Decomposing USDA Ending Stocks Forecast Errors

Raghav Goyal, Michael K. Adjemian, Joseph Glauber & Seth Meyer[1]

Paper prepared for the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management, 2021.

[1] Raghav Goyal is a PhD candidate at the University of Georgia. Michael K. Adjemian is an Associate Professor in the department of Agricultural and Applied Economics at the University of Georgia. Joseph Glauber is a Senior Research Fellow at IFPRI. Seth Meyer is a Chief Economist at the USDA.

Decomposing USDA Ending Stocks Forecast Errors

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 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, Extreme Gradient Boost Tree (EGBT), that uses a gradient boosting framework and is robust to multicollinearity. Our results indicate export and production misses to be the major contributors to ending stocks projection errors. 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 and export forecast misses, country-by-country, and show that misses on production and export levels in China, Mexico, Brazil, and Europe cost USDA the most.

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, specifically on their beginning stocks, production, domestic consumption, exports, imports, and ending stocks. The 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 lower implied volatilities and elevate realized volatilities in futures markets, indicating that market participants react to the government's otherwise unavailable information.

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. Higher ending stocks forecasts signal excess supply over demand, leading to lower farm prices. On the contrary, lower ending stocks projections signal scarcity in the market, leading to higher farm prices.

As with any forecast, errors affect the informational value of these reports. Given the significance of the information, 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. Significant forecast errors could negatively impact the welfare of the market participants. Overestimates (underestimates) of ending stocks could lead to lower (higher) producer prices and distort storage and production signals, and ultimately inefficient intertemporal resource allocation decisions. Federal government budgets, payments, and expenditures, based on these estimates, will also be inefficient (Xiao et al., 2017).

In this paper, 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 this sort of analysis in the literature 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 weighted boosting model, with weights defined as a function of expected information uncertainty at the time of making forecasts. The model provides each regressor's relative importance/ contribution to the dependent variable 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 address this gap by decomposing USDA's ending stocks projection errors for four commodities: corn, cotton, soybeans, and wheat. The United States is a key producer of all of these, but recent structural changes in these markets complicate the task of forecasters, including the impacts of 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).

According to our results, for all four commodities U.S. export and production forecast errors are the major 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, Europe, 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, Mexico, Brazil, Europe, and Southeast Asia cost USDA the most. We interpret these results to mean that better information about production expectations, both domestically and worldwide, will contribute to more efficient agricultural balance sheet forecasts.

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.

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 the 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 (like producer price and exchange rate indices) and indications of behavioral sources of error (such as lagged forecast errors and percent change in forecast level) in the ending stocks misses for all three commodities.

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 these projections—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. More recently, 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 soybeans 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.

3. Methodology

Most of the research studies in the previous section use linear regression methods to gauge the importance of demand and supply elements to ending stocks projection errors. This approach has several limitations, including that (1) regression coefficients capture the average effect of a regressor, 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. 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. 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. Although not often found in the economics literature, researchers in other fields have used GBTs to rank independent variables by their relative importance. Pan et al. (2017) use this method to compute the importance of different independent variables in predicting the association of single amino acid variations to a disease; Deng et al. (2019) use these trees to find the most critical variables for insider trading identification in China's stock market to better regulate insider trading activities; Wu et al. (2017) use this method to identify nonlinearities in traffic variation and find key variables to improve short-term traffic prediction. We identified two economics studies use gradient boosting trees. In a 2020 working paper by the Organization for Economic Cooperation and Development, Woloszko notes the importance of GBTs for macroeconomic modeling, particularly in the presence of structural breaks, and extolls the easy interpretability of these models. In a 2017 Bank of England working paper, Chakraborty and Joseph study the feature importance values obtained from decision trees. We follow these research studies and report the relative importance of each regressor obtained from GBTs.

We also report Shapley additive explanations (SHAP) values (Lundberg and Lee, 2017) in this study. SHAP values provide a localized contribution of each independent variable to the variable of interest, where local attributions assign importance scores to each independent variable by decomposing every observation. In a 2019 working paper by the Bank of England, Joseph uses SHAP values to calculate local attributions 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.

As we will note from our analysis on the decomposition of ending stocks, exports play a crucial role in determining ending stocks errors. We, therefore, study how USDA's country-by-country import projections contribute to U.S. export 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.[2] Lastly, we run a relative importance simulation exercise to gauge the performance of EGBT relative to linear regression. EGBTs, compared to GBTs, leverage computational abilities of computer hardware to speed up the calculations.

3.1 Forecast Errors

We compute forecast errors as follows, depending on which side of the balance sheet the variable resides. If their positive shocks increase the commodity market supply, as in the case of beginning stocks, production, and imports, we calculate their 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 their positive shocks contribute to commodity disappearance, as in the case of exports and domestic use, we calculate their forecast errors as:

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

where the other subscripts are the same and ***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 ***s***,***i***,***t***, with ending stocks forecast errors calculated according to equation (1).

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

3.2 Methods to compute Relative Importance:

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.

3.3 Comparison of EGBT and Linear Regression

We run two controlled simulation experiments to demonstrate the effectiveness of the EGBT model.

3.3.1 Simulation exercise 1: Baseline

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

$$Y = \beta_1 X_1 + \beta_2 X_2^2 + \beta_3 \mathrm{X}_3 + \epsilon, \quad X_3 = X_1 * X_2 \tag{3}$$

$$X_1 \ \sim N(0, 1) \tag{4}$$

$$X_2 \ \sim N(0, 1) \tag{5}$$

$$\beta_1 = \beta_2 = \beta_3 = 0.33 \tag{6}$$

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,[3] 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.

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

Figure 1a 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 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.

Figure 1a: Monte-Carlo simulation results –Baseline

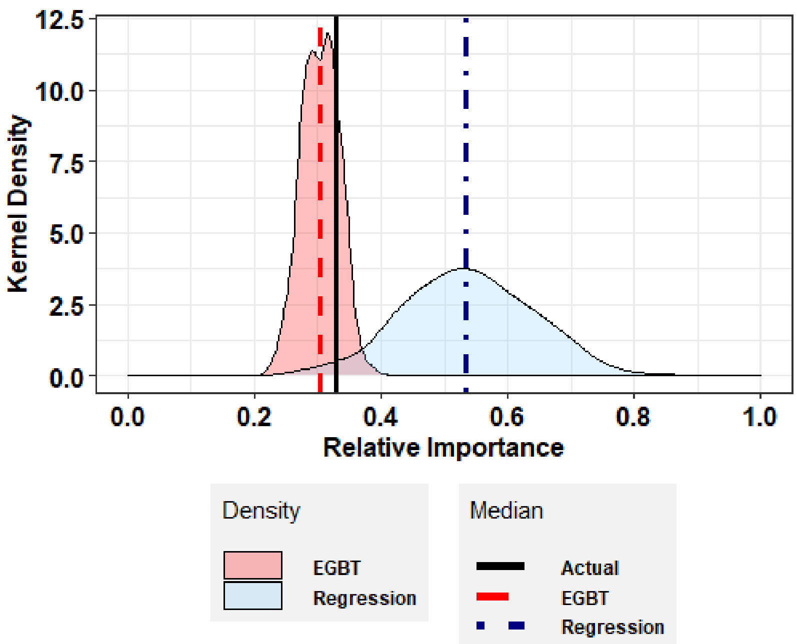

(a) Kernel density plot X_1

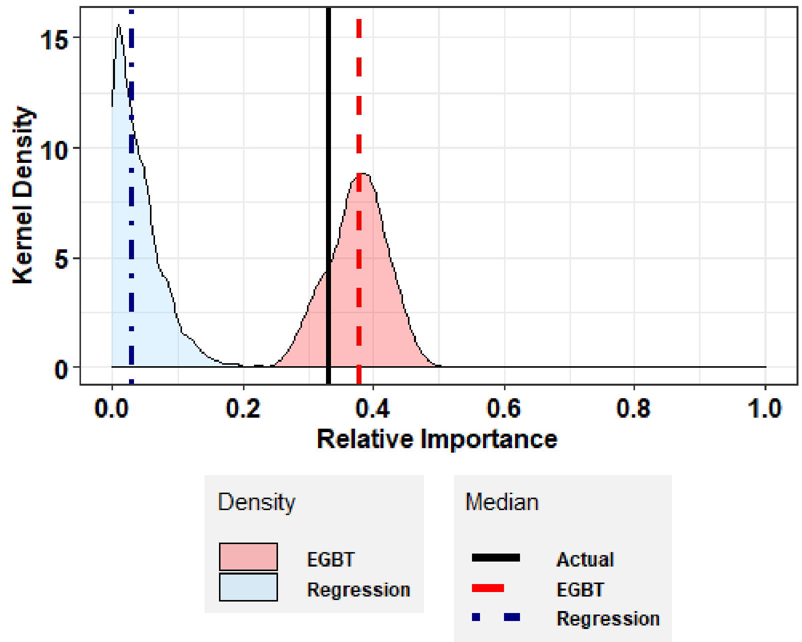

(b) Kernel density plot X_2

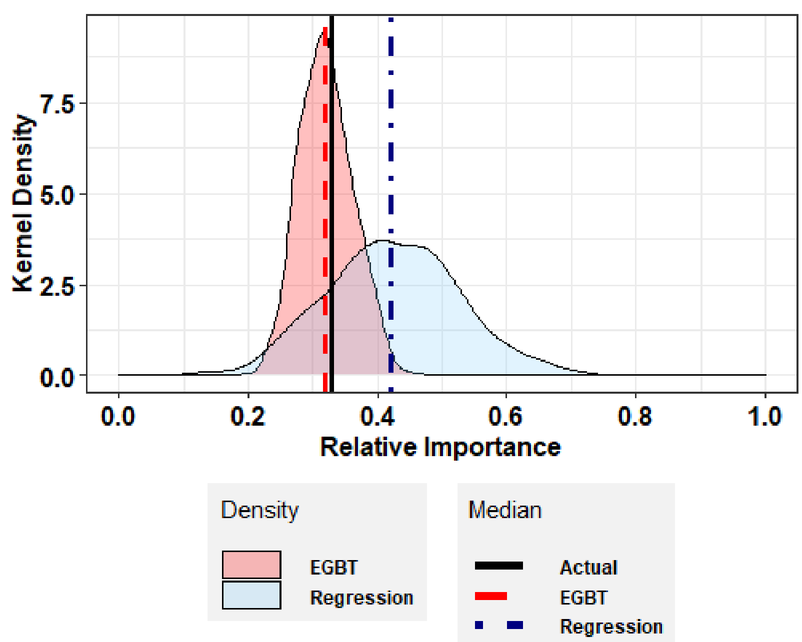

(c) Kernel density plot X_3

Source: Authors' calculations

3.3.2 Simulation exercise 2: Multicollinearities

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

$$Y = \beta_1 X_1 + \beta_2 X_2 + \beta_3 \mathrm{X}_3 + \epsilon \tag{7}$$

$$\begin{bmatrix} X_1 \\ X_2 \end{bmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix}\right] \tag{8}$$

$$X_3 \sim N(0,1) \tag{9}$$

$$\beta_1 = \beta_2 = \beta_3 = 0.33 \tag{10}$$

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 1b 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. Therefore, given the superior performance of EGBTs in both the simulations, we report all the relative importance measures based on the EGBT model.

Figure 1b: Monte-Carlo simulation results –Multicollinearities

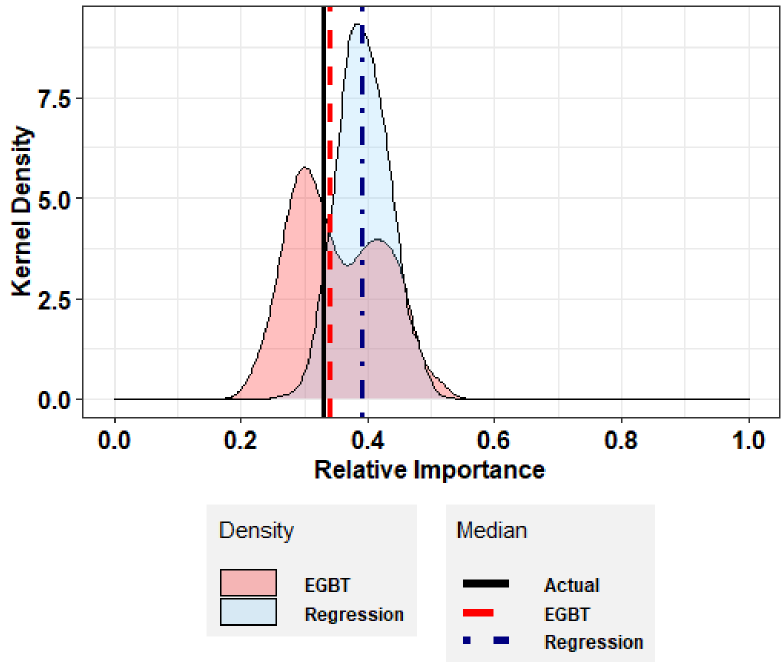

(a) Kernel density plot X_1

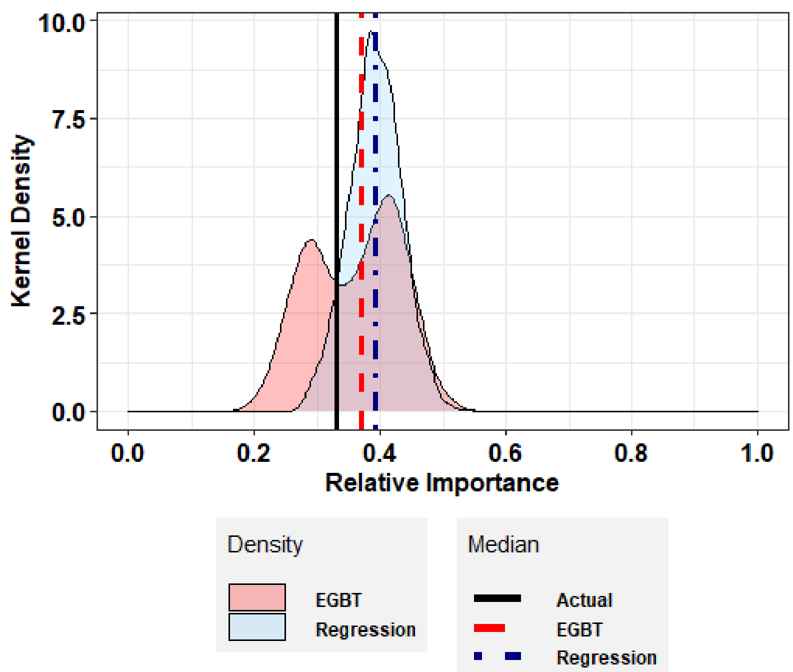

(b) Kernel density plot X_2

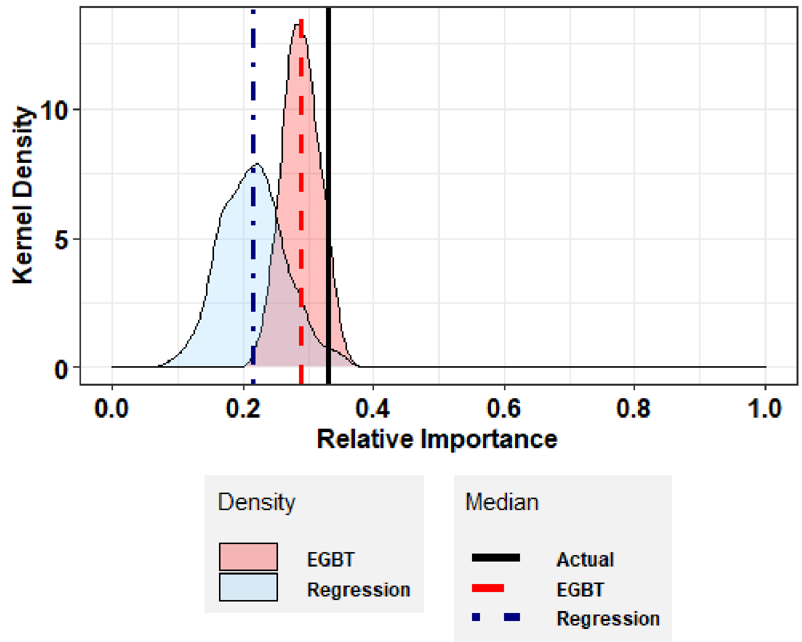

(c) Kernel density plot X_3

Source: Authors' calculations

3.4 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.[4] 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 11) 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,

[4] 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.

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:

$$Q^h = \sum_{i=1}^{n}\left(y^i - (\hat{y}^{i,(h-1)} + g^h\left(x^i\right))\right)^2 + \Omega(g^h) \tag{11}$$

where Ω is the regularization term, imposing a penalty to the loss function and preventing overfitting. $x_i$is the vector of independent variables for observation 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 (11). 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 (11), 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.

USDA makes forecasts for balance sheet elements over a horizon of 19 months, beginning in May and ending in November of the following year. During the initial months before harvest, the uncertainty about many of these elements is higher, resulting in high forecast errors. Therefore, the initial months can dominate the model and hence the importance of independent variables. To account for this asymmetry in uncertainty, we modify the loss function in (11) by using a weighted loss function. For each forecast horizon, we use the average absolute ending stocks misses (for that horizon) as weights:

$$W_f^i = \frac{\sum_{i=1}^{n}(y_{i,f})}{\sum_f \sum_{i=1}^{n}(y_{i,f})} \qquad f = \{1, \ldots\ldots, 19\} \tag{12}$$

W_F^i is the weight for loss function for observation *i* for forecast horizon *f*, while the corresponding value of the dependent variable is $y_{i,f}$. Higher W_F^i will result in higher loss due to that observation, resulting in the higher first and second derivatives of (11).

3.5 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 contributions for each forecast horizon (e.g., for *f* = 1,2, 3,…,19) for each time period, as well, 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:

$$v(z') = \phi_0 + \sum_{j=0}^{M} \phi_j z'_j \tag{13}$$

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 (13) simplifies to:

$$v(5) = \phi_{BS,5} + \phi_{Exp,5} + \phi_{Imp,5} + \phi_{FSI,5} + \phi_{FR,5} + \phi_{E,5} + \phi_{P,5} \tag{14}$$

Where *v (5)* is the SHAP value for the fifth observation. In effect, equation (14) 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. SHAP values tell how to distribute the payout equitably among each independent variable. These 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 (13), 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.

4. Data

We use 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, for 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. There are 510 observations for corn, soybeans, wheat and 434 observations for cotton (since those data are available beginning in 1996/97).[5]

Table 1 provides summary statistics for USDA forecast errors of these elements for the four commodities. We split the elements into (1) variables increasing commodity supply (2) variables decreasing commodity supply.

In table 1, units differ across commodities so we restrict our comparisons within each commodity’s balance sheet. Export forecast errors are highest for cotton, while feed & residual dominate corn and wheat. Average export projection errors are negative for all commodities, indicating the

[5] There were three missing report days, two corresponding to the October 2013 WASDE, and one corresponding to the January 2019 WASDE.

USDA's tendency to underestimate U.S. exports; ending stocks projections for corn, wheat, and cotton exhibit a similar pattern. In contrast, the Department tends to overestimate ending stocks for soybeans.

Table 1: Average forecast errors for USDA 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	7.48 (2.66)	-4.9 (1.23)	1.65 (0.45)	-0.05 (0.01)
Production	-10.7 (18.92)	9.93 (4.5)	3.57 (2.28)	0.07 (0.04)
Imports	3.97 (0.7)	1.53 (0.39)	2.38 (0.53)	0.004 (0.002)
Variables decreasing commodity supply				
Exports	-6.04 (10.2)	-33.51 (4.51)	-3.85 (3.06)	-0.25 (0.06)
Feed & Residual	33.47 (9.34)		11.22 (2.12)	
Food, Seed & Industrial*	-1.07 (2.08)	4.28 (1.29)	3.76 (16.47)	
Ethanol	-9.45 (4.99)			
Crushings		-20.48 (2.09)		
Domestic Use				0.09 (0.02)
Ending Stocks	17.64 (12.68)	-43 (4.11)	17.79 (3.07)	-0.13 (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 A1 to A4 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 mistakes are highest for cotton and wheat.

5. Results

5.1 Corn

Figure 2 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.

Figure 3 presents EGBT results for the relative importance of each balance sheet element forecast to 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 58% to ending stocks errors, while production contributes as much as 40%. Feed & Residual, on the other hand, contribute as much as 28%. Ethanol's importance increases beginning in the early 2000s, rising to 20% of ending stocks forecast errors, on average, during the last three five-year cycles.

Figure 4 provides the SHAP plot for corn, to account for variations in relative importance at the observational level. The vertical axis of the left panel labels the errors in descending order of their relative importance in determining ending stocks forecast errors. SHAP scores in figure 4 attribute the highest importance to export projections, followed by feed & residual and production. These conclusions confirm our findings in figure 3 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 such errors raise ending stocks projection errors between 250-300 million bushels. Averaging over all the SHAP values related to exports produces a value of 72.1 million bushels, 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 4. The nonlinear relationship depicted in the figure indicates that export prediction make a decreasing marginal contribution to ending stocks errors as they grow larger because other variables exert more influence on ending stocks errors.

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 for each country. Figure 5 shows forecast errors on corn imports for each country alongside U.S. corn exports. For most of the sample, USDA underestimated imports for countries around the world, consistent with underestimating U.S. exports. At the country level, USDA generally underestimated exports demand from Europe (especially in the post-2008 period of the sample) and the Rest of the World (ROW) catch-all category (including Columbia).

Figure 6 presents EGBT relative importance results for our international trade analysis. Matching our expectations from figure 5, ROW and European corn import errors are the most important.[6] 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 7 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, China, and Brazil misses are among the most important, although Mexico and Southeast Asian countries also appear in certain time periods.

5.2 Soybeans

Figure 8 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. 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

[6] 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.

directions. For example, in 2016/17 USDA missed very high on production and very low on exports.

Figure 9 presents our EGBT relative contribution results for soybean ending stocks errors over a rolling five-year period. In the figure, production (22% contribution, on average) and export misses (21%, on average) are the lead contributors. Crushings and seed residual errors each contribute 15% to ending stocks errors. Figure 10 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.

Figure 11 documents how USDA's country-level soybean import forecasts relate to its export forecast errors. 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.

Figure 12 presents EGBT relative importance results for international trade. It confirms the graphical intuition of figure 11: that USDA's Chinese import misses are responsible for the bulk of its export misses, especially as the latter grow larger in size 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 (53%), and during the ongoing trade war (48%). 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. Alternatively, figure 13 shows how the errors in the world production estimates relate to U.S. soybeans export misses. Like figure 12, China plays an important role; USDA consistently expected China to produce many more soybeans than it did, leading to underestimates of Chinese imports and U.S. exports. But, likely contributing to other import misses around the globe in figure 12, USDA also missed on production with respect to the world's other large soybean producers: Argentina and Brazil.

5.3 Wheat

Figure 14 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 15 explains how these balance sheet elements contribute to ending stocks misses. Production contributes as much as 74% to ending stocks errors, food & seed contribute as much as 41%, and exports' contribution peaks at 37%. SHAP results in figure 16, confirm the important role played by production, exports and food & seed elements misses to ending stocks errors.

Figure 17 exhibits the relation between USDA's wheat export errors and its country-by-country import projection errors. The chart demonstrates that, historically, USDA tends to underestimate both total and country-level foreign wheat imports. Figure 18 presents EGBT results for how international trade contributes to U.S. export errors. Imports by China (16%, on average), ROW (16%), North Africa (14%) and Brazil (14%) are the most influential. Likewise, figure 19 relates USDA's country-level foreign production forecasts to its U.S. wheat export errors. Clearly, Brazil, China, and North Africa feature in both the production and import stories, but Europe's production misses play a more prominent role than do USDA misses on its forecasted imports.

5.4 Cotton

Figure 20 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, 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 21 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 50% 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. SHAP summary plots in figure 22 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.

Figure 23 documents USDA’s domestic cotton export misses and its country-by-country import forecast errors. Since 2011, USDA mostly underestimated China’s imports, but this hasn’t always weighed down its U.S. export forecasts: note how small its domestic forecast errors were in 2011/12 and 2012/13, even though it underestimated Chinese imports quite significantly. At the same time, it actually overestimated imports by both India and Mexico, contributing to, at times, aggregate U.S. export overestimates. Figure 24 confirms that intuition: USDA import misses on the part of China, India, and Mexico play significant parts in U.S. export misses. Figure 25 explores how U.S. export errors are linked to misses in USDA's country-level foreign production forecasts. Europe, India, and China’s influence is notable, with China growing in importance over the latter half of the period of observation.

6. Conclusion

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 lead to distortive 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 the task. Common approaches to address multicollinearity limit result interpretability on the part of the researcher, so we address this problem by using EGBTs, and we demonstrate their value using a simulation analysis.

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 (the lone exception being cotton); their importance in generating ending stocks misses falls as export levels get very large. Because U.S. exports are clearly linked to foreign production and import demand, we also use EGBTs to explore the relationships between those forecast misses. Our findings document the role of large international buyers and competitors, and some links between country-by-country imports and production. For almost every commodity, USDA forecasts for China's imports and production play an important role, 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. We conclude that, ultimately, improvements in production forecasts would 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. Another source of forecast errors is documents in the literature; 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 making fully efficient forecasts. In ongoing work, we test USDA forecasts for the presence of information rigidities; if they are present, technological gains like those we discuss above would be helpful solutions.

References

Adjemian, M.K. 2012. "Quantifying the WASDE Announcement Effect." *American Journal of Agricultural Economics* 94 (1): 238–56. Accessed at: https://academic.oup.com/ajae/article-abstract/94/1/238/67362?redirectedFrom=fulltext

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: https://prism.ucalgary.ca/handle/1880/50265

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

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://arxiv.org/abs/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

Deng, S., C. Wang, M. Wang, and Z. Shun. 2019. "A gradient boosting decision tree approach for insider trading identification: An empirical model evaluation of China stock market." *Applied Soft Computing* 83. Accessed at: https://doi.org/10.1016/j.asoc.2019.105652

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

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?seq=1#metadata_info_tab_contents

Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008. "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., 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

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

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: https://ieeexplore.ieee.org/document/8767915

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://arxiv.org/abs/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://arxiv.org/abs/1705.07874

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

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://link.springer.com/article/10.1007/s11294-009-9228-5

Pan, Y., D. Liu, and L. Deng. 2017. "Accurate prediction of functional effects for variants by combining gradient tree boosting with optimal neighborhood properties." *PLoS One* 12(6). Accessed at: https://doi.org/10.1371/journal.pone.0179314

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

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., 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://academic.oup.com/ajae/article/101/4/1228/5485173

"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

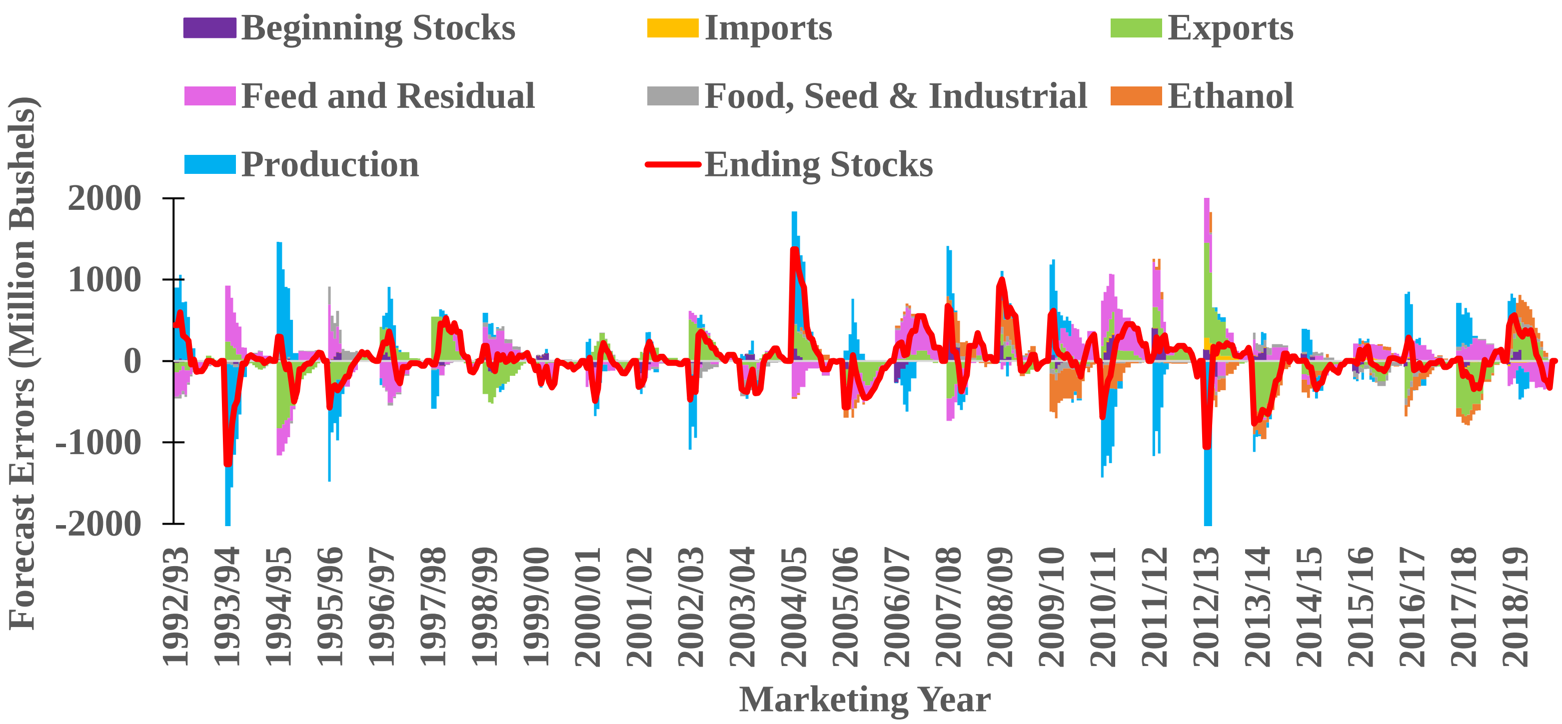

Figure 2: USDA Corn Balance Sheet Element Forecast Errors, 1992-2019

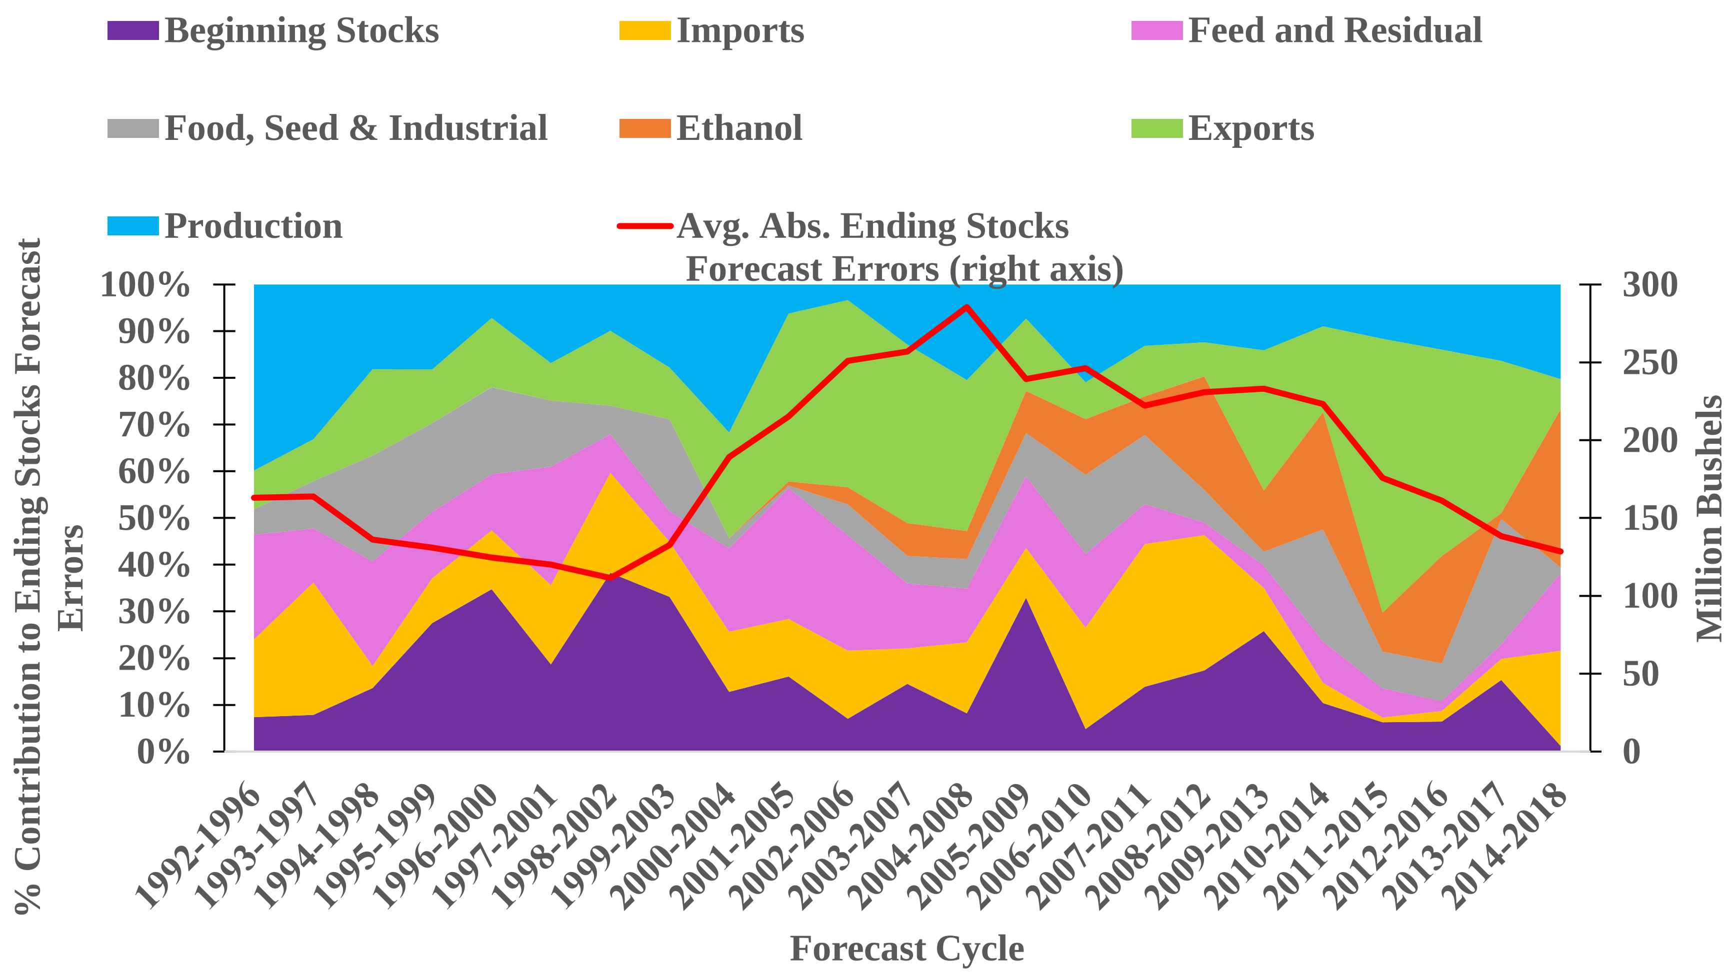

Figure 3: Contribution of Balance Sheet Elements to USDA's Corn Ending Stock Forecast Errors

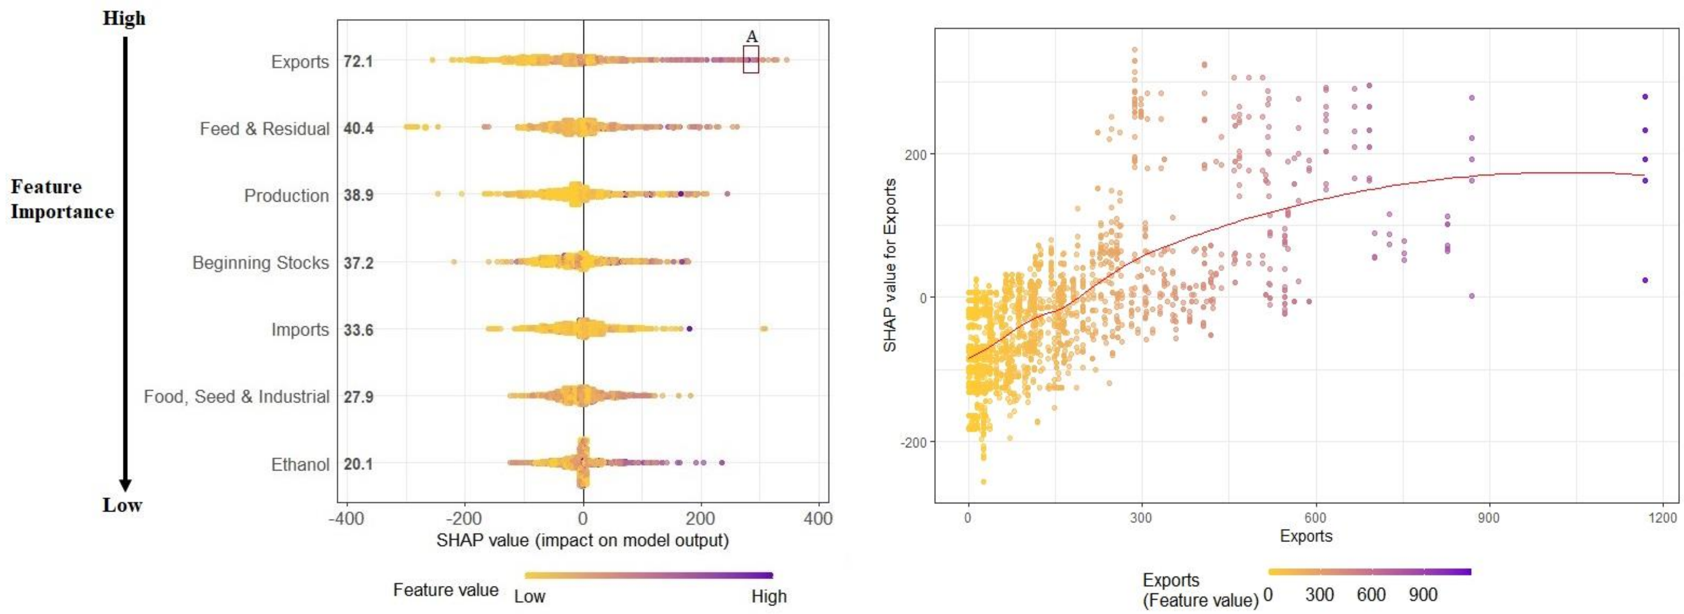

Figure 4: Corn – Exports SHAP Dependence Plot

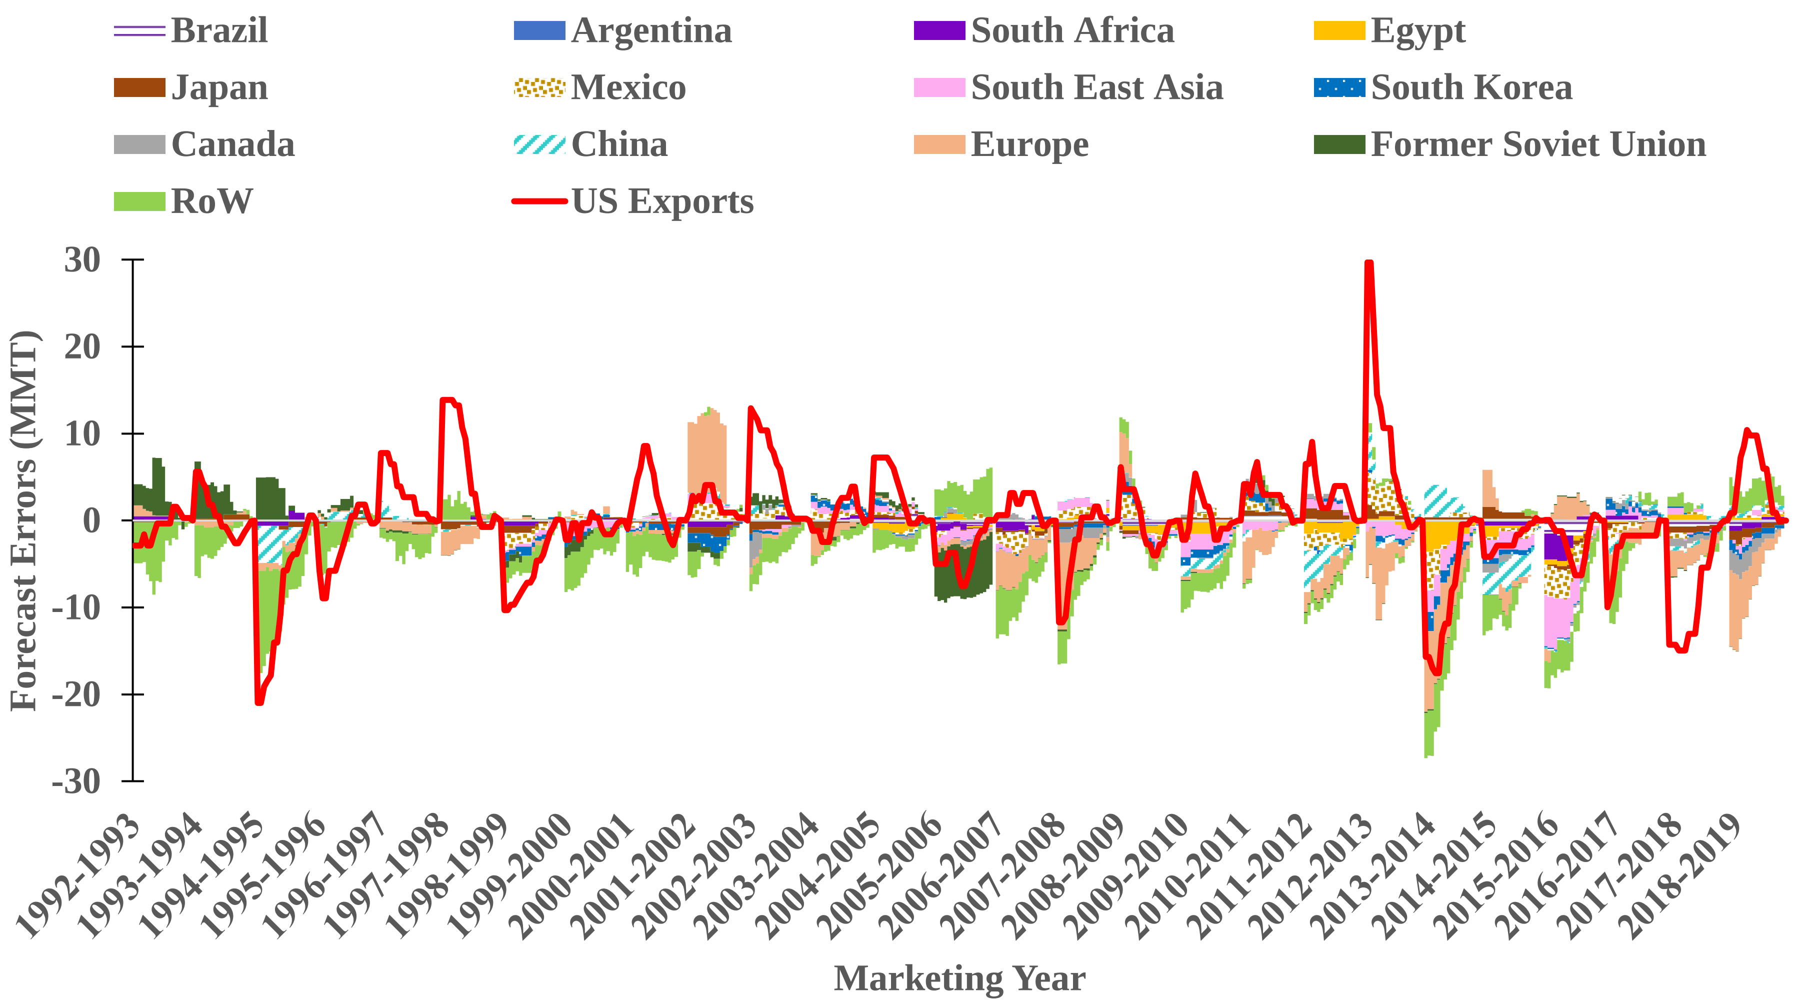

Figure 5: USDA Corn World Import/ U.S. Export Forecast Errors, 1992-2019

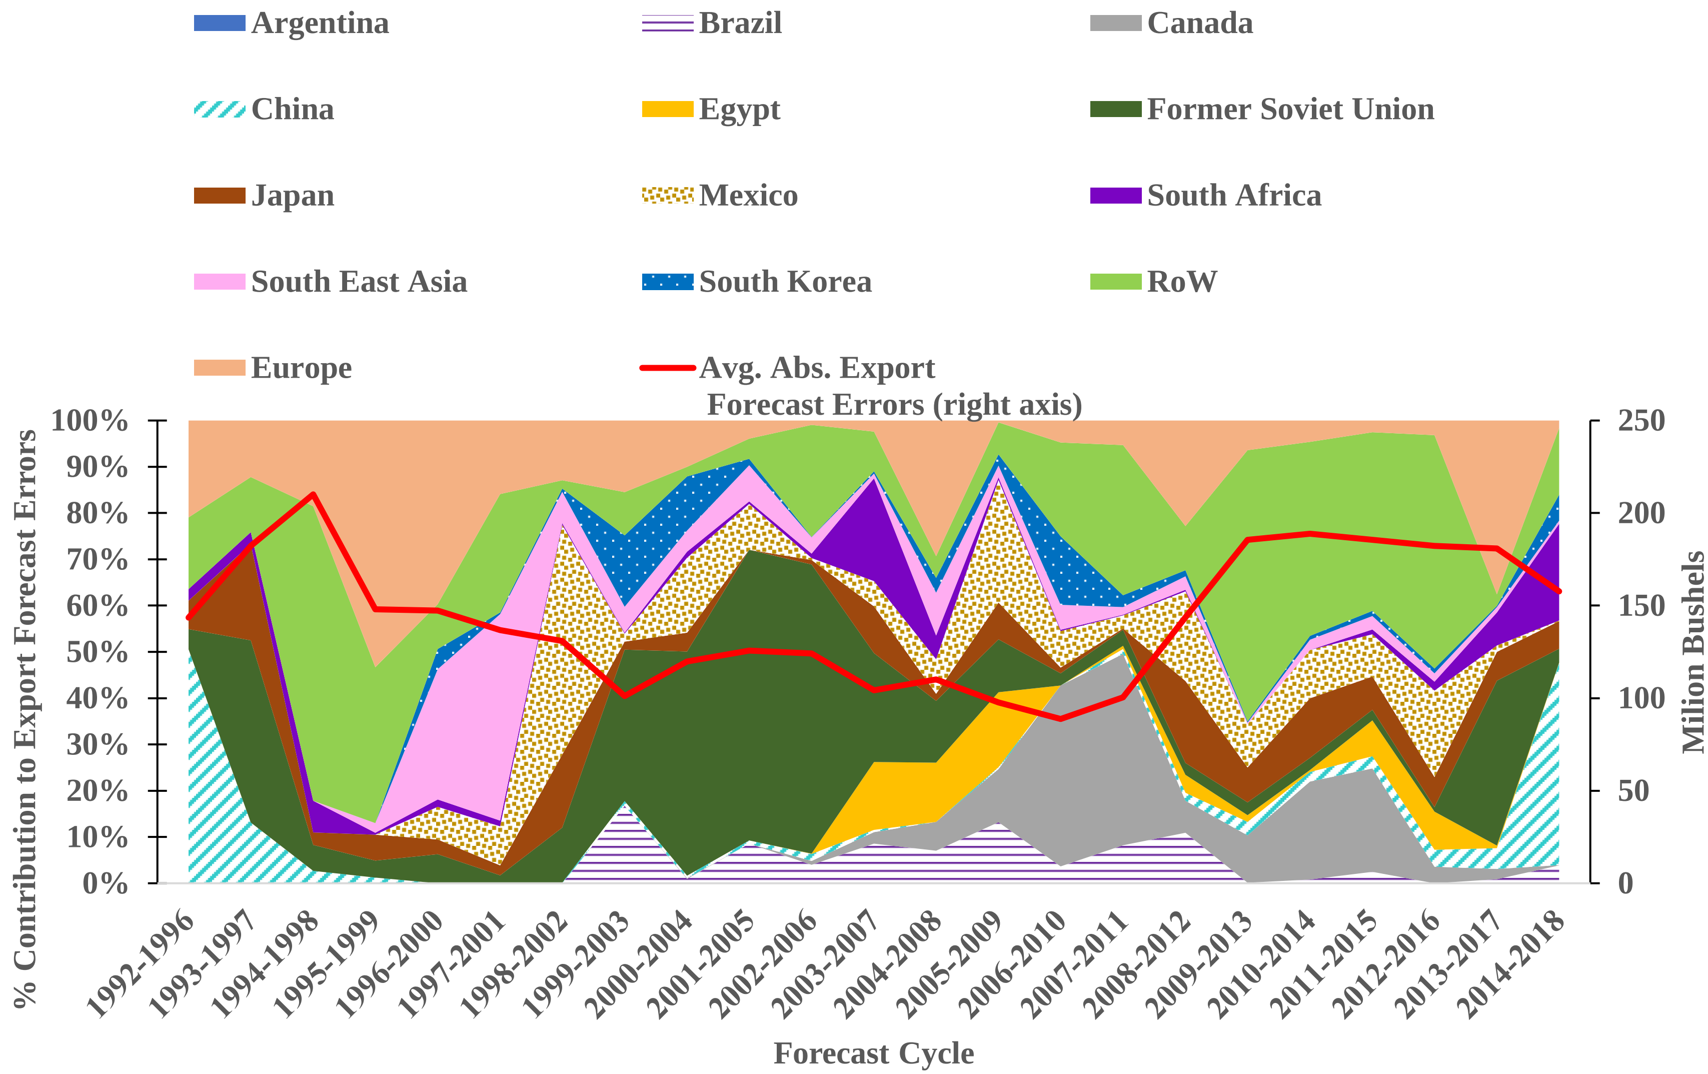

Figure 6: Contribution of World Corn Import Errors to U.S. Corn Export Projection Errors

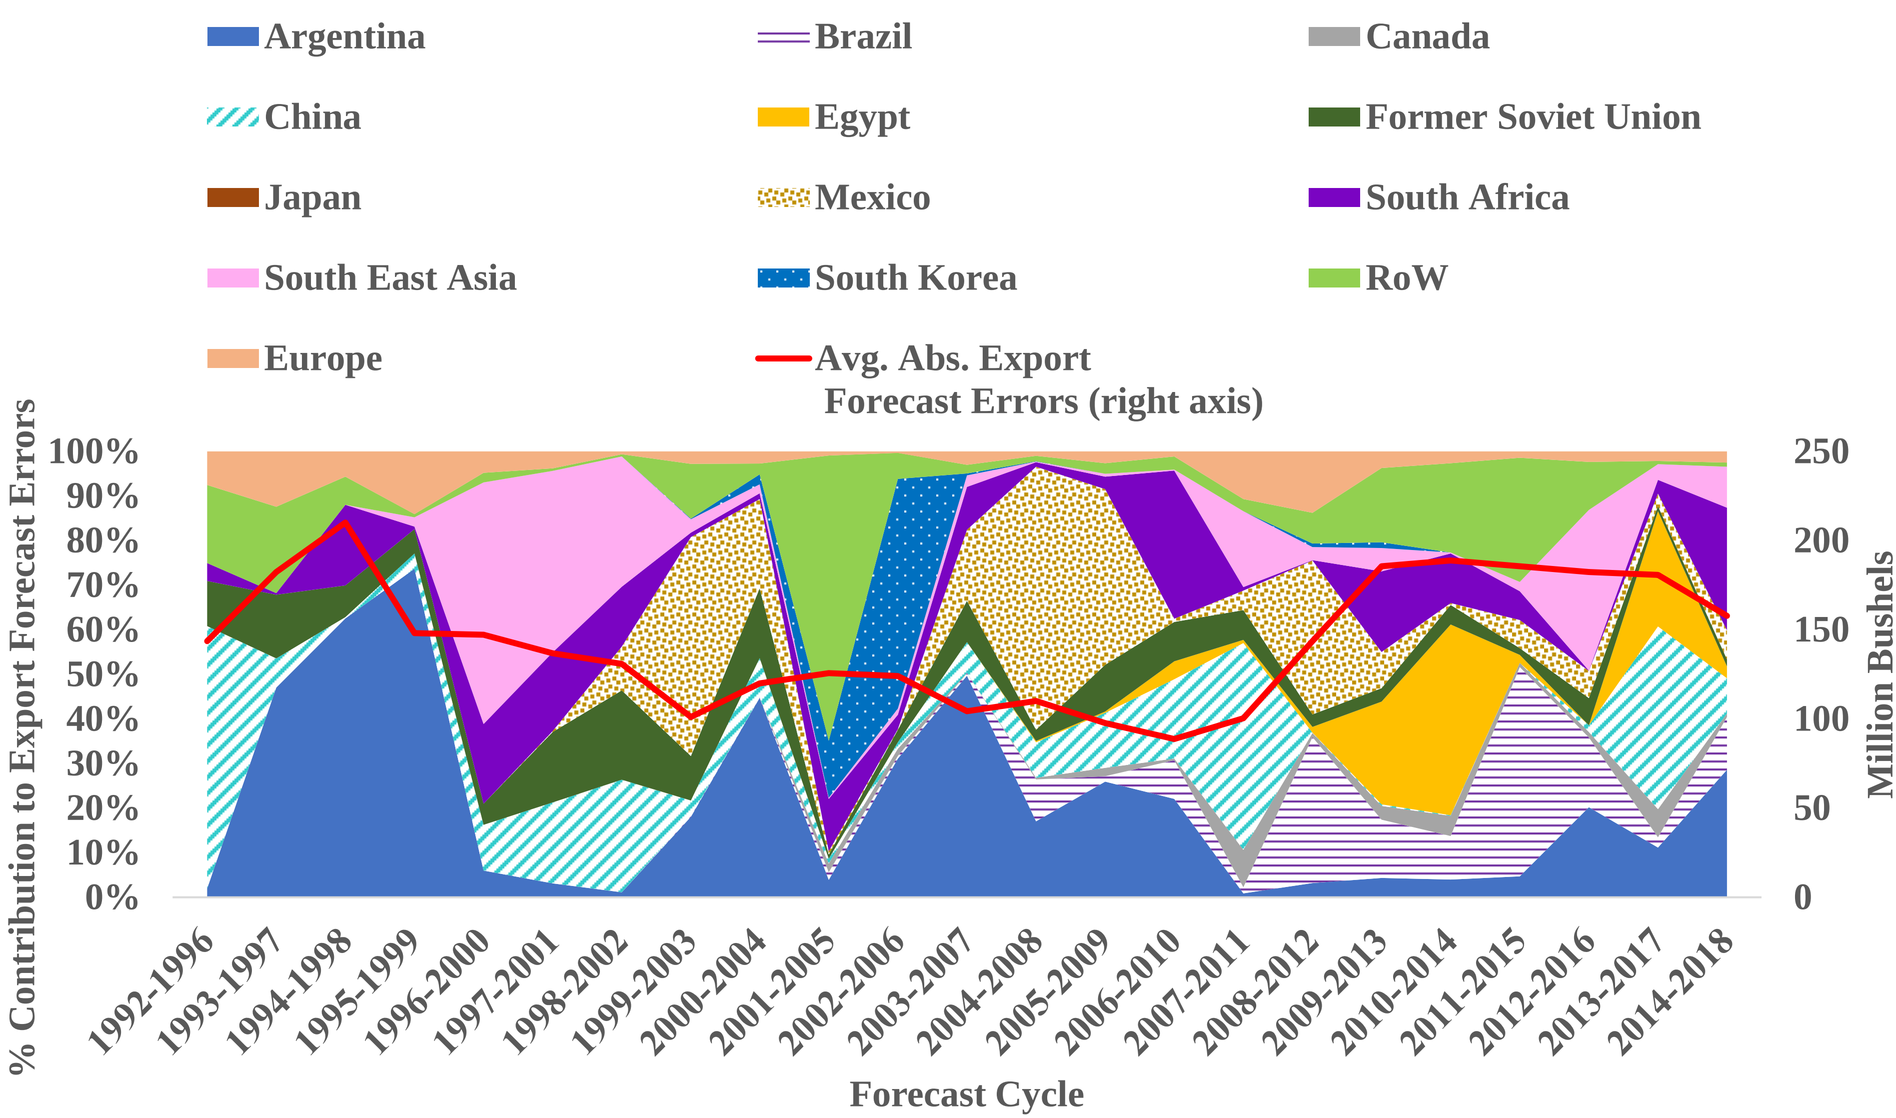

Figure 7: Contribution of World Corn Production Errors to U.S. Corn Export Errors

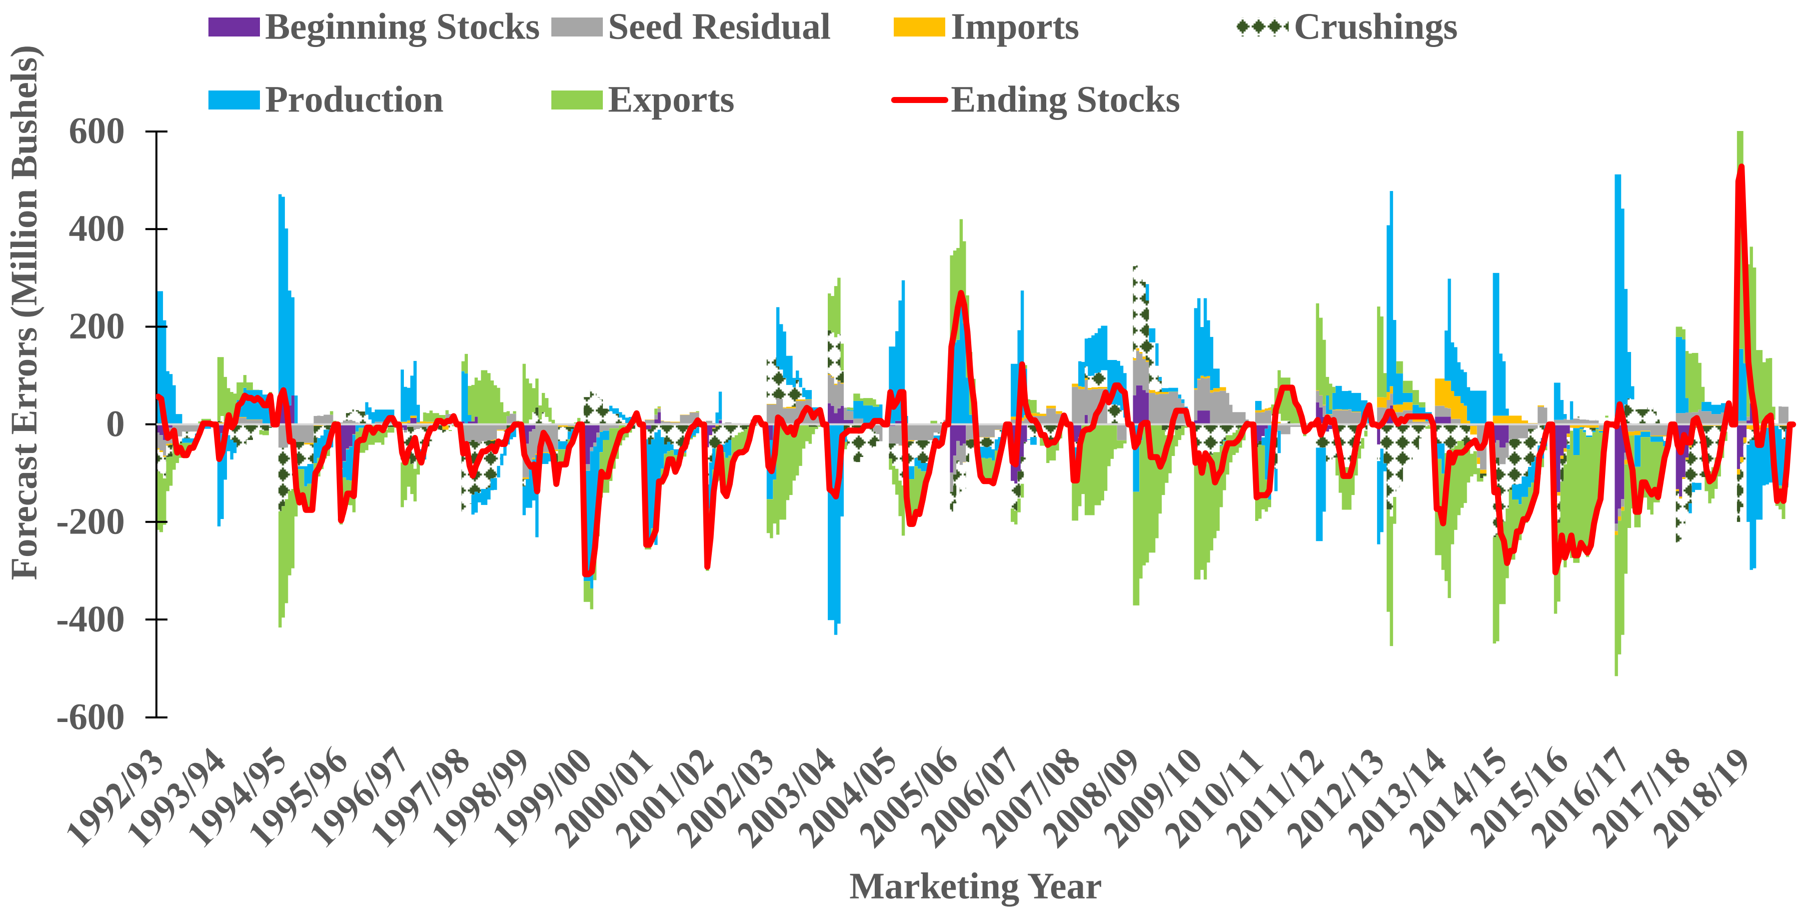

Figure 8: USDA Soybeans Balance Sheet Element Forecast Errors, 1992-2019

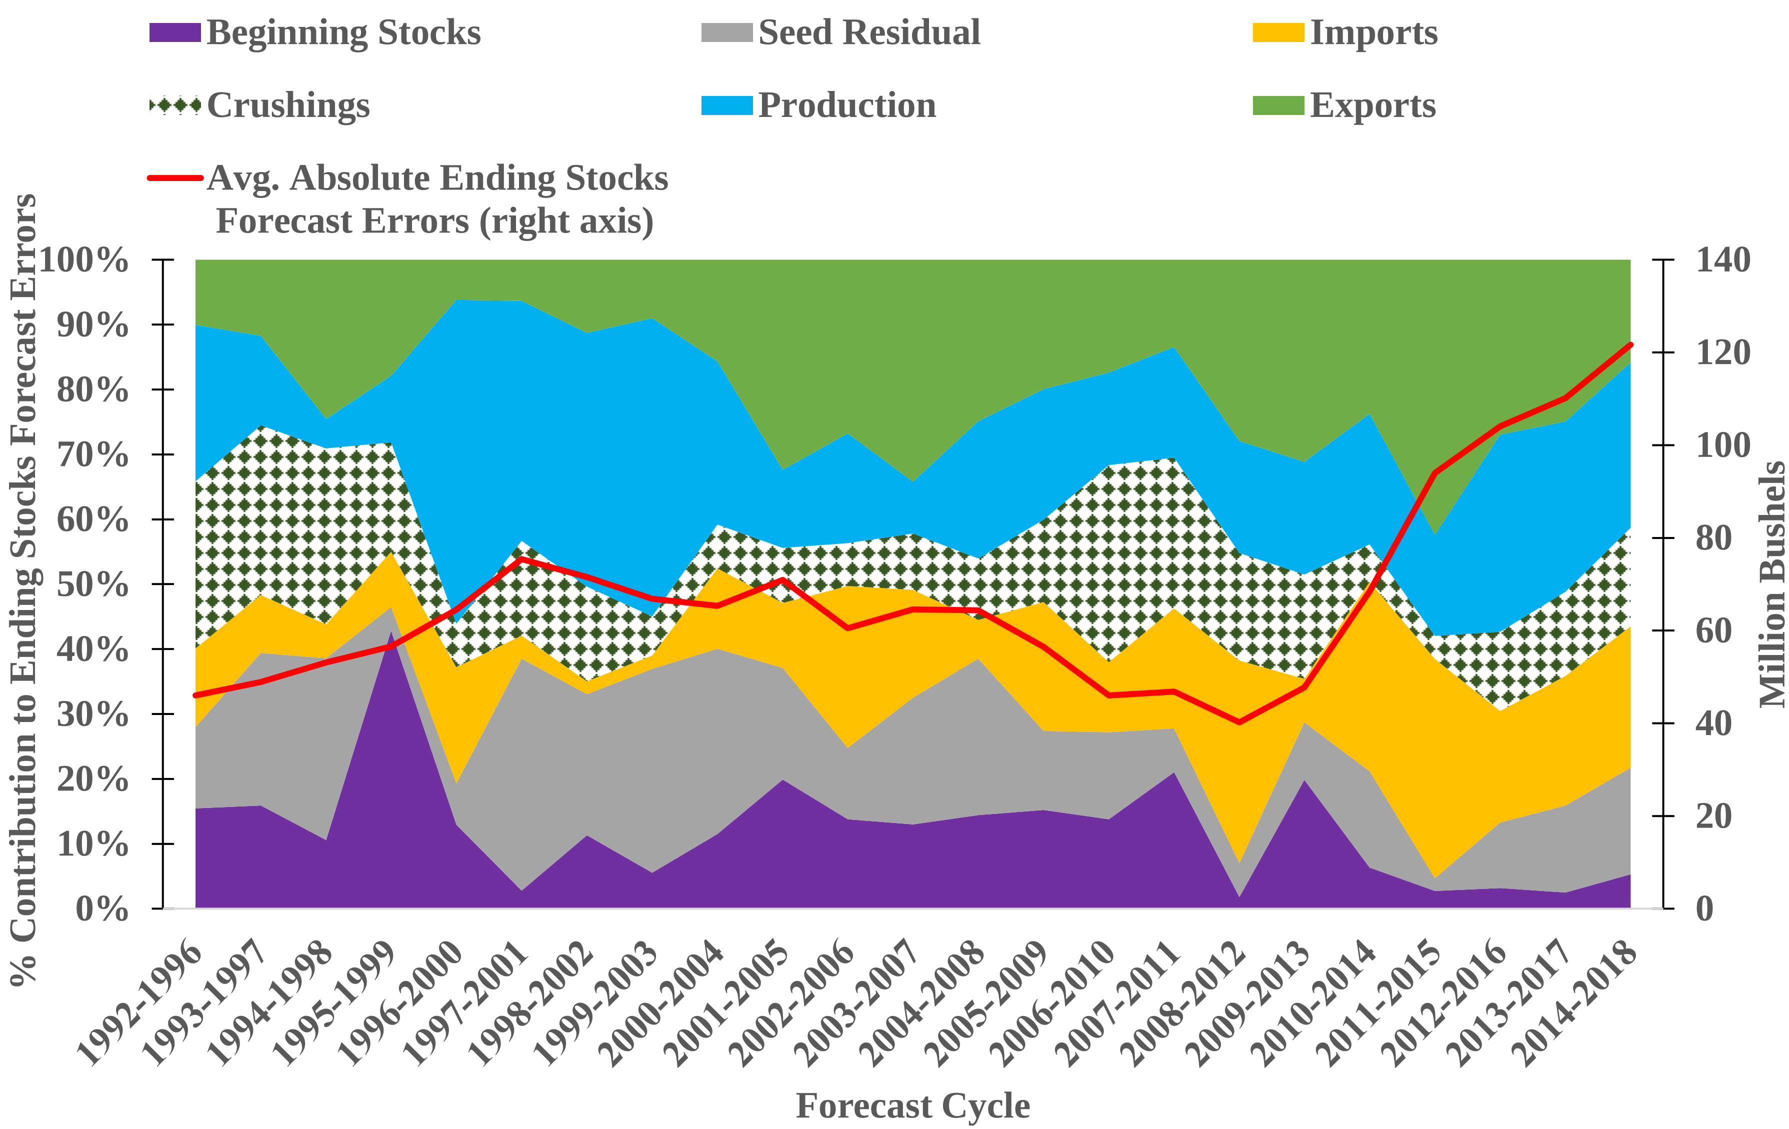

Figure 9: Contribution of Balance Sheet Elements to USDA's Soybeans Ending Stock Forecast Errors

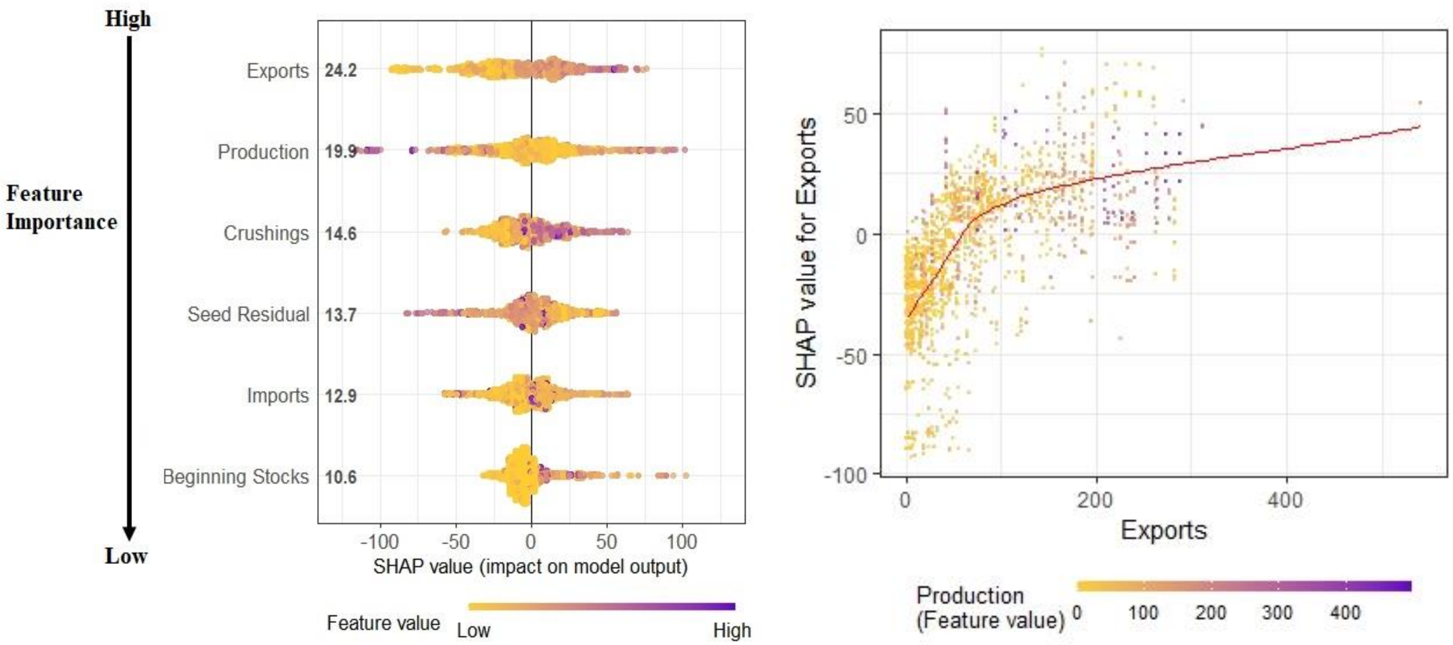

Figure 10: Soybeans – Exports SHAP Dependence Plot

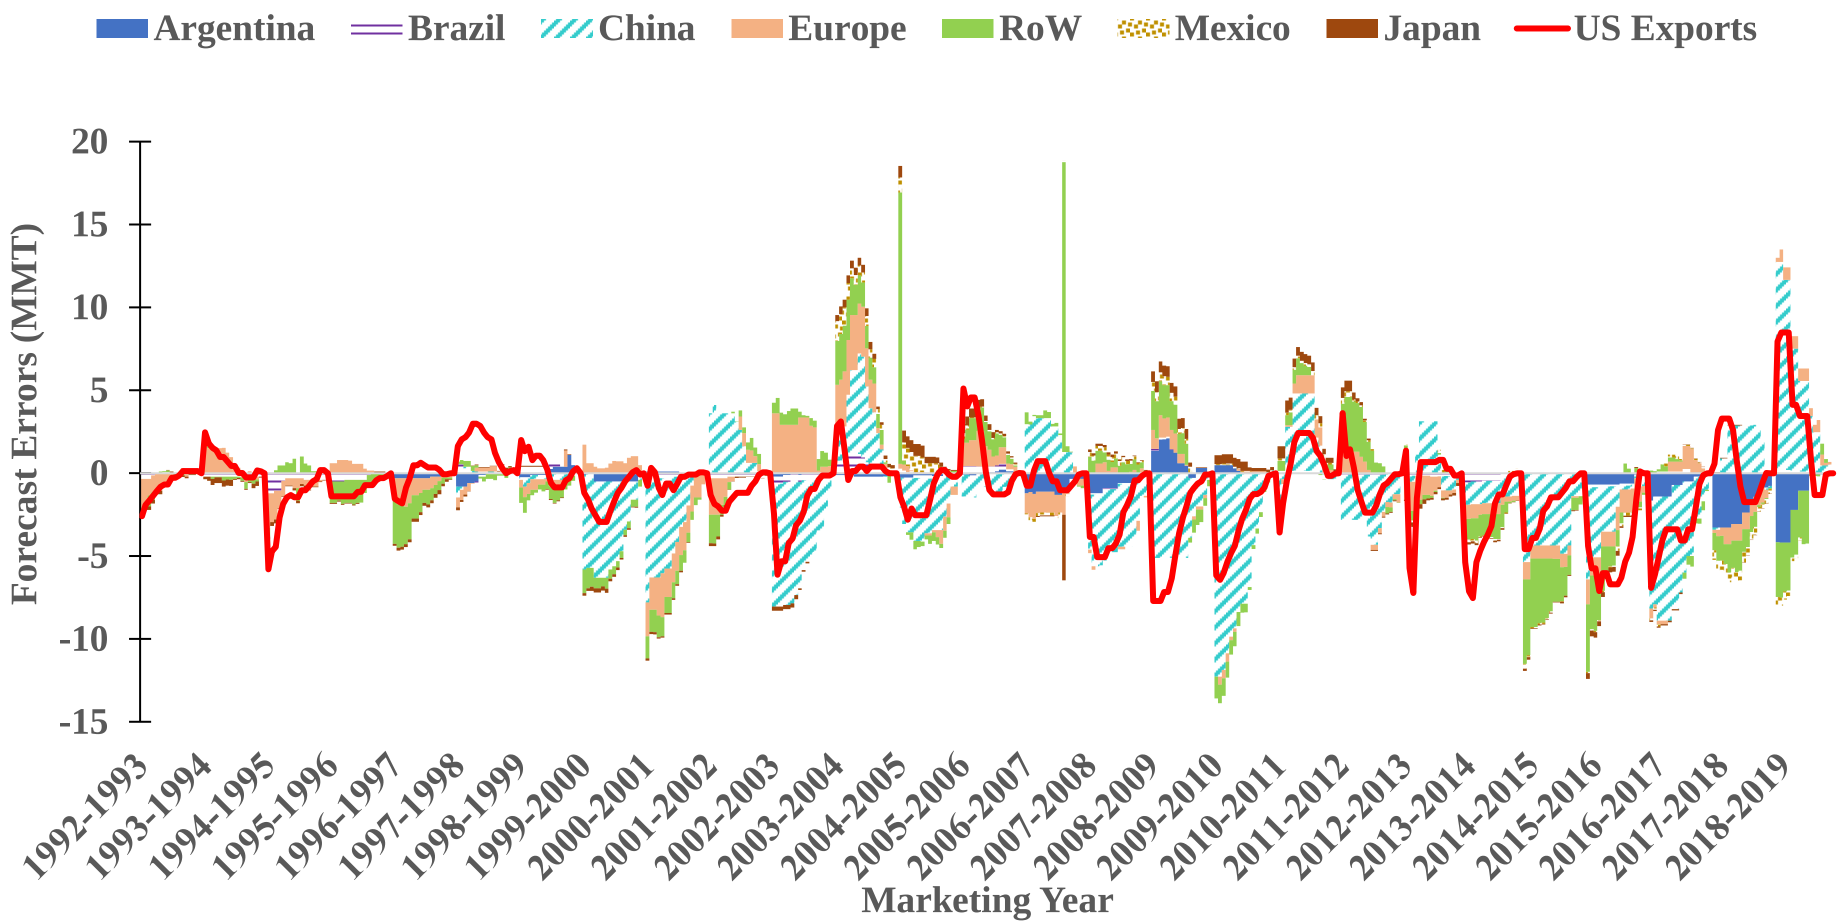

Figure 11: USDA Soybeans World Import/ U.S. Export Forecast Errors, 1992-2019

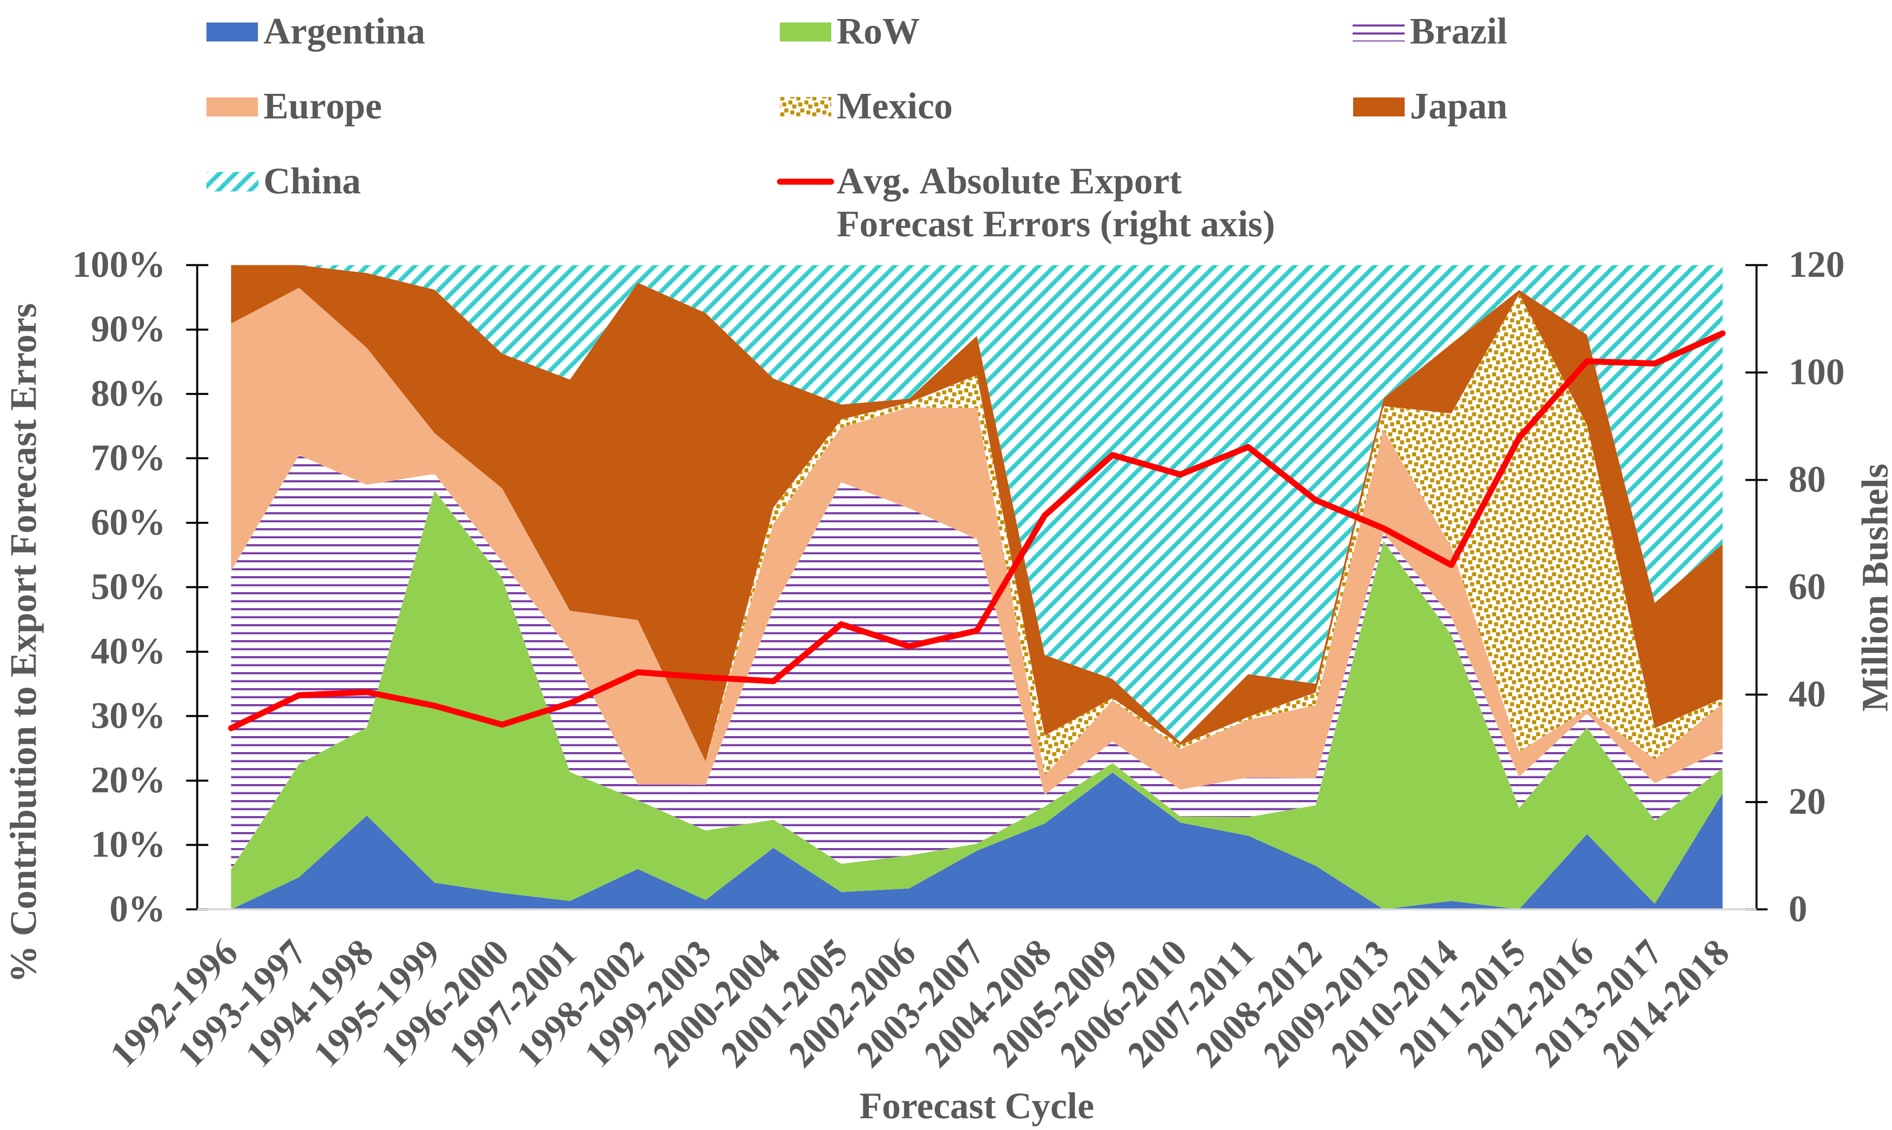

Figure 12: Contribution of World Soybeans Import Errors to U.S. Soybeans Export Projection Errors

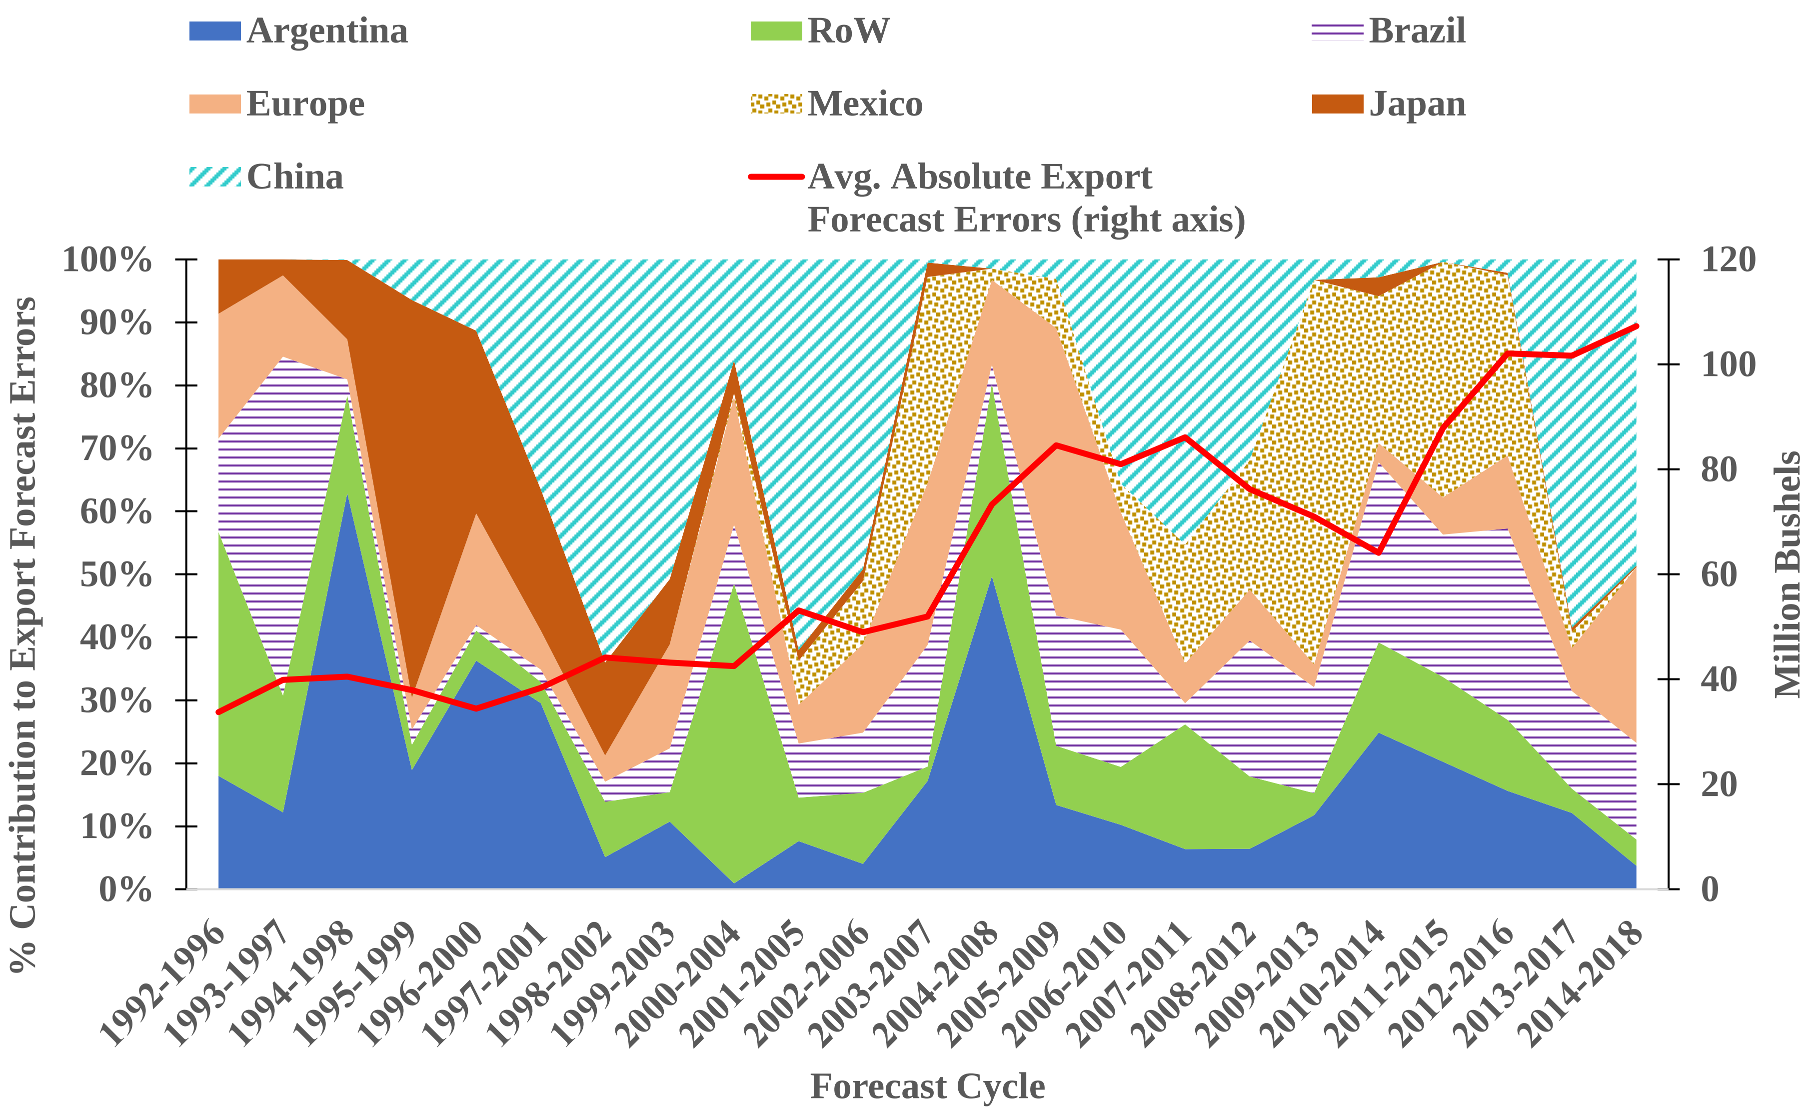

Figure 13: Contribution of World Soybeans Production Errors to U.S. Soybeans Export Projection Errors

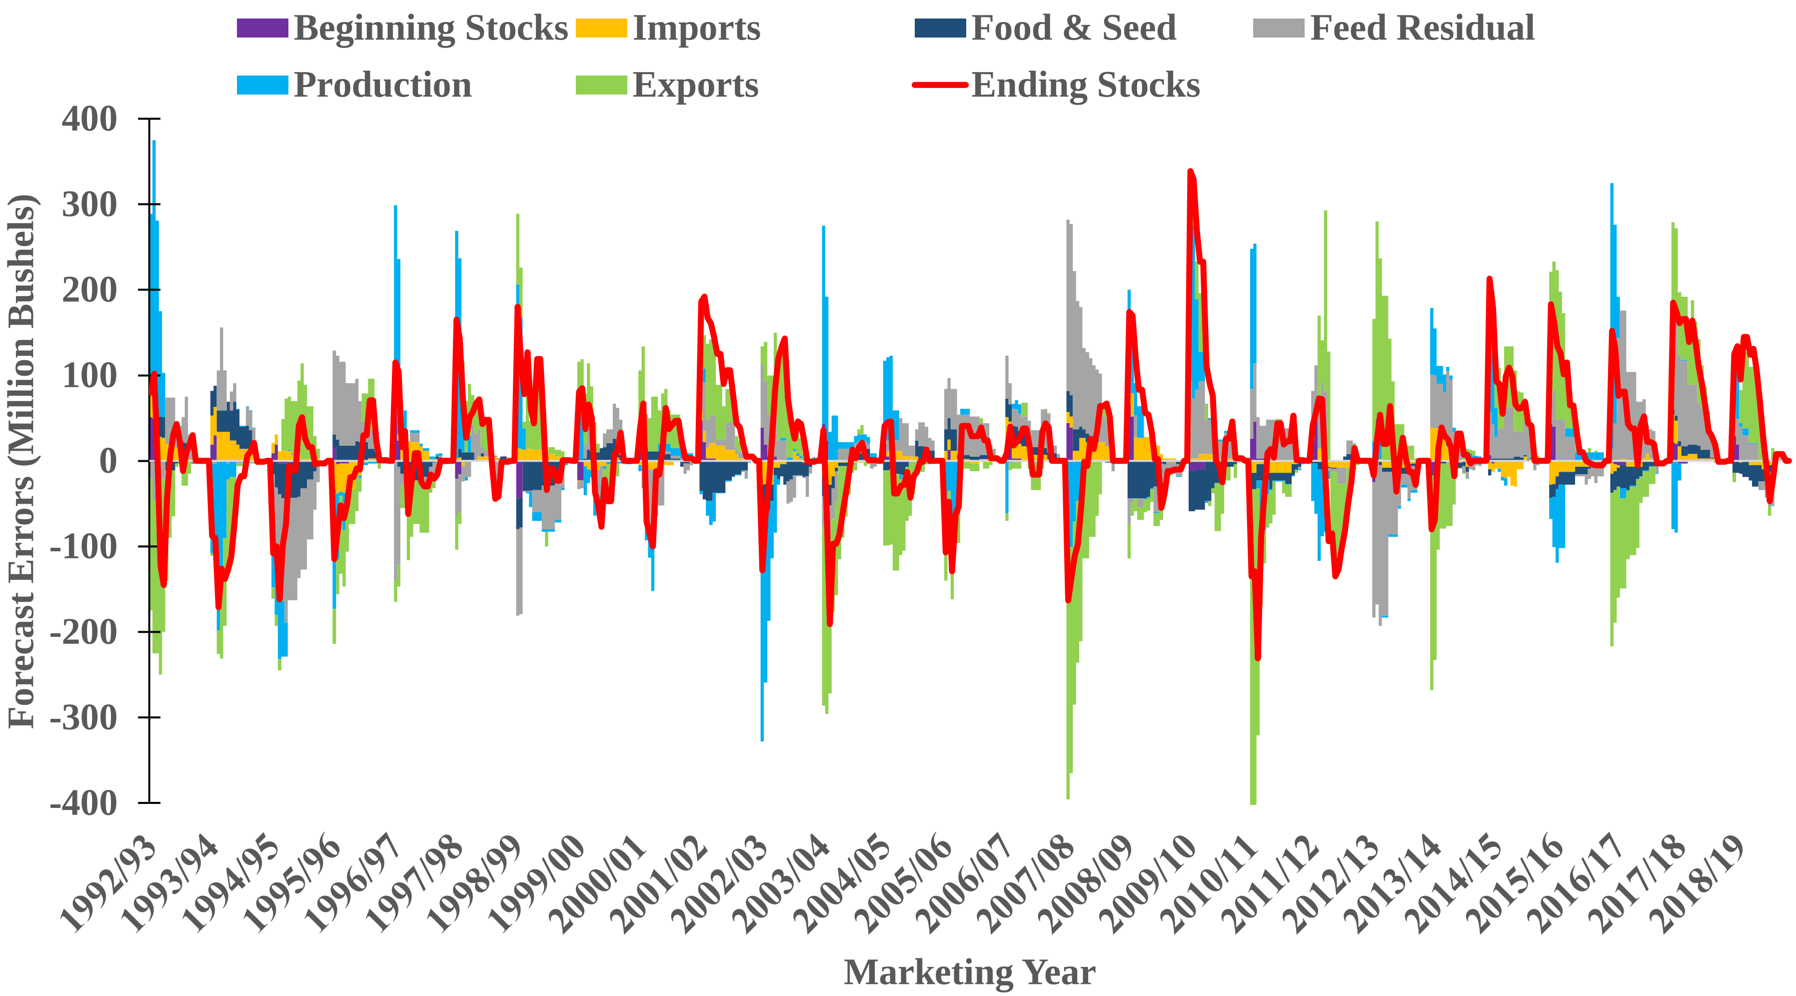

Figure 14: USDA Wheat Balance Sheet Element Forecast Errors, 1992-2019

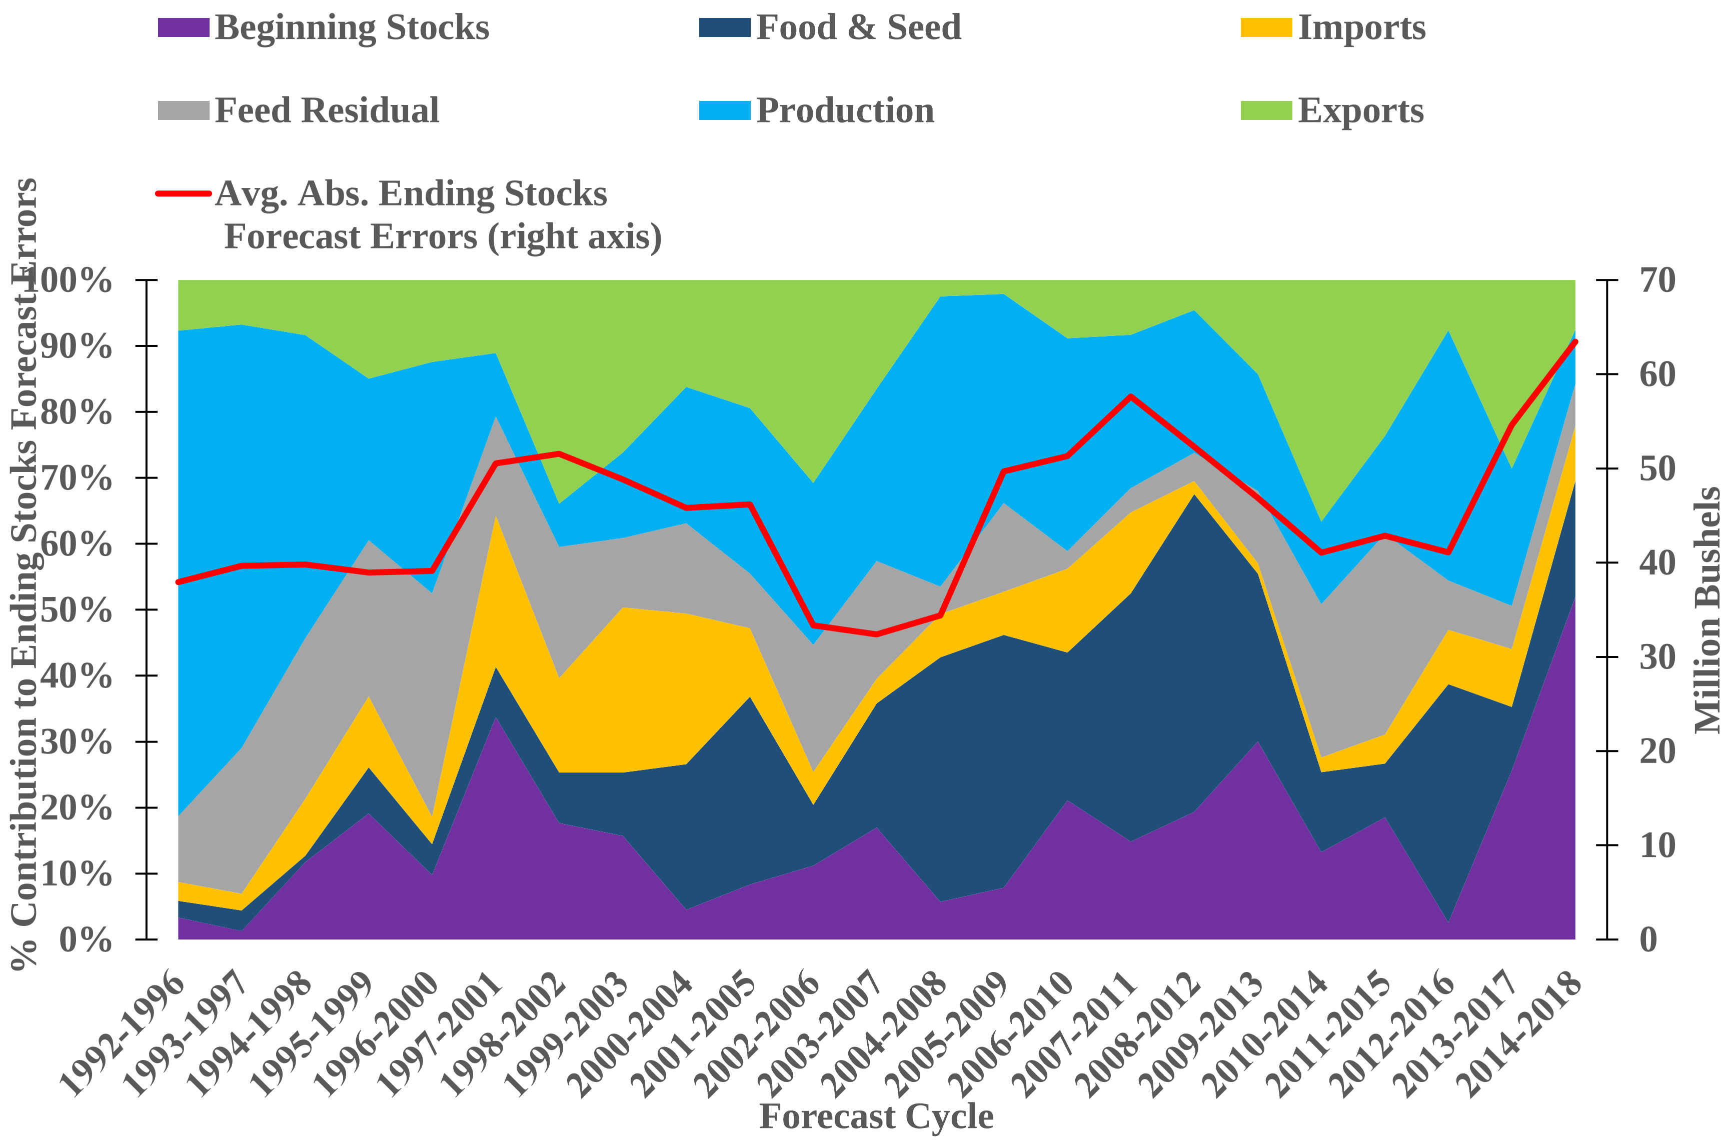

Figure 15: Contribution of Balance Sheet Elements to USDA's Wheat Ending Stock Forecast Errors

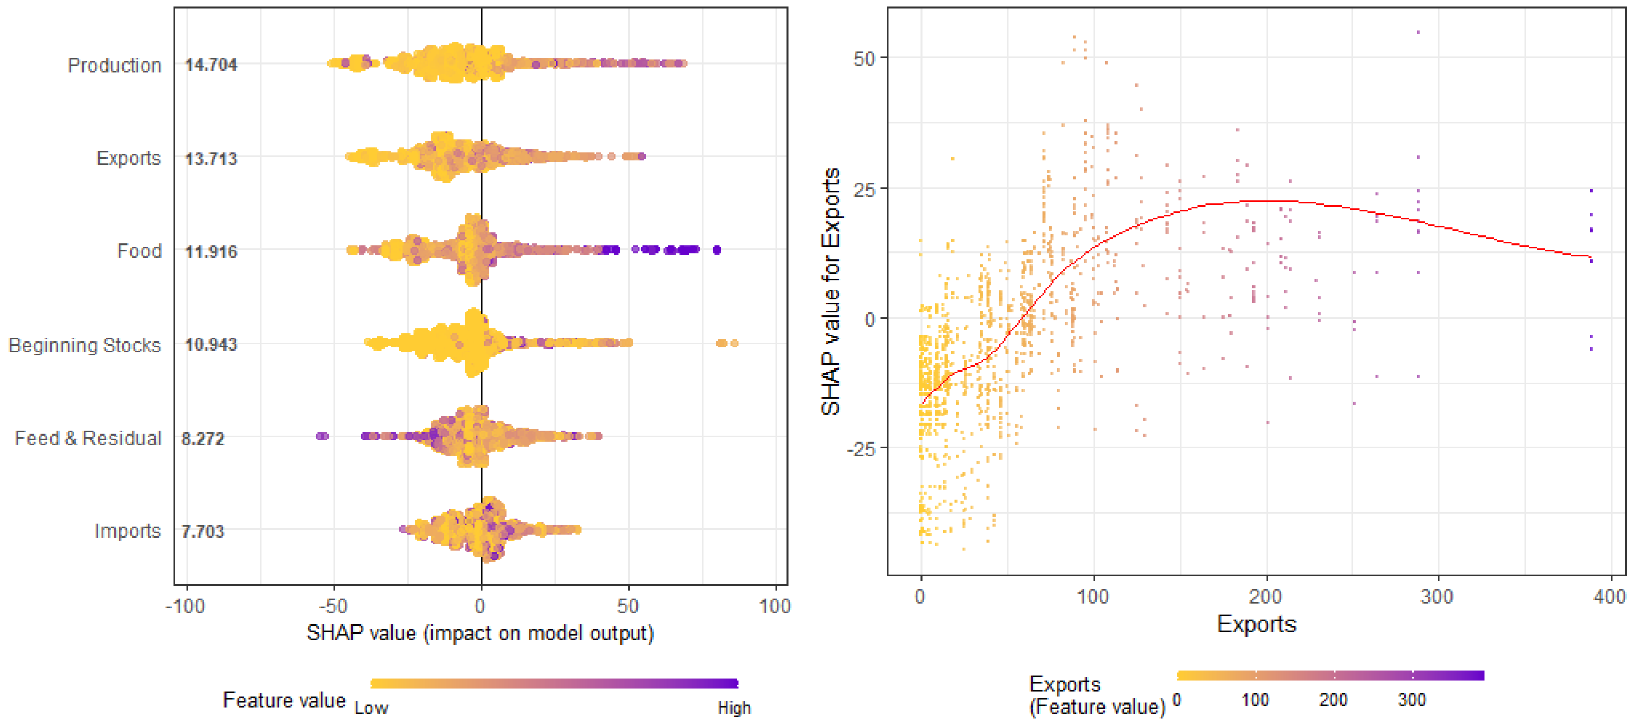

Figure 16: Wheat – Exports SHAP Dependence Plot

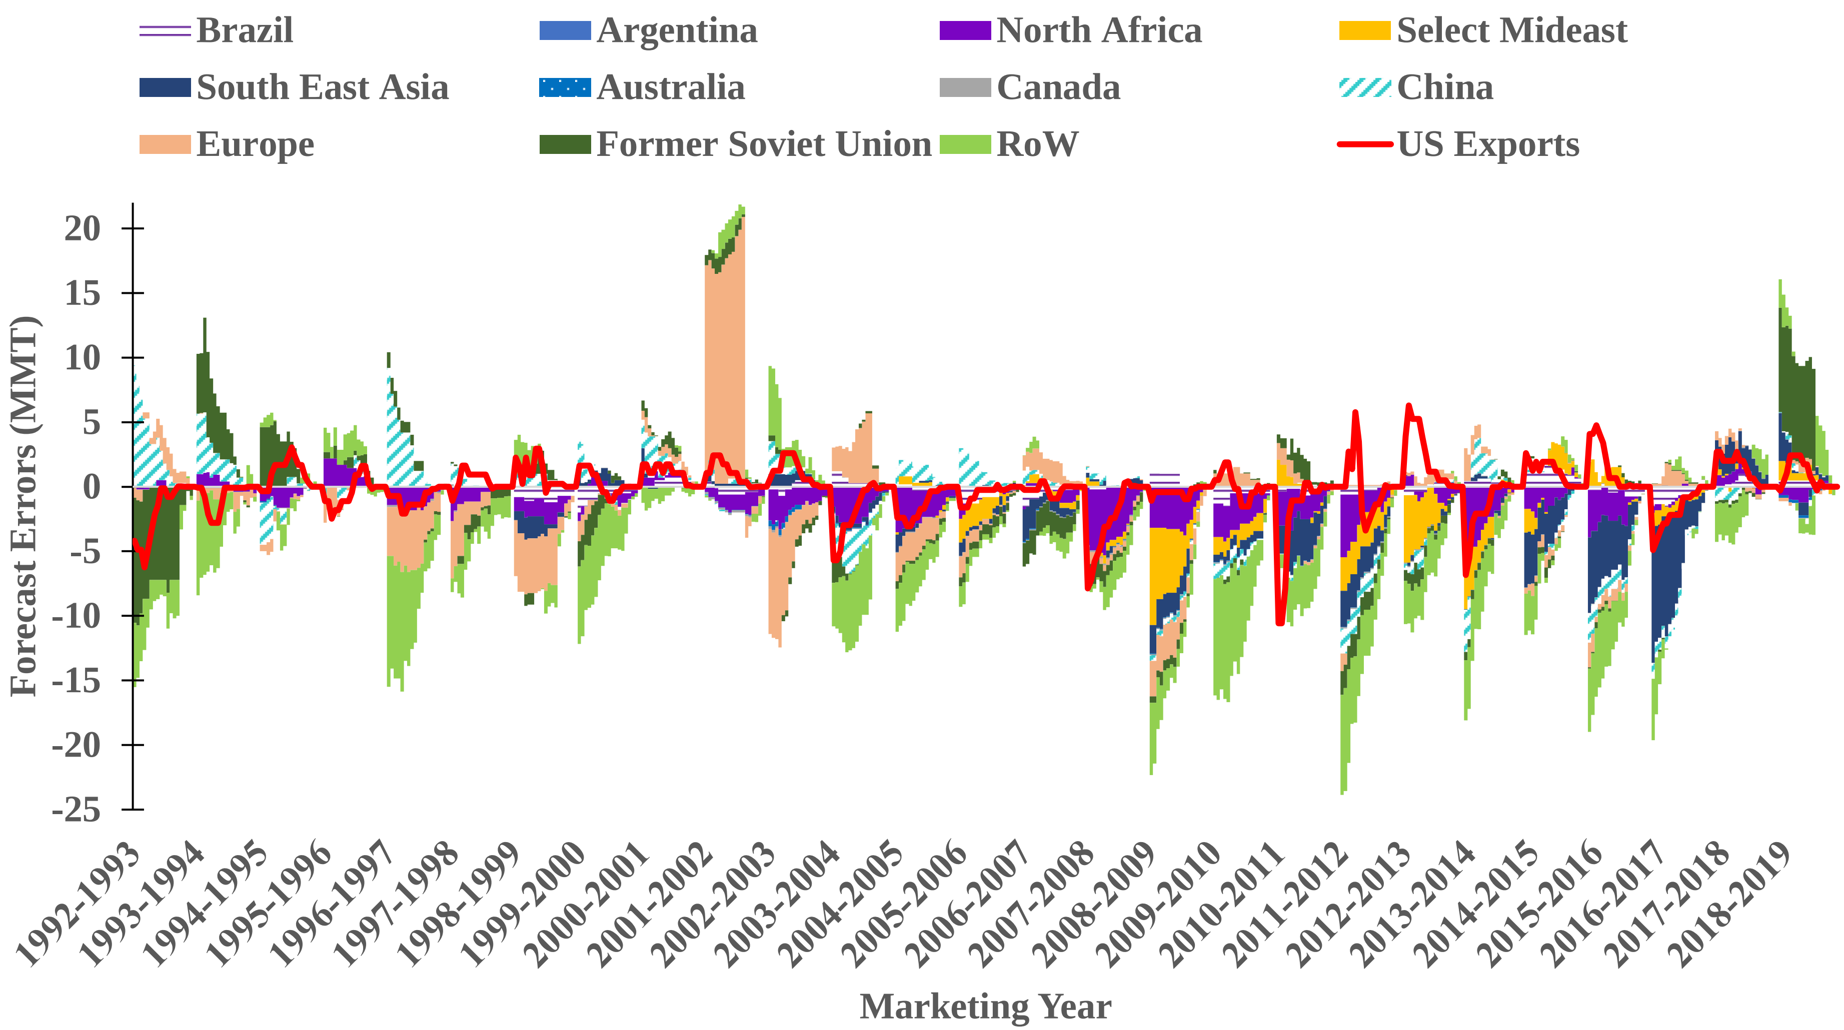

Figure 17: USDA Wheat World Import/ U.S. Export Forecast Errors, 1992-2019

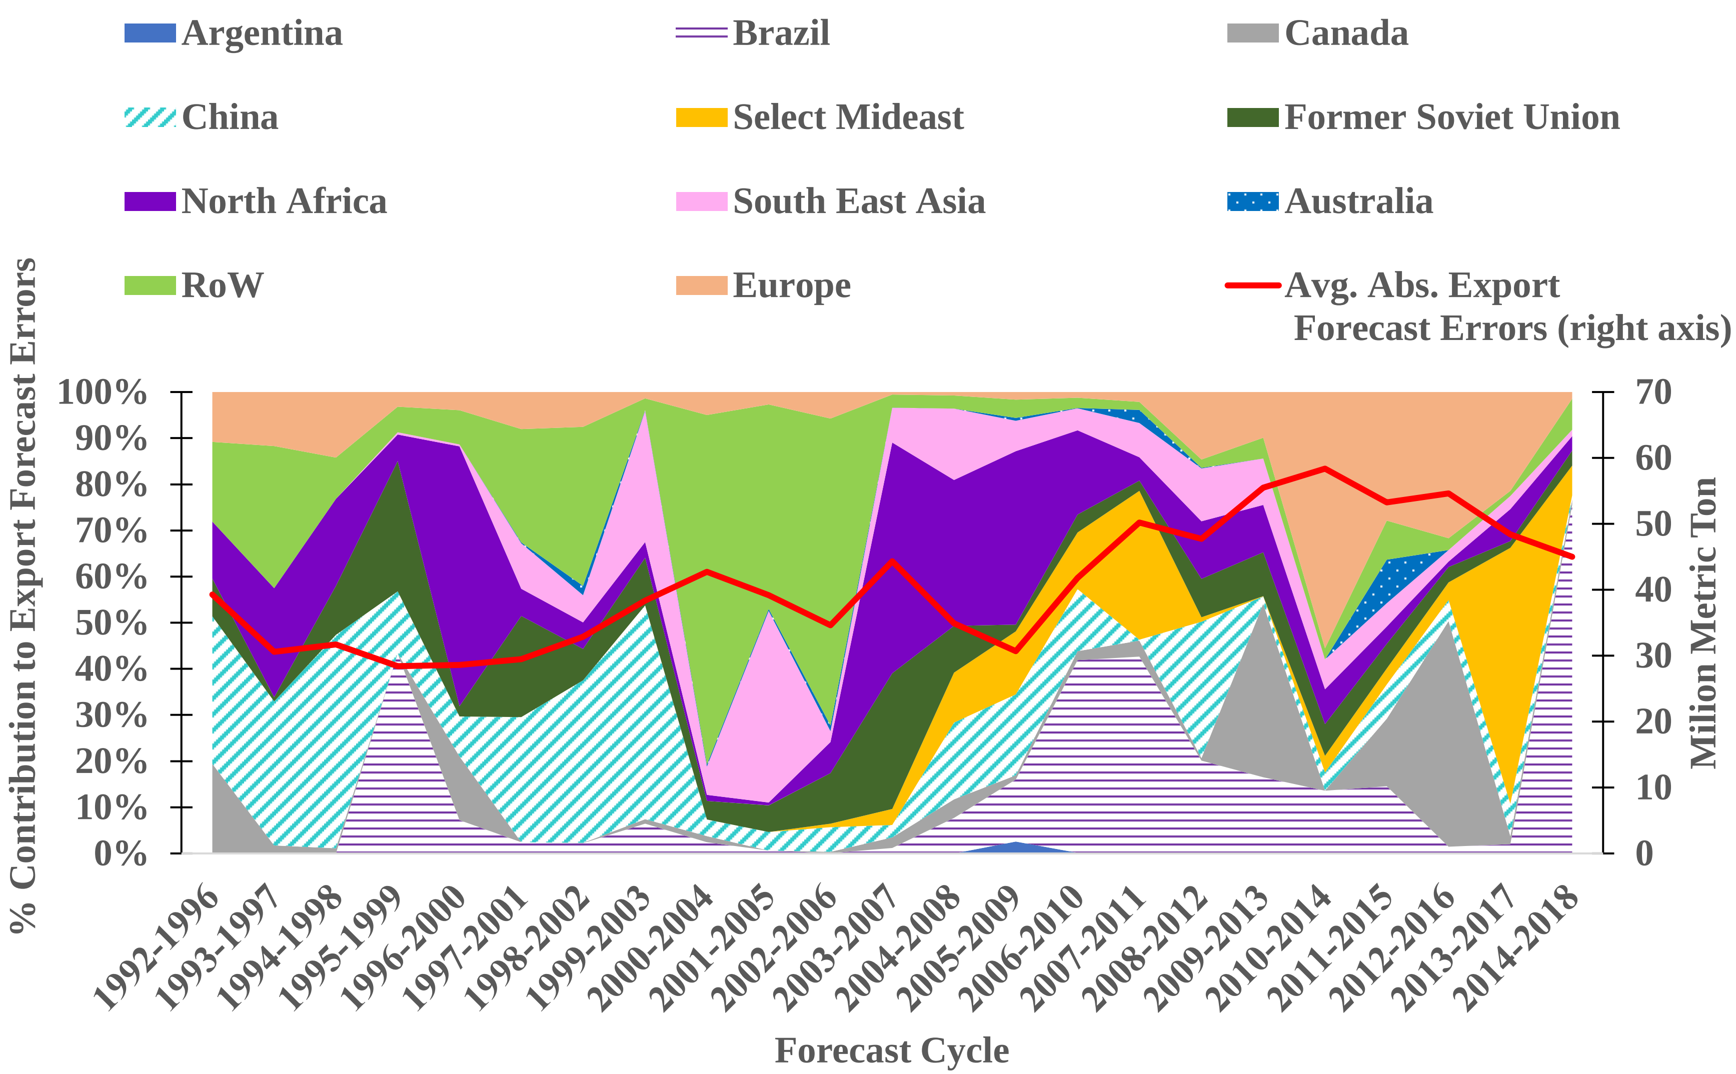

Figure 18: Contribution of World Wheat Import Errors to U.S. Wheat Export Projection Errors

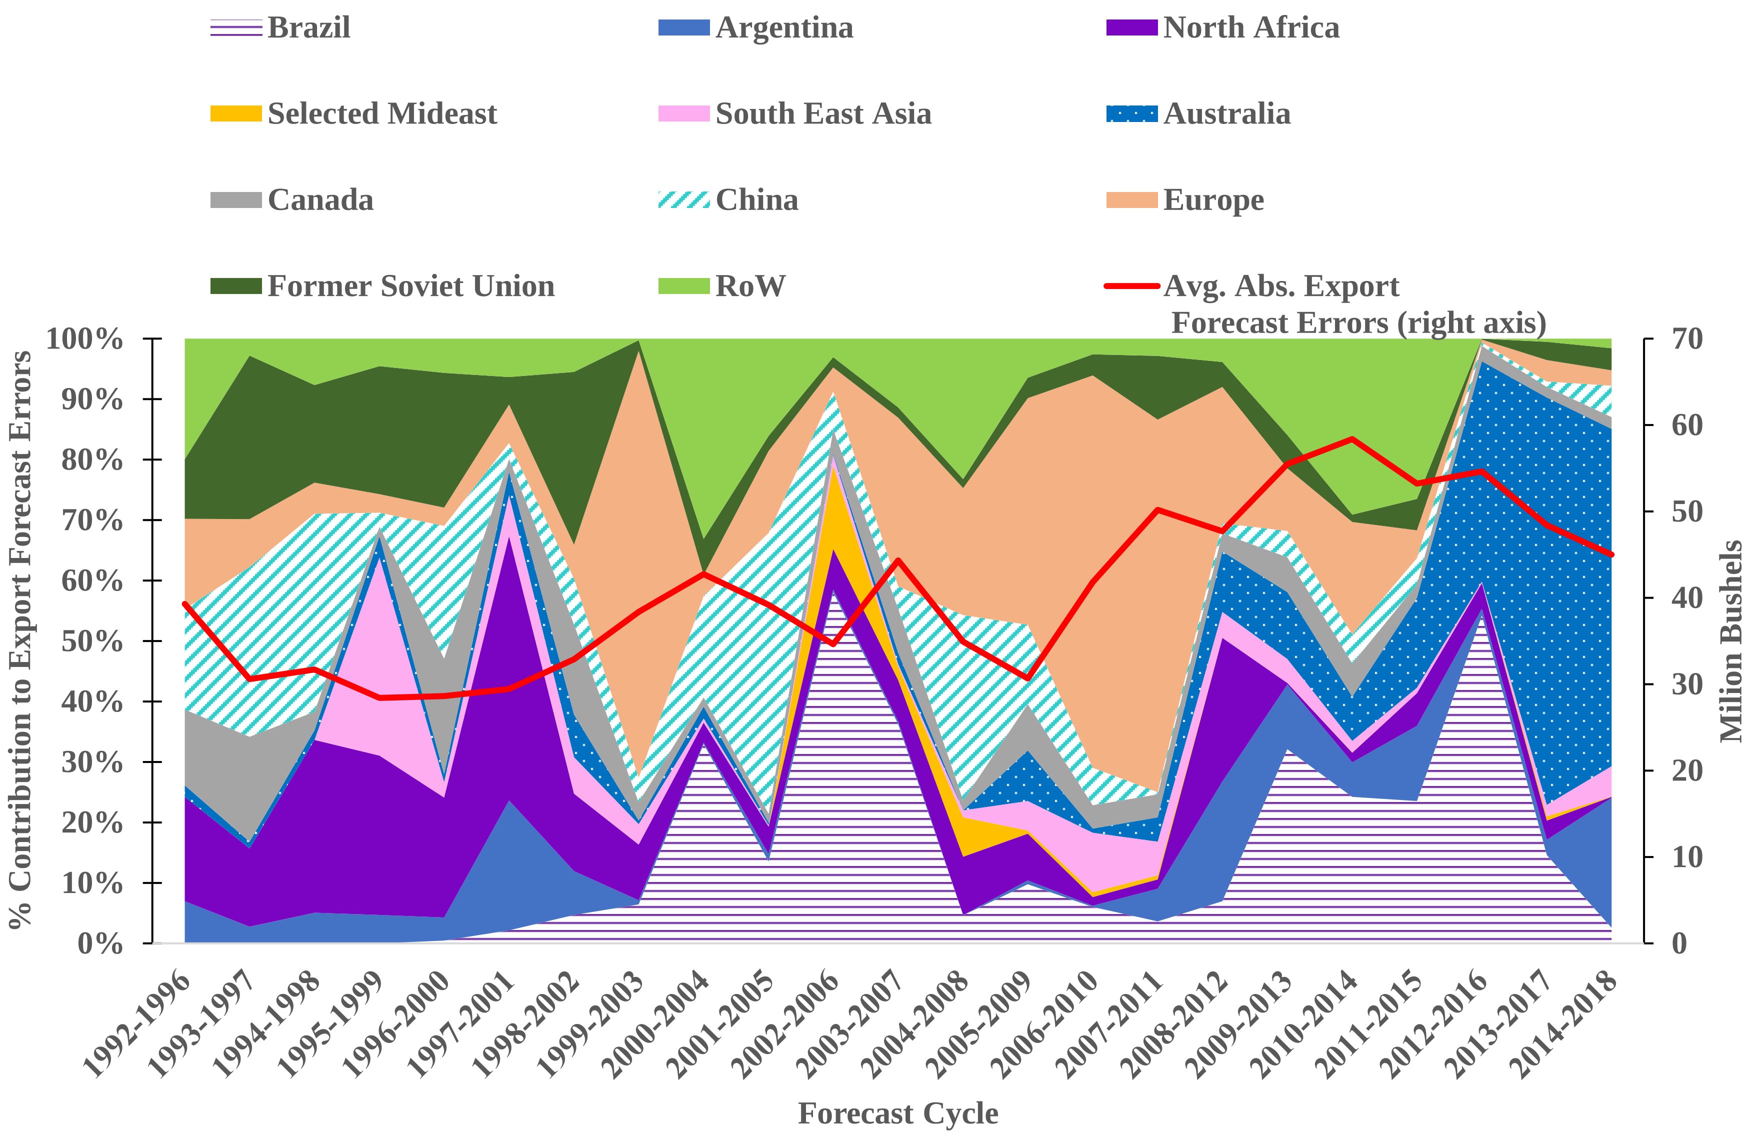

Figure 19: Contribution of World Wheat Production Errors to U.S. Wheat Export Projection Errors

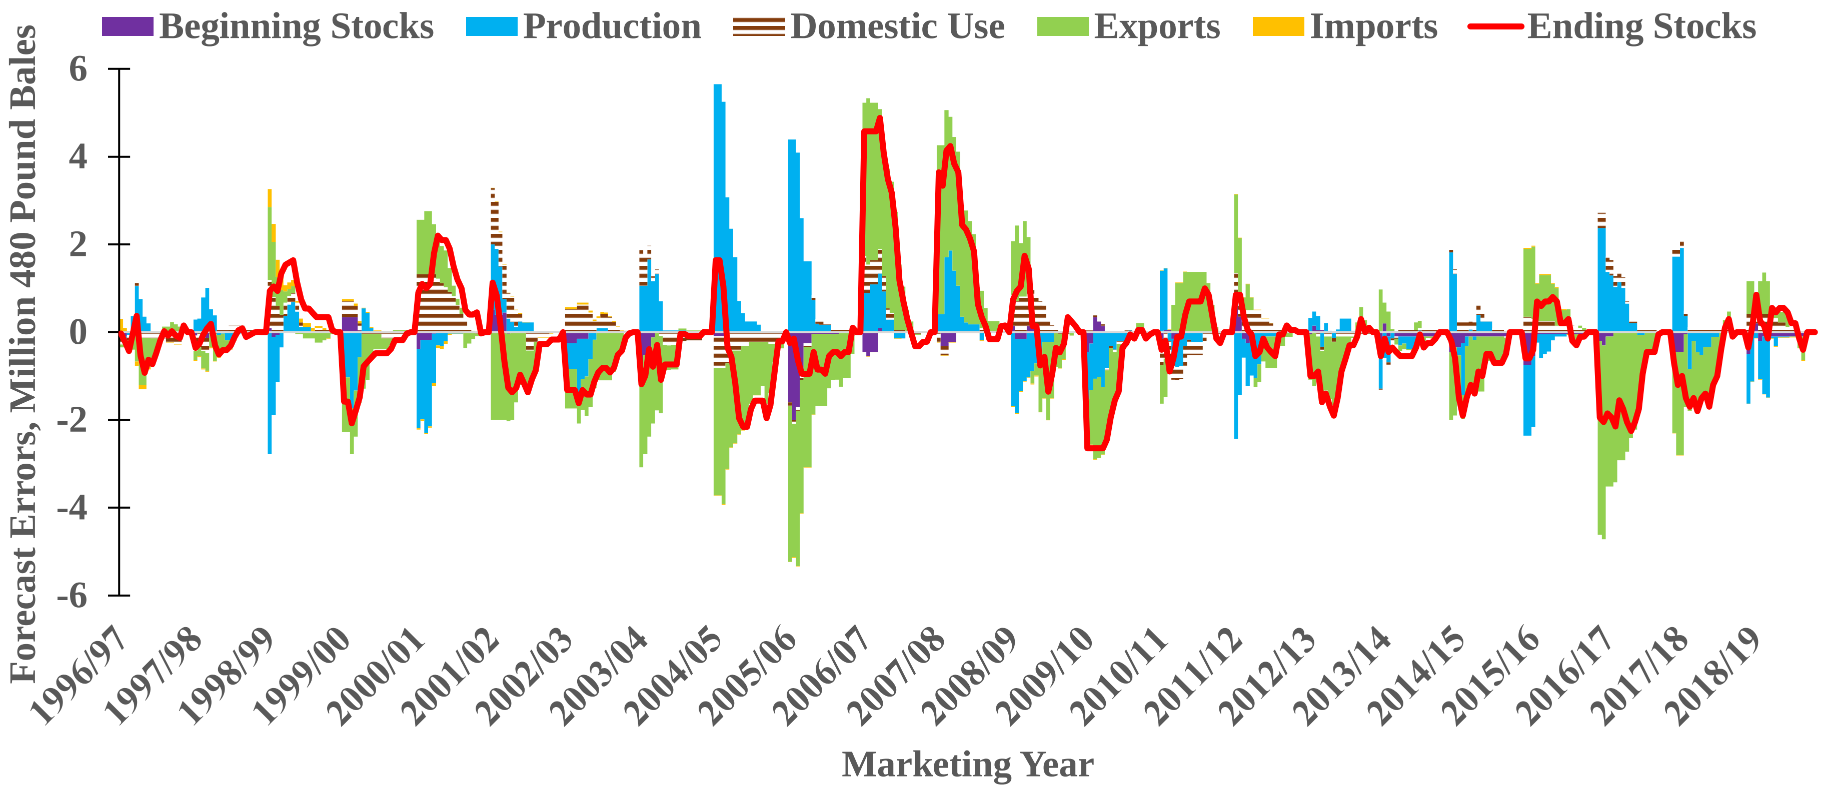

Figure 20: USDA Cotton Balance Sheet Element Forecast Errors, 1996-2019

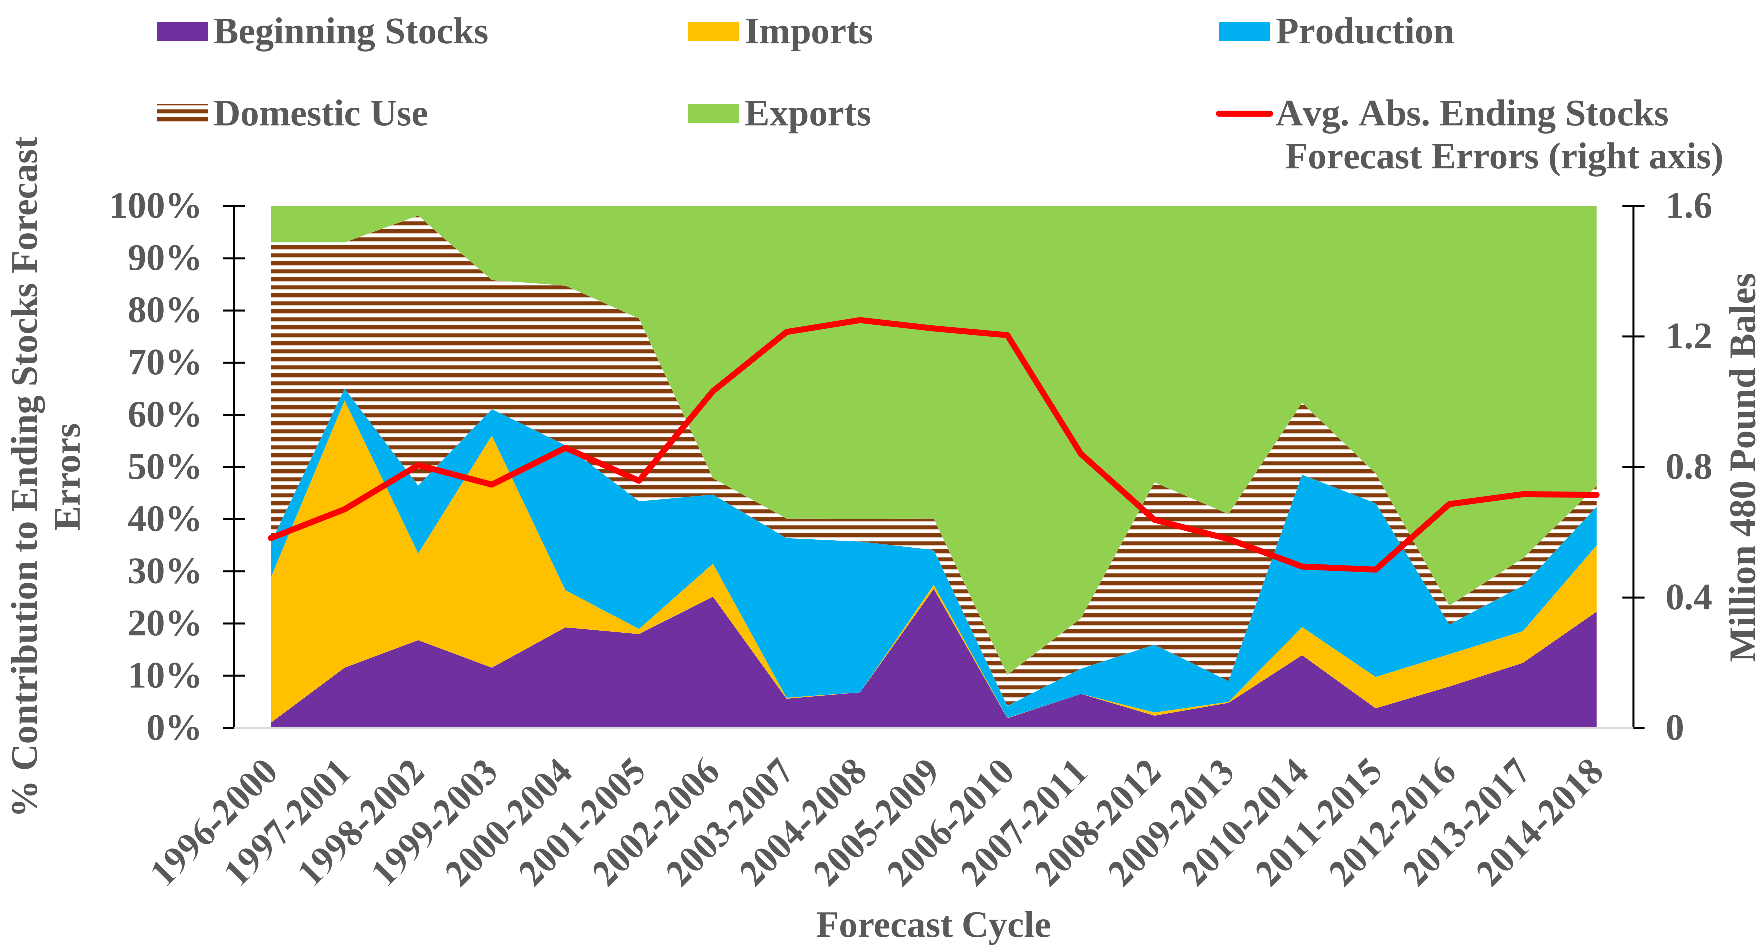

Figure 21: Contribution of Balance Sheet Elements to USDA's Cotton Ending Stock Forecast Errors

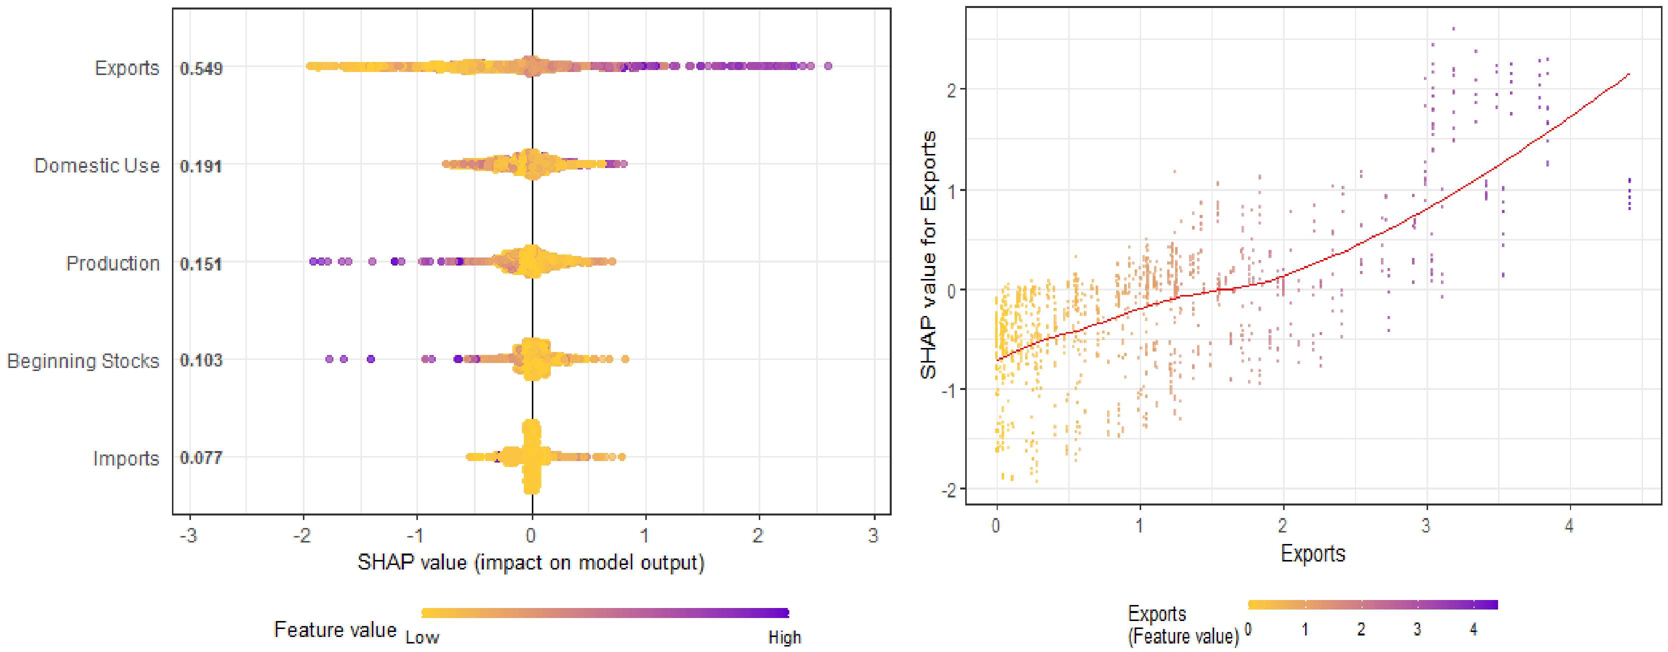

Figure 22: Cotton – Exports SHAP Dependence Plot

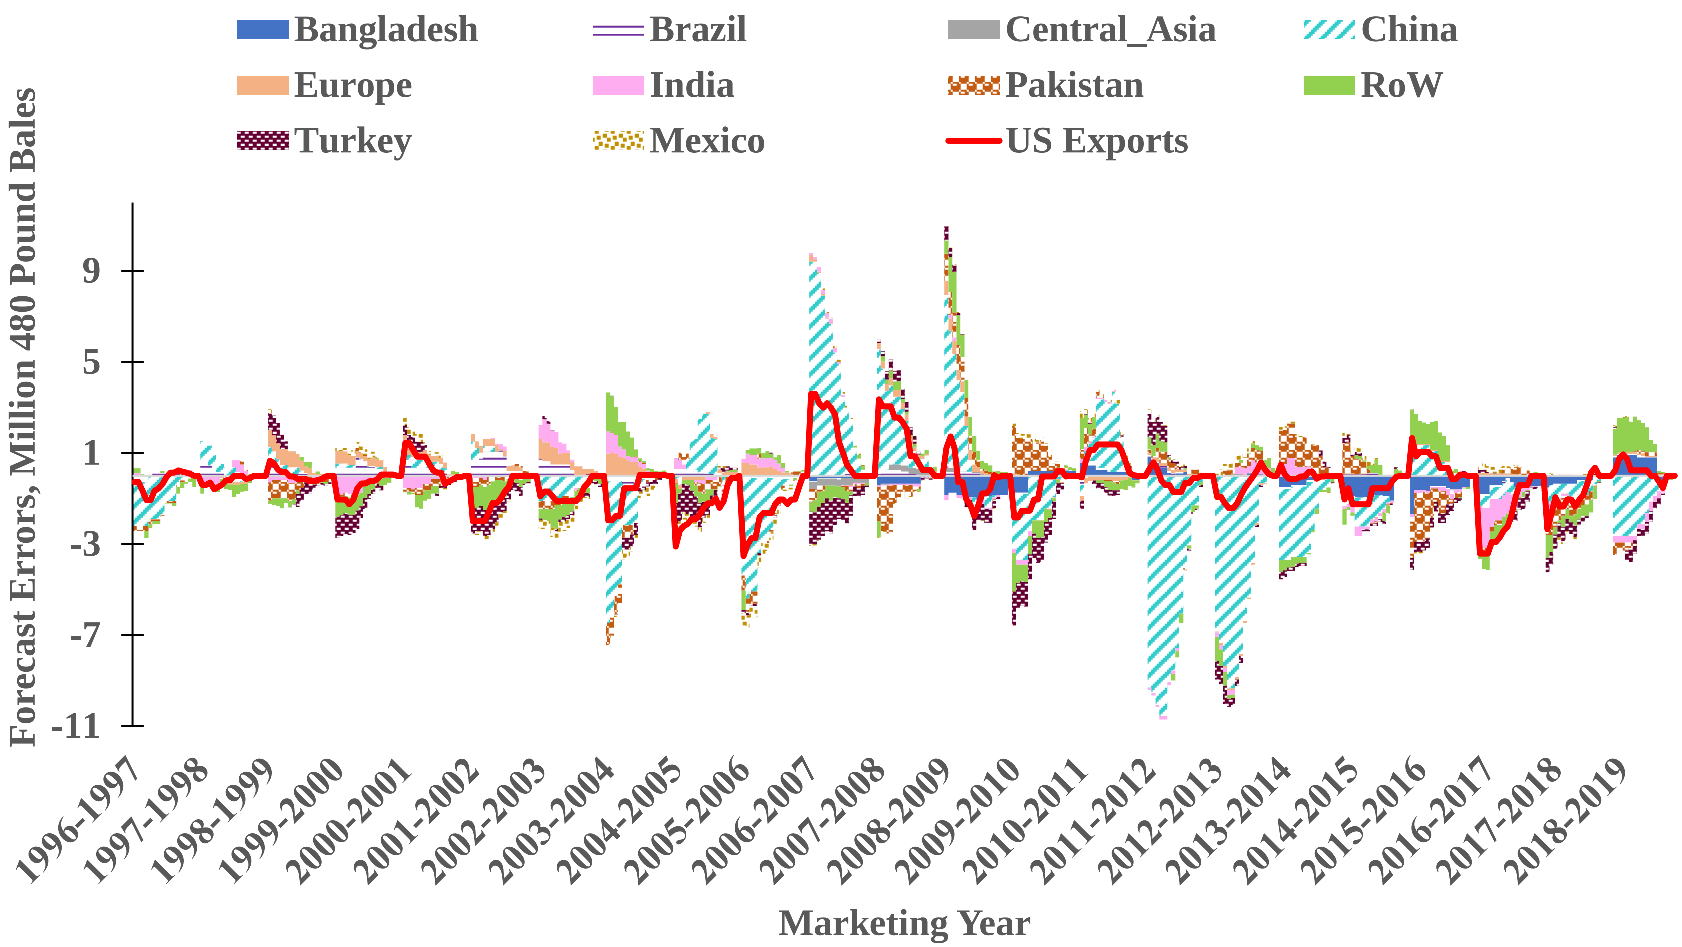

Figure 23: USDA Cotton World Import/ U.S. Export Forecast Errors, 1996-2019

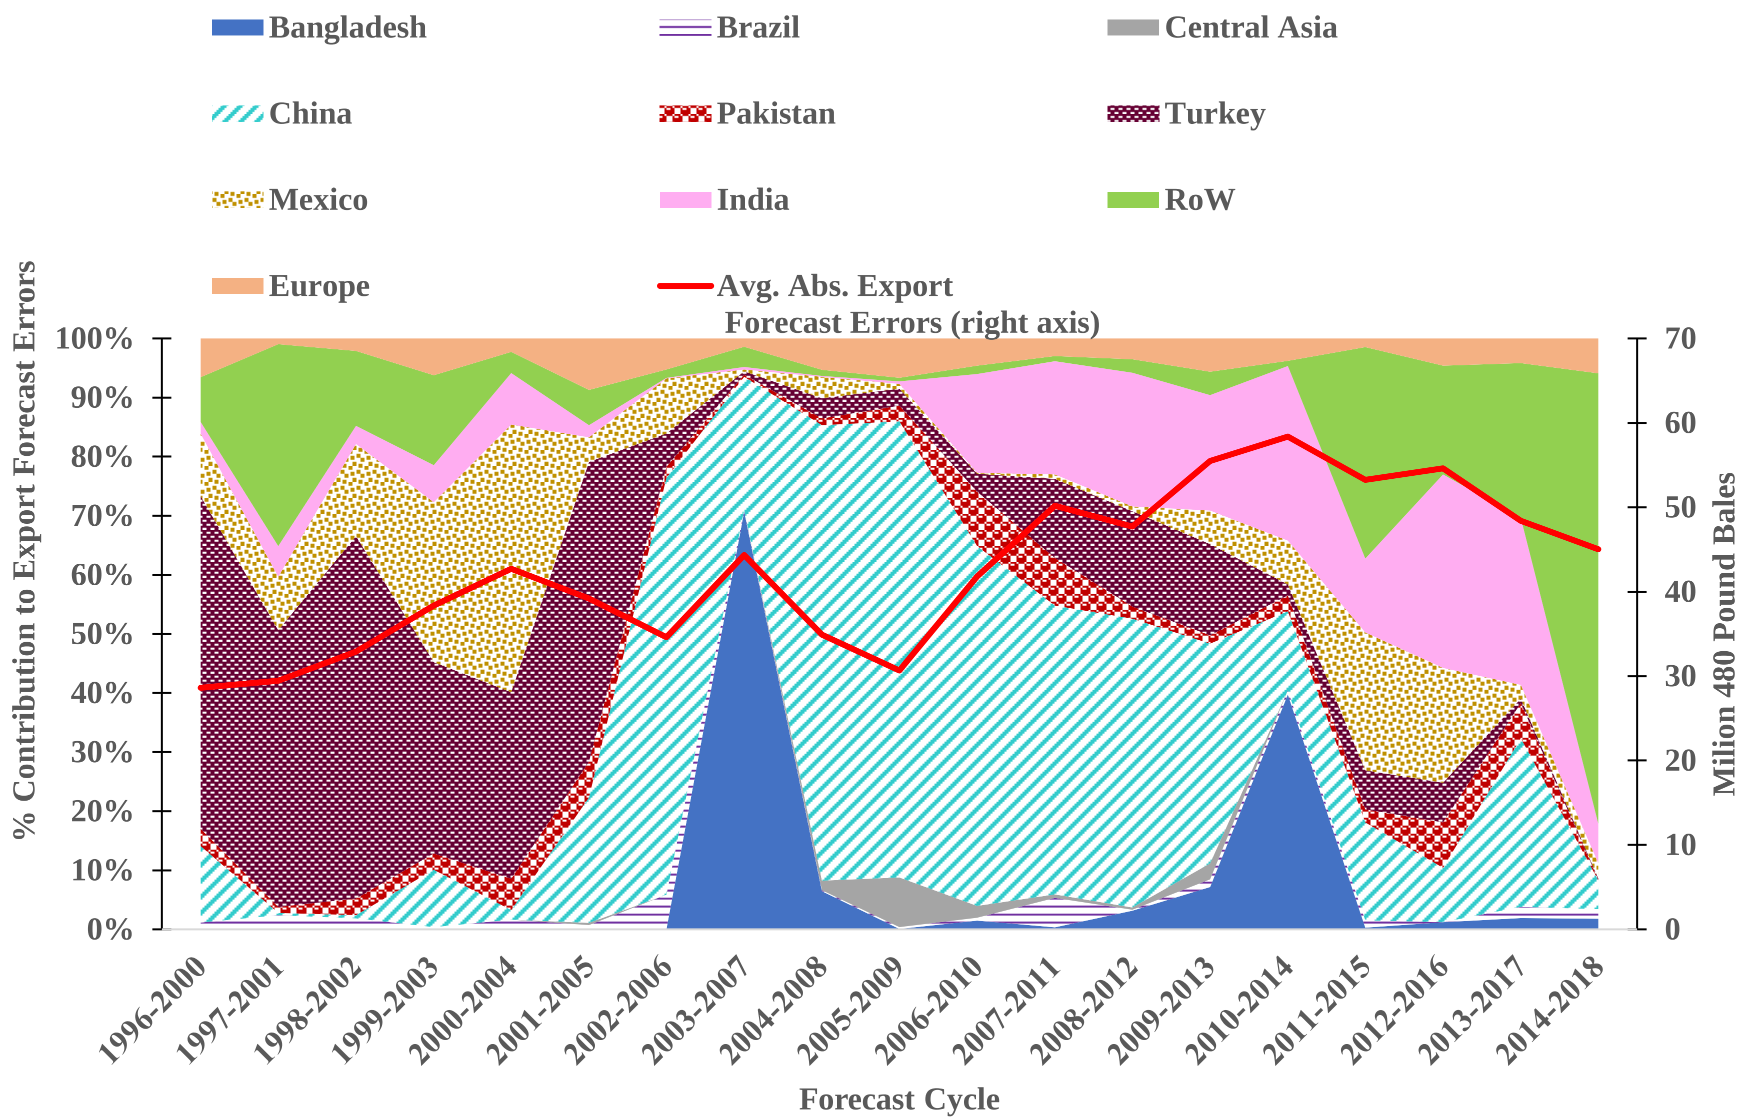

Figure 24: Contribution of Cotton Wheat Import Errors to U.S. Cotton Export Projection Errors

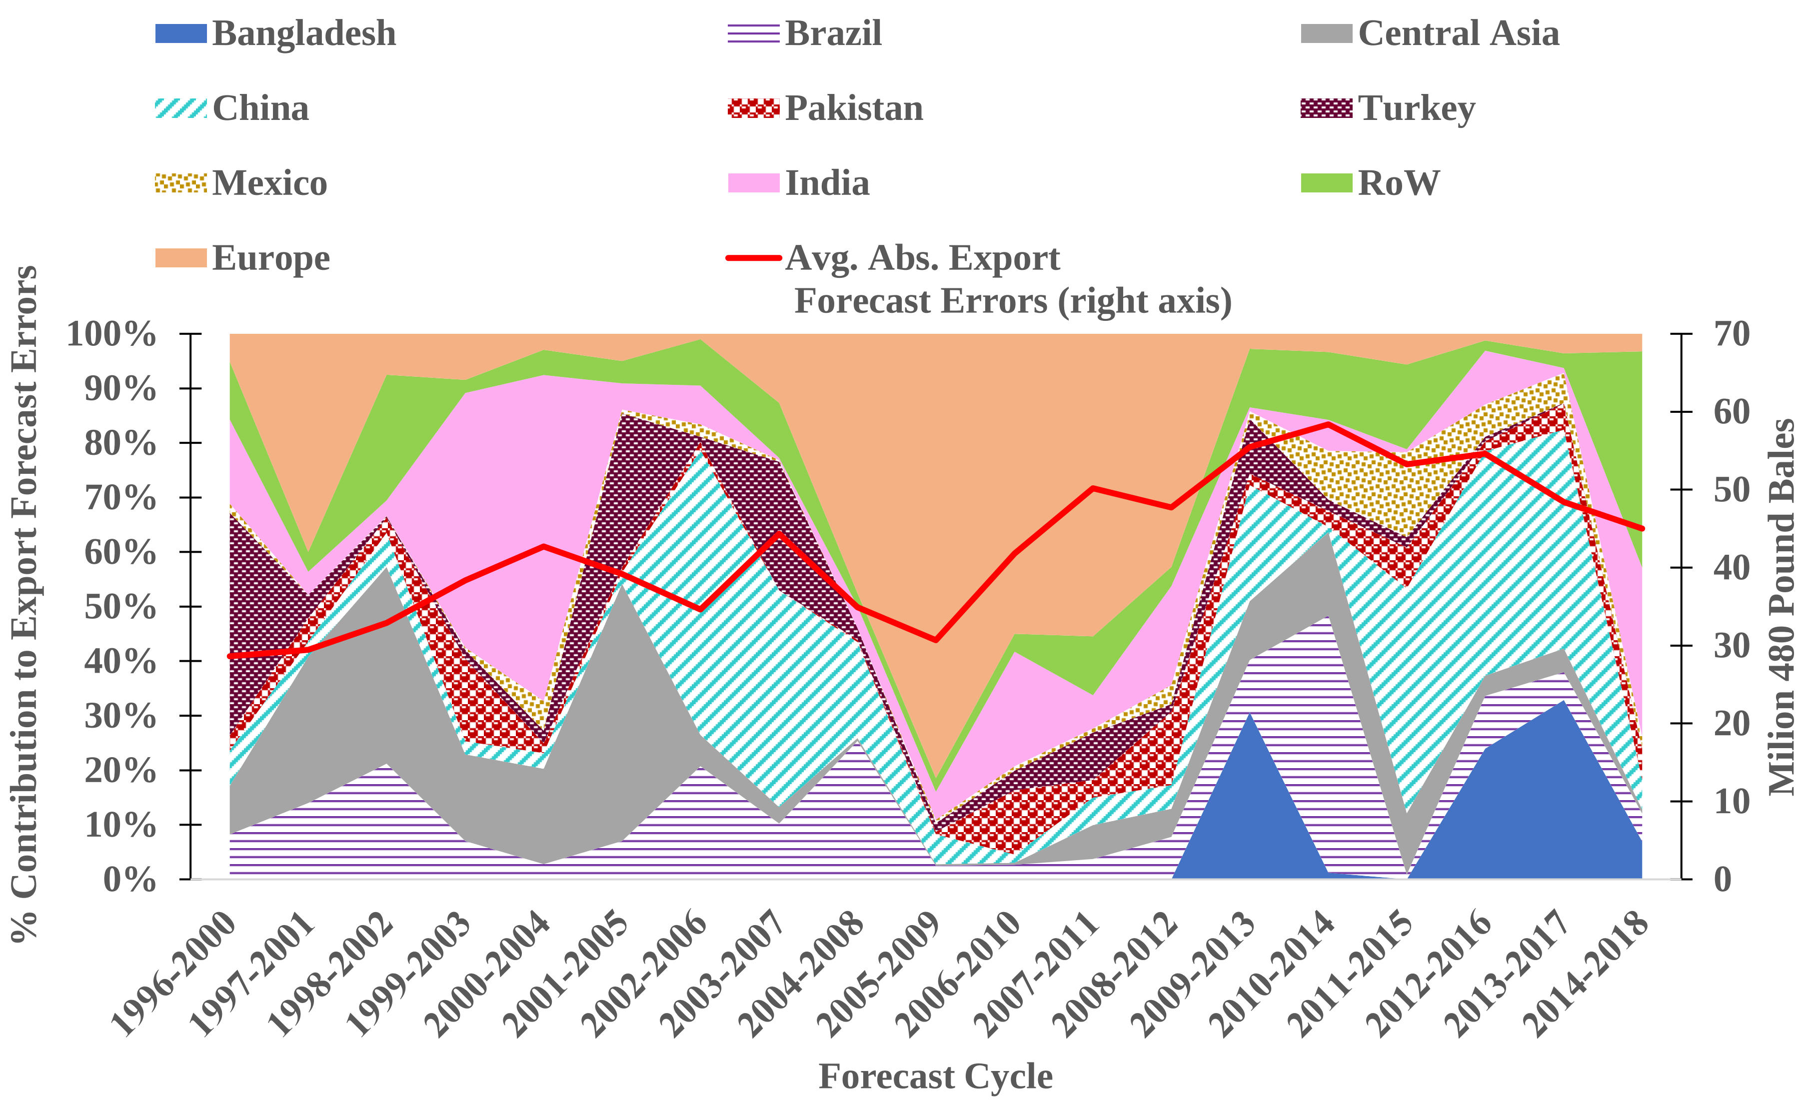

Figure 25: Contribution of World Cotton Production Errors to U.S. Cotton Export Projection Errors

Appendix

Table A1: 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 A2: 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: 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 A4: 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' calculations