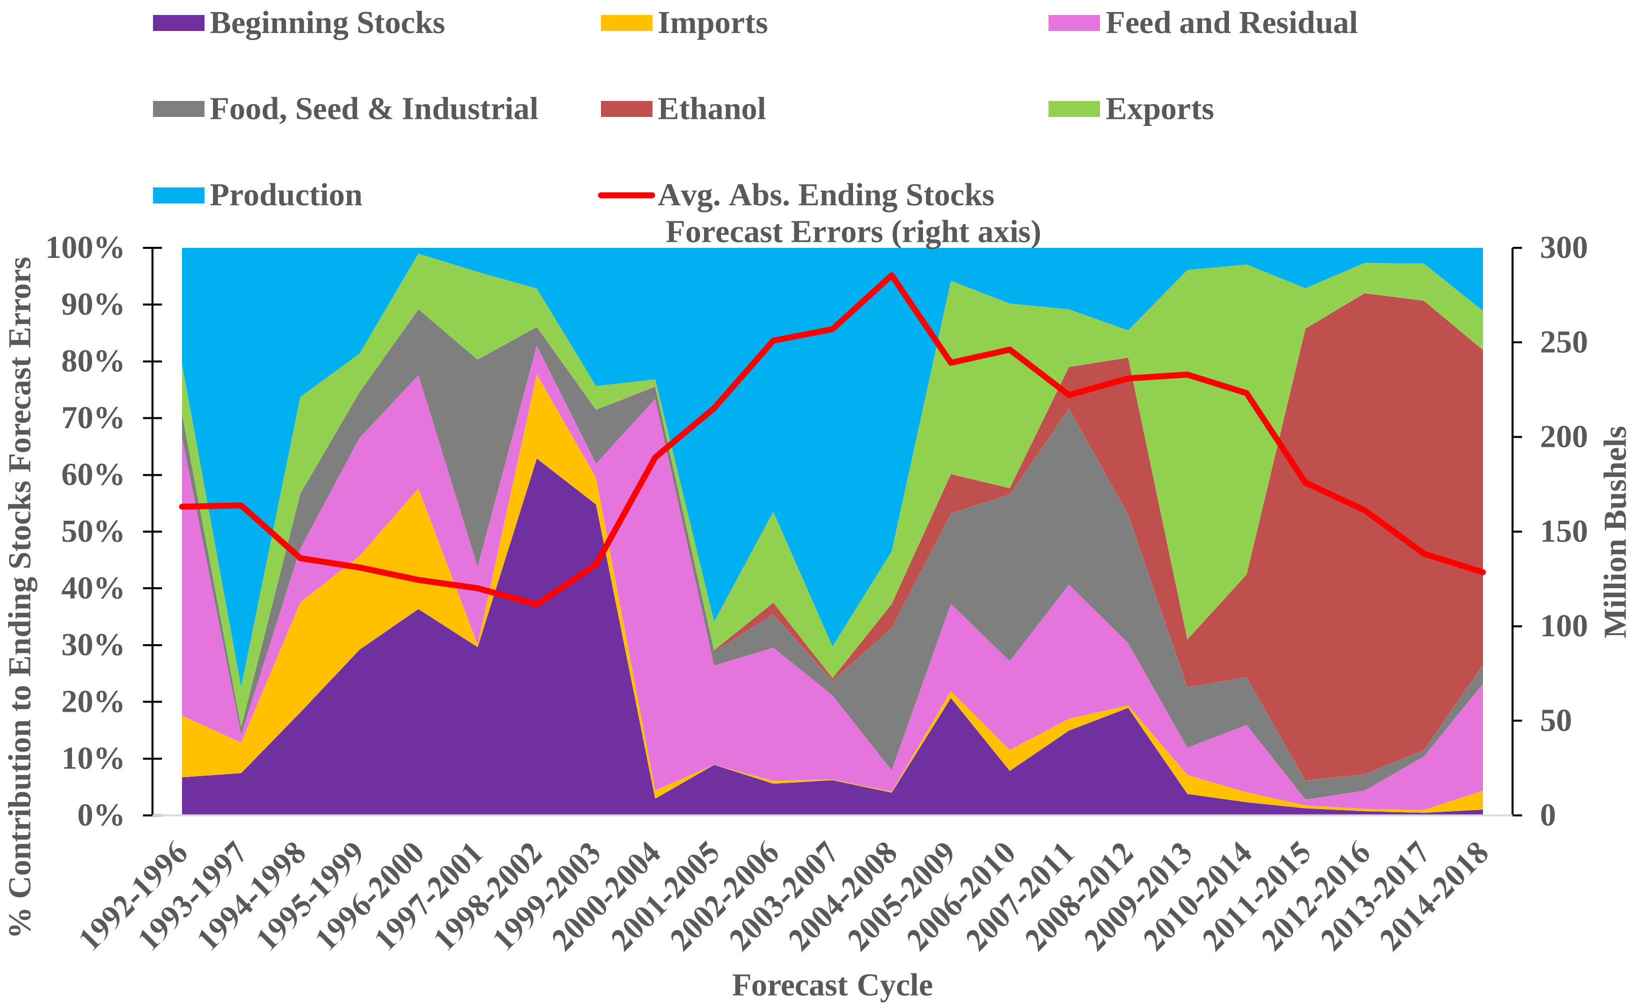

Figure 3.2: Contribution of balance sheet elements to USDA's Corn ending stock forecast errors

Note: The first year quoted on X-axis "1992-1996" represent the 5-year period from marketing year 1992/93 to 1996/97.

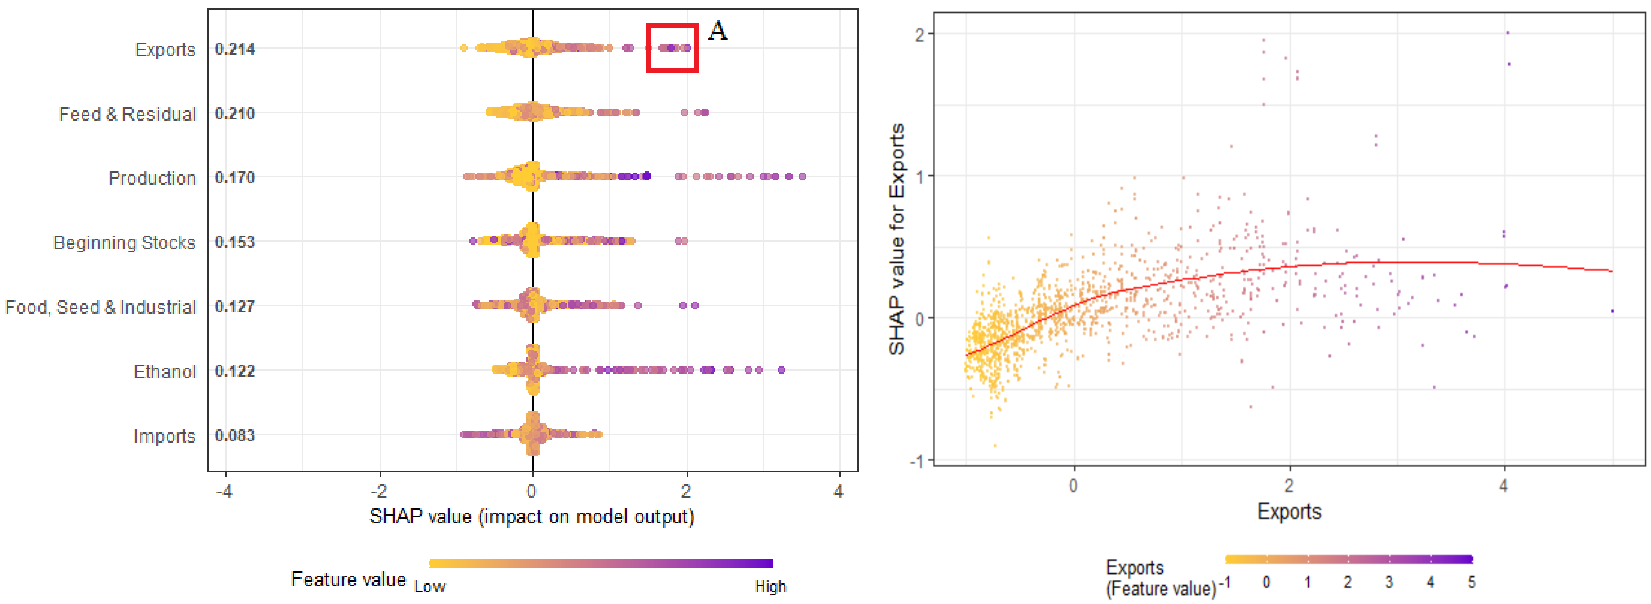

Figure 3.3: Corn – exports SHAP dependence plot

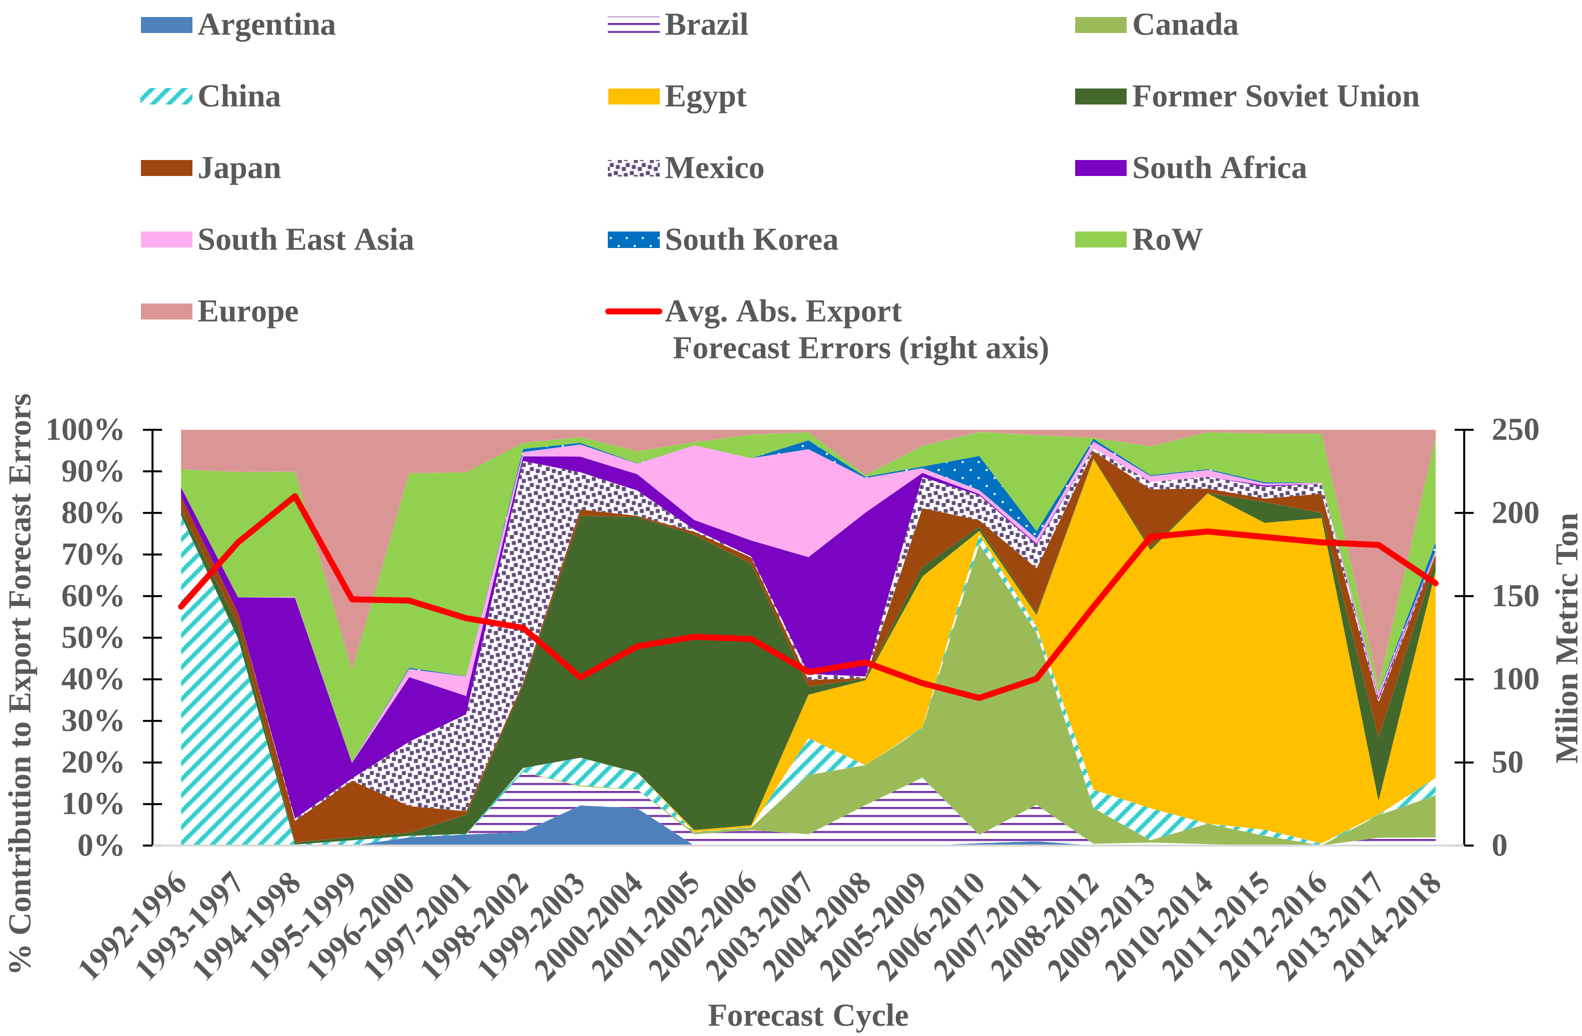

Figure 3.4: Contribution of world Corn import errors to U.S. Corn export projection errors

Note: The first year quoted on X-axis “1992-1996” represent the 5-year period from marketing year 1992/93 to 1996/97.

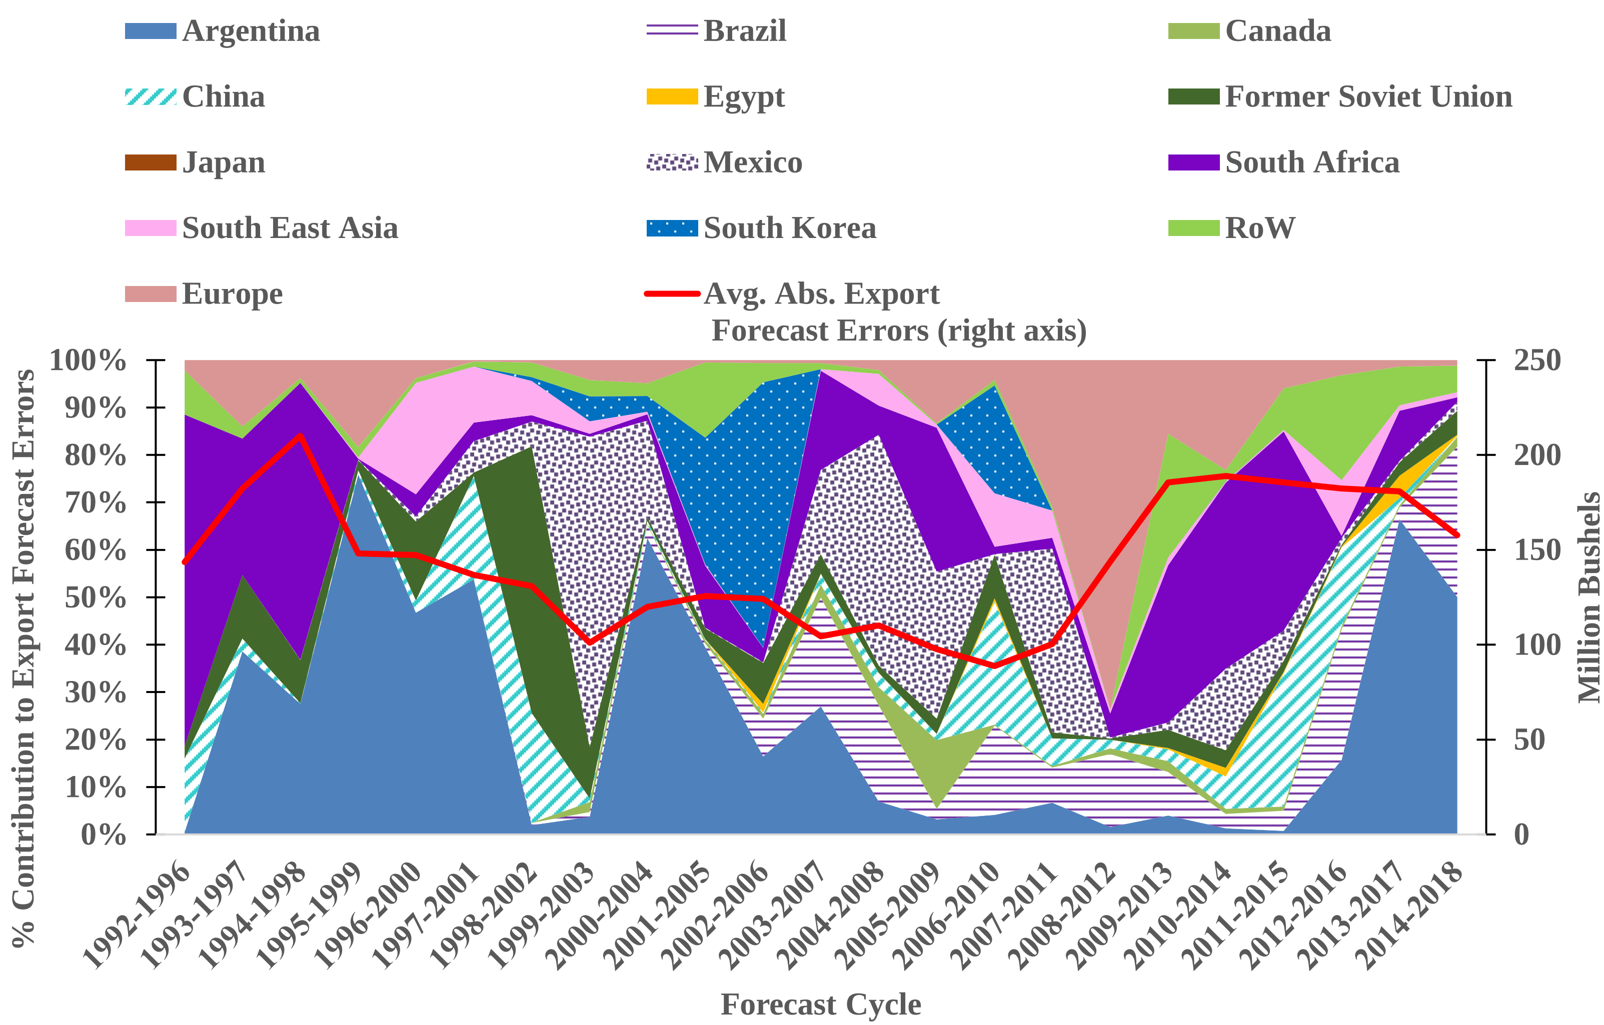

Figure 3.5: Contribution of world Corn production errors to U.S. Corn export errors

Note: The first year quoted on X-axis "1992-1996" represent the 5-year period from marketing year 1992/93 to 1996/97

3.7. Appendix

In a regression framework, researchers use LMG (Lindeman, Merenda, and Gold, 1980) and PMVD (Feldman, 2005) to estimate the relative contributions of regressors. LMG is the average R^2 contribution of the independent variables averaged over all possible orderings. PMVD is a weighted average contribution of regressors to the R^2 with data-dependent weights. Therefore, a regression of ending stocks projection errors on balance sheet forecast misses introduces multicollinearity in the system—contaminating the LMG and PMVD measures. Moreover, when constructing forecasts, USDA consciously attempts to equalize the balance sheet elements. If the agency misses high on one variable, it will effectively miss low on at least one other element. We therefore use absolute values of forecast errors, which eliminates the possibility of a perfectly collinear relationship. However, as indicated by variance inflation factors, regression models still suffer from high collinearity. Furthermore, a curvilinear relationship may exist between different demand and supply element forecast misses and ending stocks forecast errors. Therefore, relative importance analysis based on a linear regression model may be inefficient because multicollinearity reduces the precision of regression coefficients. To resolve this problem, we use EGBTs, since they perform better in the presence of multicollinearity and are robust to nonlinearities. Since we do not note any conscious efforts to balance foreign imports or production to U.S. exports, we use signed errors for our trade analysis.

In the following subsections, we provide a detailed discussion of how we set up the EGBT. But first, we conduct a simulation exercise comparing the EGBT approach to a linear regression modes in order to better understand the difference between the two approaches in computing relative importance.

3.7.1. Comparison of EGBT and Linear Regression

We run a simulation with the following data generating process (DGP)

$$Y = \beta_1 X_1 + \beta_2 X_2^2 + \beta_3 \mathrm{X}_3 + \epsilon, \quad X_3 = X_1 * X_2 \tag{1}$$

$$X_1 \sim N(0,1) \tag{2}$$

$$X_2 \sim N(0,1) \tag{3}$$

$$\beta_1 = \beta_2 = \beta_3 = 0.33 \tag{4}$$

Regressors (X_1, X_2, X_3) are independently and identically distributed. We add a second-order polynomial (X_2^2), and interaction effects in the form of variable X_3. We set the values of parameter estimates $(\beta_1, \beta_2, \beta_3)$ as 0.33. This implies that all the three regressors are equally important in determining the dependent variable, *y*. We repeatedly generate data sets of 500 draws from the prespecified probability distributions for each variable, about the same size as our data set of USDA forecast errors. For both models, we conduct 1000 Monte Carlo simulations. We represent the probability distribution of simulated relative importance scores found via linear regression using LMG,[14] and report EGBT's relative importance values based on each regressor's contribution to reducing the residual sum of squares (see more on this in the next section). To evaluate the two methods, we compare the median of simulated relative importance values to the actual value, also termed the "truth". The model closer to the truth is preferred.

Figure A3.1 compares the kernel density plots for the relative importance of each variable across all simulation trials, for the two proposed models. For all three variables in the figures, EGBT clearly performs much better compared to the regression model. The median relative

[14] PMVD is currently patented in the U.S and, hence, not available for public use.

importance and the actual importance coincide for the EGBT model for X_1. Specifically, for variables $X_2$and X_3, regression models result in inferior performance. Furthermore, regression models result in much higher variance, causing lower precision and efficiency.

Next, to check which model performs better in the presence of multicollinearities, we run a simulation with the following data generating process (DGP):

$$Y = \beta_1 X_1 + \beta_2 X_2 + \beta_3 \mathrm{X}_3 + \epsilon \tag{5}$$

$$\begin{bmatrix} X_1 \\ X_2 \end{bmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix}\right] \tag{6}$$

$$X_3 \sim N(0, 1) \tag{7}$$

$$\beta_1 = \beta_2 = \beta_3 = 0.33 \tag{8}$$

Regressors $(X_1, X_2,)$ are highly correlated with a covariance of 0.8. As with the first simulation, we set the values of parameter estimates $(\beta_1, \beta_2, \beta_3)$ as 0.33, so all three variables make equal contribution to the outcome. We follow all the steps similar to the first simulation – (1) We compute 500 draws for each variable and run 1000 Monte Carlo simulations (2) We report LMG measures for linear regression model and the simulated relative importance values for EGBT (3) We plot kernel densities in figure A3.2 and compare the median of simulated relative importance values to the truth. For all three variables in the figures, EGBT measured contributions are much closer to the truth. Specifically, for non-linear, regression models result in inferior performance. There is only a marginal difference for variable X_2. However, EGBT performs much better for the other two variables, $X_1$and X_3. Please note that the simulation results are for illustration purposes only and are dependent on the model parameters and may change if the number of variables is increased. Therefore, given the superior performance of EGBTs in both the simulations, we report all the relative importance measures based on the EGBT model.

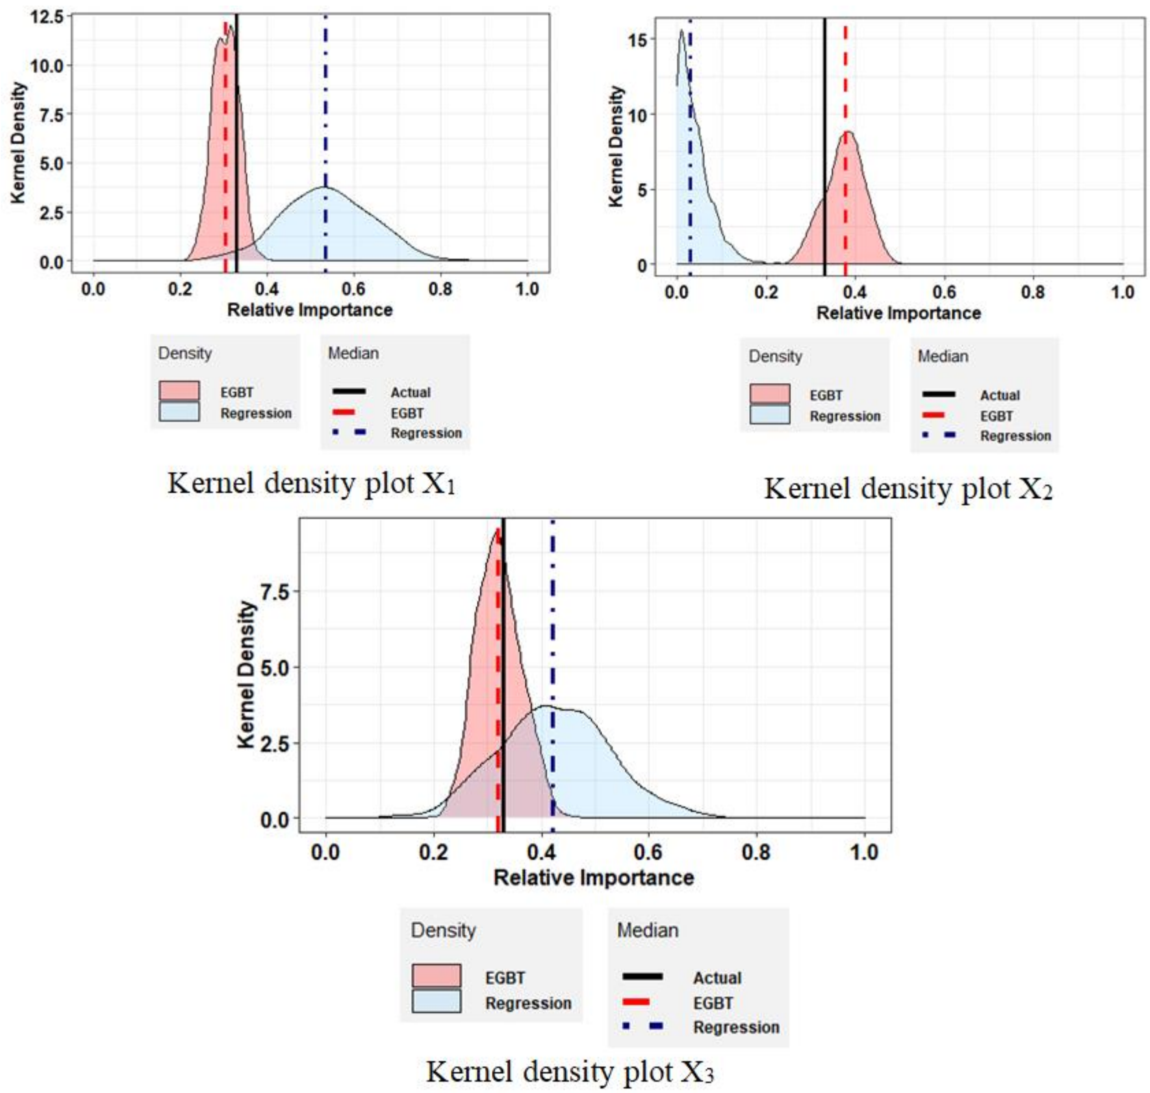

Figure A3.1: Monte-Carlo simulation results –baseline

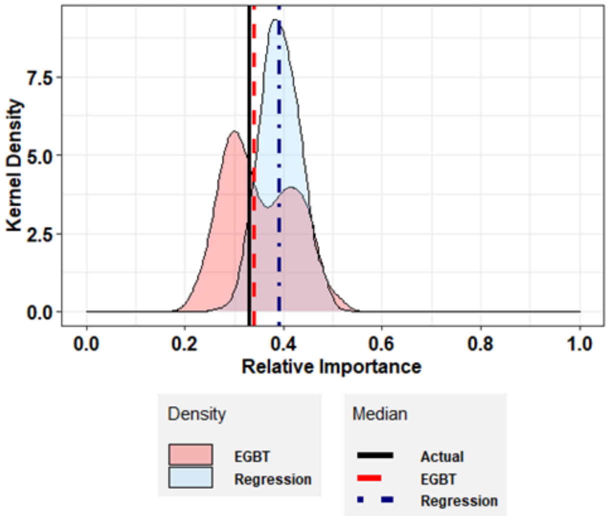

Kernel density plot X_1

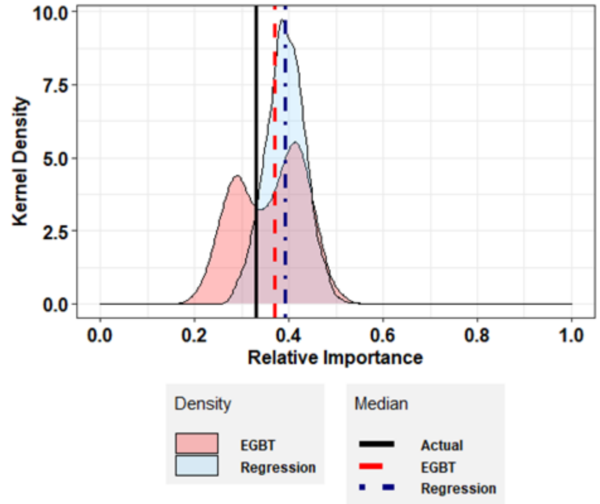

Kernel density plot X_2

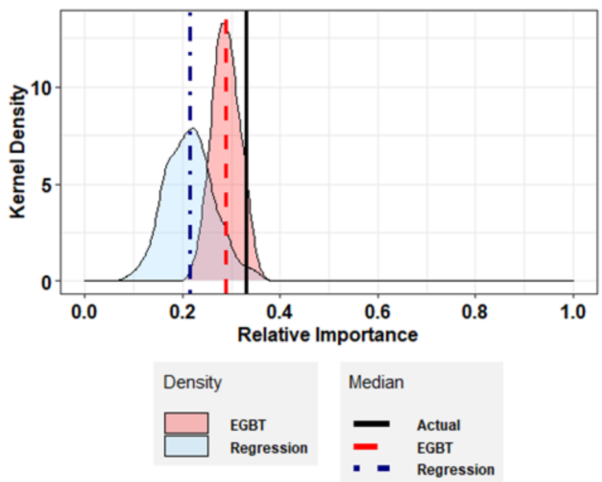

Kernel density plot X_3

Figure A3.2: Monte-Carlo simulation results –multicollinearities

Tables A1.1 to A1.4 provide summary statistics for USDA’s projection errors for country-by-country imports and production. For corn and wheat, exports to ROW countries have highest (in absolute terms) forecast errors. China’s import projection misses are highest for soybeans, while exports to Turkey dominate for cotton.

Table A3.1: Average corn forecast errors for USDA foreign balance sheet elements

Corn	Argentina	Brazil	South Africa	Egypt	Japan	Mexico	Southeast Asia	South Korea	Canada	China	Europe	Former Soviet Union	ROW
							Imports						
Mean	-0.002	-0.094	-0.090	-0.216	-0.088	-0.003	-0.325	-0.139	-0.098	-0.099	-0.816	0.070	-1.359
SE	0.001	0.016	0.022	0.027	0.023	0.048	0.046	0.023	0.024	0.048	0.113	0.071	0.106
							Production						
Mean	0.348	1.566	0.176	0.035	0.000	0.236	0.041	-0.001	0.053	4.094	0.441	0.067	2.576
SE	0.140	0.253	0.076	0.012	0.000	0.057	0.049	0.000	0.017	0.435	0.133	0.086	0.191

Source: Authors' calculations

Table A3.2: Average soybeans forecast errors for USDA foreign balance sheet elements

Soybeans	Argentina	Brazil	Japan	China	Europe	Mexico	ROW
				Imports			
Mean	-0.241	-0.022	0.032	-1.033	0.003	0.046	-0.159
SE	0.034	0.012	0.014	0.155	0.050	0.011	0.076
				Production			
Mean	-0.477	0.543	-0.002	0.086	-0.079	0.004	0.049
SE	0.198	0.185	0.001	0.026	0.084	0.001	0.107

Source: Authors' calculations

Table A3.3: Average wheat forecast errors for USDA foreign balance sheet elements

Wheat	Argentina	Australia	Brazil	Canada	China	Europe	Former Soviet Union	North Africa	Southeast Asia	Select Mideast	ROW
						Imports					
Mean	0.000	-0.024	-0.090	-0.029	0.265	-0.006	-0.039	-0.933	-0.558	-0.332	-1.482
SE	0.000	0.003	0.024	0.002	0.066	0.140	0.093	0.060	0.071	0.053	0.108
						Production					
Mean	-0.394	-0.014	0.061	0.248	0.689	0.028	0.349	0.059	0.110	0.084	1.384
SE	0.147	0.146	0.020	0.068	0.110	0.153	0.256	0.030	0.059	0.023	0.191

Source: Authors' calculations

Table A3.4: Average cotton forecast errors for USDA foreign balance sheet elements

Cotton	Bangladesh	Brazil	Central Asia	China	Europe	India	Mexico	Pakistan	Turkey	ROW
Imports										
Mean	-0.106	0.069	-0.001	-0.445	0.121	-0.042	0.000	-0.038	-0.191	-0.048
SE	0.018	0.011	0.004	0.133	0.011	0.016	0.008	0.029	0.021	0.027
Production										
Mean	0.014	0.877	-0.039	2.759	1.208	-0.559	0.003	-0.396	-0.098	-5.668
SE	0.004	0.174	0.015	0.558	0.304	0.229	0.005	0.089	0.019	1.409

Source: Authors' calculation

Table A3.5: Average forecast errors for log balance sheet elements

Balance Sheet Elements	Commodity Mean Forecast Errors			
	Corn (Million Bu.)	Soybeans (Million Bu.)	Wheat (Million Bu.)	Cotton (Million 480-pound bales)
Variables increasing commodity supply				
Beginning Stocks	0.009 (0.002)	-0.01 (0.005)	0.003 (0.0009)	-0.01 (0.002)
Production	-0.002 (0.002)	0.004 (0.002)	0.001 (0.0011)	0.0016 (0.002)
Imports	0.126 (0.016)	0.073 (0.02)	0.027 (0.0057)	-0.0001 (0.023)
Variables decreasing commodity supply				
Exports	0.002 (0.006)	-0.03 (0.004)	-0.002 (0.0029)	-0.022 (0.005)
Feed & Residual	0.007 (0.002)		0.111 (0.0147)	
Food, Seed & Industrial*	0.0002 (0.001)	0.028 (0.013)	0.025 (0.0024)	
Ethanol	-0.002 (0.002)			
Crushings		-0.013 (0.001)		
Domestic Use				0.018 (0.003)
Ending Stocks	0.025 (0.01)	-0.152 (0.014)	0.022 (0.0047)	-0.04 (0.01)

Standard errors are reported in parentheses. Mean forecast errors for corn, soybeans, and wheat are in million bushels. Whereas for cotton, they are in million 480-pound bales.

* We provide seed & residual for soybeans and food & seed for wheat in the food, seed & industrial category.

Source: Authors' calculations based on USDA data.

Table A3.5 reports summary statistics for the logarithm of balance sheet forecast errors. Export forecast errors are highest for cotton, while imports errors are largest for corn and soybeans. Feed and residual errors dominate for wheat. Table A3.6 reports the standardized root mean squared error observed for the test dataset for the two models. We find the performance of both models to be quite similar, with Random Forests performing slightly better compared to EGBT for some commodities, while EGBT does for others.

Table A3.6: Standardized RMSE for EGBT and Random Forests

	Random Forest	Extreme Gradient Boosting Tree
Domestic Forecasts		
Corn	-6.07	-6.46
Soybeans	-0.71	-0.94
Wheat	-1.41	-0.14
Cotton	10.53	14.74
International Forecasts (Exports)		
Corn	3.03	1.07
Soybeans	-2.25	-1.6
Wheat	3.04	3.66
Cotton	1.94	1.19
International Forecasts (Production)		
Corn	-1.21	-1.61
Soybeans	3.16	2.46
Wheat	6.44	7.96
Cotton	-1.39	-0.13

Source: Authors' Calculations

We report the parameter tuning figures for EGBT in the appendix table A3.7. We fit between 100 to 200 trees, and the change in parameter values does not impact the model output.

Table A3.7: Parameter Tuning Values

Parameter	Range
No of Trees	[100, 200]
Depth of the tree	[2,4,6]
Learning Rate	[0.1, 1]

Source: Authors' Calculations

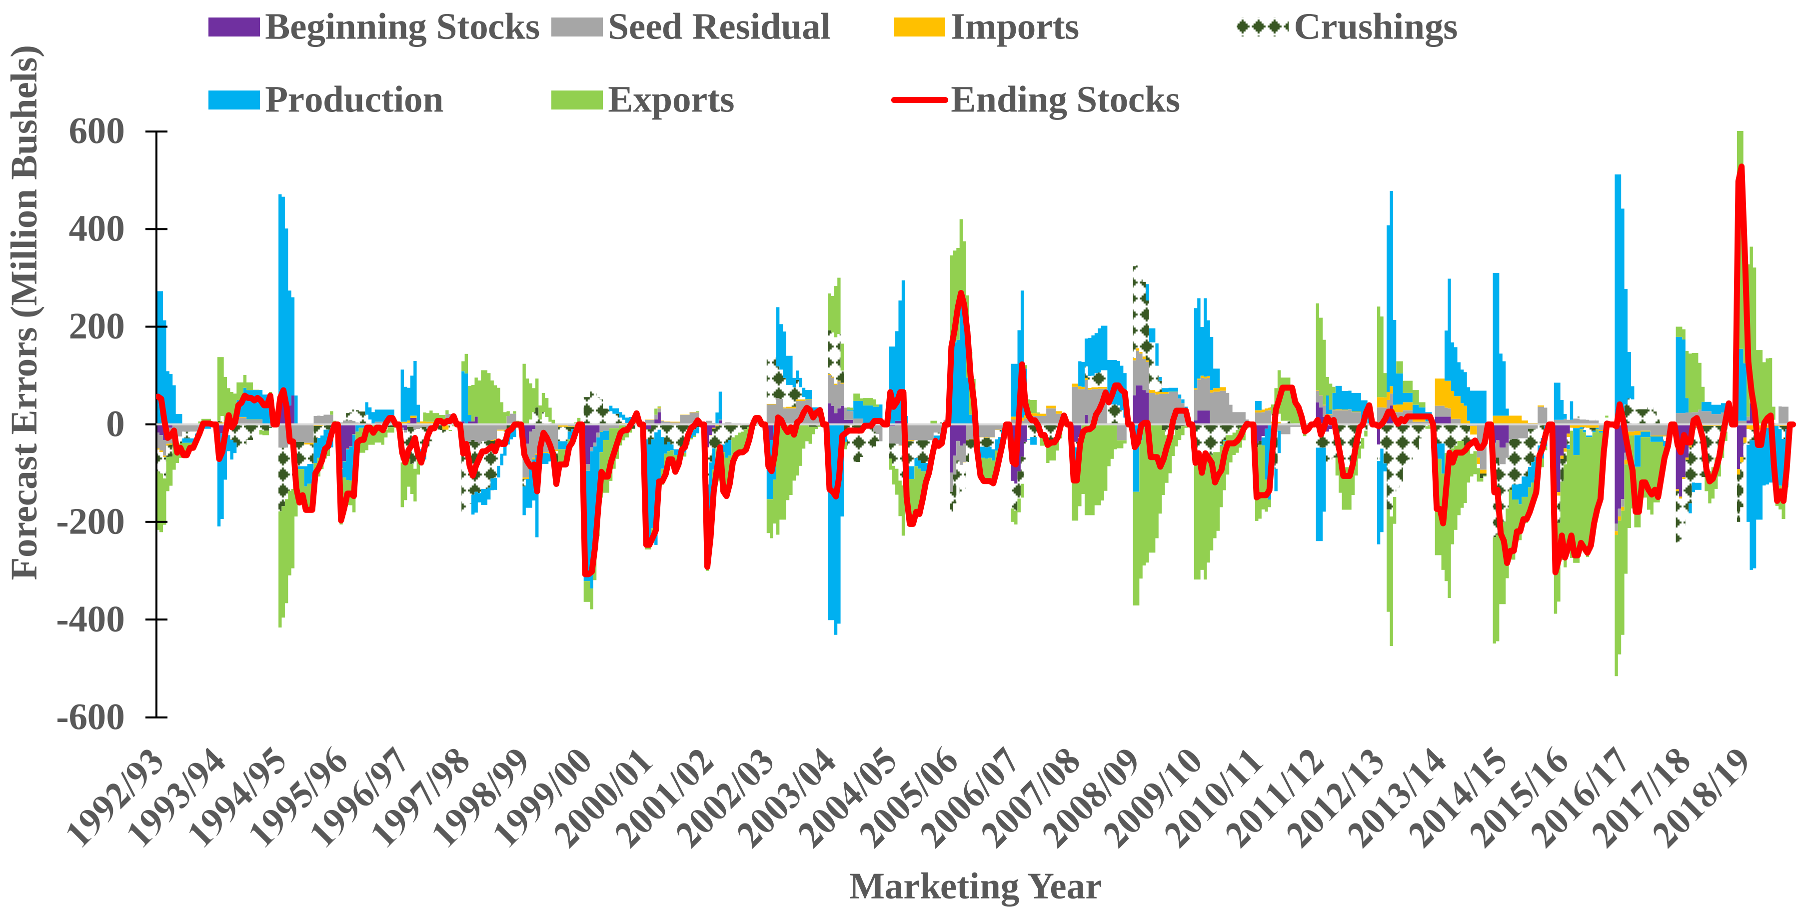

Figure A3.3: USDA Soybeans balance sheet element forecast errors, 1992-2019

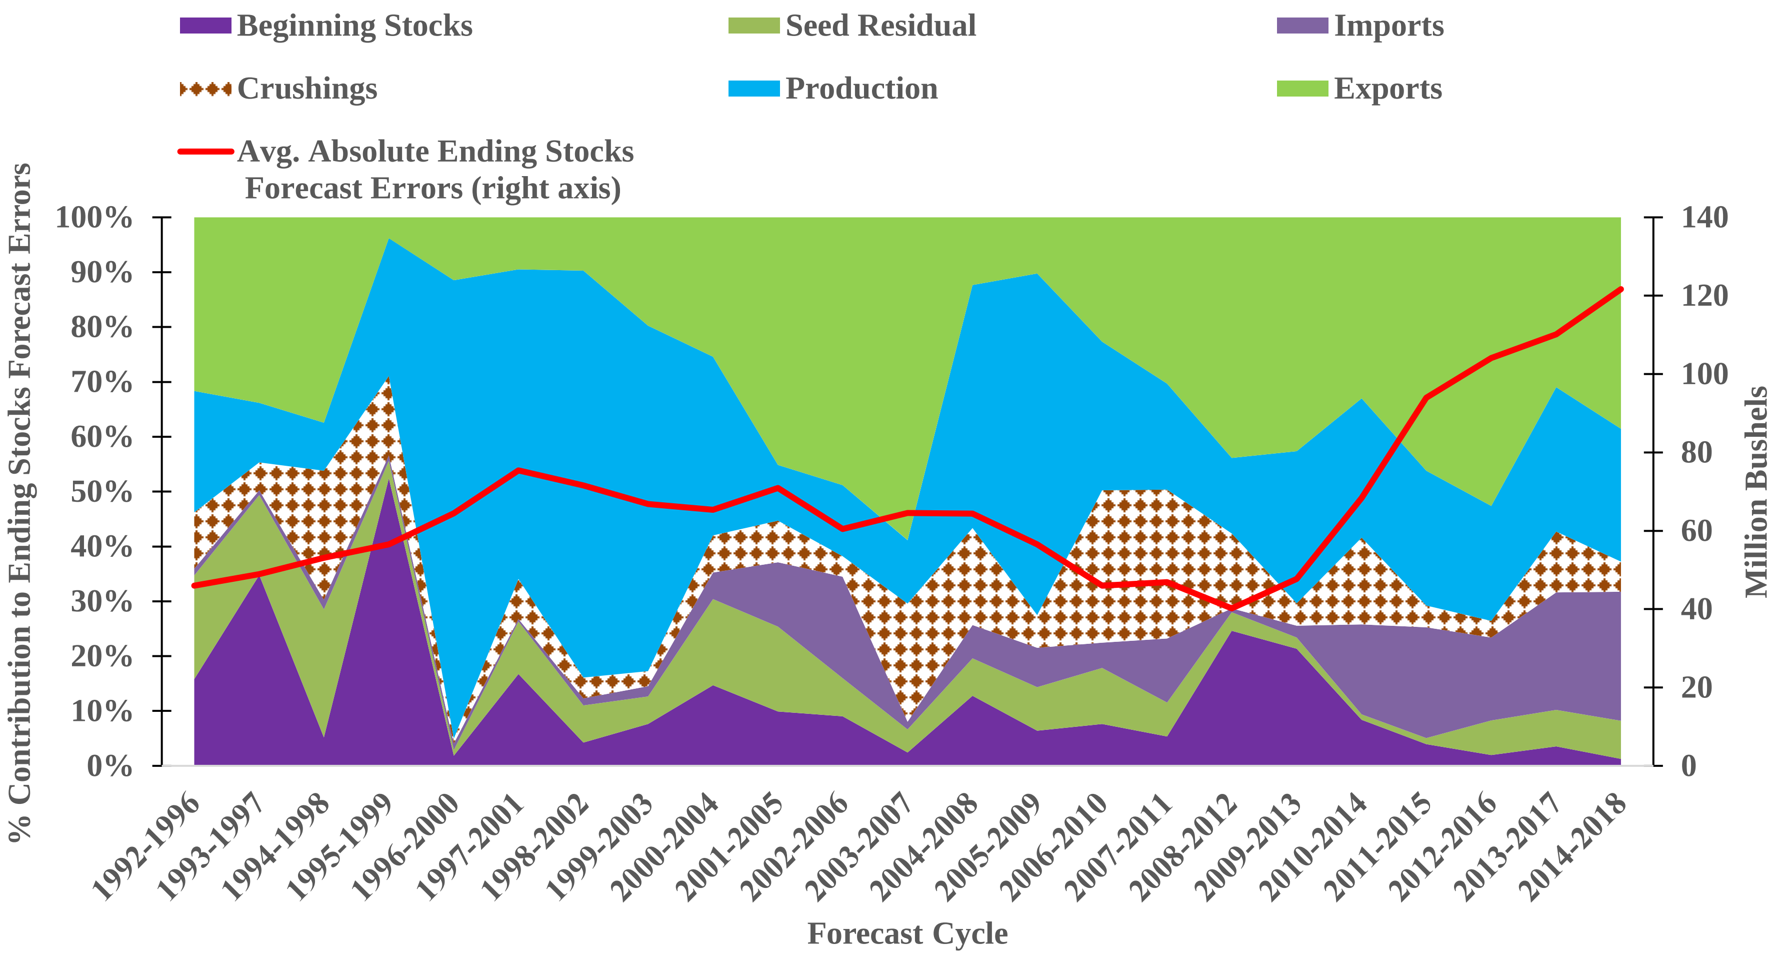

Figure A3.4: Contribution of balance sheet elements to USDA's Soybeans ending stock forecast errors

Note: The first year quoted on X-axis "1992-1996" represent the 5-year period from marketing year 1992/93 to 1996/97.

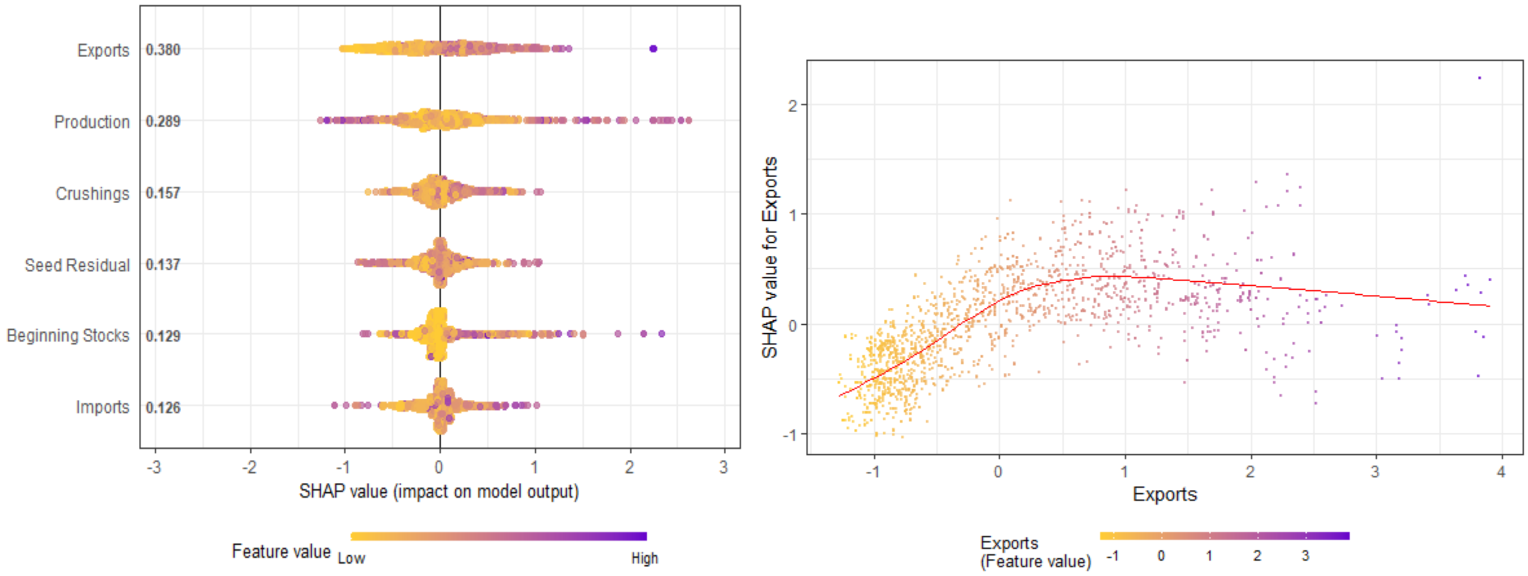

Figure A3.5: Soybeans – exports SHAP dependence plot

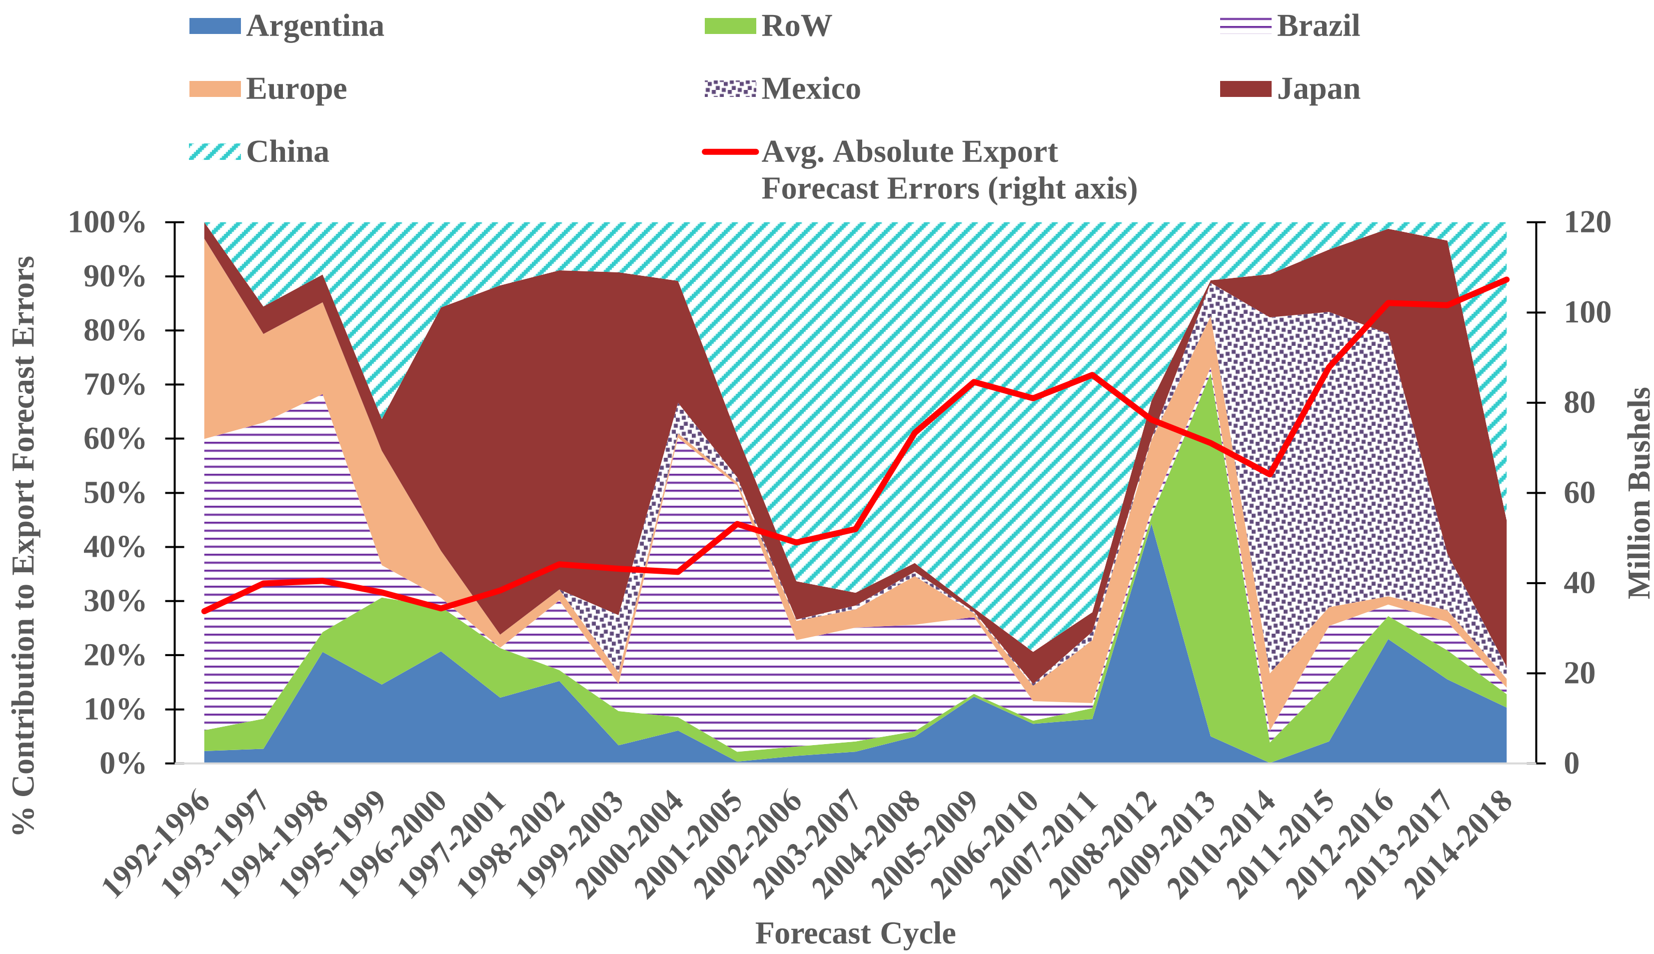

Figure A3.6: Contribution of world Soybeans import errors to U.S. Soybeans export projection errors

Note: The first year quoted on X-axis "1992-1996" represent the 5-year period from marketing year 1992/93 to 1996/97.

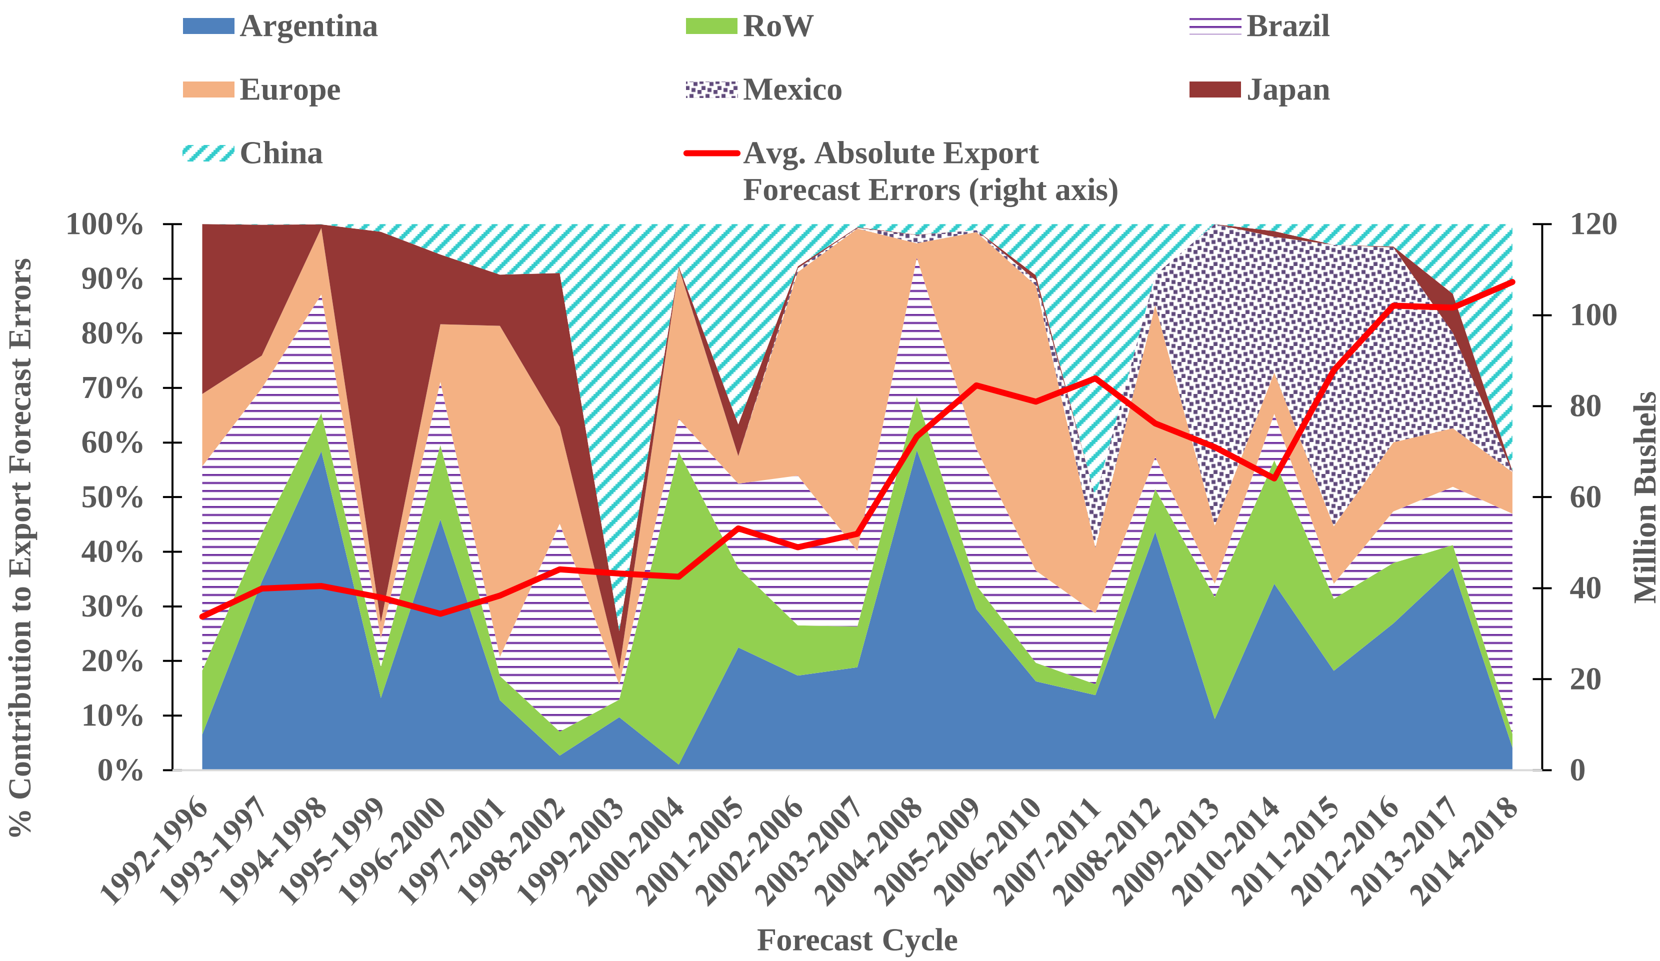

Figure A3.7: Contribution of world Soybeans production errors to U.S. Soybeans export projection errors

Note: The first year quoted on X-axis “1992-1996” represent the 5-year period from marketing year 1992/93 to 1996/97.

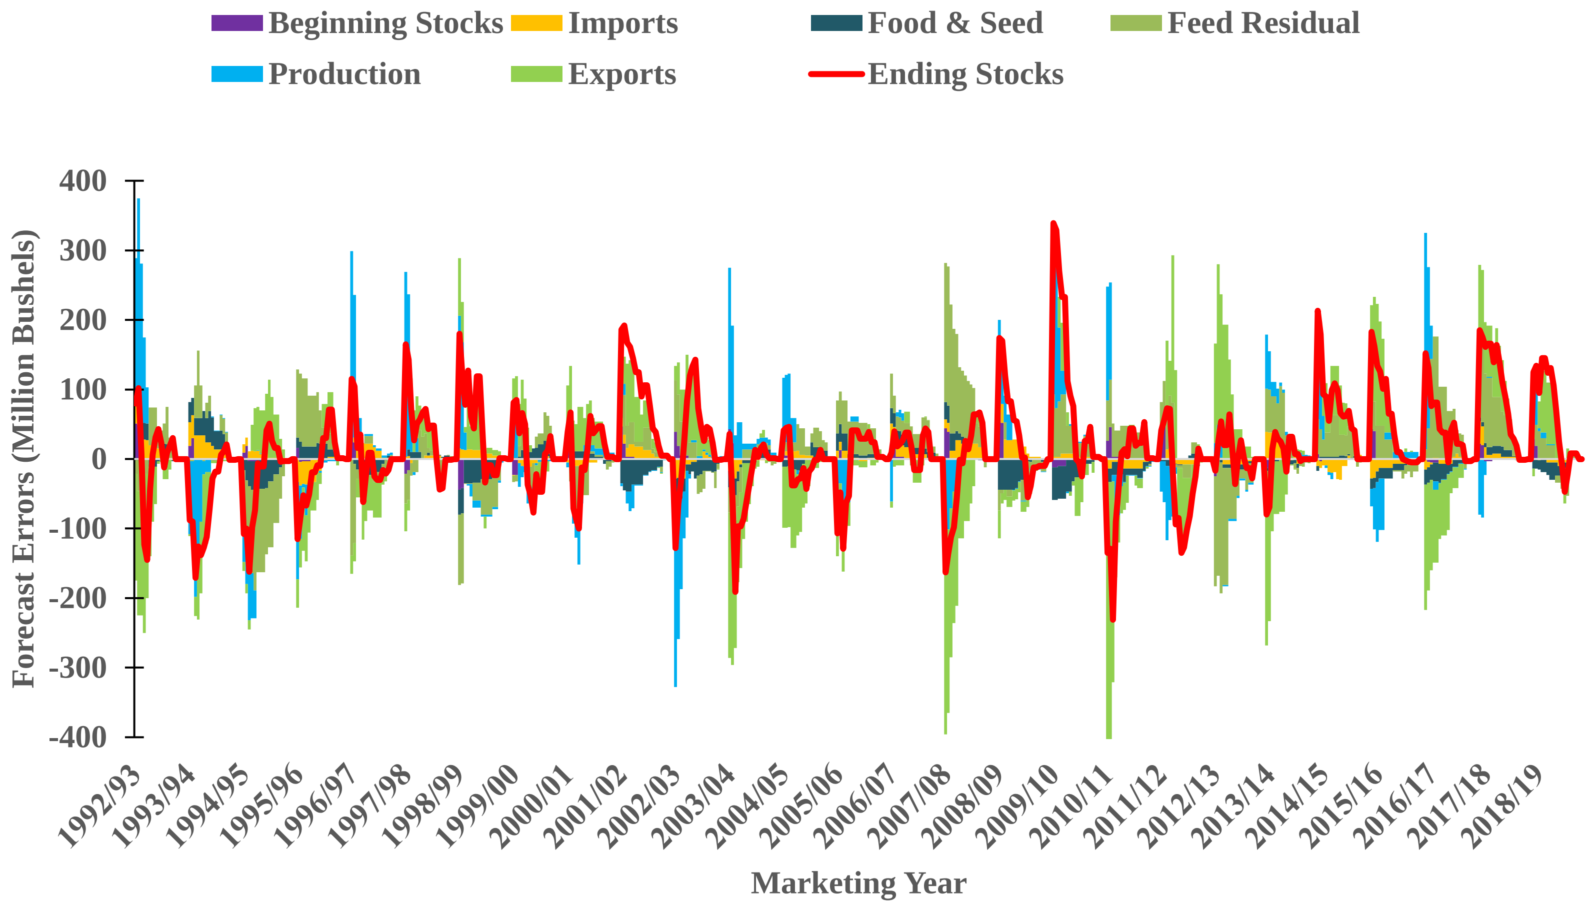

Figure A3.8: USDA Wheat balance sheet element forecast errors, 1992-2019

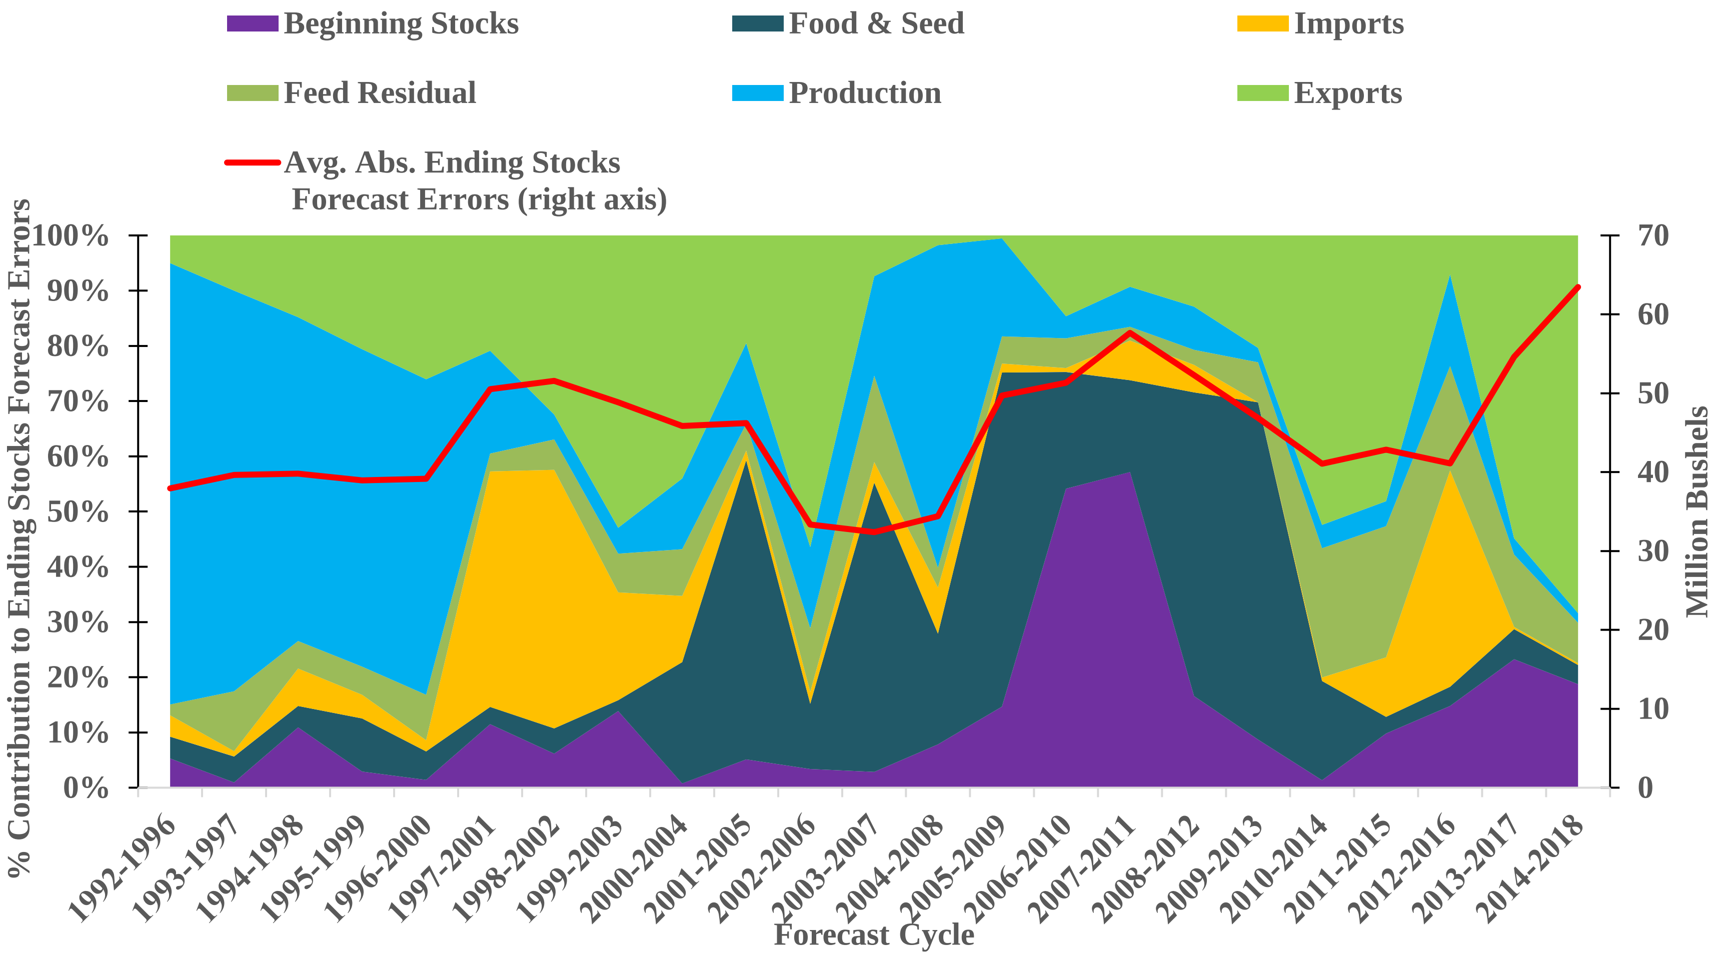

Figure A3.9: Contribution of balance sheet elements to USDA's Wheat ending stock forecast errors

Note: The first year quoted on X-axis “1992-1996” represent the 5-year period from marketing year 1992/93 to 1996/97.

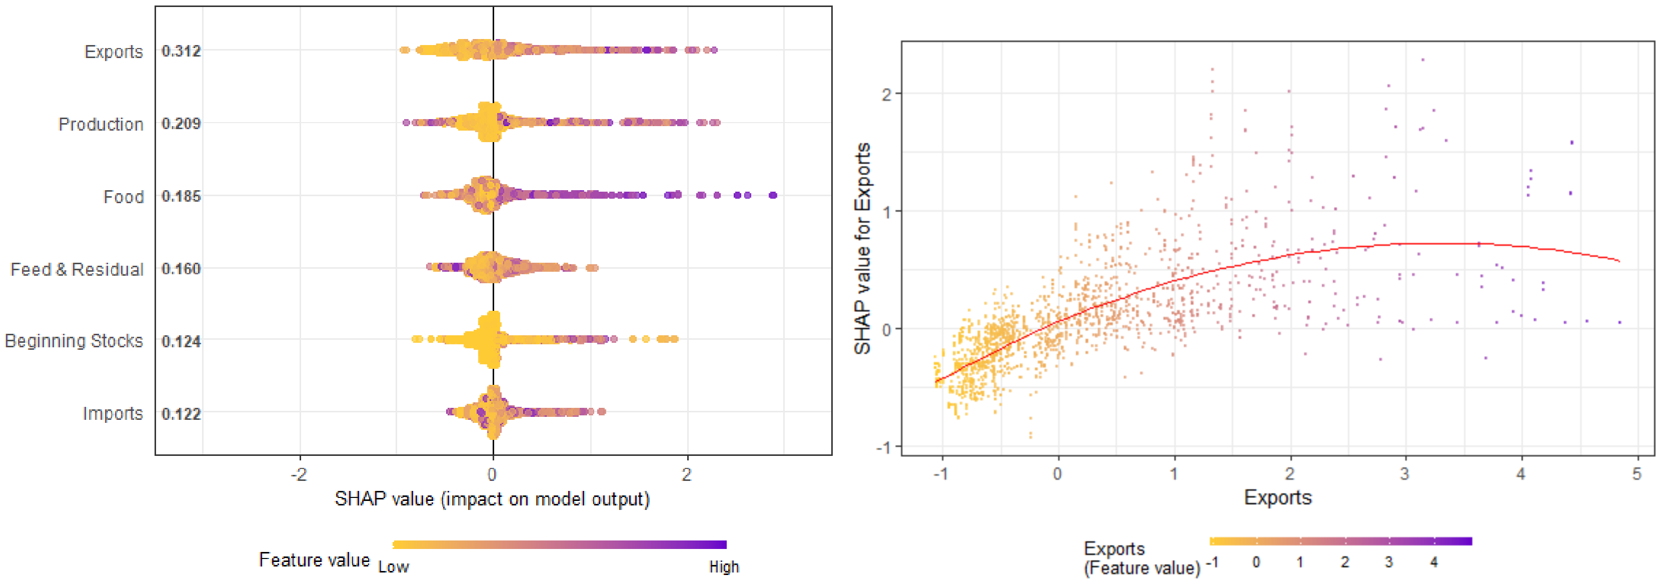

Figure A3.10: Wheat – exports SHAP dependence plot

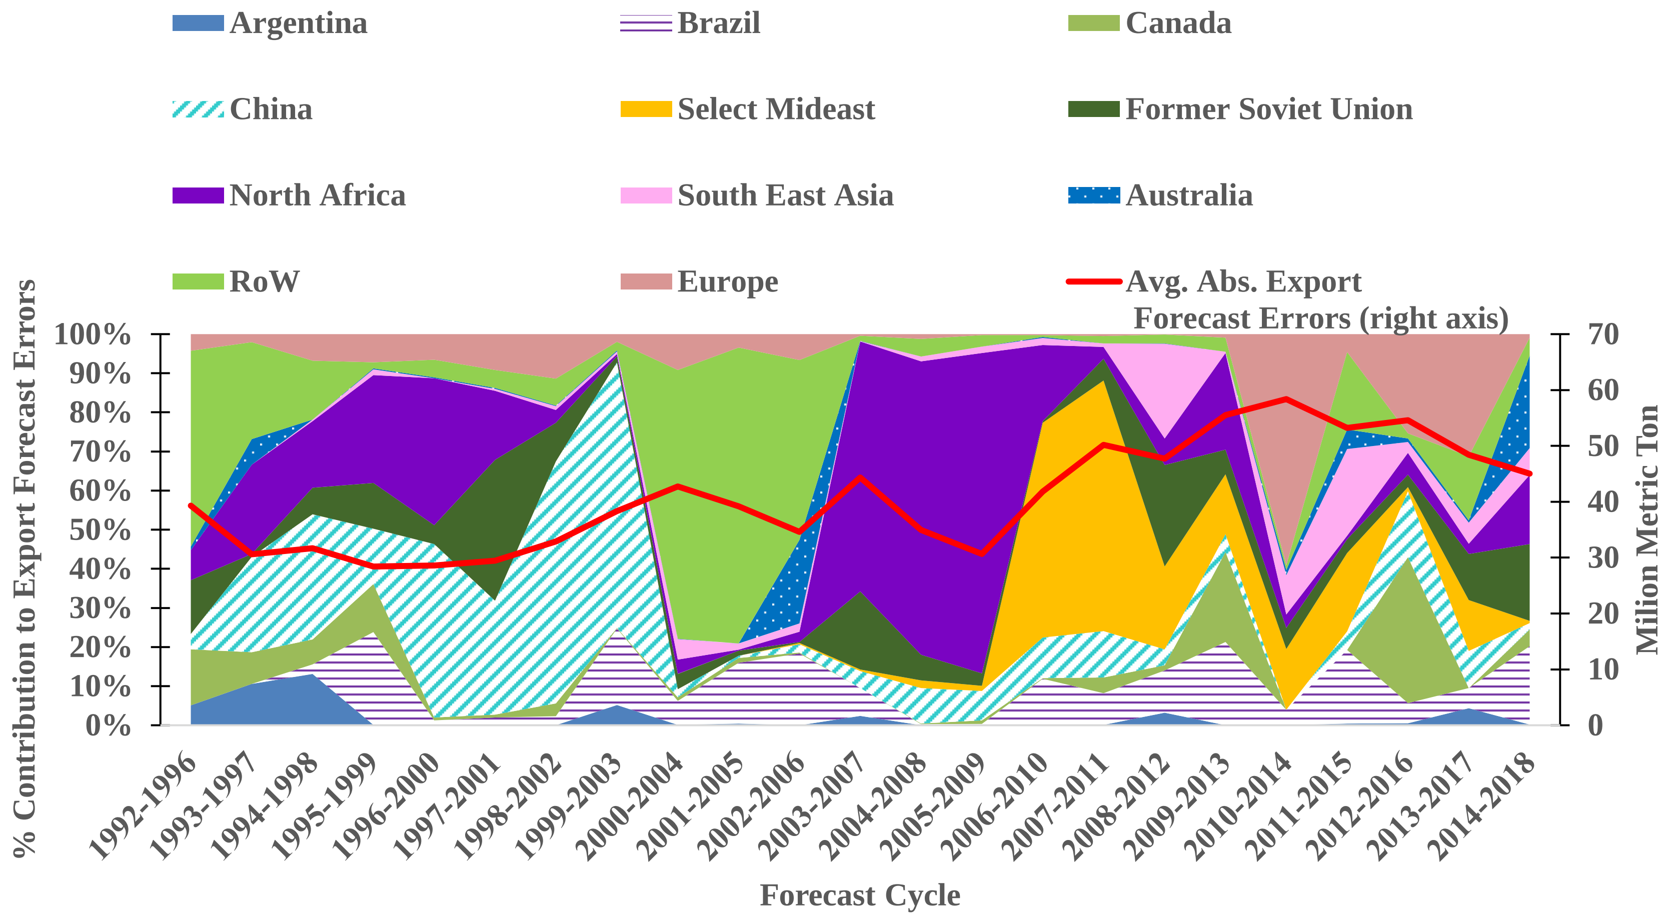

Figure A3.11: Contribution of world Wheat import errors to U.S. Wheat export projection errors

Note: The first year quoted on X-axis "1992-1996" represent the 5-year period from marketing year 1992/93 to 1996/97.

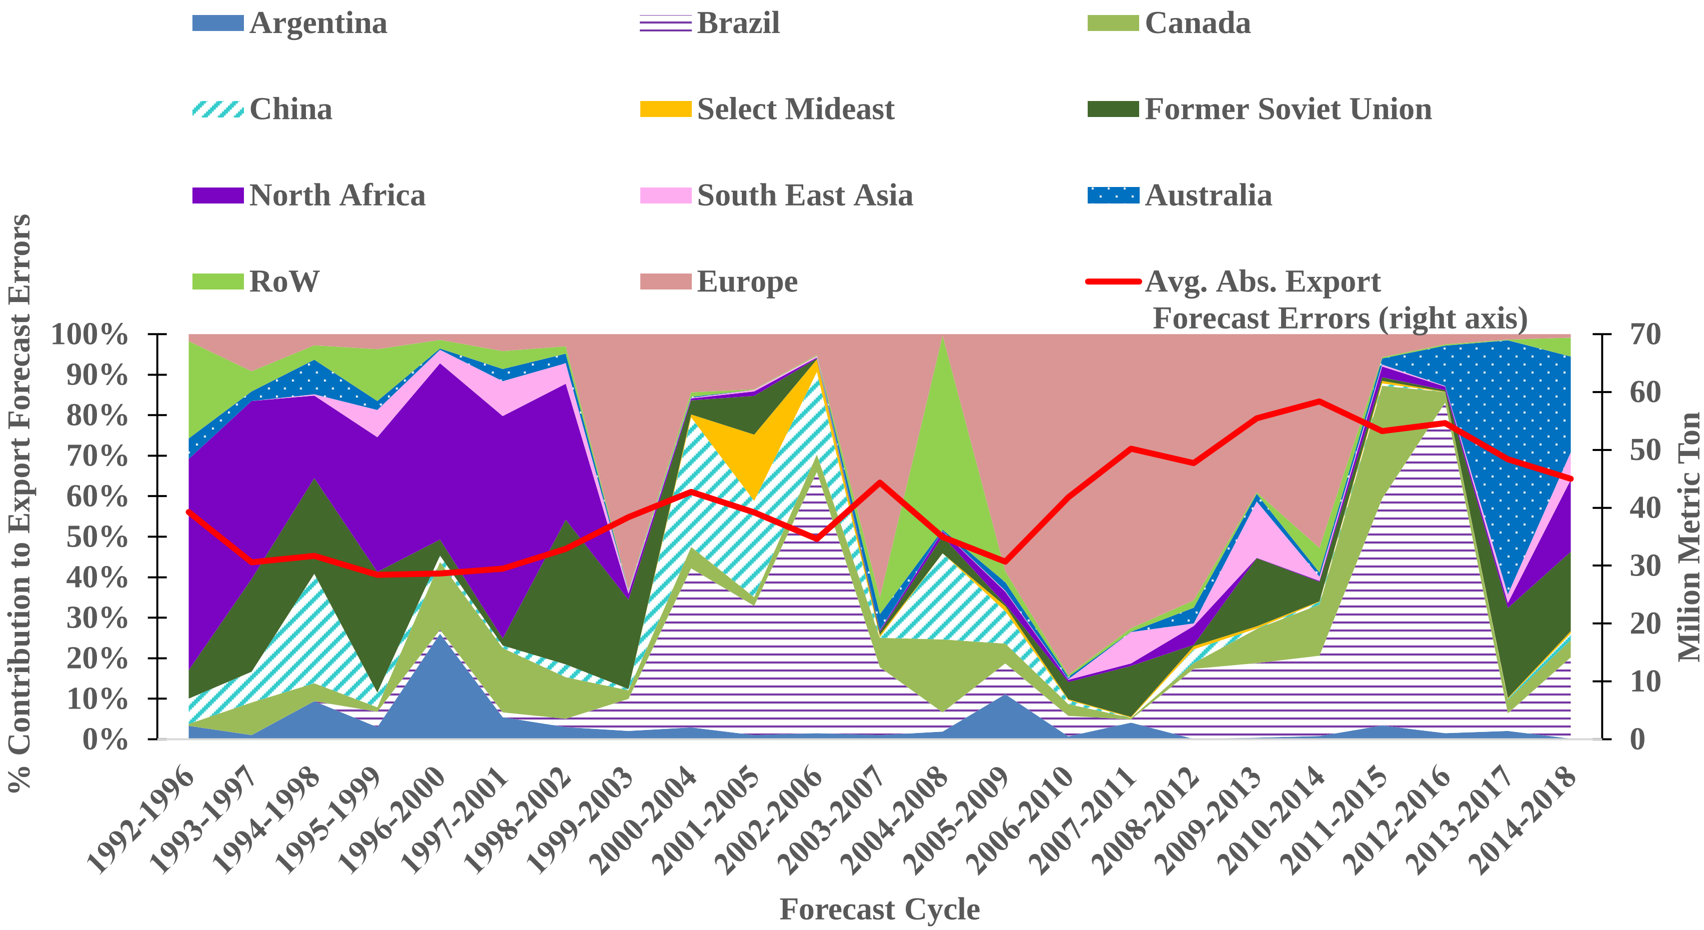

Figure A3.12: Contribution of world Wheat production errors to U.S. Wheat export projection errors

Note: The first year quoted on X-axis "1992-1996" represent the 5-year period from marketing year 1992/93 to 1996/97.

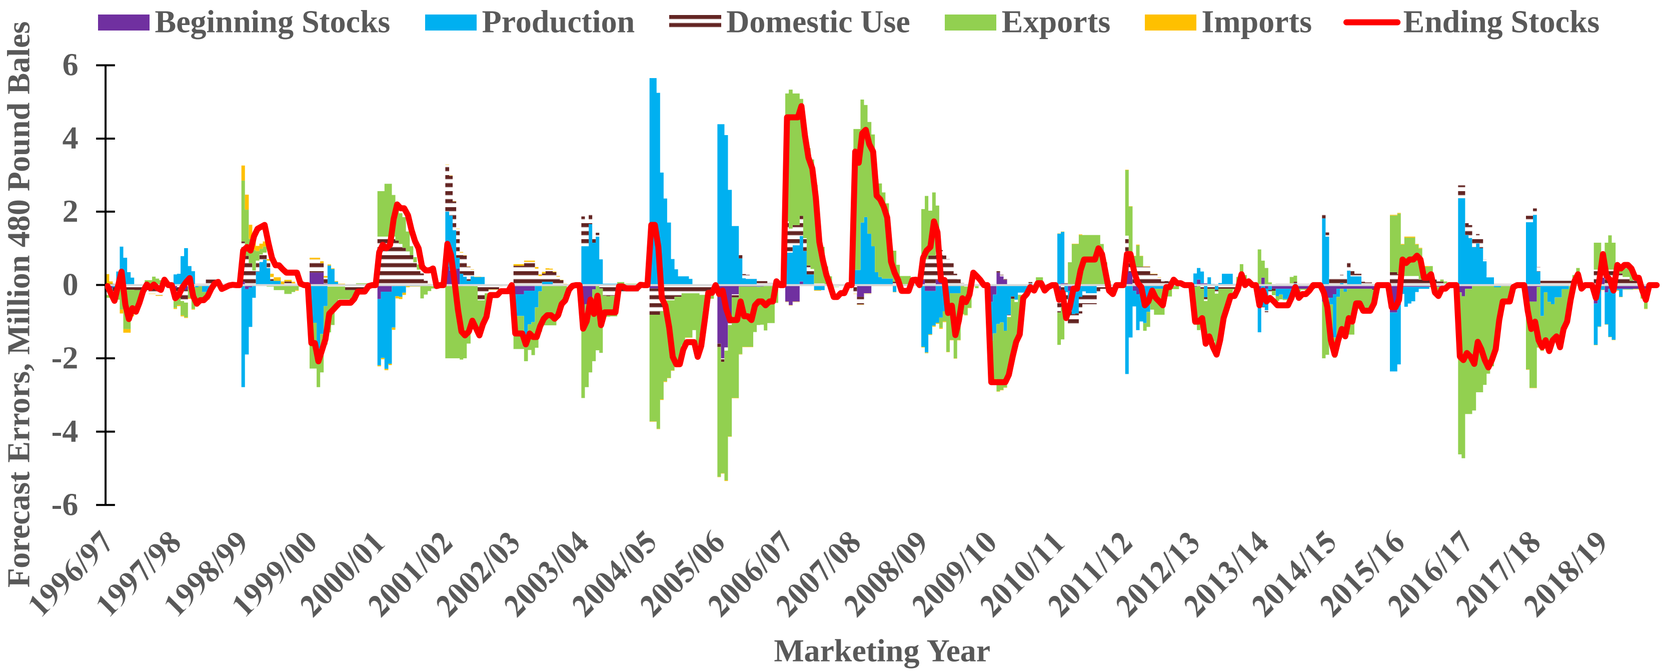

Figure A3.13: USDA Cotton balance sheet element forecast errors, 1996-2019

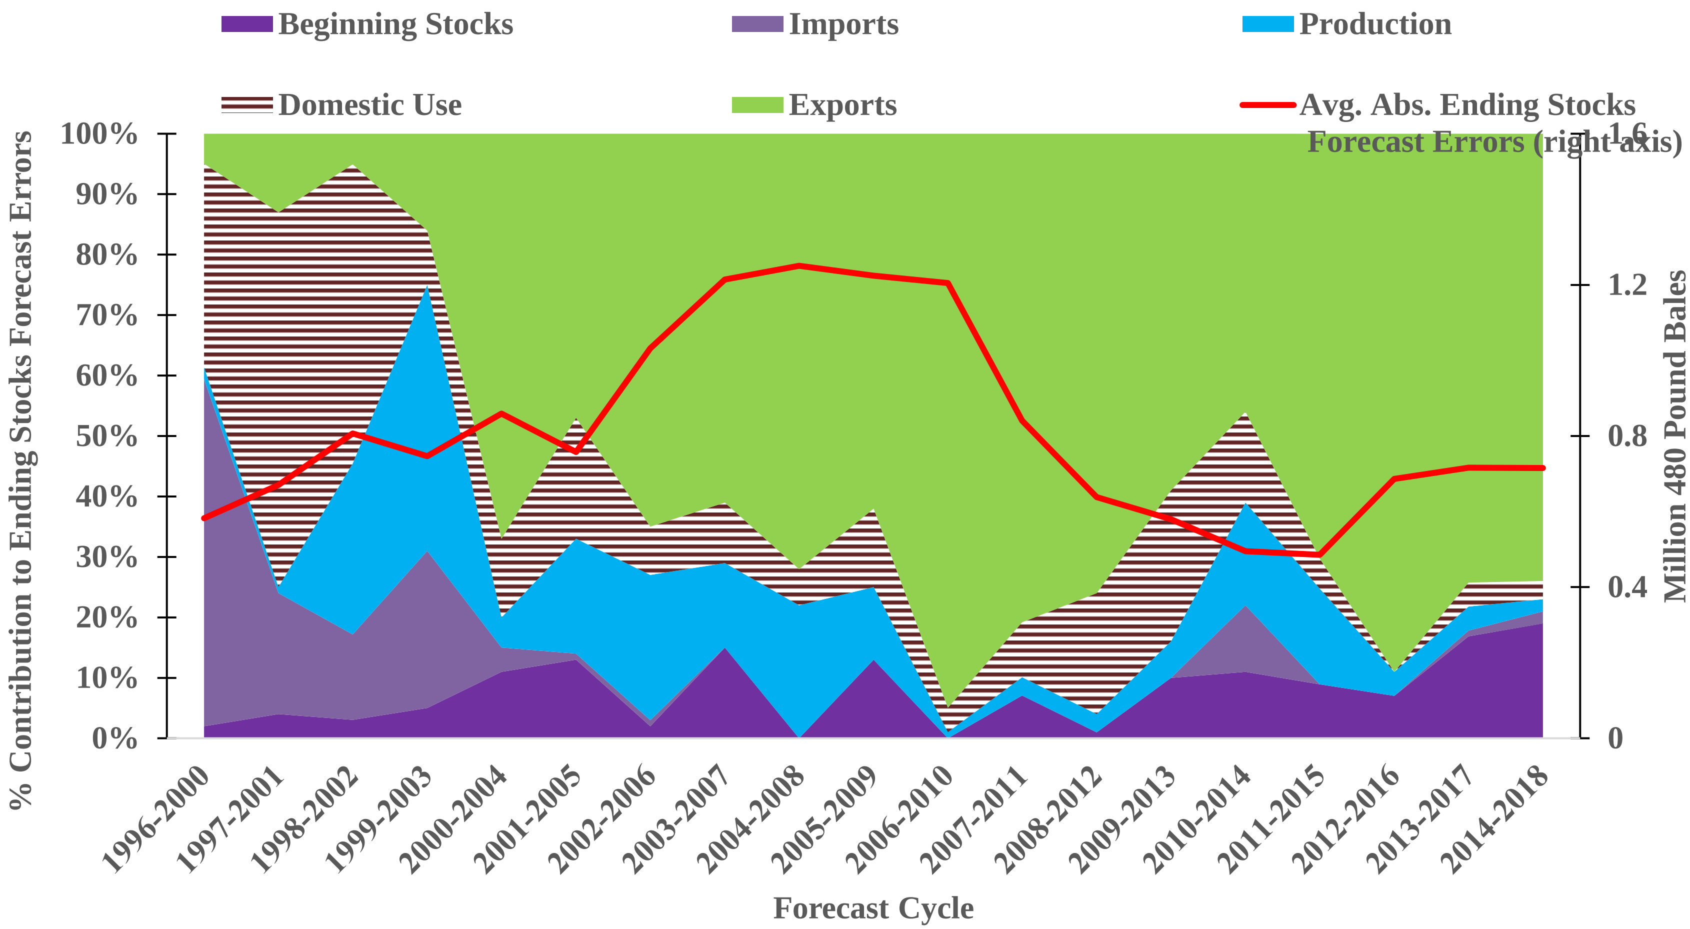

Figure A3.14: Contribution of balance sheet elements to USDA's Cotton ending stock forecast errors

Note: The first year quoted on X-axis “1992-1996” represent the 5-year period from marketing year 1996/97 to 2000/01.

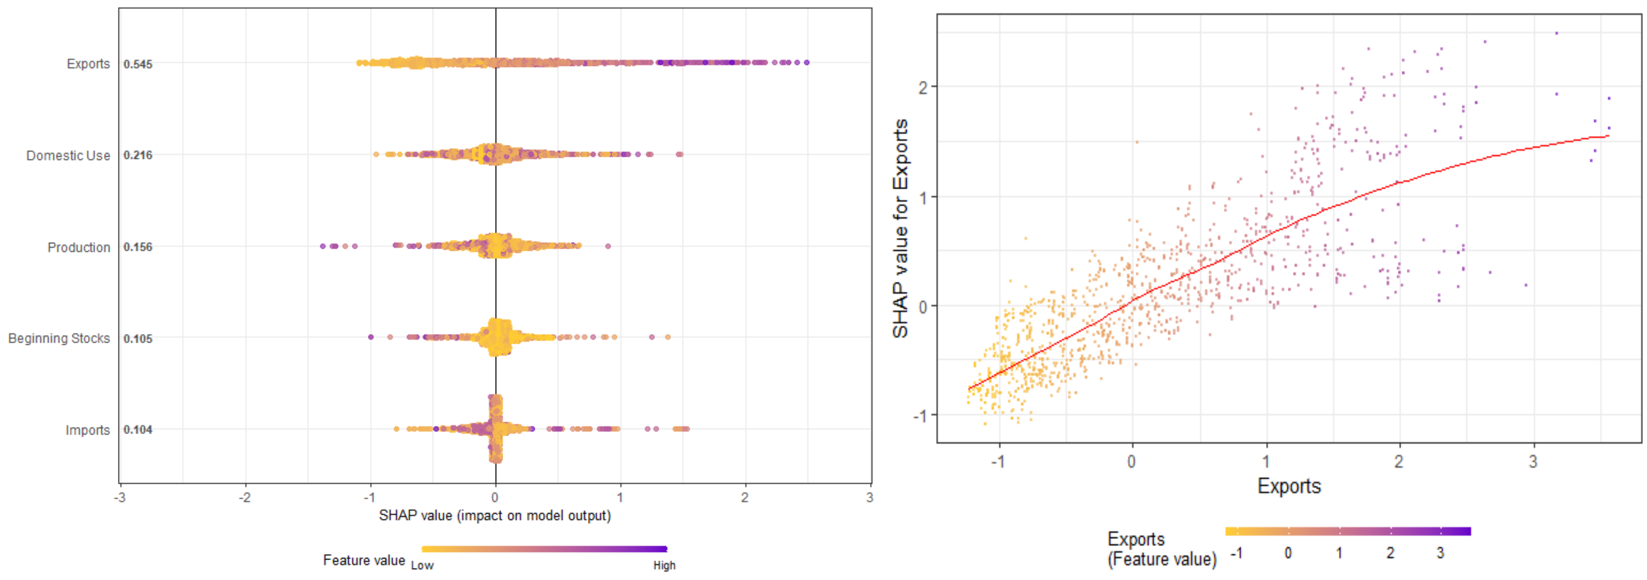

Figure A3.15: Cotton – exports SHAP dependence plot

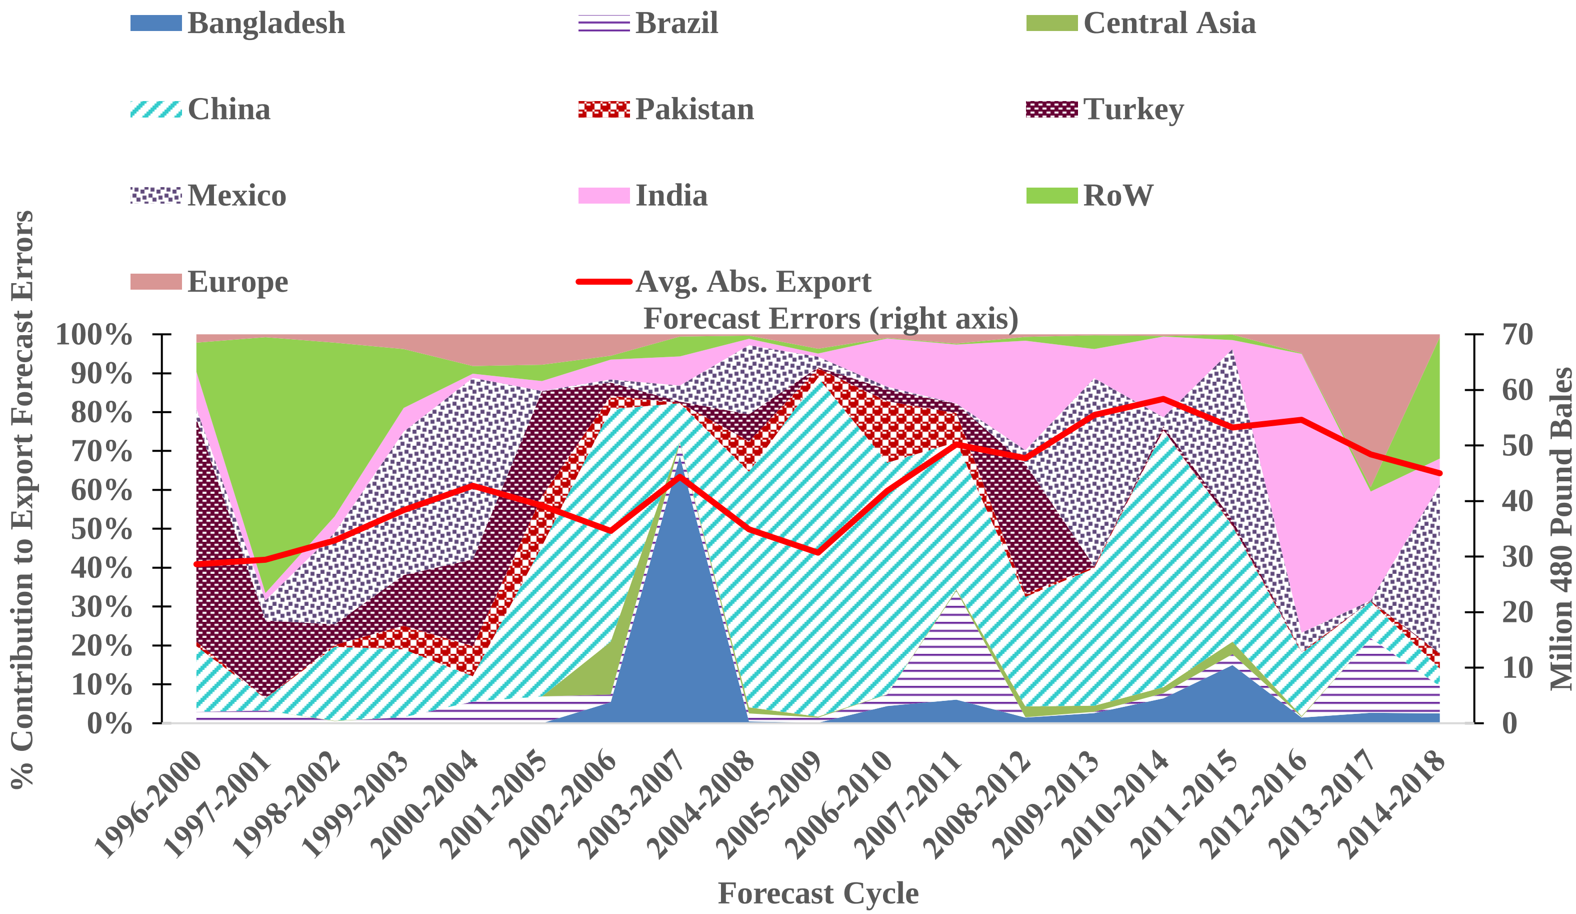

Figure A3.16: Contribution of Cotton import errors to U.S. cotton export projection errors

Note: The first year quoted on X-axis "1996-1997" represent the 5-year period from marketing year 1996/97 to 2000/01.

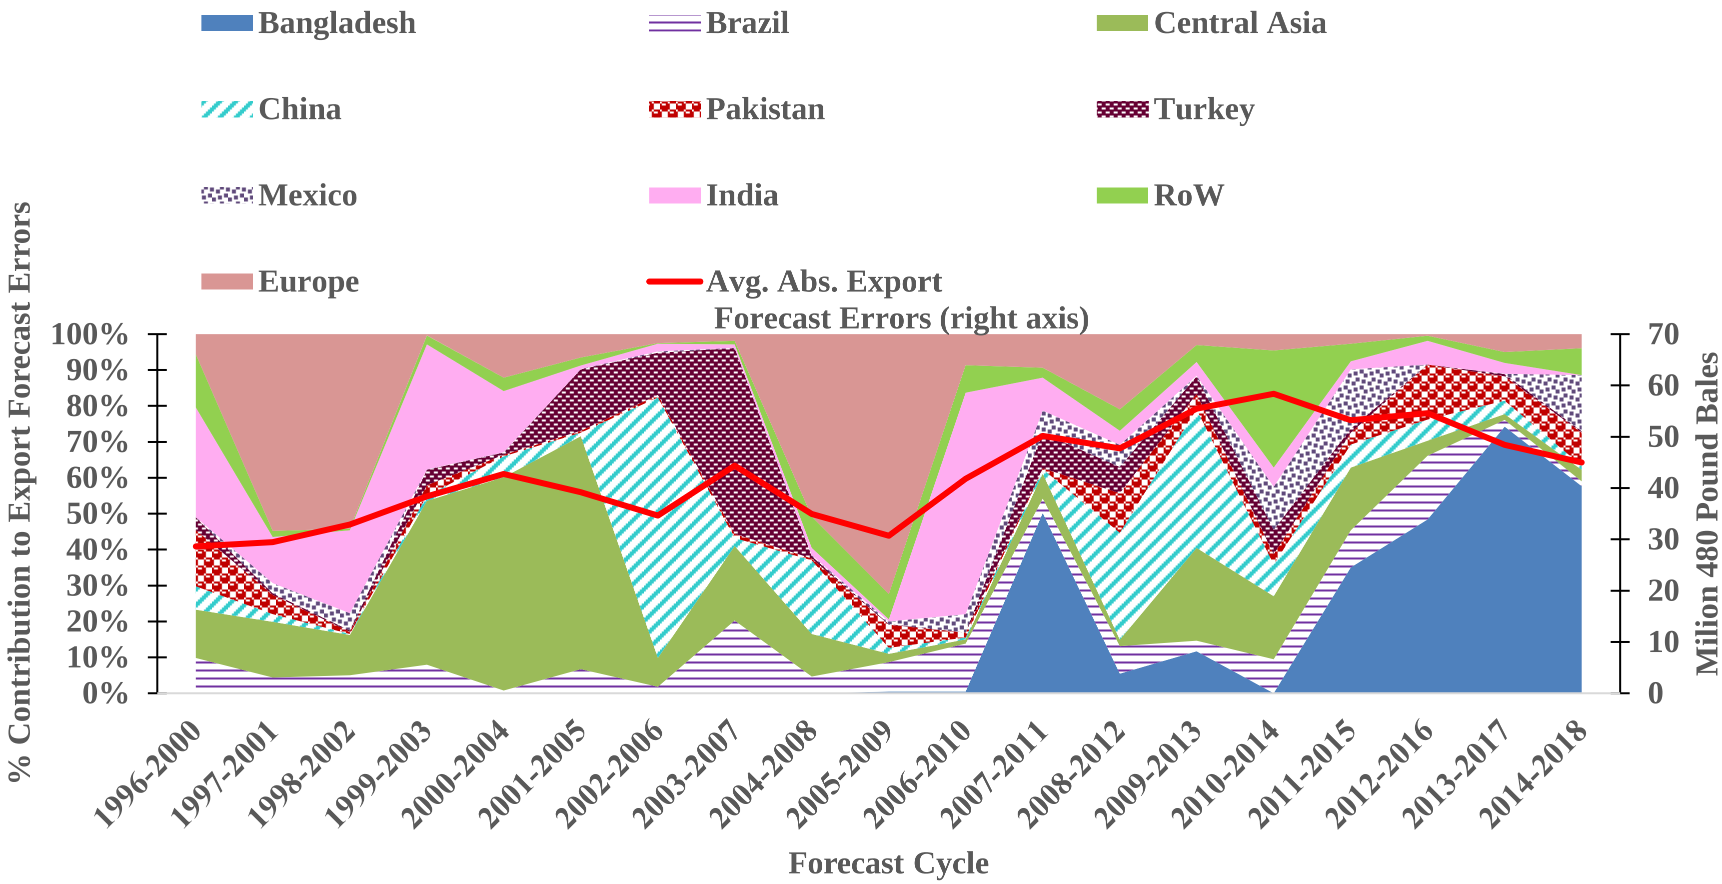

Figure A3.17: Contribution of world Cotton production errors to U.S. Cotton export projection errors

Note: The first year quoted on X-axis “1996-2000” represent the 5-year period from marketing year 1996/97 to 2000/01.

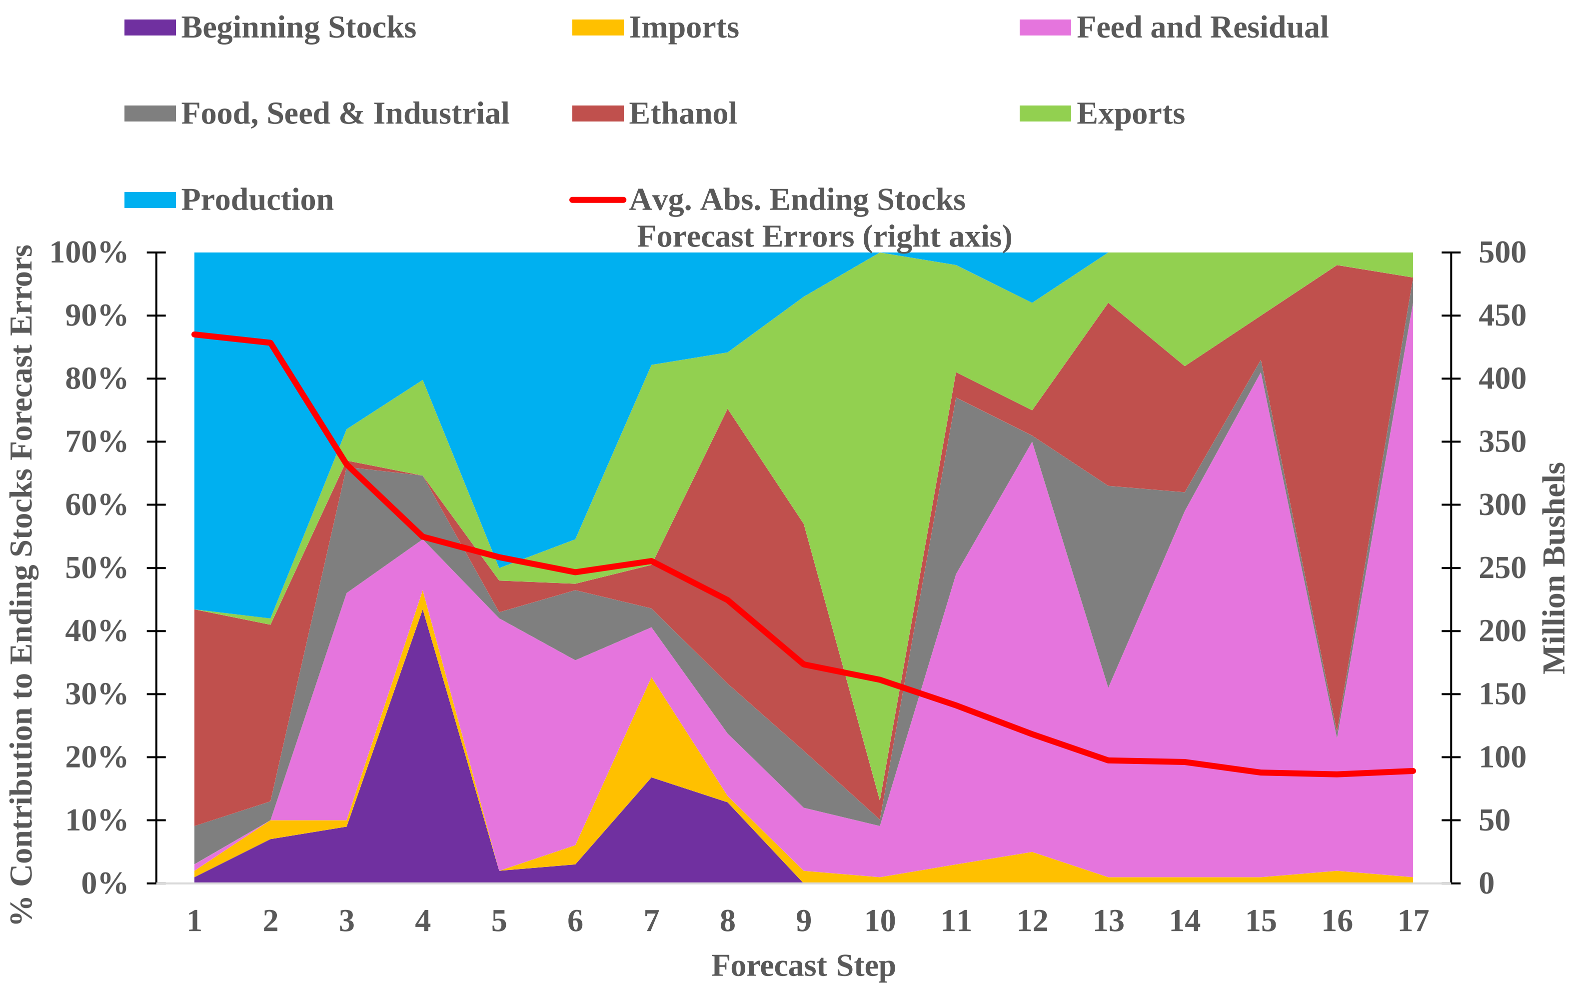

Figure A3.18: Contribution of balance sheet elements to USDA's Corn ending stock forecast errors

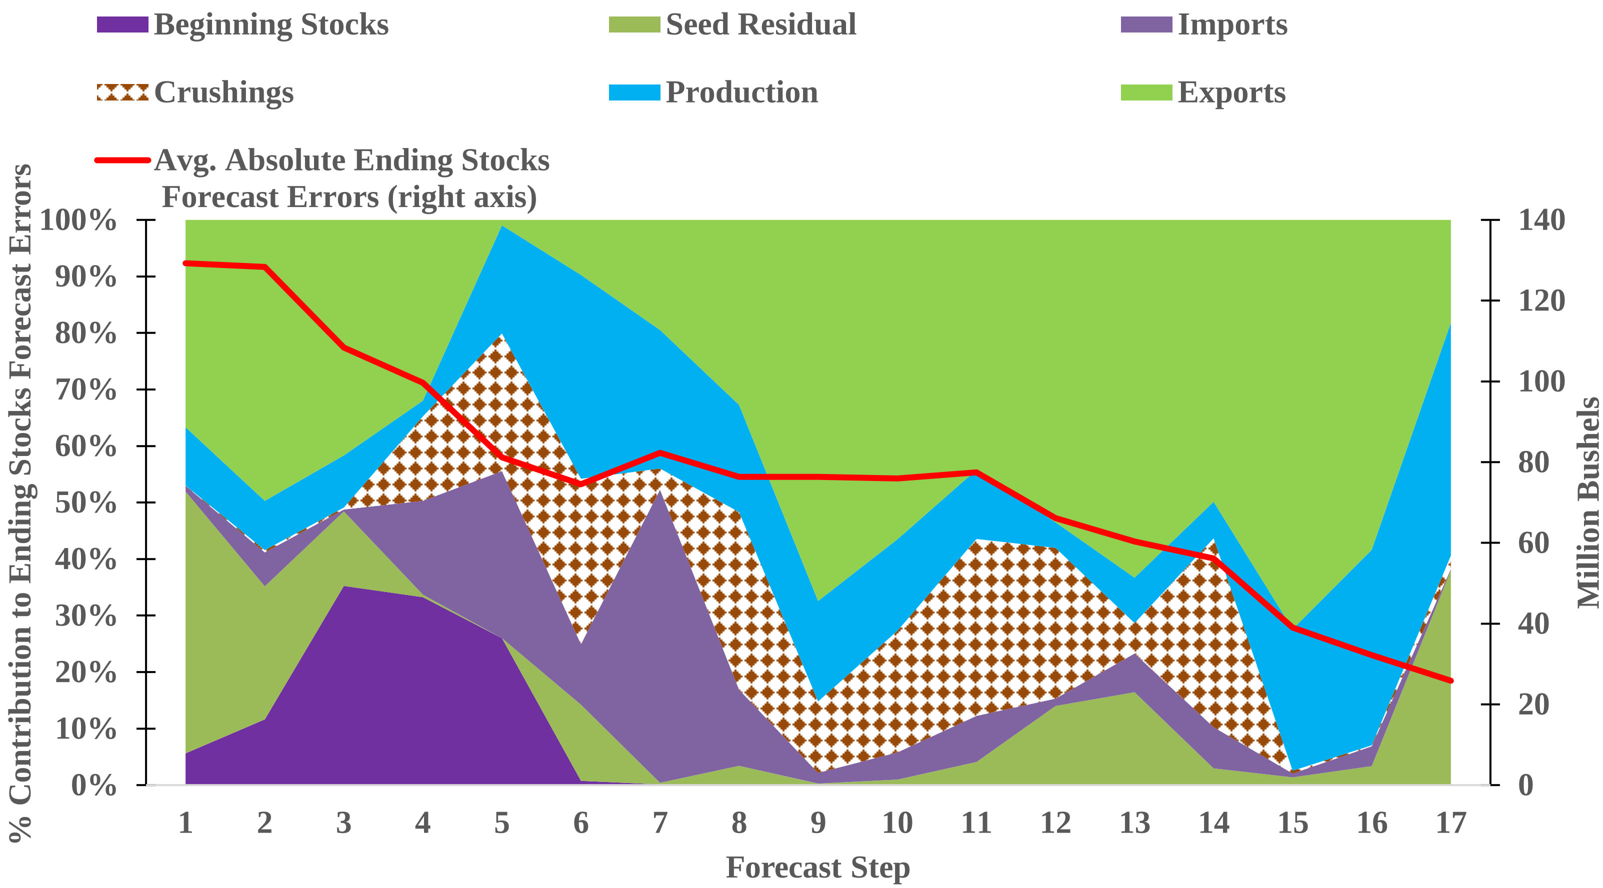

Figure A3.19: Contribution of balance sheet elements to USDA's Soybeans ending stock forecast errors

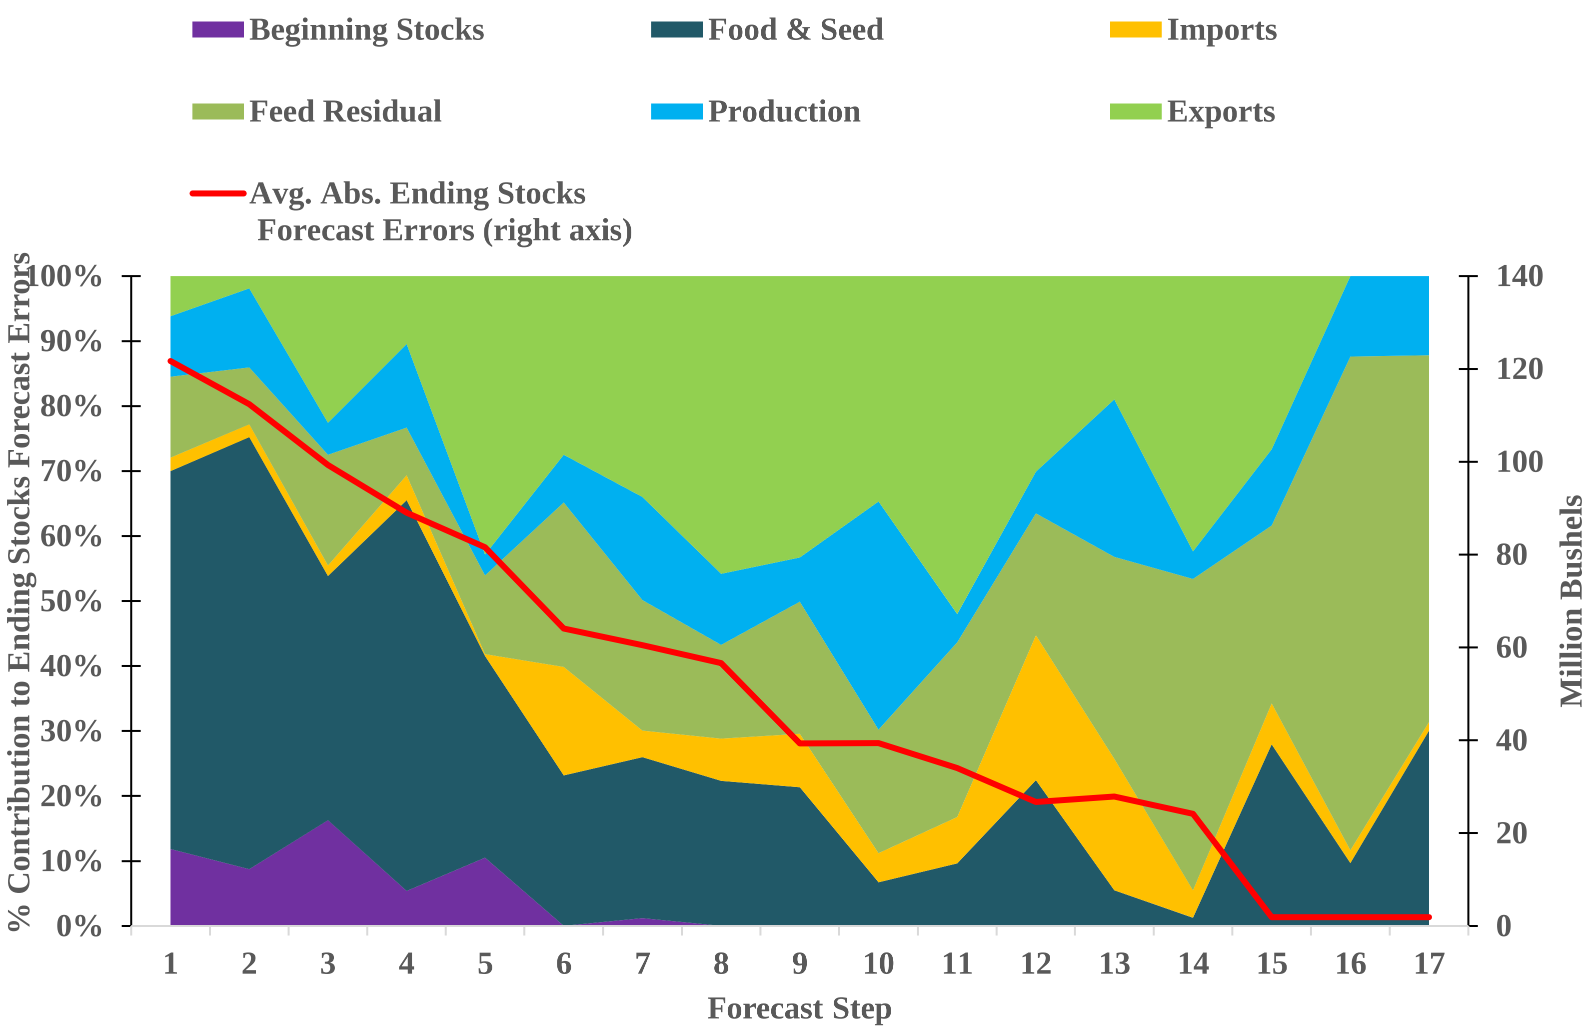

Figure A3.20: Contribution of balance sheet elements to USDA's Wheat ending stock forecast errors

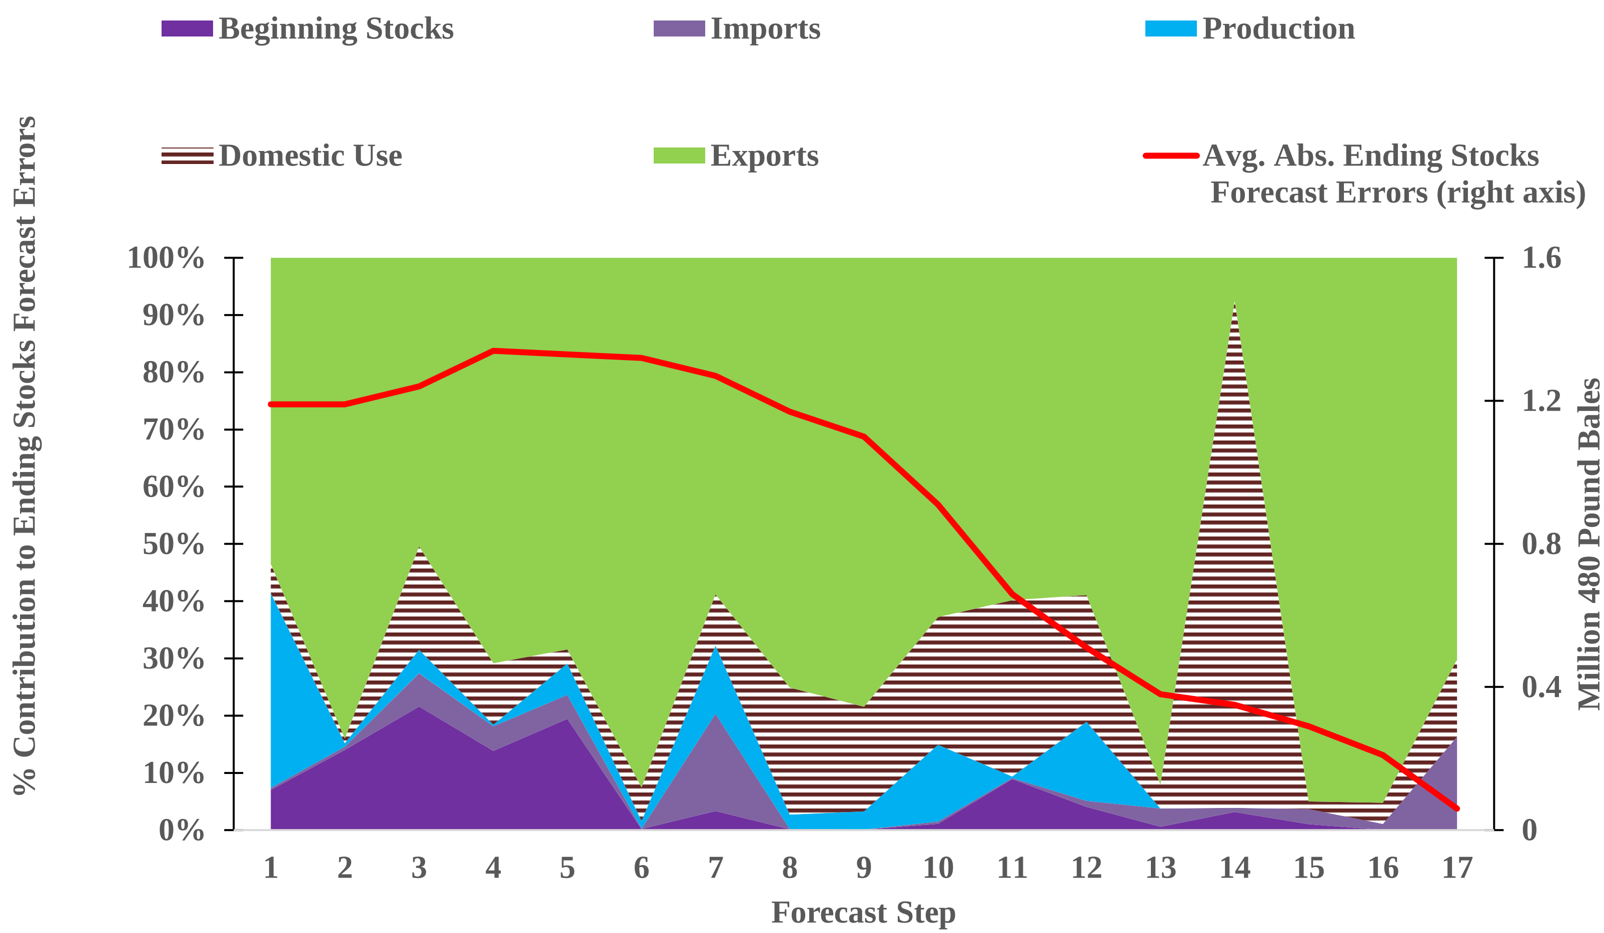

Figure A3.21: Contribution of balance sheet elements to USDA's Cotton ending stock forecast errors

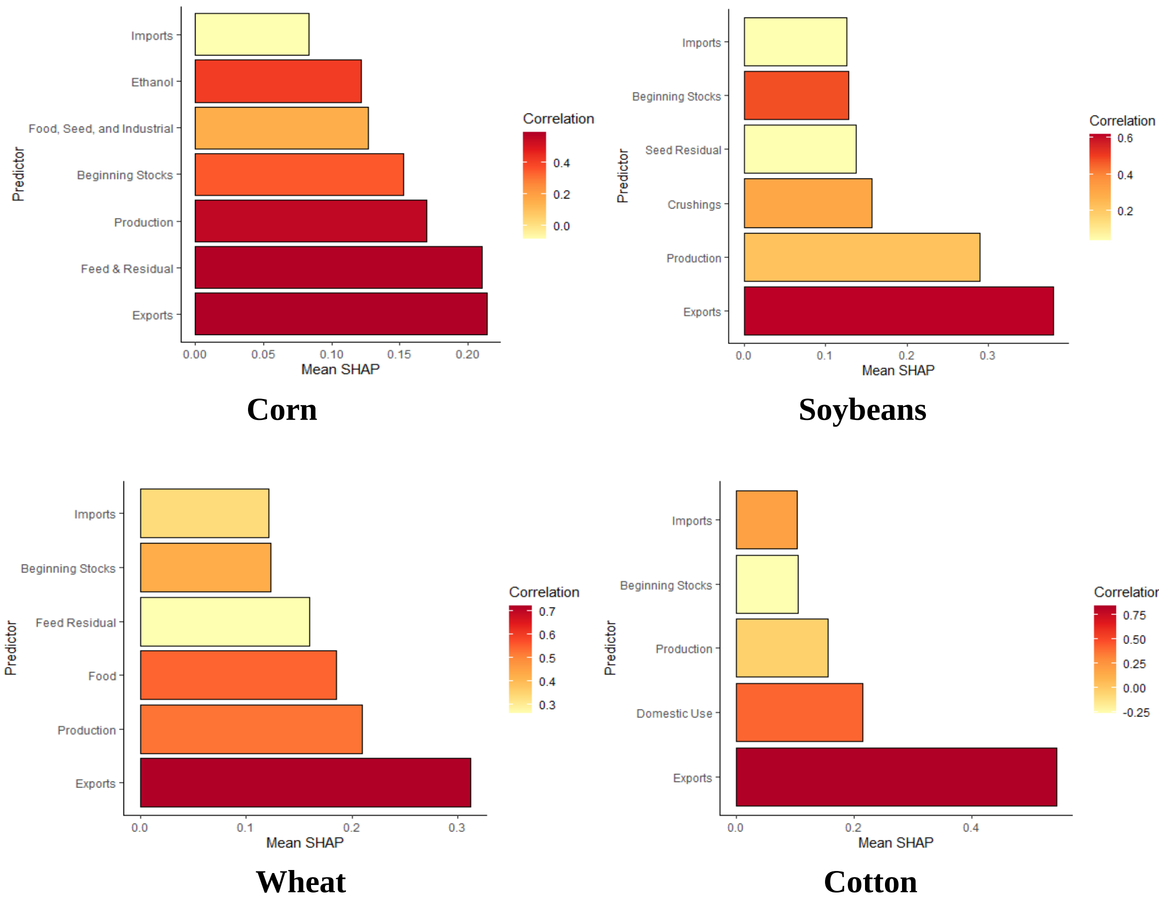

Figure A3.22: Top Predictors of Ending Stocks Errors

CHAPTER 4

ESTIMATING SUPPLY ELASTICITIES FOR CORN IN THE UNITED STATES: ACCOUNTING FOR PROSPECTIVE PLANTINGS[15]

[15] To be submitted to *AEPP*.

Abstract:

We propose adding USDA-reported prospective plantings to conventional acreage observations to estimate agricultural supply elasticities. Using past yield shocks as instruments to endogenous futures prices, we estimate a supply elasticity for corn at 0.26 percent (with a 95 percent confidence interval that ranges from 0.14 to 0.38 percent). Using three closely-related methods, we show that this elasticity *implies the COVID-19 pandemic generated welfare losses to domestic corn producers of about $5.4 billion. Through its Coronavirus Food Assistance Program (CFAP) the federal government compensated producers with payments of $ 6.9 billion.*

4.1. Introduction

Practitioners use commodity demand and supply elasticities to conduct agricultural policy analysis because policy makers and firms often request an understanding of the implications of sudden and unexpected supply (e.g., weather changes, natural disasters, pest infestations, and crop diseases) and demand shocks (recession, terrorist attacks, epidemics, and pandemics). In most years, commodity price variation is dominated by supply shifts. Consequently, Tomek (1979), Gray (1974), Chua and Tomek (2010), Adjemian and Smith (2012), and several other researchers have estimated demand elasticities and (their related) price flexibilities for agricultural commodities. Econometricians, however, need demand-induced price variations to study the shape of the commodity supply curve. Hendricks et al. (2014) show that estimating supply elasticities for agricultural commodities is complicated by the fact that demand shifters are difficult to observe, and their explanatory power with respect to price variation is limited, relative to supply shocks like weather events.

Recent volatility in commodity prices invited the attention of economists (Kim and Moschini, 2018; Carter, Rausser, and Smith, 2016; Roberts and Schlenker, 2013; Mallory, Hayes, and Babcock, 2011), who link the increases in commodity prices to the adoption of the Renewable Fuel Standard (RFS). RFS was established under the 2005 Energy Policy Act and expanded under the 2007 Energy Independence and Security Act. Post RFS adoption, ethanol consumption increased from 0.5% of total oil consumption in 2008 to 5.37% in 2018 (Schmitz, Moss, and Schmitz, 2020). Moreover, ethanol production accounts for 40% of total corn usage (2021) in the United States (USDA, ERS).

The onset of the COVID-19 pandemic and the ensuing lockdowns and closure of non-essential business reduced the demand for fuel (and therefore the corn used to generate ethanol). From March through November 2020, ethanol production declined by 2 billion gallons (a year-on-year decrease of 2% from 2018 to 2019 and 12% from 2019 to 2020), decreasing corn usage by 700 million bushels (RFA, 2020). Therefore, the COVID-19 pandemic is a large-scale demand shock, raising the question of its impact on farmers' land allocation decisions. This becomes even more important to study because the United States Department of Agriculture (USDA) launched the Coronavirus Food Assistance Program (CFAP) to compensate farmers based on their expected losses.

Roberts and Schlenker (2013) measure supply elasticities by instrumenting prices with the prior-year yield shock. We apply this method to a data set that includes the covid-19 event to estimate and compare supply elasticities used by USDA for policy analysis. Every year USDA releases (1) the Prospective Plantings report containing farmers' planting intentions and (2) the

Actual Acreage report. The Prospective Plantings report helps set production expectations and guide market decisions for the coming year. To this point, this information has been left out in the literature when assessing acreage response to price fluctuations. We use the data on prospective plantings, actual plantings, and futures prices to study farmers' responses to price shocks. We also connect this work to policy by using our elasticities to estimate the welfare damage generated by the COVID shock and comparing this to the payments USDA made to domestic corn and soybean producers under CFAP. To do so, we estimate the implied shift in the quantity via three related methods during the year 2020: (1) using the observed price difference from January to April (used by USDA in their calculations), (2) the observed price difference from January to August (harvest month for corn), and (3) the observed price difference from January to November (one month before contract expiration). We find that the supply elasticity for corn is 0.26% (estimates in the literature range between 0.40% and 0.50%). Under all three scenarios, our analysis shows that the USDA was able to entirely compensate farmers for their losses because of the COVID-19 pandemic.

We expand on the literature by adding information on farmers' intended plantings to estimate agricultural supply elasticities using the data on farmers' intended plantings and actual plantings. This offers us more data points, increasing statistical power and improving the precision of our estimates. Ultimately, our analysis will help government and private organizations form effective strategies when facing future demand shocks. In addition, the information on farmers' responses to price changes will assist traders and market participants inform their production and risk management decisions, improving the efficiency of the agricultural supply chain.

4.2. Background and Literature Review

Supply Elasticities

Nerlove (1958), in a seminal work, studies agricultural supply elasticity analysis using lagged prices as control variables. Roberts and Schlenker (2013) estimate demand and supply elasticities for corn, soybeans, wheat, and rice by aggregating these commodities into calories. The authors construct weather-induced yield shocks as deviations from yield trends and use them as instruments to estimate supply elasticities. Their elasticity estimate lies between 0.08% and 0.1%. Further, they use these results to assess the US biofuel mandate's impact on food prices between 2005 and 2008, showing that the mandate increased food prices by 30% (~ loss in consumer surplus of 180 billion). On the contrary, Hendricks et al. (2014a) propose using OLS to estimate supply elasticities by regressing planting area on futures prices with the current-year realized yield shock as a control variable. They decompose supply elasticities into changes in deviations from yield trend, the composition of growing area, and total growing area to show that adding the current year yield shock eliminates any sources of endogeneity, making OLS a suitable model choice.

Hendricks et al. (2014b) use satellite data on land use in three corn belt states, Illinois, Indiana, and Iowa, to estimate short and long-run elasticities for corn and soybeans. With roughly over 8 million observations from 2000 to 2010, the authors use first-order Markov transition probabilities and find an elasticity estimate of 0.29 for corn and 0.26 for soybean. Kim and Moschini (2018) argue that price changes from 2005 to 2015 were due to demand factors following RFS adoption. Using price changes (in this period) as exogenous variables, they estimate corn and soybean supply elasticities for 12 midwestern states from 2005 to 2015. The authors report own-price elasticities of 0.38 for soybeans and 0.50 for corn.

In a recent article, Thompson et al. (2021) study how agricultural supply responds to extreme events such as the global price surge of 2005-07, the African swine fever, the trade war, and the COVID-19 pandemic. These authors evaluate how elasticities from the literature perform under these demand shocks for corn and soybeans. Specifically, they consider the performance of elasticities estimated by Goodwin and Mishra (2006), Hendricks et al. (2014a), and Kim and Moschini (2018). The authors find that the price changes were associated with small area effects; for example, price changes as high as 50% resulted in area changes as low as 3%. They also show that cross-price elasticities are crucial in determining area response to price changes.

Thompson et al. (2021), however, do not estimate supply elasticities; they only assess the values measured in the literature. This article fills this gap by estimating supply elasticities for corn and soybeans using recently advanced techniques and methods. Furthermore, although most supply work relies on a single, annual observation of price and plantings, we expand the dataset by accessing one additional source of data: information on planting expectations that the USDA releases in its Prospective Plantings report in March before planting is complete.

Coronavirus Food Assistance Program

The United States government announced Coronavirus Food Assistance Program in April 2020. USDA allocated $16 billion to assist farmers and ranchers based on pandemic-associated losses. Producers who have suffered a price decline of 5% or greater were eligible for the program. Applications were accepted from May 26, 2020, to September 11, 2020. Later USDA offered a second round of payments (CFAP2) and received applications from April 2021 through October

2021. USDA made a total payment of $11.8 billion under CFAP1, and $19.2 billion under CFAP2, summing to $31 billion. Figure 4.1 shows the total CFAP payments made for corn (~$6.86 billion) across different states. As expected, disbursements are higher in corn belt states (Iowa, Illinois, Nebraska) compared to other regions.

Figure 4.1: Total CFAP Payments for Corn

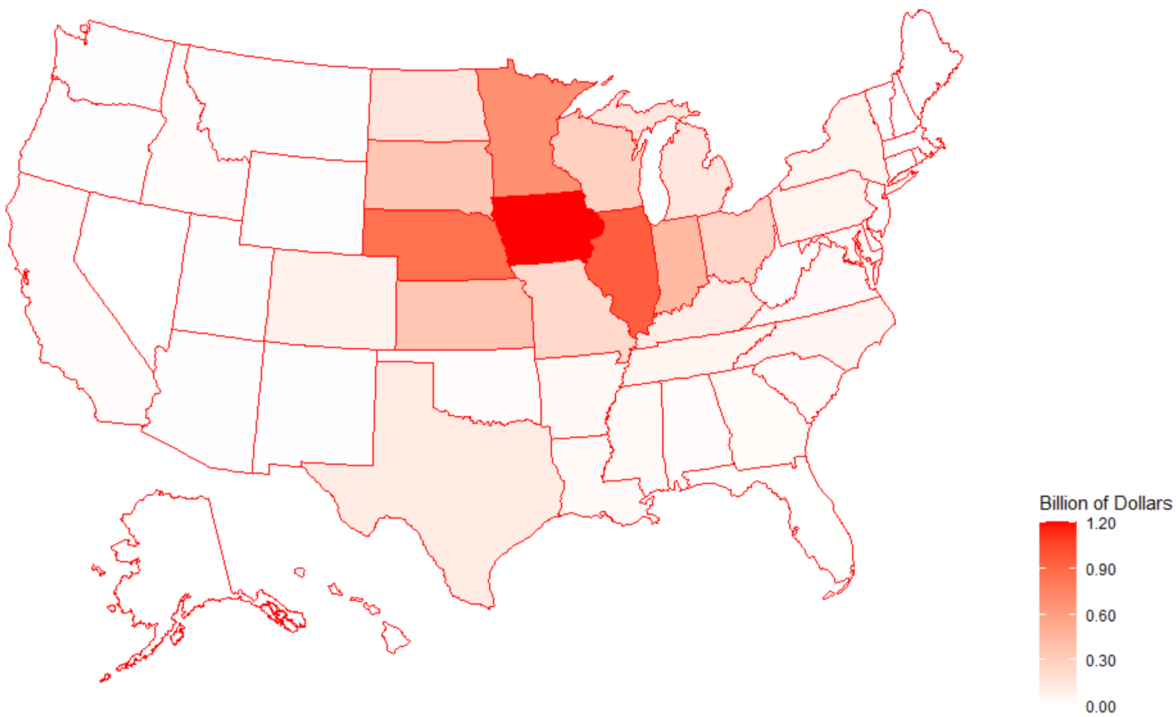

Source: USDA Data

4.3. Data

End of March every year, USDA releases the Prospective Plantings report containing information on farmers' planting intentions for that marketing year. The organization conducts farmer surveys towards the end of February and early March (March Agricultural Survey) and provides markets with an indication of the expected crop size (Good and Irwin, 2011). A total of 78,900 farm operators were included in this survey in 2021 (Lofthus, 2021). In addition, every year in June, USDA publishes the Crop Acreage report presenting information on planted acreage for major crops across different states. The data on yield is available in the annual crop production summary published every year in January. The Cornell University Library maintains these data from 1980 through 2020. In addition, we also collect data on ethanol production from the U.S. Department of Energy. Table 4.1 provides summary statistics for each of the above series we use in our analysis.

Table 4.1: Summary Statistics

Variable	Time Period	Mean	Standard Error
Prospective Plantings (1,000 acres)	1980-2020	82635.5	1191.4
Planted Acreage (1,000 acres)		82340.6	1300.6
Yield (per acre)		136.58	4.06
Ethanol Production (Million gallons)	2000-2020	9870	1206.02

Source: USDA Reports

Domestic prospective and planted corn acreage increased from 1980 to 2020, as shown in figure 4.2. In most years, prospective and actual plantings are fairly close. Their mean within-year difference over the period of observation is 0.36%. In 1983 prospective plantings exceeded actual acreage by 9.4 million acres—the largest difference observed in our dataset—potentially due to a drought in the mid-west that year. The next largest difference occurred in 2020, where planted acreage was about 5 million fewer acres than expected. This difference is most likely due to the reduced demand for ethanol and therefore corn as a result of the COVID-19 pandemic. Moreover, the pandemic also disrupted an increasing ethanol production trend in 2020 (a decrease of 1.8 billion gallons). For the analysis we use harvest contract futures closing prices traded on the day prospective and acreage plantings reports are released. We draw these prices from Bloomberg.

Figure 4.2: Corn Acreage over Time

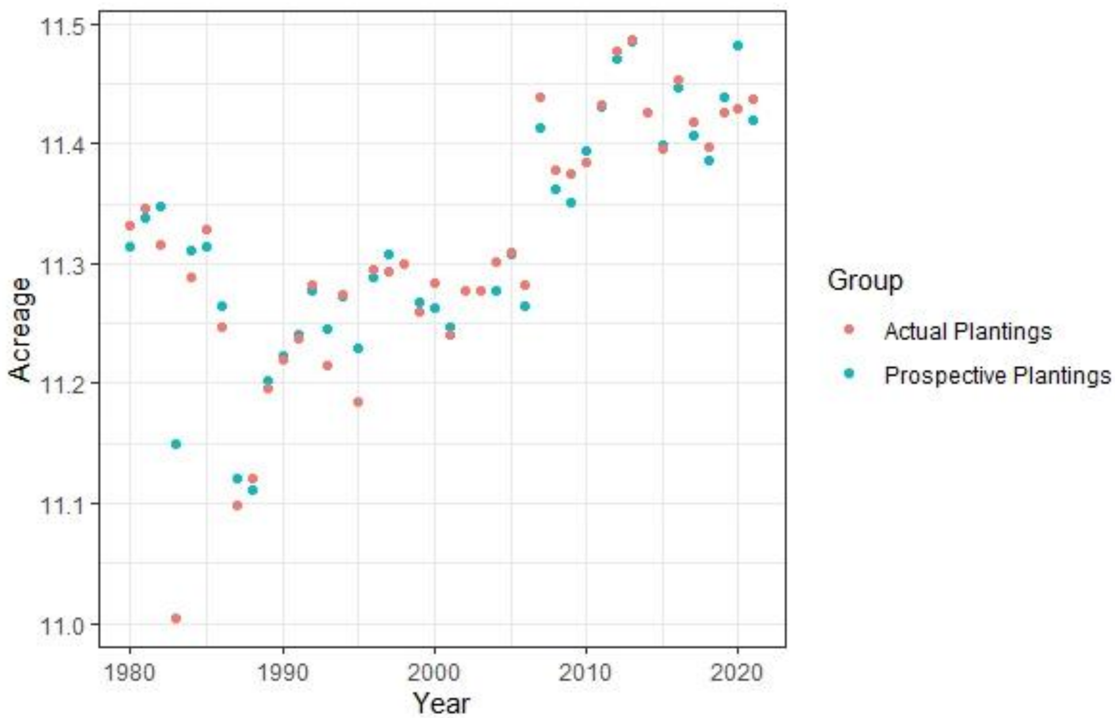

Source: USDA reports

4.4. Methodology

We follow the approach of Roberts and Schlenker (2013) to estimate agricultural supply elasticities. In their model, quantity is specified as a function of price and other covariates:

$$q_{st} = \alpha_s + \beta_s p_{st} + \gamma_s \omega_t + f_s(t) + u_t \tag{1}$$

where q_{st} is the log corn supply (prospective and actuals) at time *t*. The log futures price is denoted by p_{st}, ω_t refers to the yield shock at time *t*, and the time trend is captured by $f_s(t)$. However, farmers might make acreage decisions based on futures prices, and simultaneously acreage might determine futures prices, resulting in endogenous prices negatively biasing the elasticity coefficient. We follow Roberts and Schlenker (2013) and use the lagged yield shock as an instrument to address the concern of endogenous prices. Yield shocks are unexpected deviations from yield trends that occur due to random weather shocks. We approximate these shocks by calculating Jackknife residuals by fitting a yield trend using a restricted cubic spline with three knots (Roberts and Schlenker, 2013; Hendricks et al., 2014a). We also perform a robustness test by repeating the procedure using four and five knots. More specifically, the first stage of the 2-SLS method is given by:

$$p_{st} = \delta_s + \mu_{s0}\omega_t + \mu_{s1}\omega_{t-1} + f_s(t) + \epsilon_t \qquad (2)$$

Lagged yield shock (ω_{t-1}) is an excluded instrument. Before conducting our analysis, we deflate futures prices with the seasonally adjusted consumer price index (CPI).

Figure 4.3: Demand and Supply Movements due to COVID-19

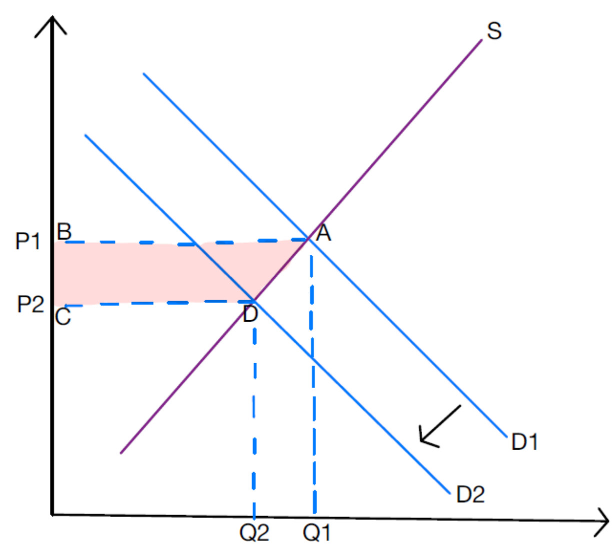

As shown in figure 4.3, because of the COVID-19 pandemic, the demand curve shifted inward, resulting in losses for farmers and ranchers; in the figure, it moved from D1 to D2. In the short run, because the supply technology hasn't changed the supply curve remains constant.[16]

[16] Planting harvest of field crops isn't affected by covid lockdown restrictions

Therefore, we use our estimated elasticities from equation (1) to approximate the welfare damage—or the loss in producer welfare—resulting from this shift in the demand curve (trapezoid ABCD in the figure).[17] To assess this, we use the following equation:

$$\text{Welfare Loss} = \frac{1}{2}\left(P_{January-P'}\right)\left(\frac{Q_{Acreage\ 2020}}{1+\beta_s\frac{\left(P'-P_{January}\right)}{P_{January}}}+Q_{Acreage\ 2020}\right) \qquad (3)$$

Where, $P_{January}$ is the average futures price for the December 2020 corn contract in January 2020, and $Q_{Acreage\ 2020}$ is the actual corn acreage planted in 2020 (from the June Acreage report). The supply elasticity estimate is given by β_s (from equation (1)). We use three different estimates for P': (1) the harvest futures contract price in April (used by USDA in their calculations), (2) in August (harvest month), and (3) in November (one month before contract expiration).

4.5. Results

In Table 4.2, we report the 2-stage least squares (2-SLS) model estimated using equations (1) and (2) using the data from 1980 to 2020. Column (A) corresponds to our base specification with yield shock and time trend (fitted using a cubic spline with three knots) as controls and lagged yield shock as an instrument. Our result implies that, on average, a 1% increase in corn's harvest futures prices, increases acreage by 0.27%. In columns (B) and (C), we fit the time trend using a cubic spline with four and five knots, respectively. The supply elasticity varies between 0.26% – 0.28% with overlapping confidence intervals, indicating the robustness of our results.

[17] We assume that the demand and supply curves are linear.

Elasticity coefficients are statistically significant at a 5% significance level across all regression models. Column (B) is our preferred model because it has the lowest Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) values.[18] Figure 4.4 compares our estimated corn elasticity to the results in the literature. It is important to note that Hendricks et al. (2014) use data from 2000 to 2010, and Kim and Moschini use data from 2005 to 2015; our results are based on the information from 1980 to 2020. Moreover, none of these papers instruments for futures prices.

Next, we assess the welfare damages to corn producers because of the pandemic by estimating equation (3) with $\beta_s = 0.264\%$. Based on price changes from January to April (the period used by USDA in their calculations), we estimate that the pandemic generated $5.4 billion (95% CI: [$5.35 bn, $5.45 bn]) in damages for corn producers due to an inward shift in the demand curve. In addition, our results project an estimated loss of $5.76 billion (95%CI: [$5.71 bn, $ 5.82 bn]) when using price changes from January to August. However, when using price changes from January to November (one month before the 2020 harvest contract expires), our analysis projects that corn producers would have instead gained during the pandemic. However, it is important to note that increasing our period (January – November) might be adding noise and diluting the pandemic's impact since the country was adapting and moving towards normalcy. Figure 4.5 documents the welfare damages in three different scenarios and compares them to the USDA's payments under CFAP 1 and 2 – all our estimates indicate that the USDA was able to fully compensate farmers for their damages due to the pandemic.

[18] AIC measures the relative distance between the fitted and the true (unknown) likelihood function, whereas BIC estimates the probability of the estimated model being true; lower AIC and BIC values are preferred.

Table 4.2: 2-SLS Regression Results for Corn

	Dependent variable:		
	Log Acreage		
	3 knots (A)	**4 knots (B)**	**5 knots (C)**
Log Price (β_s)	0.276^{***}	0.264^{***}	0.284^{***}
	(0.068)	(0.062)	(0.070)
Yield Shock	0.025^{***}	0.026^{***}	0.024^{***}
	(0.008)	(0.008)	(0.008)
			(0.098)
Constant	-6.61	-10.19	-1.9
	(6.957)	(12.329)	(12.178)
Weak instruments	0	0	0
Wu-Hausman	0.11	0.04	0.03
Observations	82	82	82
R^2	0.696	0.716	0.702

*Note: Robust standard errors are reported in parentheses. Asterisks *,**, and *** denote statistical significance at the 10%, 5%, and 1%, respectively.*
Source: Authors

4.6. Conclusions

Previous literature addresses the concern of endogenous futures prices by using lagged yield shocks as instruments. However, these papers only use planted acreage as a dependent variable. USDA publishes prospective plantings and acreage in two separate reports published three months apart. We take advantage of this information and use both expected plantings and actual plantings to estimate agricultural supply elasticity for corn from 1980 to 2020. Our results show consistency across different model specifications, with a mean elasticity estimate of 0.26%.

Figure 4.4: Comparison of Corn Elasticities

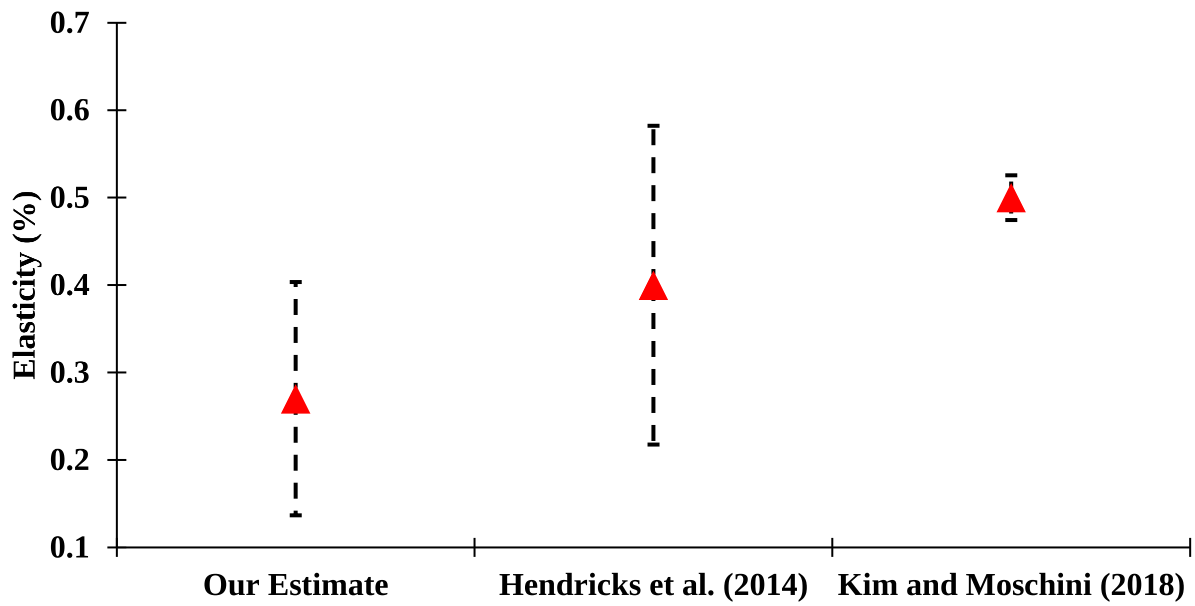

Source: Authors

Recently, due to the COVID-19 pandemic, demand for ethanol plummeted – resulting in an inward shift in the demand curve for corn. Consequently, our analysis shows that corn producers suffered losses worth \$5.4 billion (using price changes from January to April), 27% lower compared to the actual payments USDA made under CFAP. Moreover, all our estimates suggest that the USDA was able to fully compensate producers for any losses they faced due to demand reduction.

Figure 4.5: Estimated Corn Producer Surplus Effects of COVID-19

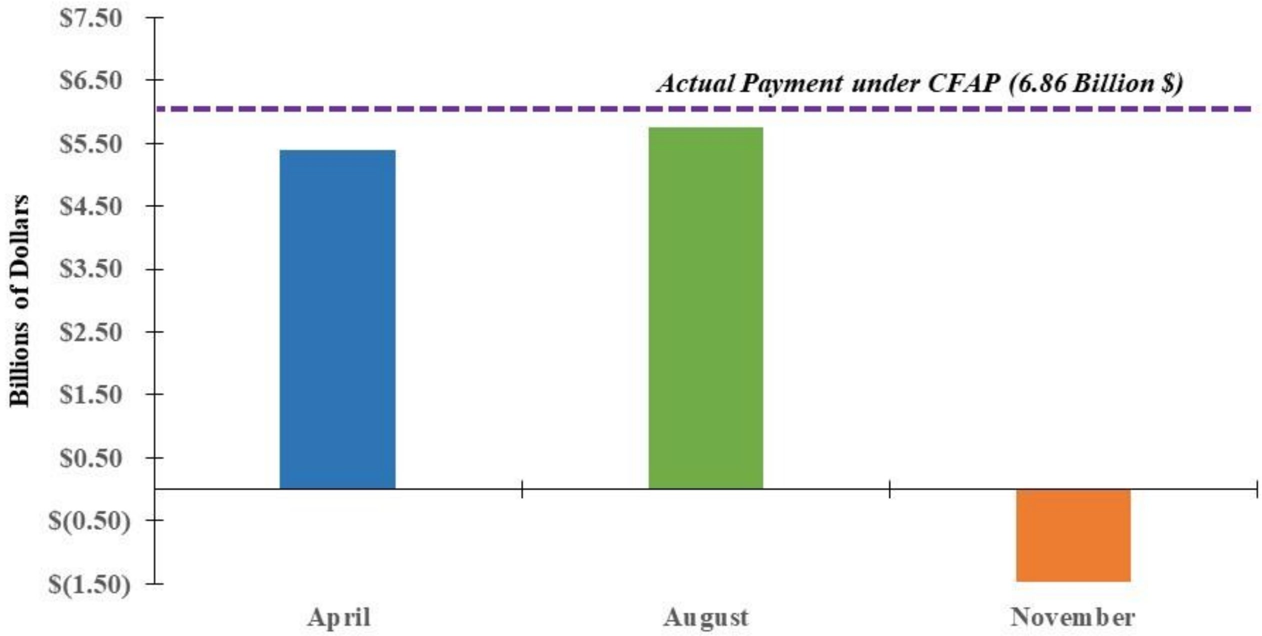

Source: Authors

CHAPTER 5

CONCLUSIONS

In conclusion, my work underscores the importance of government information in reducing uncertainty in agricultural commodity markets. The absence of such information increases implied volatility resulting in losses for market participants. For example, the 2019 government shutdown and the subsequent curtailment of the WASDE generated damages worth 7 million dollars for corn and soybean producers alone. Given the importance of these reports, the government forecasts must be highly efficient and accurate. In another chapter, I show that production and export misses are the significant contributors to endings stocks forecast errors. Errors in predicting US exports to China, Mexico, Brazil, and Southeast Asian countries are the leading contributors to U.S. export errors. In the last chapter, I show that the COVID-19 pandemic resulted in approximately $5.5 billion loss for corn farmers. My findings will help governmental organizations understand the impact of adverse events on farmers and help them develop policies accordingly.

REFERENCES

Adjemian, M. K., and A. Smith. “Using USDA Forecasts to Estimate the Price Flexibility of Demand for Agricultural Commodities” *American Journal of Agricultural Economics*. 94(4), 978-995. Accessed at: https://files.asmith.ucdavis.edu/2012_AJAE_AS_flexibility.pdf

Adjemian, M.K. 2012. “Quantifying the WASDE Announcement Effect.” *American Journal of Agricultural Economics* 94 (1): 238–56. Accessed at: https://doi.org/10.1093/ajae/aar131

Adjemian, M.K., R. Johansson, A. McKenzie, and M. Thomsen. 2018. “Was the Missing 2013 WASDE Missed?” *Applied Economic Perspectives and Policy* 40(4): 653-671. Accessed at: https://doi.org/10.1093/aepp/ppx049

Adjemian, M.K., A. Smith, and W. He. 2021. “Estimating the Market Effect of a Trade War: The Case of Soybean Tariffs.” *Food Policy* 105. Accessed at: https://doi.org/10.1016/j.foodpol.2021.102152

Adjemian, M.K. and S.H. Irwin, 2018. "USDA Announcement Effects in Real-Time.” *American Journal of Agricultural Economics, Agricultural and Applied Economics Association*. 100(4): 1151-1171. Accessed at https://onlinelibrary.wiley.com/doi/full/10.1093/ajae/aay018

Agarwal, J., and T., Wu. 2004. "China's entry to WTO: global marketing issues, impact, and implications for China." *International Marketing Review* 21(3): 279-300. Accessed at: http://dx.doi.org/10.11575/PRISM/34081

Amin, M. D., S. Badruddoza, and J. J. McCluskey. 2021. "Predicting access to healthful food retailers with machine learning." *Food Policy* 99. Accessed at: https://doi.org/10.1016/j.foodpol.2020.101985

Baker, S.R., N. Bloom, and S.J. Davis, 2016. "Measuring Economic Policy Uncertainty." *The Quarterly Journal of Economics* 131(4): 1593 - 1636. Accessed at: https://academic.oup.com/qje/article/131/4/1593/2468873

Beckers, S., 1980. "The Constant Elasticity of Variance Model and Its Implications For Option Pricing." *The Journal of Finance* 35(3): 661 – 673. Accessed at: https://www.jstor.org/stable/2327490

Barone-Adesi, G. and R.E. Whaley, 1987. "Efficient Analytic Approximation of American Option Values." *The Journal of Finance* 42(2): 301-320. Accessed at: https://doi.org/10.1111/j.1540-6261.1987.tb02569.x

Bollerslev, T., 1986. "Generalized autoregressive conditional heteroskedasticity." *Journal of Econometrics* 31(3): 307–327. Accessed at: https://doi.org/10.1016/0304-4076(86)90063-1

Benami, E., Z. Jin, M. R. Carter, A. Ghosh, R. J. Hijmans, A. Hobbs, B. Kenduiywo, and D. B. Lobell. 2021. "Uniting remote sensing, crop modelling and economics for agricultural risk management." *Nat Rev Earth Environ* 2: 140 – 159. Accessed at: https://doi.org/10.1038/s43017-020-00122-y

Botto, A.C., O., Isengildina-Massa, S.H. Irwin, and D.L. Good. 2006. "Accuracy trends and sources of forecast errors in WASDE balance sheet categories for corn and soybeans.", Paper presented at the *2006 American Agricultural Economics Association Annual Meeting, Long Beach, California, July 23-26.* Accessed at: https://core.ac.uk/download/pdf/7055392.pdf

Badruddoza, S., A. C. Carlson, and J. J. McCluskey. 2021. "Long-term dynamics of US organic milk, eggs, and yogurt premiums." *Agribusiness* 1-28. Accessed at: https://doi.org/10.1002/agr.21723

Breiman, L. 2001. "Random Forests." *Machine Learning* 45(1): 5-32. Accessed at: https://doi.org/10.1023/A:1010933404324

Cater, C. A., G. C. Rausser, and A. Smith. 2016. "Commodity Storage and the Market Effects of Biofuel Policies." *American Journal of Agricultural Economics*. 99(4), 1027-1055. Accessed at: https://doi.org/10.1093/ajae/aaw010

Chakraborty, C., and A. Joseph 2017. "Machine learning at central banks." *Bank of England Working Papers, No. 674*. Accessed at: https://www.bankofengland.co.uk/working-paper/2017/machine-learning-at-central-banks

Chen, T., and C. Guestrin. 2016. "Xgboost: A scalable tree boosting system." In Proceedings of the *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*: 785–794. Accessed at: https://doi.org/10.48550/arXiv.1603.02754

Coibion, O., and Y. Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review*, 105 (8): 2644-78. Accessed at: https://www.aeaweb.org/articles?id=10.1257/aer.20110306

Colling P.L. and S.H. Irwin, 1990. "The Reaction of Live Hog Futures Prices to USDA Hogs and Pigs Reports." *American Journal of Agricultural Economics* 72(1): 84–94. Accessed at: https://onlinelibrary.wiley.com/doi/abs/10.2307/1243147

Chua, H.W.P., and W.G. Tomek. 2010. "On the Relationship of Expected Supply and Demand to Futures Prices." Paper presented at *NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management*. St. Louis, April 19, 2010. Accessed at: https://ageconsearch.umn.edu/record/121055/?ln=en

Davidson, W.N., J.K. Kim, E. Ors, and A. Szakmary, 2001. “Using Implied Volatility on options to measure the relation between asset returns and variability.” *Journal of Banking & Finance* 25(7): 1245 – 1269. Accessed at: https://doi.org/10.1016/S0378-4266(00)00128-X

Egelkraut T.M., P. Garcia, S.H. Irwin and D.L. Good, 2003. "An Evaluation of Crop Forecast Accuracy for Corn and Soybeans: USDA and Private Information Agencies." *Journal of Agricultural and Applied Economics* 35(1): 79–95. Accessed at: https://doi.org/10.1017/S1074070800005952

Engle, R., 2002. "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models." *Journal of Business & Economic Statistics* 20(3):339 - 350. Accessed at: http://www.jstor.org/stable/1392121

Feldman, B.E. 2005. “Relative Importance and Value.” *SSRN.* Accessed at: https://ssrn.com/abstract=2255827

Gauthier, J., Q.V. Wu, and T.A. Gooley, 2020. “Cubic splines to model relationships between continuous variables and outcomes: a guide for clinicians.” *Bone Marrow Transplant* 55: 675–680. Accessed at: https://doi.org/10.1038/s41409-019-0679-x

Goodwin, B.K., and A.K. Mishra. 2006. “Are Decoupled Farm Program Payments Really Decoupled? An Empirical Evaluation.” *American Journal of Agricultural Economics.* 88(1):

73-89. Accessed at:
https://www.jstor.org/stable/3697967?seq=1#metadata_info_tab_contents

Good, D., and S. Irwin. 2011. “USDA Corn and Soybean Acreage Estimates and Yield Forecasts: Dispelling Myths and Misunderstandings.” *Marketing & Outlook Briefs.* Accessed at: https://farmdoc.illinois.edu/publications/usda-corn-and-soybean-acreage-estimates-and-yield-forecastsdispelling-myths-and-misunderstandings

Goyal, R., and M. K. Adjemian. 2021. “The 2019 government shutdown increased uncertainty in major agricultural commodity markets.” *Food Policy* 102. Accessed at: https://doi.org/10.1016/j.foodpol.2021.102064

Goyal, R., and M. K. Adjemian. 2021. “Information Rigidities in USDA Forecasts.” *Working Paper*. Accessed at: https://www.econ.iastate.edu/files/events/files/information_rigidities.pdf

Gray, R.W. 1974. “Grain Reserve Issues.” Paper presented at National Agricultural Outlook Conference. Washington, DC, December 9–12.

Hassett, K.A., and G.E. Metcalf., 2001. "Investment with Uncertain Tax Policy: Does Random Tax Policy Discourage Investment." *The Economic Journal* 109(457): 372 - 93. Accessed at: https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0297.00453

Hendricks, N. P., J. P. Janzen, and A. Smith. 2014a. “Futures Prices in Supply Analysis: Are Instrumental Variables Necessary?” *American Journal of Agricultural Economics*. 97(1), 22-39. Accessed at: https://onlinelibrary.wiley.com/doi/full/10.1093/ajae/aau062

Hendricks, N. P., A. Smith, and D. A. Sumner. 2014b. “Crop Supply Dynamics and the Illusion of Partial Adjustment” *American Journal of Agricultural Economics*. 96(5), 1469-91. Accessed at: https://www.jstor.org/stable/24476988?seq=1

Hodrick, R.J., 1992. “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement.” *The Review of Financial Studies* 5(3): 357–86. Accessed at: https://academic.oup.com/rfs/article/5/3/357/1576229?login=true

Hossain, M., C. Mullally, and M. N. Asadullah. 2019. “Alternatives to calorie-based indicators of food security: An application of machine learning methods.” *Food Policy* 84: 77-91. Accessed at: https://doi.org/10.1016/j.foodpol.2019.03.001

Irwin, S., and D. Good. 2016. “The Relationship between Stocks-to-Use and Corn and Soybean Prices: An Alternative View.” *farmdoc daily* (6): 66. Accessed at: https://farmdocdaily.illinois.edu/2016/04/relationship-between-stock-to-use-and-prices.html

Isengildina-Massa, O., S.H. Irwin, and D.L. Good. 2006. “Are Revisions to USDA Crop Production Forecasts Smoothed?” *American Journal of Agricultural Economics* 88(4): 1091-1104. Accessed at: https://www.jstor.org/stable/4123548

Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008a. “Impact of WASDE Reports on Implied Volatility in Corn and Soybean Markets.” *Agribusiness* 24 (4): 473–490. Accessed at https://doi.org/10.1002/agr.20174

Isengildina-Massa O., S.H. Irwin, D.L. Good and J.K. Gomez, 2008b. "The Impact of Situation and Outlook Information in Corn and Soybean Futures Markets: Evidence from WASDE Reports." *Journal of Agricultural and Applied Economics* 40(1): 89–103. Accessed at: http://ageconsearch.umn.edu/record/45048/files/jaae-40-01-89.pdf

Isengildina-Massa, O., B. Karali, and S.H. Irwin. 2013. “When do the USDA forecasters make mistakes?” *Applied Economics* 45:36, 5086-5103. Accessed at: https://doi.org/10.1080/00036846.2013.818213

Isengildina-Massa, O., B. Karali, T. Kuethe, and A. Katchova. 2020. “Joint Evaluation of the System of USDA’s Farm Income Forecasts.” Applied Economic Perspectives and Policy 1-21. Accessed at: https://doi.org/10.1002/aepp.13064

Isengildina-Massa O., B. Karali, and S.H. Irwin., 2020. "Can Private Forecasters Beat the USDA? Analysis of Relative Accuracy of Crop Acreage and Production Forecasts" *Journal of Agricultural and Applied Economics* 52(14): 545–61. Accessed at: https://doi.org/10.1017/aae.2020.18

Jankova, Z., 2018. "Drawbacks and Limitations of Black-Scholes Model for Options Pricing." *Journal of Financial Studies & Research*. Accessed at: https://ibimapublishing.com/articles/JFSR/2018/179814/

Joseph, A. 2019. "Parametric inference with universal function approximators." *Bank of England Working Papers, No. 784*. Accessed at: https://www.bankofengland.co.uk/working-paper/2019/shapley-regressions-a-framework-for-statistical-inference-on-machine-learning-models

Karali, B., O.I. Massa, S.H. Irwin, M.K. Adjemian and R. Johansson, 2019. "Are USDA reports still news to changing crop markets?" *Food Policy* 84: 66-76. Accessed at https://doi.org/10.1016/j.foodpol.2019.02.005

Kim, H., and G. Moschini. 2018. "The Dynamics of Supply: U.S. Corn and Soybeans in the Biofuel Era." *Land Economic*s. 94(4):593-613. Accessed at: http://le.uwpress.org/cgi/reprint/94/4/593

Leiva, R. C., A. F. Anta, V. Mancuso, and P. Casari. 2019. "A Novel Hyperparameter-free Approach to Decision Tree Construction that Avoids Overfitting by Design." *IEEE Access* 7: 99978-99987. Accessed at: 10.1109/ACCESS.2019.2930235

Liaw, A., and M. Wiener. 2002. "Classification and Regression by randomForest." *R News* 2:18-22. Accessed at: https://cogns.northwestern.edu/cbmg/LiawAndWiener2002.pdf

Lindeman, R.H., P.F. Miranda, and R.Z. Gold. 1980. *Introduction to Bivariate and Multivariate* Analysis, Scott Foresman, Glenview, IL.

Lundberg, S.M., G. G. Erion, and S.-I. Lee. 2018. "Consistent Individualized Feature Attribution for Tree Ensembles." Accessed at: https://doi.org/10.48550/arXiv.1802.03888

Lundberg, S.M., and S.-I. Lee. 2017. "A unified approach to interpreting model predictions." *Advances in Neural Information Processing Systems* 4765-4774. Accessed at: https://doi.org/10.48550/arXiv.1705.07874

Lofthus, D. 2021. "Understanding the Prospective Plantings Report." *United States Department of Agriculture*. Accessed at: https://www.nass.usda.gov/Statistics_by_State/Minnesota/Publications/Other_Press_Releases/2021/USDA-NASS-Transparency-1-March-2021-PP.pdf

MacDonald, S., F. Gale, and J. Hansen. 2015. "Cotton Policy in China.", *U.S. Department of Agriculture.* Accessed at: https://www.ers.usda.gov/webdocs/outlooks/36244/52550_cws-15c-01.pdf?v=6603.4

MacDonald, S., M. Ash, and B. Cooke. 2017. "The Evolution of Inefficiency in USDA's Forecasts of U.S. and World Soybean Markets.", *Munich Personal RePEc Archive.* Accessed at: https://mpra.ub.uni-muenchen.de/id/eprint/87545

MacDonald, S., and M. Ash. 2016. "Detecting the Sources of Information Rigidity: Analyzing Forecast Bias and Smoothing in USDA's Soybean Forecasts.", Paper presented at the *2016 Agricultural & Applied Economics Association Annual Meeting, Boston, Massachusetts, July 31-August 2.* Accessed at: https://ageconsearch.umn.edu/record/235349

Mallory, M. L., D. J. Hayes, and B. A. Babcock. 2011. "Crop-Based Biofuel Production with Acreage Competition and Uncertainty." *Land Economics* 87(4): 610-627. Accessed at: http://le.uwpress.org/content/87/4/610.short

McNew, K.P. and J.A. Espinosa, 1994. "The Informational Content of USDA Crop Reports: Impacts on Uncertainty and Expectations in Grain Futures Markets." *Journal of Futures Markets* 14:475 – 492. Accessed at https://doi.org/10.1002/fut.3990140408

Nelson, D. B., 1991. "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica* 59 (2), 347–370. Accessed at: https://www.jstor.org/stable/2938260

Nerlove, M. 1958. "The Dynamics of Supply: Estimation of Farmer's Response to Price." Baltimore: Johns Hopkins University Press.

No, C. S., and M.E. Salassi. 2009. "A sequential rationality test of USDA preliminary price estimates for selected program crops: Rice, Soybeans, and Wheat." *International Advances In Economic Research* 15: 470-482. Accessed at: https://doi.org/10.1007/s11294-009-9228-5

Nodari, G., 2014. "Financial regulation policy uncertainty and credit spreads in the US." *Journal of Macroeconomics* 41(C): 122- 32. Accessed at: https://econpapers.repec.org/article/eeejmacro/v_3a41_3ay_3a2014_3ai_3ac_3ap_3a122-132.htm

Quinlan, J. R. 1993. “Programs for Machine Learning.” *Morgan Kaufmann Publishers, San Mateo, California.*

Roberts, M. J., and W. Schlenker. 2013. “Identifying Supply and Demand Elasticities of Agricultural Commodities: Implications for the US Ethanol Mandate.”, *American Economic Review.* 103(6), 2265– 2295. Accessed at: https://www.aeaweb.org/articles?id=10.1257/aer.103.6.2265

RFA. 2020. “Ethanol Industry Needs Support, as COVID-19 Losses Near $4 Billion.” Accessed at: https://ethanolrfa.org/media-and-news/category/news-releases/article/2020/12/rfa-ethanol-industry-needs-support-as-covid-19-losses-near-4-billion

Ruppert, D., 2002. “Selecting the number of knots for penalized splines.” *Journal of computational and graphical statistics* 11(4):735–757. Accessed at: https://www.jstor.org/stable/1391159

Schmitz, A., C. B. Moss, and T. G. Schmitz. 2020. "The Economic Effects of COVID-19 on the Producers of Ethanol, Corn, Gasoline, and Oil." *Journal of Agricultural & Food Industrial Organization* 18(2). Accessed at: https://doi.org/10.1515/jafio-2020-0025

Sumner D.A. and R.A.E. Mueller, 1989. "Are Harvest Forecasts News? USDA Announcements and Futures Market Reactions." *American Journal of Agricultural Economics* 71(1): 1–8. Accessed at: http://dx.doi.org/10.2307/1241769

Thompson, W., S. Arita, J. Cooper, P. Westhoff, S. Gerlt, and S. Meyer. 2021. "Agricultural Supply Response under Extreme Market Events and Policy Shocks." *Selected Paper 2021 AAEA & WAEA Joint Annual Meeting, Austin, TX, August 1-3, 2021*. Accessed at: https://ageconsearch.umn.edu/record/313930/references

Tomek, W.G. 1979. "Estimating Demand Relations Using Futures Prices." *Agricultural Economics Research* 31(4):30–39. Accessed at: https://ideas.repec.org/a/ags/uersja/148574.html

Vogel F.A. and G.A.Bange, 1999. "Understanding USDA Crop Forecasts." Washington DC: U.S. Department of Agriculture, National Agricultural Statistics Service and Office of the Chief Economist, World Agricultural Outlook Board. Accessed at: https://www.nass.usda.gov/Education_and_Outreach/Understanding_Statistics/pub1554.pdf

Whaley, R.E. 1986. "Valuation of American Futures Options: Theory and Empirical Tests." *The Journal of Finance* 41(1): 127 – 150. Accessed at: https://www.jstor.org/stable/2328348?seq=1#metadata_info_tab_contents

Wiggins, J.B., 1987. "Option values under stochastic volatility: Theory and empirical estimates." *Journal of Financial Economics* 19(2): 351-372. Accessed at: https://doi.org/10.1016/0304-405X(87)90009-2

Woloszko, N. 2020. "Adaptive Trees: A New Approach to Economics Forecasting." *OECD Economics Department Working Papers, No. 1593, OECD Publishing*. Accessed at: https://doi.org/10.1787/5569a0aa-en

Xiao J., S.H.Lence and C.Hart, 2014. "USDA and Private Analysts Forecasts of Ending Stocks: How Good Are They?" Selected paper presented at 2014 AAEA Annual Meeting, Minneapolis, MN, July 27-29, 2014.

Xiao, J., C. E. Hart, and S. H. Lence. 2017 "USDA Forecasts of Crop Ending Stocks: How Well Have They Performed?", *Applied Economic Perspectives and Policy* 39(2), 220–241. Accessed at: https://doi.org/10.1093/aepp/ppx023

Yang, S., J. Wu, Y. Du, Y. He, and X. Chen. 2017. "Ensemble Learning for Short-Term Traffic Prediction Based on Gradient Boosting Machine." *Journal of Sensors*. Accessed at: https://www.hindawi.com/journals/js/2017/7074143/

Ying, J., Y. Chen, and J.H. Dorfman. 2019. "Flexible Tests for USDA Report Announcement Effects in Futures Markets." *American Journal of Agricultural Economics* 101(4): 1228-46. Accessed at: https://doi.org/10.1093/ajae/aaz013

"Colombia 2019 Export Highlights." U.S. Department of Agriculture, Foreign Agricultural Service. Accessed at: https://www.fas.usda.gov/colombia-2019-export-highlights

"Focus on Prices and Spending: Import and Export Prices, Second Quarter 2011." U.S. Bureau of Labor Statistics 2(5). Accessed at: https://www.bls.gov/opub/btn/archive/the-impact-of-soaring-cotton-prices-on-consumer-apparel-prices.pdf

"Interdependence of China, United States, and Brazil in Soybean Trade." Economic Research Service, U.S. Department of Agriculture. 2019. Accessed at: https://www.ers.usda.gov/webdocs/outlooks/93390/ocs-19f-01.pdf?v=5860.4

"USDA Coexistence Fact Sheets, Soybeans." U.S. Department of Agriculture. 2017. Accessed at: https://www.usda.gov/sites/default/files/documents/coexistence-soybeans-factsheet.pdf

"World Agricultural Supply and Demand Estimates, December" U.S. Department of Agriculture. 2021. Accessed at: https://www.usda.gov/oce/commodity/wasde/wasde1221.pdf

"Soybean prices, economic growth and poverty in Argentina and Brazil" Commodities and Development Report, Food and Agriculture Organization of the United Nations. 2017. Accessed at: http://www.fao.org/3/I8316EN/i8316en.pdg

USDA, 2019. "USDA Reschedules Reports Affected by Lapse in Federal Funding." National Agricultural Statistics Service, Agricultural Statistics Board Notice. January 30th. Accessed at: https://downloads.usda.library.cornell.edu/usda-esmis/files/tm70mv177/dz010w33d/j9602589k/01-30-19_Report_Reschedule_ASB_Notice_Final.pdf

USDA. 2021. "Feedgrains Sector at a Glance." Accessed at: https://www.ers.usda.gov/topics/crops/corn-and-other-feedgrains/feedgrains-sector-at-a-glance/