Sources of Bias in the USDA International Baseline Projections

Rabail Chandio[*] Ani Katchova[†]

Working Paper

November 10, 2022

*Please find the latest version at this [**link**].*

Dept. of Agricultural, Environmental, and Development Economics
The Ohio State University.

[*]Ph.D. Candidate. Email: chandio.1@osu.edu
[†]Professor and Farm Income Enhancement Chair. Email: katchova.1@osu.edu

Abstract

USDA's annual Agricultural Baseline Projections contribute significantly to agricultural policy in the United States, and hence their accuracy is vital. The baseline projections present a neutral policy scenario assuming a specific macroeconomic situation and allow the analyses of alternative policies and their micro and macroeconomic impacts in the United States. We investigate the trends and heterogeneity in the incidence of bias in the USDA International Baseline Projection reports from 2002 to 2022 to find a potential source of the bias. First, we use the dynamic time warping algorithm to examine whether experts tend to group together the projections for certain crops across different countries, defined as herding, producing similar projection trends and find strong evidence for herding of projection trends toward the United States. Then, we compute the bias in projections, decompose it by projection horizon, and estimate whether the bias is higher across crops or across countries with more substantial evidence for grouping behavior. Results show that for corn yield, soybeans exports, and imports, herding the projections toward the United States lowers the bias while for almost all other crop variable combinations, herding is associated with higher bias. This suggests that the projections for our sample countries may irrationally be herded toward the United States, which is bias inducing.

Keywords: agricultural baselines, USDA baselines, commodity projections, time series similarity, forecast evaluation.

JEL Codes: C53, E17, Q11, Q13, Q14, Q18.

1 Introduction

From long-term agricultural policies that are incorporated into the law through quinquennial Farm Bills and the annual Presidential budgeting, USDA's annual Agricultural Baseline Projections have substantial implications for various stakeholders. The baseline projections present a neutral policy scenario, which assume a specific macroeconomic situation and allow for the analyses of alternative policies and their micro and macroeconomic impacts in the United States. They help evaluate local and foreign policy scenario changes and their subsequent implications for United States' farmers (Skorbiansky, Childs, and Hansen, 2018; Langholtz et al., 2012). Therefore, any policy evaluations utilizing the baselines projections will be as useful and informative as the projections are accurate. Recently, academic research has started evaluating the accuracy of these projections, and has made initial discoveries about the incidence of bias and limited informativeness of the various USDA projections and forecasts (Bora, Katchova, and Kuethe, 2022, 2021; Regmi et al., 2021; Isengildina-Massa et al., 2021; Kuethe, Hubbs, and Sanders, 2018). However, from a policy perspective, understanding the source of bias is essential for minimizing it and improving the projections, which has not received much attention in the literature.

Our study addresses this gap in the literature and investigates the level of herding and its contribution towards bias in the USDA international baseline projections, where herding is defined as significant similarity in two projected series. To answer this question, we have a two-step empirical strategy. First, we examine whether experts tend to group together the projections for certain crops across different countries, defined as herding behavior, producing similar projection trends. Since projections are produced for a decade in any given year, and may be correlated in both leads and lags, we employ time-series techniques novel to our field to accommodate this time-varying correlation and quantify herding. Utilizing rich time-series data where each time period nests an entire projection series (nested time series), we use a Dynamic Time Warping (DTW) algorithm to assess the degree of herding in the baselines projections. The DTW algorithm has been adapted to various fields in recent applications (Berndt and Clifford, 1994; Müller, 2007; Jeong, Jeong, and Omitaomu, 2011; Varatharajan et al., 2018) and is a robust way to measure similarity in two time series.

Second, we compute the bias in projections, where bias is defined as the difference between projected values and realized values, and assess how the bias varies across crops and across countries with differing level of herding behavior. Quite reasonably, some similarities in projections may be stemming from the fact that the true realized data for those regions follow a similar trajectory. In that case, herding in the projections may occur from an informed choice, and it's relation to bias should be evaluated with that in mind. Since we only observe the final values and cannot observe any adjustments made to the initially proposed projected values, we account for this phenomenon by defining the rationality of the projections on an ex post facto basis. If herding the projections of other countries towards the United States is associated with lower projection errors, we classify this as a "rational" herding. Similarly, if herding the projections toward the United States is associated with higher projection errors, it is irrational to align the projection trends of other regions this way. This allows us to observe the heterogeneity in the distribution of bias, determine the contribution of herding towards it, and improve the accuracy of projections.

Herding is a behavioral phenomenon often observed in financial markets, when investors and experts with private information align their choices and decisions with others as a risk management strategy. It can be rational if the individuals make the choice to align their decisions with others based on superior private information, or it can be irrational if individuals ignore their private information in order to adopt similarity with others (Devenow and Welch, 1996). Behavioral finance research suggests that propensity to herd is a response to a private cost minimization strategy by individuals. As long as there is no significant cost to a single agent of agreeing with the group opinion, the majority of people choose to follow the group consensus, regardless of their individual prior beliefs (Huang et al., 2017). Moreover, when an institution (such as USDA) or a specific forecasting team is considered a single entity, the forecast behavior of experts within the institution and/or a team is affected by and aligned with the overall beliefs of the institution/team, and hence their forecasts herd together (Benchimol et al., 2020; Van Campenhout and Verhestraeten, 2010). Whether herding is rational or irrational, it increases volatility in stock and commodity markets and is suboptimal for the market (Blasco, Corredor, and Ferreruela, 2012; Wang and Wang, 2018).

Moreover, to understand how herding may occur in the baseline projections, we recapitulate the process of baseline projections preparation. Released each year by the

USDA Interagency Agricultural Projections Committee, baseline projections combine model-based values and judgment-based adjustments to these values (USDA Agricultural Projections to 2030). Experts from various committees in USDA, including the Economic Research Service, World Agricultural Outlook Board, and the Office of Chief Economist, evaluate the region specific projected values and adjust them until a point of consensus is reached among the committee. Therefore, the projections process involves a first stage where region specific projections are prepared by individual teams, and a second stage where all the regional projections are considered in unison for a global model, and the region specific values may be adjusted. We, however, only observe the finalized projection values after both stages are completed and cannot observe any adjustment. Conceptually, for the baseline projections, similarity in projections can be introduced at either of the two stages.

We find that the baseline projections for all countries are statistically significantly aligned with the projections of the United States in their trends for all crops and variables. Moreover, for most crops and variables including corn, soybeans, and wheat total consumption, this correlation is associated with significantly higher errors in projections. In addition, corn yield, soybeans exports, and wheat imports are the only variables for which herding the projections towards the United States is associated with more accurate projections for the other countries. We also find that other globally leading producers of certain crops can also provide useful for adjusting the projection trends of countries with little or inaccurate information. Our findings have implications for the USDA baseline experts as well as government agencies and users of the baselines reports.

We make three main contributions to the literature. First, our study identifies that the projections for all countries included in the USDA International Baseline Projection reports are correlated with the United States beyond the correlation in the realized values for each country, which may inform USDA on another criterion that needs to be examined prior to releasing their projection reports. Second, we provide conclusive evidence that for other countries' projections of most crop variables, correlation with the United States is associated with higher bias and lower accuracy of these countries' projections. This informs the USDA baseline experts to assess the accuracy of the baseline models. If the highly correlated projections were a result of the model values being correlated, then the model input or models themselves may need to be updated. However, if the models resulted in dissimilar or uncorrelated projections that were later smoothed by the baseline experts to look similar, this would suggest that the baseline committee need to reconsider their herding behavior. Third, we recognize the heterogeneity in the relationship between projections' bias and projections' correlation. By highlighting the variables for which correlation reduces the error in projections, a higher accuracy in the projections can be achieved if the projections are improved. Overall, these insights can be incorporated by the team preparing the USDA baselines projections to minimize the bias in the projections, and prepare more accurate projections allowing for better policy discussion.

The remainder of this paper is organized as follows. Section 2 describes the USDA International Baseline Projections and the variables included in our study. Section 3 details the empirical strategy, which is followed by presentation and discussion of the results in section 4. Section 5 contains the concluding remarks.

2 Data

The baseline projections are one of various economic forecasts produced by several agencies in the US. We use the official USDA International Baseline Projections data from 2002 to 2022 which include 10-year domestic (United States) and international (other countries') projections for several crops each year. We limit our analysis to three crops — corn, soybeans, and wheat — and six variables — area harvested, yield, imports, exports, ending stocks, and total consumption. Balance sheet equation dictating the relationship between the variables we study is as follows

$$\begin{aligned} &BeginningStocks + Production + Imports \\ &= Exports + TotalConsumption + EndingStocks \end{aligned} \tag{1}$$

where the $BeginningStocks_t = EndingStocks_{t-1}$, making it a redundant variable, and $Production = AreaHarvested \times Yield$. We focus only on the variables that are identified independently, thus, not considering beginning stocks and production.

The available baseline data also includes the realized values for up to three years before the release date of the reports. We utilize these limited historical data in each year's report to construct an annualized panel data set for realized values that is used for bias calculations and accuracy evaluations of the projections.

The baseline projections have a structure which is statistically referred to as nested time-series data, where each year nests the series of ten incremental horizon projections for 10 years or horizons into the future. A representative projection $\hat{Y}_{rcvt}$ is the projection series for country r (belonging to an unbalanced panel of 34 countries observed annually over the study period), for crop $c \in \{corn, soybeans, wheat\}$, variable $v \in \{yield, area\ harvested, imports, exports, total\ consumption, ending\ stocks\}$, and report year $t \in \{2002, ..., 2022\}$. $\hat{Y}_{rcvt}$ is a series that has a length of $H = 10$, where h represents the different projection horizons such that $\hat{Y}_{rcvt} = (\hat{Y}_{h_0}, \hat{Y}_{h_1}, ..., \hat{Y}_{h_9})$.

The bias in the baseline projections is defined as the difference between the projection and the realized value. We employ the logged error measure for assessing projections accuracy, which can be interpreted in percentage terms:

$$LoggedError_{rcvth} = \Big(log(\hat{Y}_{rcvth}) - log(Y_{rcvth})\Big) \tag{2}$$

where Y_{rcvth} is the actual value realized for the projection $\hat{Y}_{rcvth}$ and $LoggedError_{rcvth}$ is the error for each single projection made. $LoggedError_{rcvh}$ is the average error calculated over the report years t in the projections for country r, crop c, variable v, and horizon h.

3 Methods

There are two main components of our empirical analysis. First, we estimate the degree of similarity and correlation among various countries' baseline projections using the dynamic time warping algorithm. We also compute the similarity in realized data over the study period to provide a reference to understand whether the similarity in projections is coming from the true realized data.

Second, we use regression analysis to study the relationship between the degree of herding and the size of bias in the projections. To fully understand the heterogeneity in this relationship, we estimate it separately for all crops and variables, and we also vary the benchmark country. That allows us to observe whether the top producers of a commodity are used as benchmarks in practice when these projections reports are produced. If this is the case, then we should observe similarity in projections with the top producer of a commodity.

3.1 Evaluating the Degree of Similarity

We begin our analysis by evaluating the differences in projections of specific countries for each crop, variable, report year, and projection horizon to estimate the degree of similarity or herding in the projections. We use a dynamic time warping algorithm to compute the distance between each set of projection series for each crop-variable-year-horizon combination (for instance, corn yield projections for all 10 horizons in the future that are included in the report year 2010) and evaluate whether the projections exhibit similar trends among all countries. The algorithm finds the minimum distance needed to make two time-series as similar as possible. We use the DTW algorithm to compute the distances between all country pairs for each crop-variable-year-horizon to determine the closest projection "neighbors" of the top country producers for each crop that have the smallest distance among all countries.

We suppress the indices cvt since they remain the same for each pair of countries whose projections are being compared. To determine the distance between the projections for any two countries $\hat{Y}_{r_i}$ and $\hat{Y}_{r_j}$, we define the two time-dependent series $\hat{Y}_{r_i}$ and $\hat{Y}_{r_j}$ and compute an expansive local cost matrix (LCM) between them. The LCM is populated by pairwise comparisons of each horizon's projections for $\hat{Y}_{r_i}$ with each horizon's projections of $\hat{Y}_{r_j}$, resulting in a square matrix of dimensions 10×10 since $length(\hat{Y}_{r_i}) = length(\hat{Y}_{r_j}) = 10$ for the 10 horizons. The LCM matrix is defined as:

$$LCM(\hat{Y}_{r_i}, \hat{Y}_{r_j}) = \begin{pmatrix} d_{\hat{Y}_{r_i h_0}, \hat{Y}_{r_j h_0}} & d_{\hat{Y}_{r_i h_0}, \hat{Y}_{r_j h_1}} & \cdots & d_{\hat{Y}_{r_i h_0}, \hat{Y}_{r_j h_9}} \\ d_{\hat{Y}_{r_i h_1}, \hat{Y}_{r_j h_0}} & d_{\hat{Y}_{r_i h_1}, \hat{Y}_{r_j h_1}} & \cdots & d_{\hat{Y}_{r_i h_1}, \hat{Y}_{r_j h_9}} \\ \vdots & \vdots & \ddots & \vdots \\ d_{\hat{Y}_{r_i h_9}, \hat{Y}_{r_j h_0}} & d_{\hat{Y}_{r_i h_9}, \hat{Y}_{r_j h_1}} & \cdots & d_{\hat{Y}_{r_i h_9}, \hat{Y}_{r_j h_9}} \end{pmatrix}$$

where each matrix element $d_{\hat{Y}_{r_i h_a}, \hat{Y}_{r_j h_b}} = \sqrt{(\hat{Y}_{r_i h_a} - \hat{Y}_{r_j h_b})^2}$ denotes the Euclidean distance between a^{th} and b^{th} horizon projections of series $\hat{Y}_{r_i}$ and, $\hat{Y}_{r_j}$ respectively and $a, b \in h$.

We find the distance between the two projection series by defining $\phi(k)$ to be the path from $d_{\hat{Y}_{r_i h_0}, \hat{Y}_{r_j h_0}}$ to $d_{\hat{Y}_{r_i h_9}, \hat{Y}_{r_j h_9}}$ where $k = (1, 1), ..., (H, H)$. For a given path ϕ, we compute the Euclidean distance measuring similarity between the projections for two countries $\hat{Y}_{r_i}$ and $\hat{Y}_{r_j}$ as

$d_{\phi}(\hat{Y}_{r_i}, \hat{Y}_{r_j}) = \sum_k [(LCM(k)) \times m_{\phi}(k)]$

where $m_{\phi}(k)$ is the per-step weighting coefficient, allowing to only add the distance in horizon-specific projections that falls on a given path $\phi(k)$. At this step, all the possible paths from the beginning of LCM at $k = (1, 1)$ to the end at $k = (H, H)$ are computed. Next, we find a path ϕ within the LCM matrix that gives the minimum total distance between the projections in two countries $d_{\phi}(\hat{Y}_{r_i}, \hat{Y}_{r_j})$. Hence, we use the DTW algorithm to find the minimum distance between the two countries' projections for the close-by horizons by solving the following optimization problem:

$$DTW(\hat{Y}_{r_i}, \hat{Y}_{r_j}) = min_{\phi}(d_{\phi}(\hat{Y}_{r_i}, \hat{Y}_{r_j})) \tag{3}$$

We impose two constraints on the minimization problem to avoid meaningless loops or inefficient paths. First, we require monotonicity, which restricts the direction of the path taken within the LCM to only increasing projection horizon for at least one of the projection series :

$$\phi : \phi\left(k + 1\right) \geq \phi(k) \tag{4}$$

Second, we impose the Sakoe-Chiba window constraint to restrict the path that can be taken to a window h^{window} around the principal diagonal of the LCM. We use a window h_{window} as the only valid points of the LCM where a path ϕ can be made in the range:

$$\left|\left(h_a, h_b - h_{window}\right), \left(h_a, h_b + h_{window}\right)\right| \tag{5}$$

for all horizons (h_a, h_b) along the principal LCM diagonal. This constraint allows only reasonable lead and lag relationships of the horizons to be considered while two projection series are being compared. We set $h_{window} = 2$ restricting projection $\hat{Y}_{r_i}$ for, say, horizon 5 to only be compared to projection $\hat{Y}_{r_j}$ for another country from horizon 3 to 7 allowing the algorithm to detect at max a 2-year lead or a 2-year lag relationship between the two projection series. This approach is more general than restricting comparisons between countries' projections to be only for the same horizon. To ensure comparability across different countries' projections, we scale the baseline projections data using z-score normalization for each projection series being compared.

The distance calculation detailed above is repeated for each crop c, variable v, report year t, and the projection series of each country have the length of horizon h. To compute the overall distance in projections of two countries, we average the DTW distances between projections for two countries $\hat{Y}_{r_i}$ and $\hat{Y}_{r_j}$ over all report years t:

$$distance(\hat{Y}_{r_icv}, \hat{Y}_{r_jcv}) = \frac{1}{length(t)} \sum_{t=2002}^{2021} (DTW(\hat{Y}_{r_icvt}, \hat{Y}_{r_icvt})) \tag{6}$$

where $distance(\hat{Y}_{r_icv}, \hat{Y}_{r_jcv})$ gives average difference in projections of two countries for each crop c and variable v and $DTW(\hat{Y}_{r_icvt}, \hat{Y}_{r_jcvt})$ refers to distance from the minimization problem described in equations (3), (4), and (5). We compute the standard errors for the computed distance using the bootstrapping method with 250 replications.

If a country's distance from a benchmark country is statistically indistinguishable from zero, we would infer that its projections are correlated with those of the benchmark country, in other words, the two countries are similar in trends. Repeating this process for all crops and variable combinations and all benchmark countries provides enough information to evaluate whether and where herding behavior of countries' projections occurs.

3.2 Relationship between Bias and Herding

We estimate the relationship between similarity in projection trends and the bias in the baseline projections. This allows us to answer whether herding the projections for various countries is desirable in the case of limited information. Moreover, we can evaluate the heterogeneity in the impact of herding on bias across crops, variables, horizons, and report years. Using different countries which are the top producers of these crops (United States, Brazil, or China) as benchmark countries for herding in the projections, we begin by estimating the following equation:

$$LoggedError_{r_jh} = \beta_0 + \beta_1 log(distance(\hat{Y}_{r_{US}}, \hat{Y}_{r_j})) + \epsilon_{r_jh} \tag{7}$$

for each crop and variable separately, using the logged error calculated in equation (2). $log(distance(\hat{Y}_{r_{US}}, \hat{Y}_{r_j}))$ is the log of computed DTW distance of country r_j's projections from the United States for each projection horizon. This is the distance calculated in equation (3) using the dynamic time warping algorithm.

The coefficient on $log(distance(\hat{Y}_{r_{US}}, \hat{Y}_{r_j}))$ shows the β_1% change in the projection error associated with 1% increase in the distance of country r's projections from the United States. If this value is significant and positive, it implies that herding the projections towards the United States is associated with lower errors, while a significant negative estimate suggests that herding towards the United States may be causing higher projection errors. A significant positive β_1 would signal that herding is irrational and not aiding the accuracy of the projections for the other countries, on average. Similarly, a significant negative β_1 indicate rational herding of the other countries' projections towards the United States because herding towards the United States is seemingly reducing the error in the projections of other countries.

We estimate two more versions of equation (7) with Brazil and China being the benchmark countries. With income and imperfect information, trends of other global leaders may also provide helpful insights for a given crop's projections in other regions. Therefore, we estimate

$$LoggedError_{r_jh} = \beta_0 + \beta_1 log(distance(\hat{Y}_{r_{China}}, \hat{Y}_{r_j})) + \epsilon_{r_jh} \tag{8}$$

and

$$LoggedError_{r_jh} = \beta_0 + \beta_1 log(distance(\hat{Y}_{r_{Brazil}}, \hat{Y}_{r_j})) + \epsilon_{r_jh} \tag{9}$$

4 Results

We first present the results of the dynamic time warping algorithm that measures the distance in projections of various countries compared to the projections of select benchmark countries i.e., United States, Brazil, and China. These three regions are chosen as the benchmarks because they are among the leading global producers for the crops we are focusing on: corn, soybeans, and wheat. The results are categorized by crop (corn, soybeans, and wheat) and variable (yield, area harvested, ending stocks, total consumption, imports, and exports). We also depict the errors in the projections for each country, displaying them by projection horizon. So, the top panels of 1 through 18 correspond to similarity estimates from equation 6 and the bottom panel is a visual representation of error calculations as per equation 2. The similarity measures are estimates, so we also depict the 95% confidence interval for all of them. Therefore, if the confidence intervals for estimated distance contains 0, we conclude that the projections for the country on the horizontal axis are statistically significantly similar to the projections for the benchmark country. We have estimated the similarity for all countries in our sample, but display only twelve countries in the figures, since they are among the major producers of the three crops included in our analysis. Figures 1 to 6 correspond to corn, figures 7 to 12 represent soybeans, and figures 13 to 18 contain the results for wheat.

Our results show strong correlation in the baseline projections trends of other countries with the United States. For all the countries in our sample, all confidence intervals for the distance from the United States contain 0 regardless of the magnitude of the point estimate. That implies that projections for all crops and variables are significantly similar to the United States (their distance from the United States is indistinguishable from 0). On the other hand, there is reasonable variation in similarity of projections with the other benchmark countries (Brazil and China). For instance, in figure 1, distance in projections between Brazil and United States overlaps zero (projections are similar) while distance between Brazil and China is significantly greater than zero (projections are not similar). Ukraine and Indonesia are two other countries in figure 1 with significant similarity to the United States while being significantly distant from Brazil.

The most notable result from these set of figures is that the USDA baseline projections do not significantly differ from the United States compared to other sample countries for any crop or variable considered in the analysis, while reasonable dissimilarity exists when Brazil or China are set as the benchmark country. However, the distances in realized values of different countries from the United States are often greater in magnitude than the estimated distance in projections. For select variables, the distance in realized values are outside the 95% confidence interval of estimated distance in projections (i.e., distances are irrational), the exception to this occurs for variables with very large confidence intervals (see corn area harvested—figure 2, and corn total consumption—figure 3).

Moreover, projection error trends in the bottom panels of figures 1 through 18 show the error in projections for the twelve countries depicted in the top panels. This is the average error over the study period for each projection horizon, marked on the horizontal axis. A larger projection horizon means that projections are being made further into the future. We observe that while the projection errors for the United States remain among the lowest, the errors for most other regions increase drastically as the projection horizon increases. Compared to other variables, yields are projected with the most certainty since there are no significant changes in a country's productivity from year to year. Therefore, the errors for yields are relatively smaller than other variables for all three crops across all countries. For the United States, errors in corn, soybeans, and wheat yields remain lower than 10% in magnitude for all horizons, indicative of highly accurate yield projections. Contrarily, yield projection errors for other regions (see Ukraine and Brazil, for instance) are lower than errors in their projections for other variables, but are still visibly much higher than the errors for the United States' projections.

While the overall errors for other variables are higher in magnitude compared to errors in yields, the United States consistently continues to have lower errors relative to other countries in almost all the figures. It is reasonable to wonder, then, if the projections for other regions being herded towards the United States has any relation or contribution to this phenomenon. USDA can project the values for the United States very well, but there is always room for improvement. Therefore, a higher level of accuracy in the projections of other countries will make for a better understanding of the global environment and can lead to an improvement in domestic projections. To see why it might be so, consider the balance sheet relationship depicted in equation (1). The crop specific balance sheet equation holds for each country individually in the projections, and it also holds in summation for the whole world. Therefore, we now look at the role of herding with the United States, and other commodity leaders (Brazil and China) in the determination of errors.

4.1 Relationship between Herding and Bias

Results for estimation equation (7) for each of the three benchmark countries — United States, China, and Brazil, respectively — are presented separately in tables 1 to 3. Equations (7), (8), and (9) estimate the association between error in projections and the distance in projections of a region from the benchmark country. Each column and crop panel in shows estimation results corresponding to equation (7) with the benchmark country specified in the table caption.

Table 1 shows the results with the United States as the benchmark country. The coefficient on *Distance from Base$_r$ (logged)* measures the percentage change in the error if the projections are 1% farther from the United States, which serves as the base country.

The first row implies that corn total consumption and corn ending stocks projections being more distant from the United States are associated with significantly lower projection errors among other countries (i.e., a negative significant coefficient on *Distance from Base*$_r$ *(logged)*). That is, decreasing the herding towards the United States in these projections for all countries, which on average are all significantly correlated with the United States, is associated with decreases in their errors. We observe a similar relationship for all soybeans projections except exports, and all wheat projections except imports. Herding, in these cases, is related to higher errors.

As the results depict, it seems that herding the projections towards the United States is rational for corn yields, soybeans exports, and wheat imports while it is seemingly irrational for other crop and variable combinations.

Overall, there are two main takeaways from these results. First, on average, countries whose projections are more distant from the United States are associated with lower projection errors for most variables across all three crops, even if marginally so. This suggests that the projections of other countries, that are following similar trends to the United States projections, is associated with higher bias in these countries' projections in most cases. Moreover, the three exceptions are corn yield, soybean exports, and wheat imports, where greater distance from the United States in projections is significantly associated with higher projection errors for these countries' variables. Second, herding toward the United States is seemingly more helpful for corn compared to other crops.

While limited information of the baseline team may result in projections for other countries to follow and correlate with the projections of the major producers, it does not necessarily mean that the United States trends are applicable globally. It is also reasonable to expect that for some crops and variables, herding towards the projection trends of China or Brazil in projections may reduce errors because of the massive global contribution of these countries, while for others it may not. Avileis and Mallory (2022) show that recently in corn markets, the volatility, and price information has been spilling over from Brazil towards the United States whereas this relationship used to be the opposite a few years ago. This motivates us to consider other major producers of the commodities under consideration.

We use Brazil and China as two additional benchmark countries because they are major crop producers and trade partners of the United States. We estimate equations

(8) and (9) to repeat the analysis done for the United States and assess the relationship between similarity in projections and projections error by setting China and Brazil as benchmark countries instead of the United States. Tables 2 and 3 show the estimation results from equation (7) by setting China and Brazil as base countries, respectively. Since the projections for all countries are statistically correlated with the United States in their trends, using other top producers as the benchmark countries allows us to evaluate whether the other major producer countries offer an alternate error minimizing approach to herding.

For most crops and variables in table 2, there are negative significant estimates, suggesting an irrational correlation in projections. That is because increasing a country's projection distance from China is associated with reductions in the projection errors on average. Similarly, in table 3, where Brazil is the benchmark country, most estimates are negative with some differences in sign as well as magnitude compared to the previous two tables. Corn yield, corn exports, along with wheat exports depict a significant positive relationship between other countries' correlation with Brazil and higher projection errors.

Despite the overall similar implications when China and Brazil are set as benchmark countries, there are meaningful differences when compared with table 1. The relationship between projection similarity and errors is more defined for corn when China is used as the benchmark country. Corn yield and corn imports have positive significant coefficients, which are also greater in magnitude, suggesting that projection trends of China are more insightful than the United States for countries will less information. Moreover, corn yield estimate has an even larger coefficient when Brazil is set as the benchmark country. This suggests that given incomplete and imperfect information for a region, aligning the projection trends towards Brazil for corn yield is associated with a greater error reduction than herding the projections towards the United States.

Additionally, a significant estimate with a sign opposite to another benchmark country (in tables 1, 2, and 3) presents an upper and lower bound for adjusting the projections of a country with incomplete information. For instance, the relationship between projections correlation and error is negative significant when China is set as benchmark and positive significant when Brazil is the benchmark country. That suggests if Brazil is optimistic in their outlook for corn exports while China is pessimistic in their expectation, on average it is likely that a country with less information might be experiencing a similar optimistic corn exports outlook as the Brazil — the global leader in corn production. We observe similar bounds for corn imports and soybeans exports.

5 Conclusion

While the USDA International Baseline Projections are prepared through a combination of model-based values and expert analysts' judgements, the baseline projections for all other countries are found to have an underlying correlation with the United States. Our results show that this is associated with lower bias for a few crop variables, but contributes to higher error in other cases, compromising the overall accuracy of the projections. Given the importance of the baseline projections in domestic agricultural policy, it is imperative to identify where the projection correlation among countries beyond the correlation in the realized values is increasing bias and reducing accuracy.

We employ various methods to identify the correlation in the projections of different countries, assess their degree of accuracy, and map the relationship between projections' similarity in terms of correlation and projections' error. Our results show that select variables that are grouped together in their projection trends are associated with reduced errors. Corn yield, soybeans exports, and wheat imports are the variables where herding the projections towards the United States is associated with more accurate projections for the other countries. Among other crop variables, our results show that herding the projection towards the United States is associated with significantly lower accuracy of the projections in corn total consumption, and ending stocks, soybeans yields, imports, total consumption, and ending stocks, and all wheat variables except imports.

These findings are highly useful for the team preparing the USDA baselines projections, as they present a simple and straight forward way to cater to incomplete and inaccurate information for countries included in the baseline projections. The heterogeneity observed in our results provides a reasonable starting point for the team preparing the baseline projections to get closer to the source of bias in the projections. Utilizing data on the projections made from models compared to the projections after adjustments can offer valuable insights on where and why herding occurs, and highlights the cases where it is associated with bias reductions.

References

Avileis, F.G., and M.L. Mallory. 2022. "The impact of Brazil on global grain dynamics: A study on cross-market volatility spillovers." *Agricultural Economics* 53:231–245.

Benchimol, J., M. El-Shagi, Y. Saadon, et al. 2020. *Do Expert Experience and Characteristics Affect Inflation Forecasts?*. Bank of Israel-Research Department.

Berndt, D.J., and J. Clifford. 1994. "Using dynamic time warping to find patterns in time series." In *KDD workshop*. Seattle, WA, USA:, vol. 10, pp. 359–370.

Blasco, N., P. Corredor, and S. Ferreruela. 2012. "Does herding affect volatility? Implications for the Spanish stock market." *Quantitative Finance* 12:311–327.

Bora, S.S., A.L. Katchova, and T.H. Kuethe. 2022. "The accuracy and informativeness of agricultural baselines." *American Journal of Agricultural Economics*, pp. .

—. 2021. "The Rationality of USDA Forecasts under Multivariate Asymmetric Loss." *American Journal of Agricultural Economics* 103:1006–1033.

Devenow, A., and I. Welch. 1996. "Rational herding in financial economics." *European economic review* 40:603–615.

Huang, R., M. Krishnan, J. Shon, and P. Zhou. 2017. "Who herds? Who doesn't? Estimates of analysts' herding propensity in forecasting earnings." *Contemporary Accounting Research* 34:374–399.

Isengildina-Massa, O., B. Karali, T.H. Kuethe, and A.L. Katchova. 2021. "Joint Evaluation of the System of USDA's Farm Income Forecasts." *Applied Economic Perspectives and Policy* 43:1140–1160.

Jeong, Y.S., M.K. Jeong, and O.A. Omitaomu. 2011. "Weighted dynamic time warping for time series classification." *Pattern recognition* 44:2231–2240.

Kuethe, T.H., T. Hubbs, and D.R. Sanders. 2018. "Evaluating the USDA's Net Farm Income Forecast." *Journal of Agricultural and Resource Economics* 43:457–468.

Langholtz, M., R. Graham, L. Eaton, R. Perlack, C. Hellwinkel, and G. Daniel. 2012. "Price projections of feedstocks for biofuels and biopower in the US." *Energy Policy* 41:484–493.

Müller, M. 2007. "Dynamic time warping." *Information retrieval for music and motion*, pp. 69–84.

Regmi, M., A.M. Featherstone, B.C. Briggeman, and D. Subedi. 2021. "Assessing and Improving USDA's Farm Income Baseline Projections.", pp. .

Skorbiansky, S.R., N. Childs, and J. Hansen. 2018. "Rice in Asia's feed markets." Working paper.

Van Campenhout, G., and J.F. Verhestraeten. 2010. "Herding behavior among financial analysts: A literature review.", pp. .

Varatharajan, R., G. Manogaran, M.K. Priyan, and R. Sundarasekar. 2018. "Wearable sensor devices for early detection of Alzheimer disease using dynamic time warping algorithm." *Cluster Computing* 21:681–690.

Wang, G., and Y. Wang. 2018. "Herding, social network and volatility." *Economic Modelling* 68:74–81.

6 Figures

Figure 1: Corn Yield—Correlation Estimates and Error Calculations

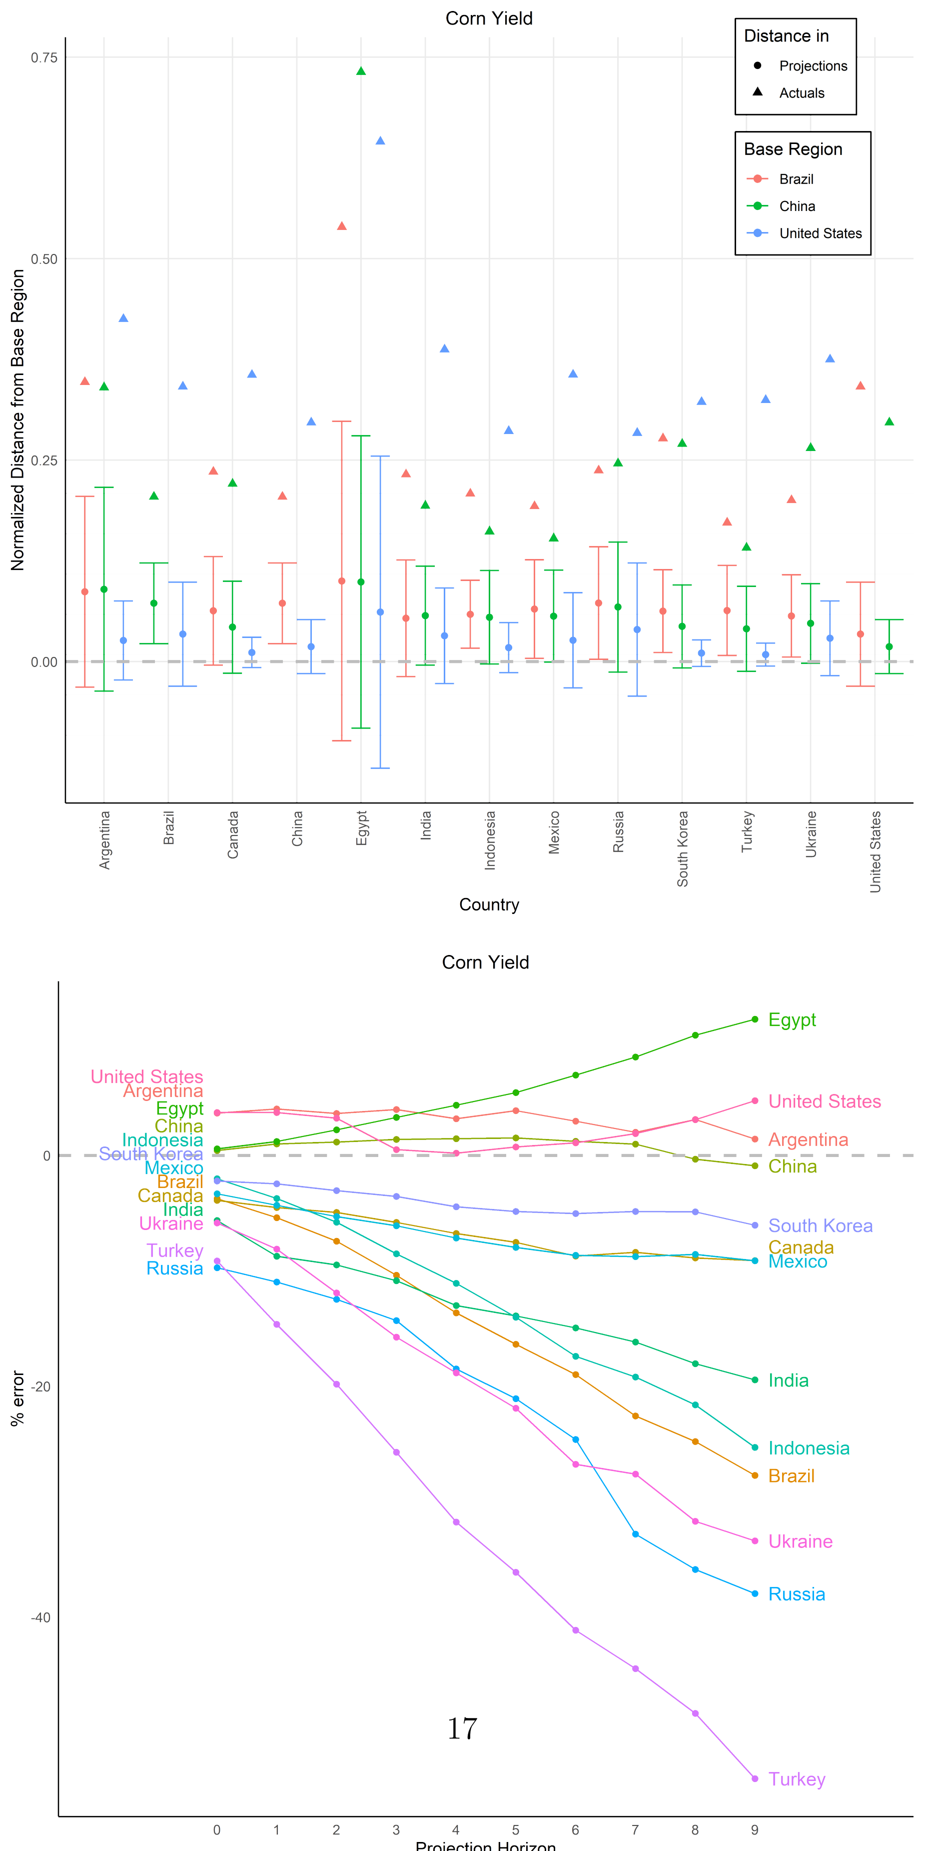

Figure 2: Corn Area Harvested—Correlation and Error Calculations

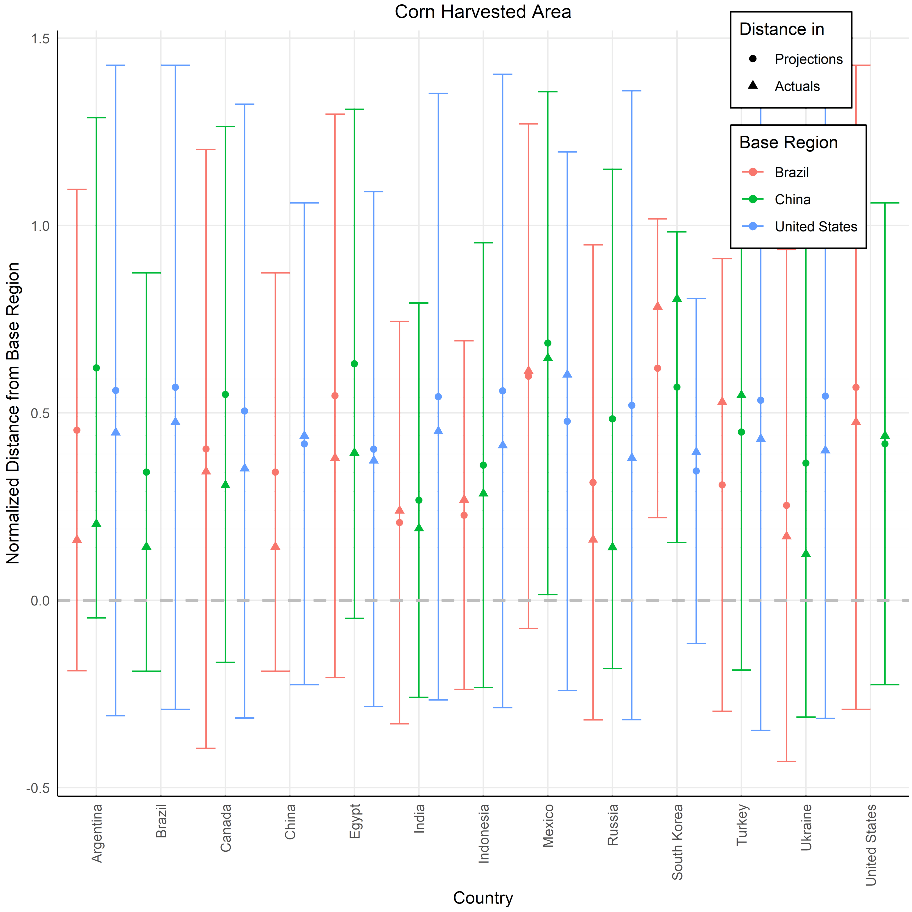

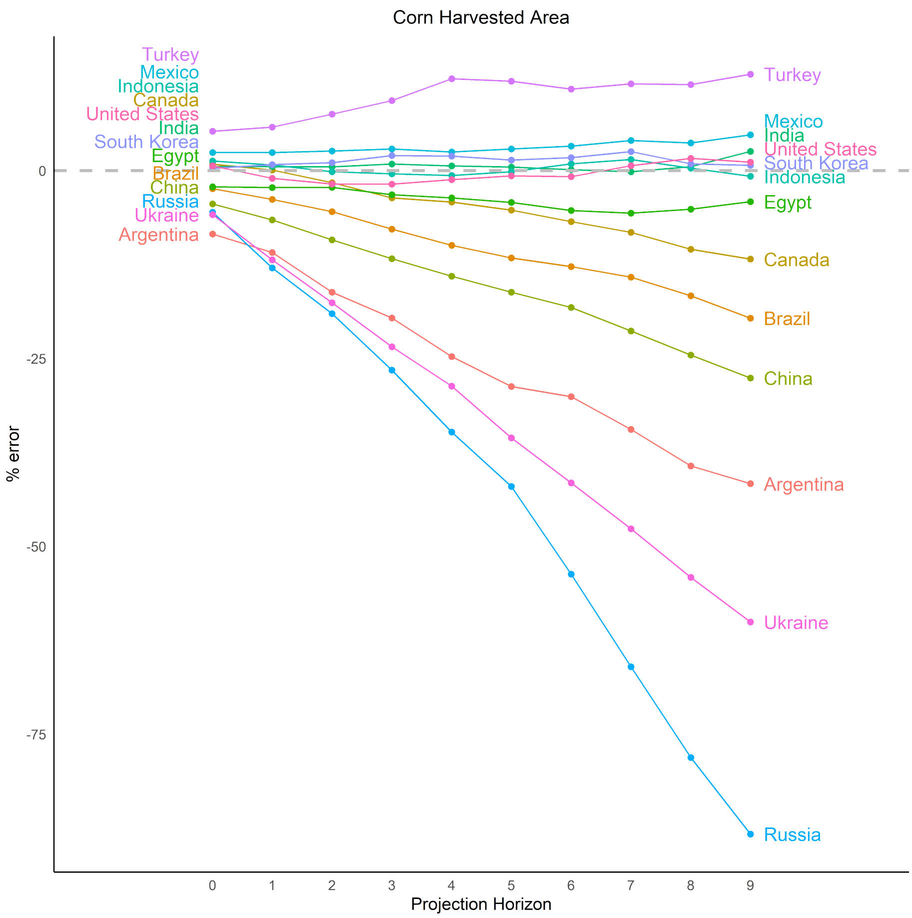

Figure 3: Corn Total Consumption—Correlation and Error Calculations

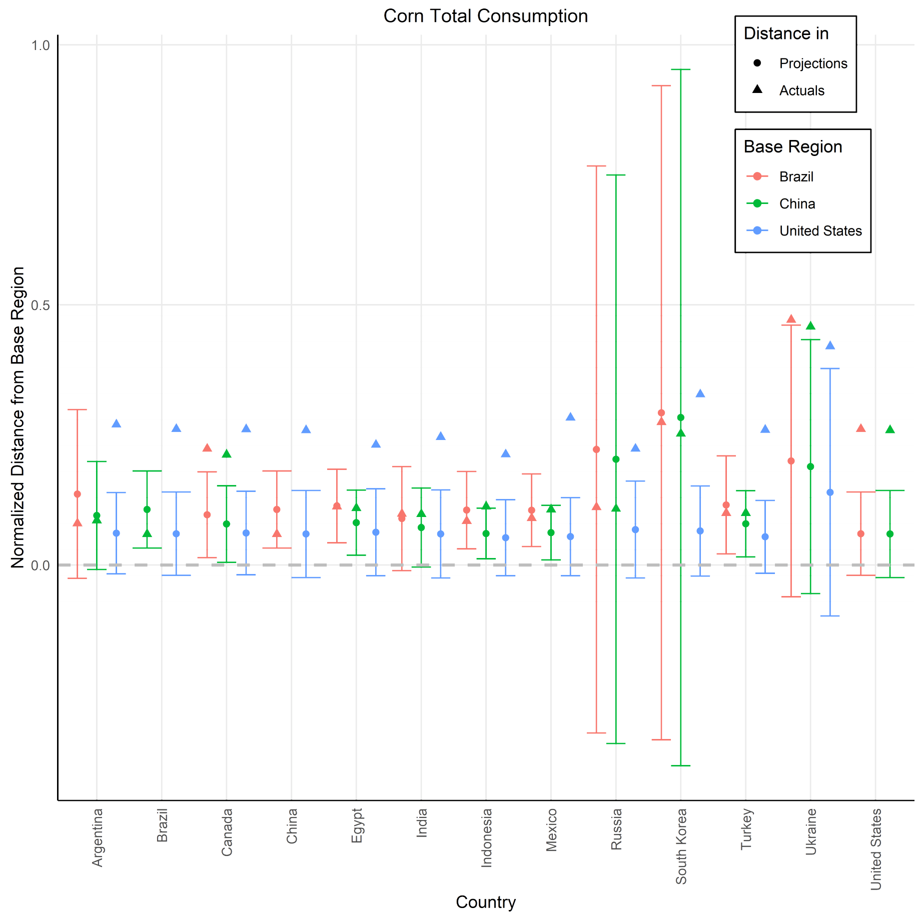

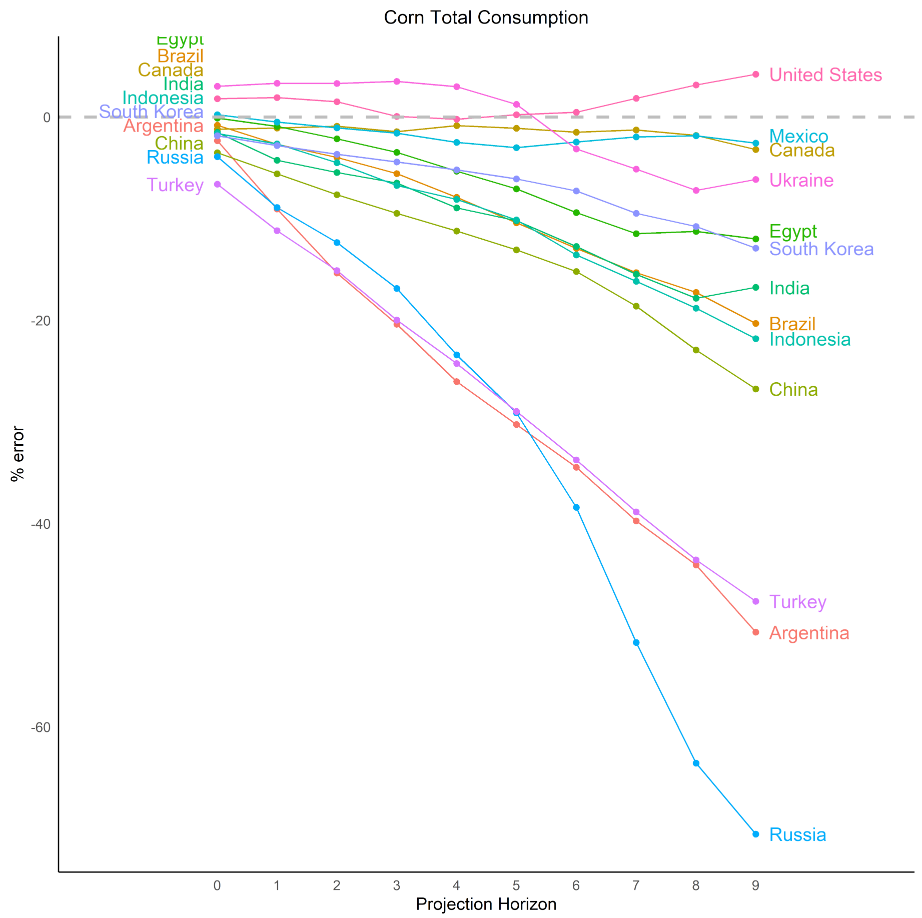

Figure 4: Corn Ending Stocks—Correlation and Error Calculations

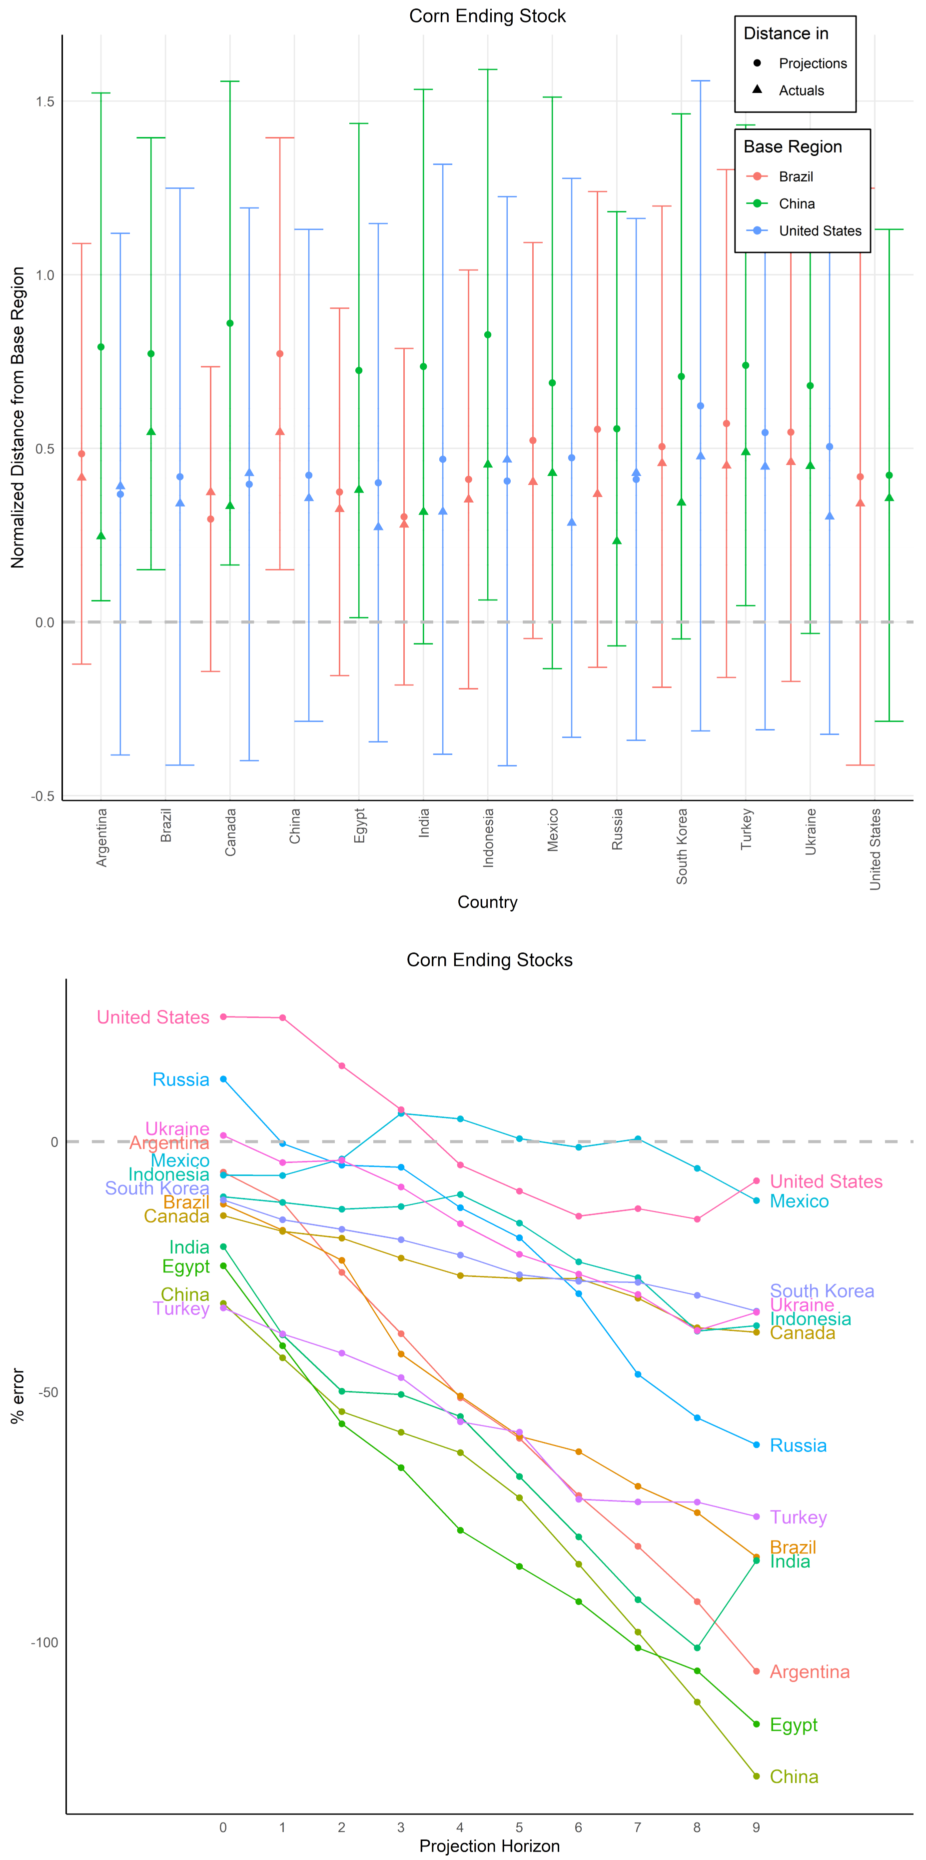

Figure 5: Corn Imports—Correlation and Error Calculations

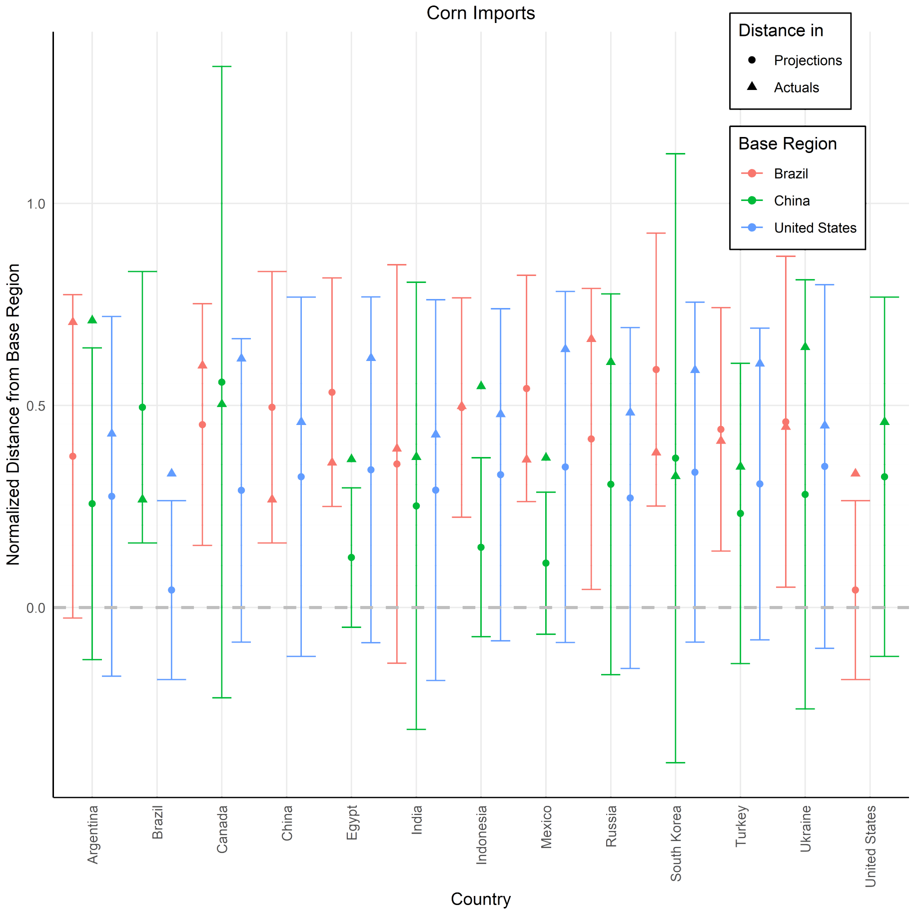

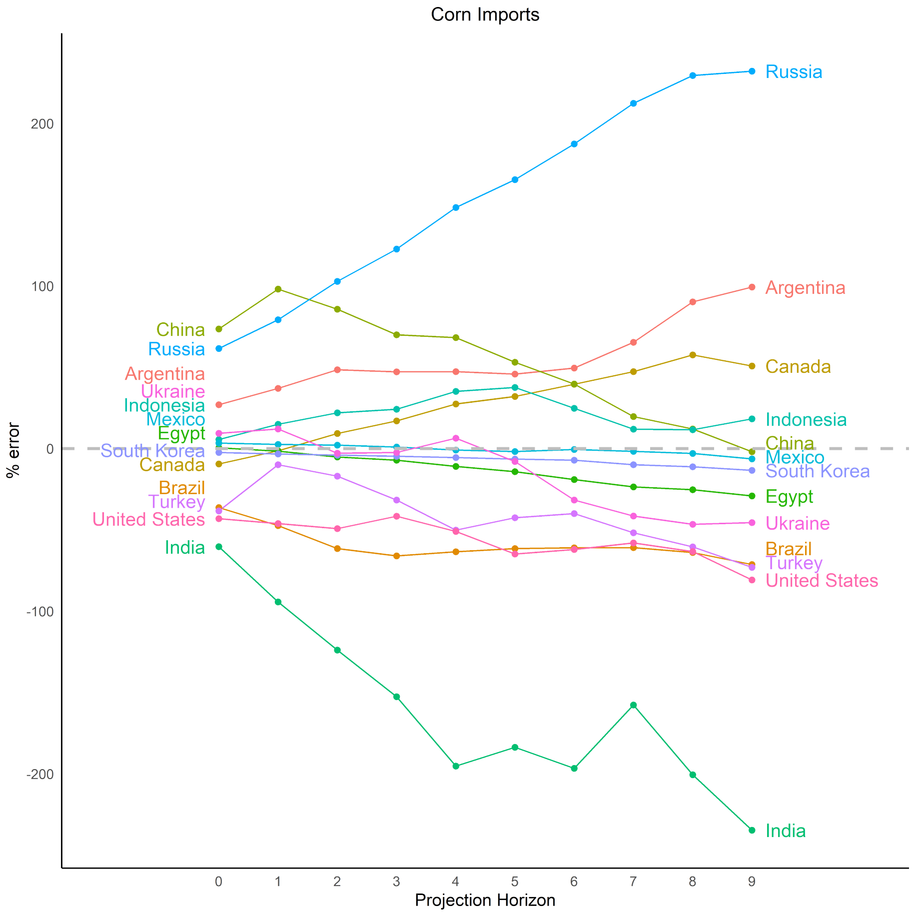

Figure 6: Corn Exports—Correlation and Error Calculations

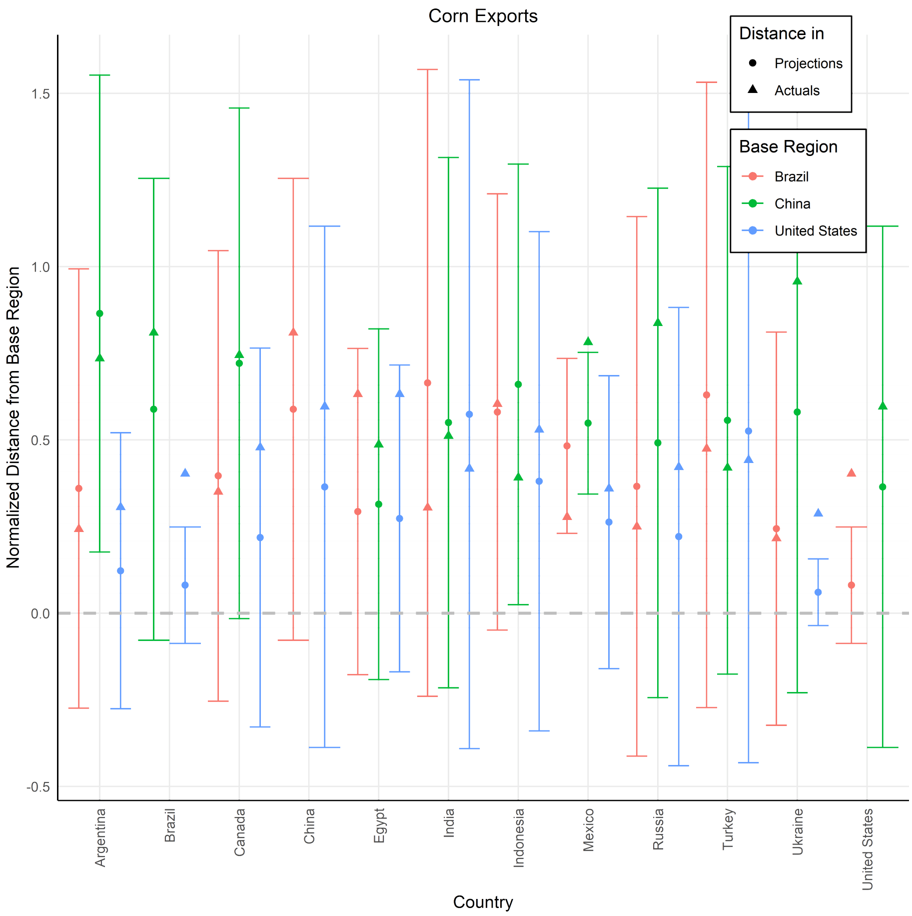

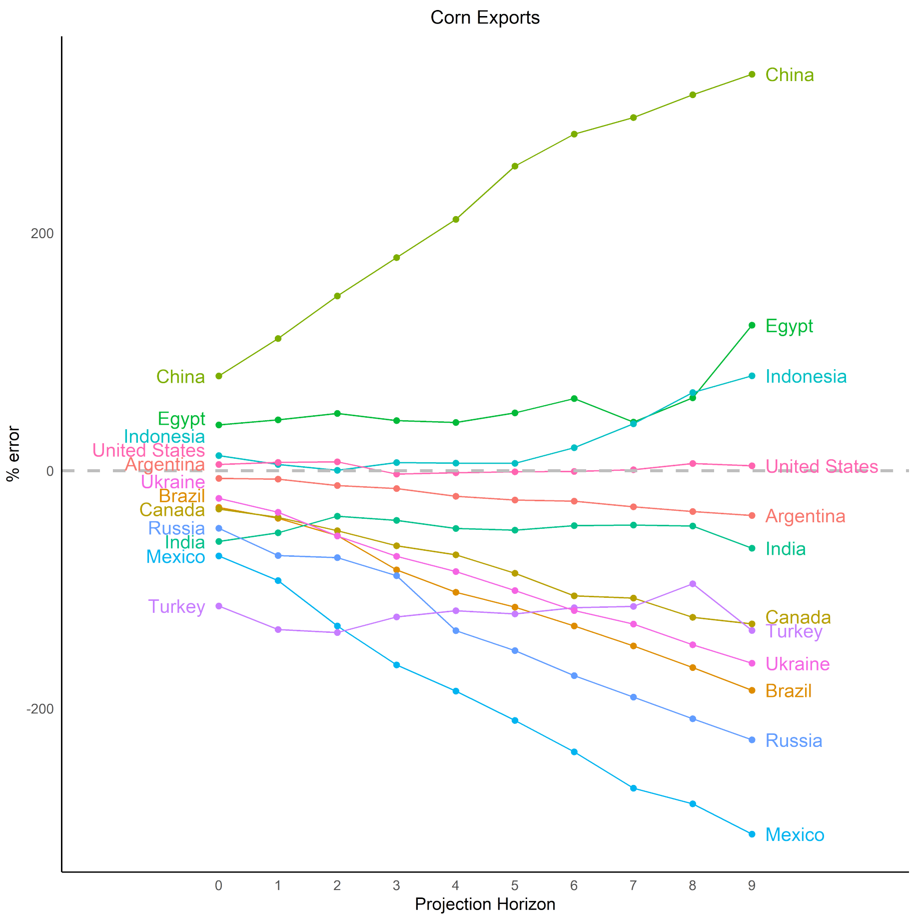

Figure 7: Soybeans Yield—Correlation and Error Calculations

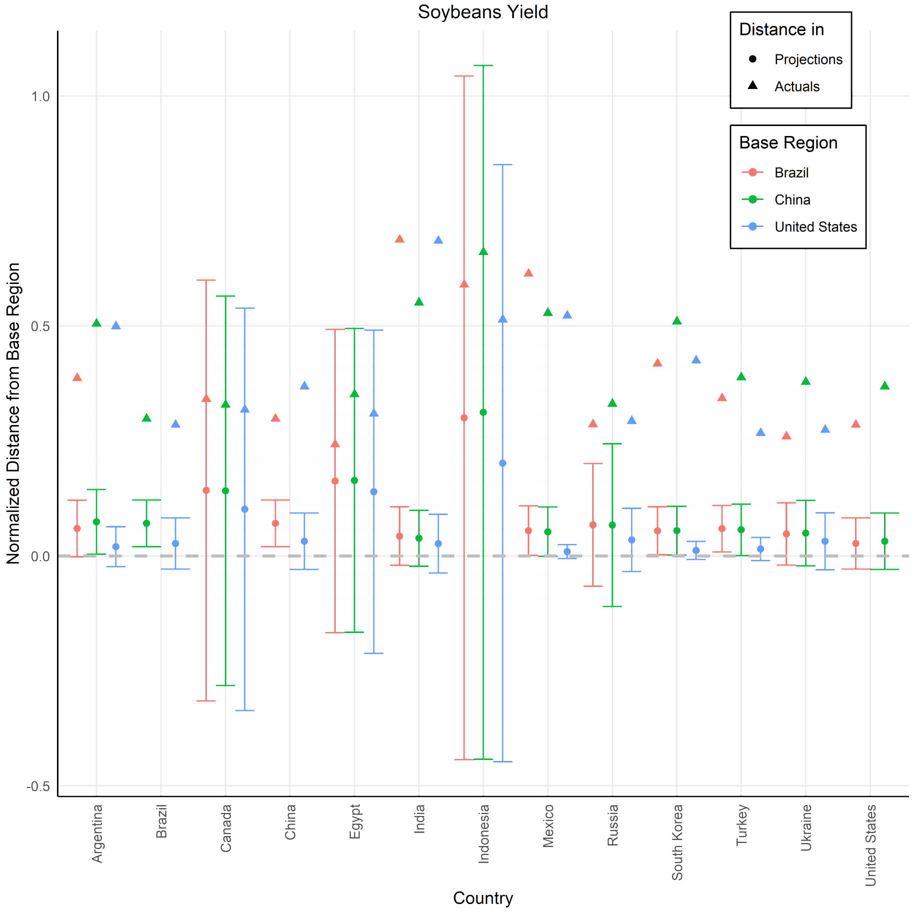

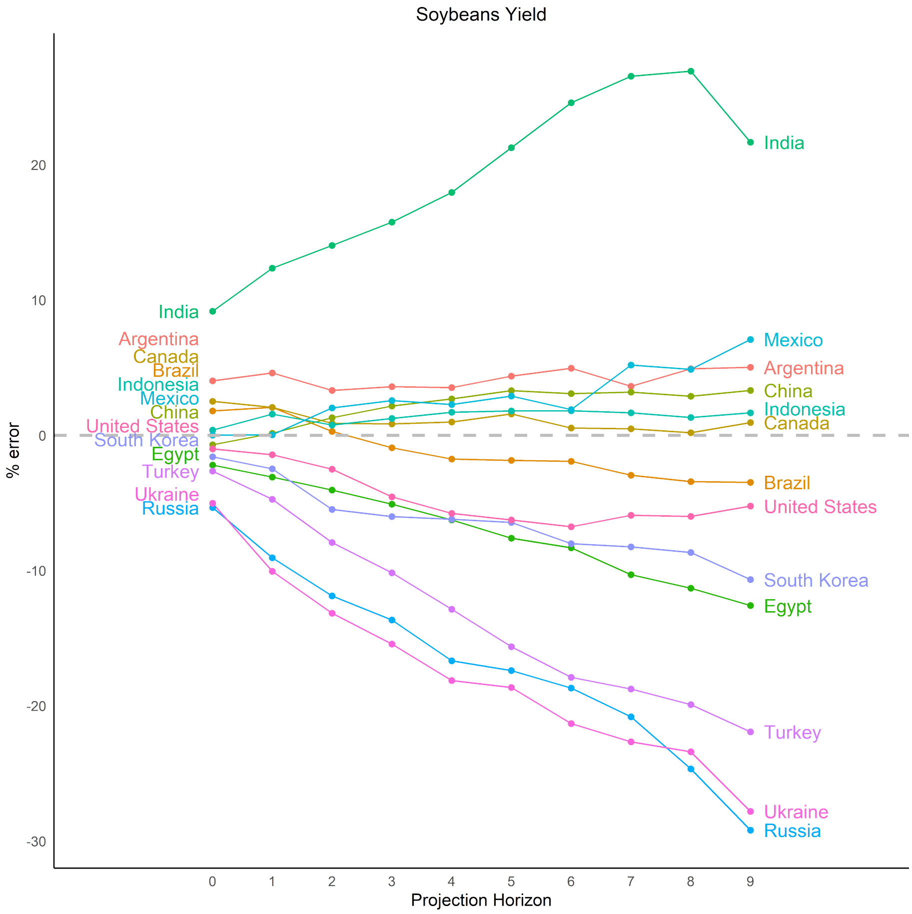

Figure 8: Soybeans Area Harvested—Correlation and Error Calculations

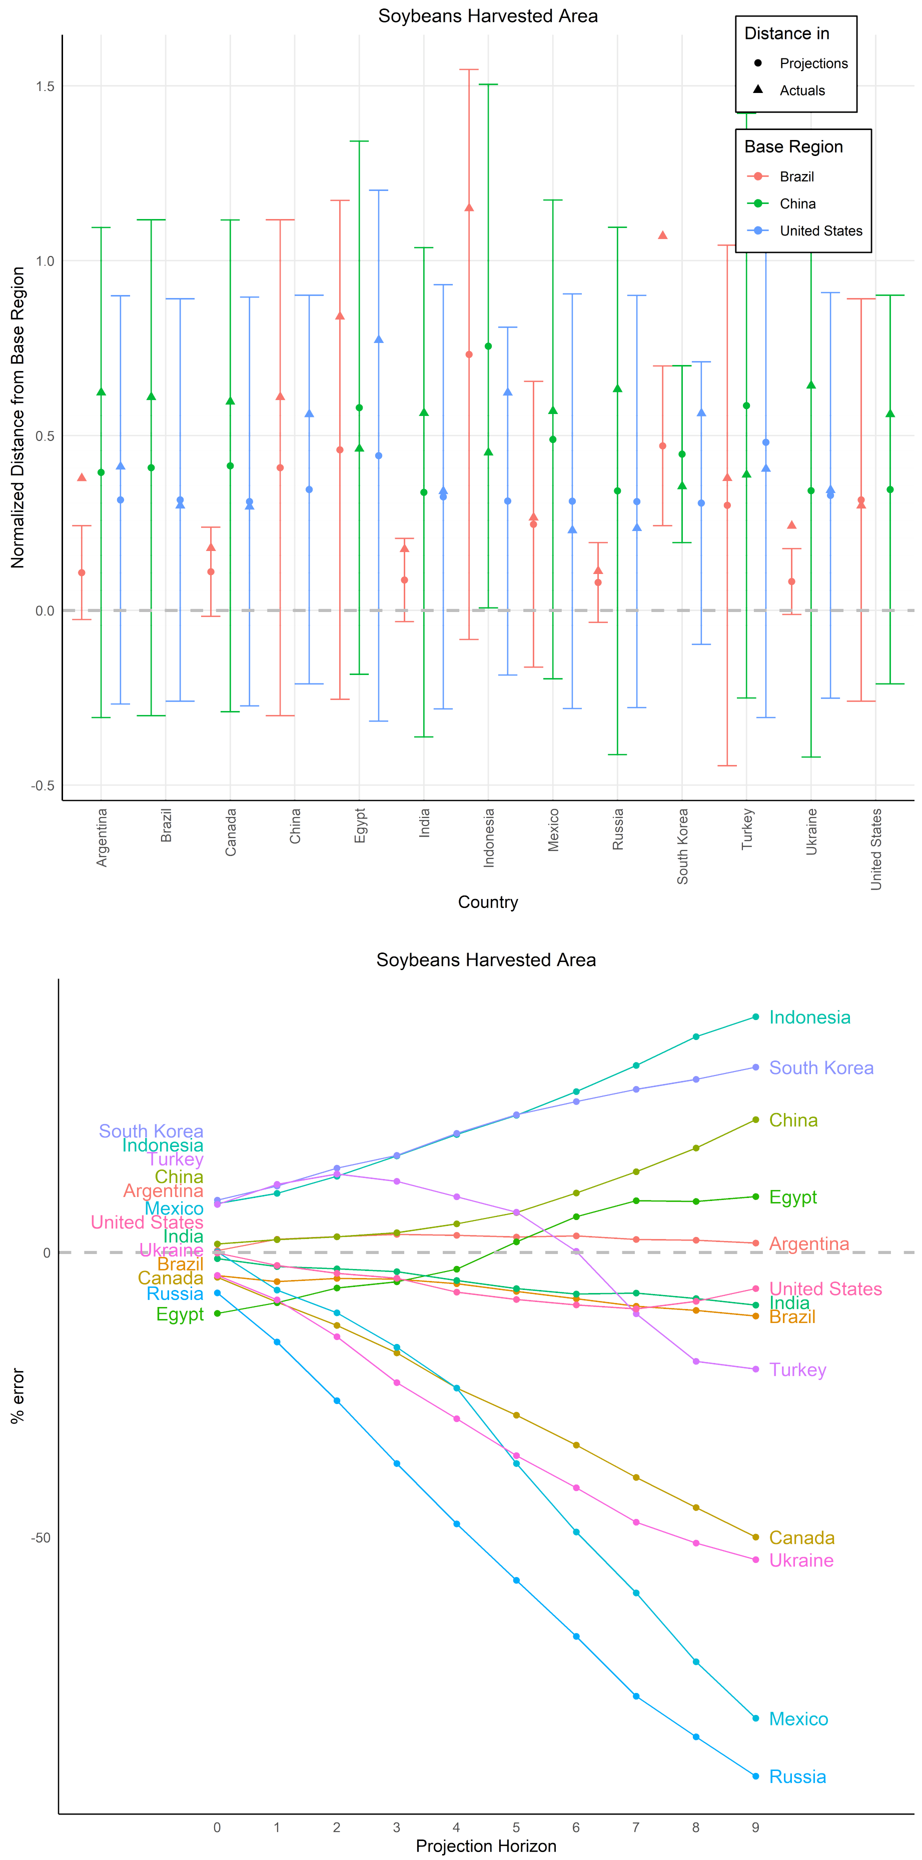

Figure 9: Soybeans Total Consumption—Correlation and Error Calculations

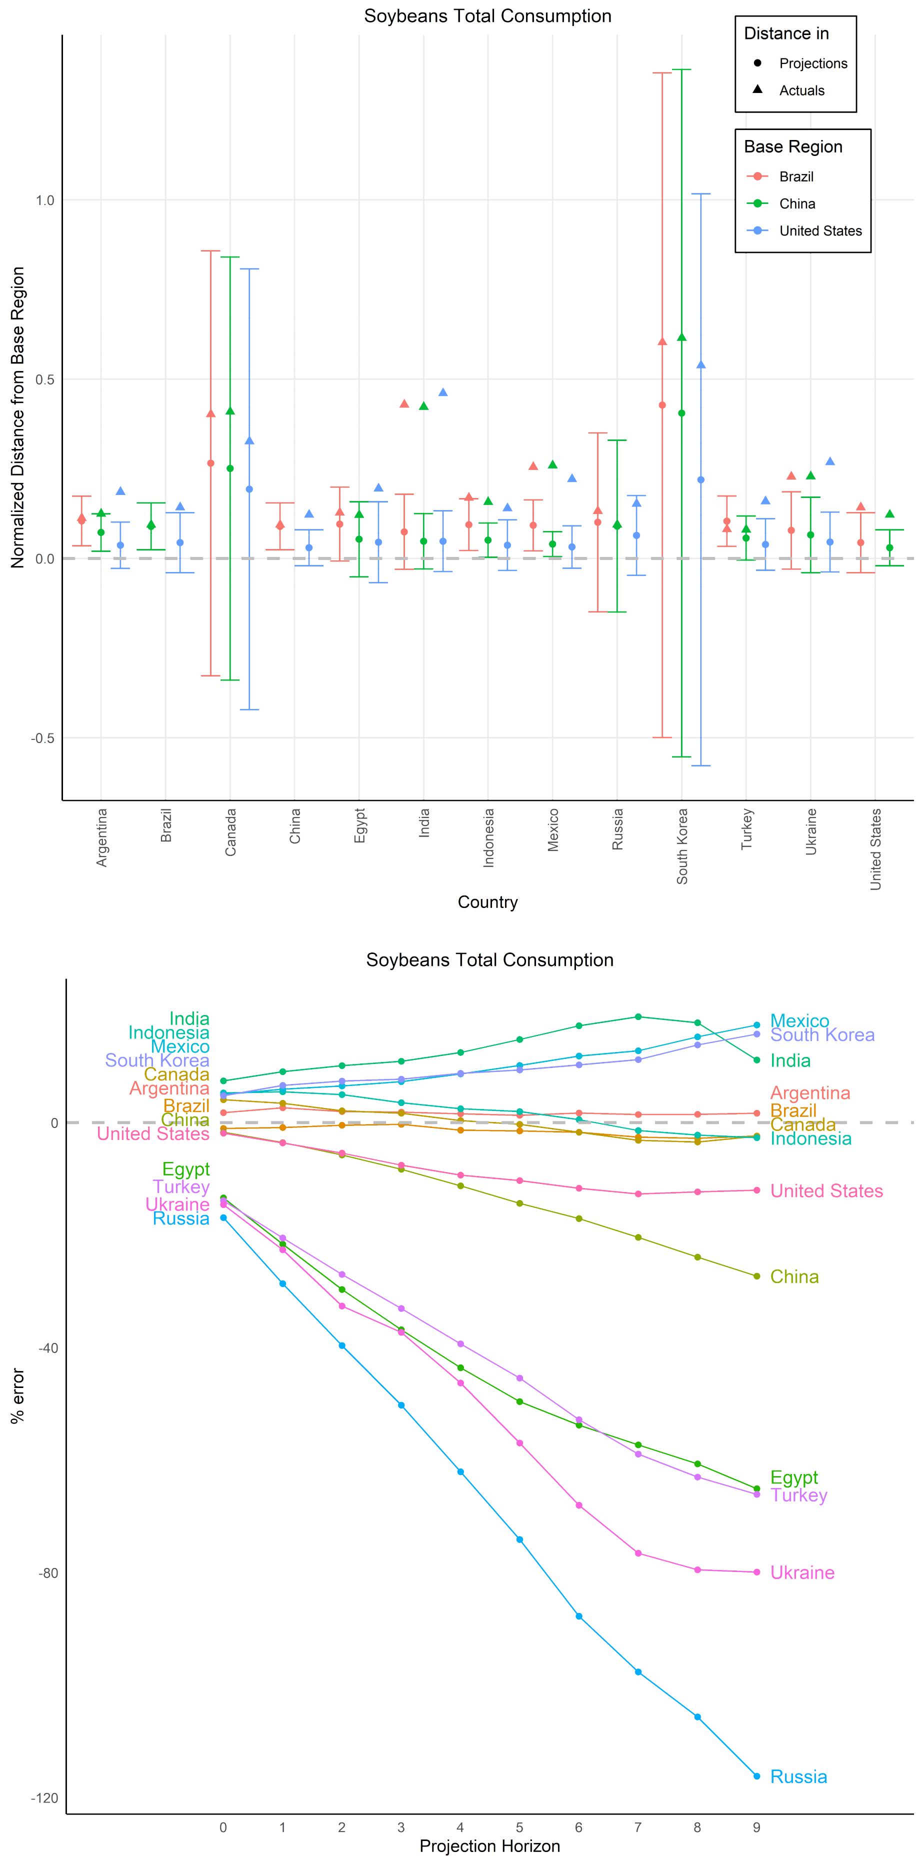

Figure 10: Soybeans Ending Stocks—Correlation and Error Calculations

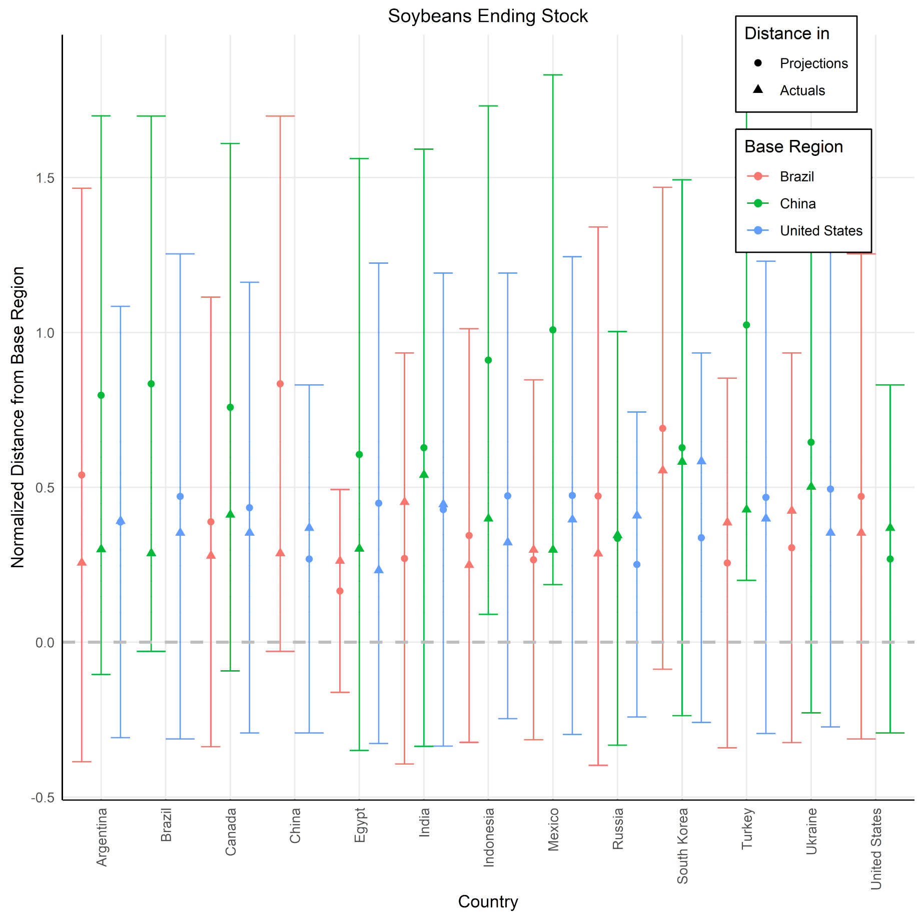

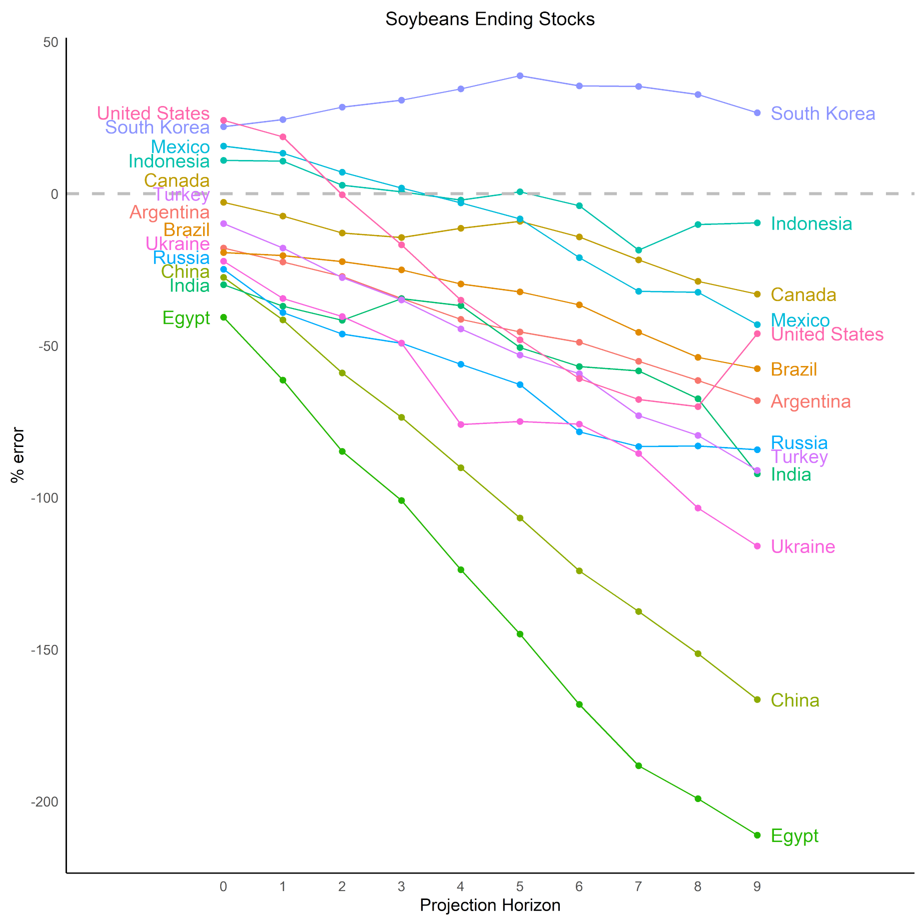

Figure 11: Soybeans Imports—Correlation and Error Calculations

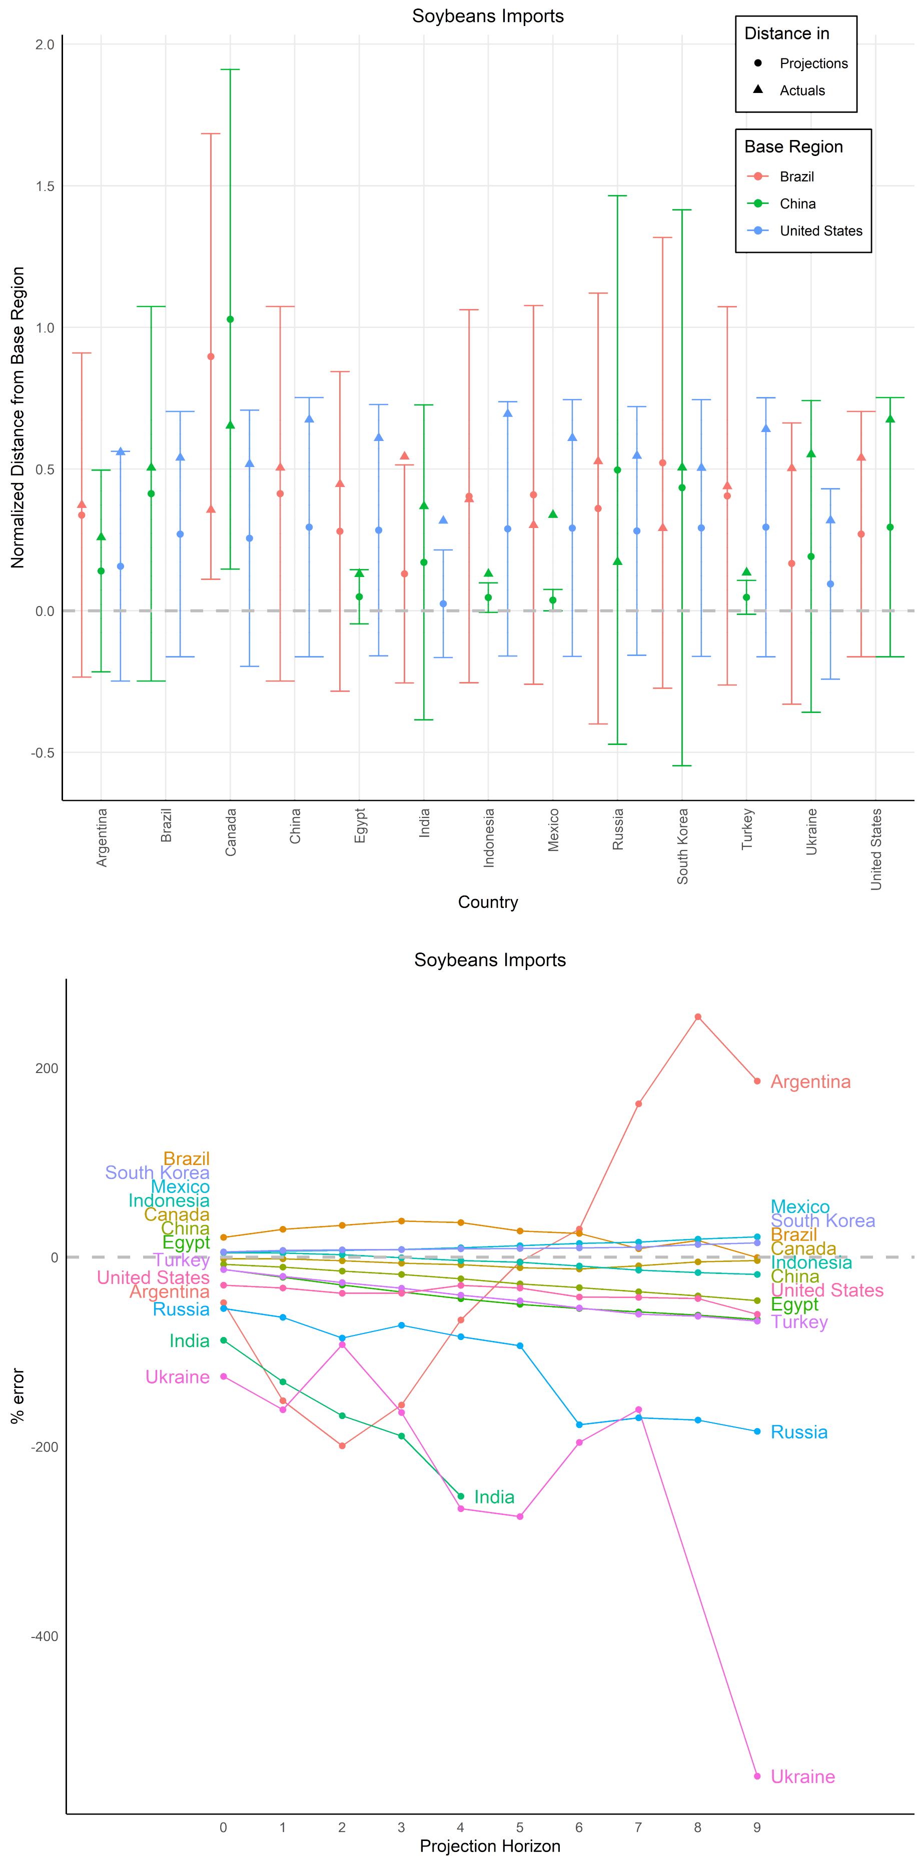

Figure 12: Soybeans Exports—Correlation and Error Calculations

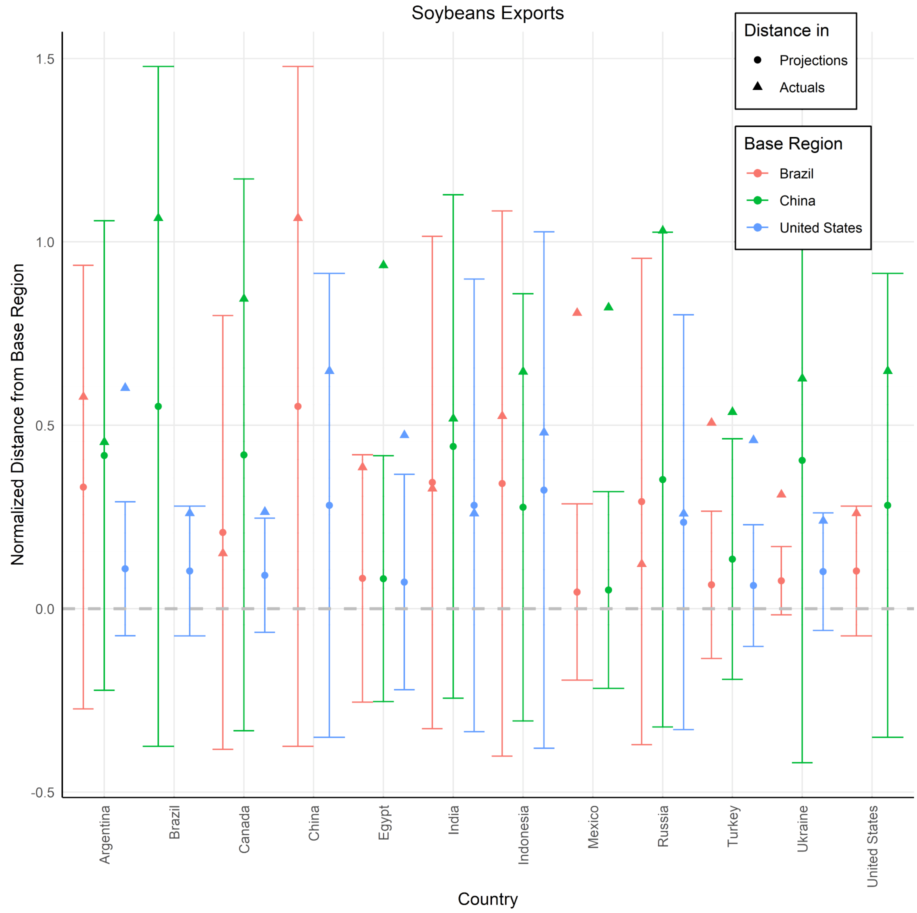

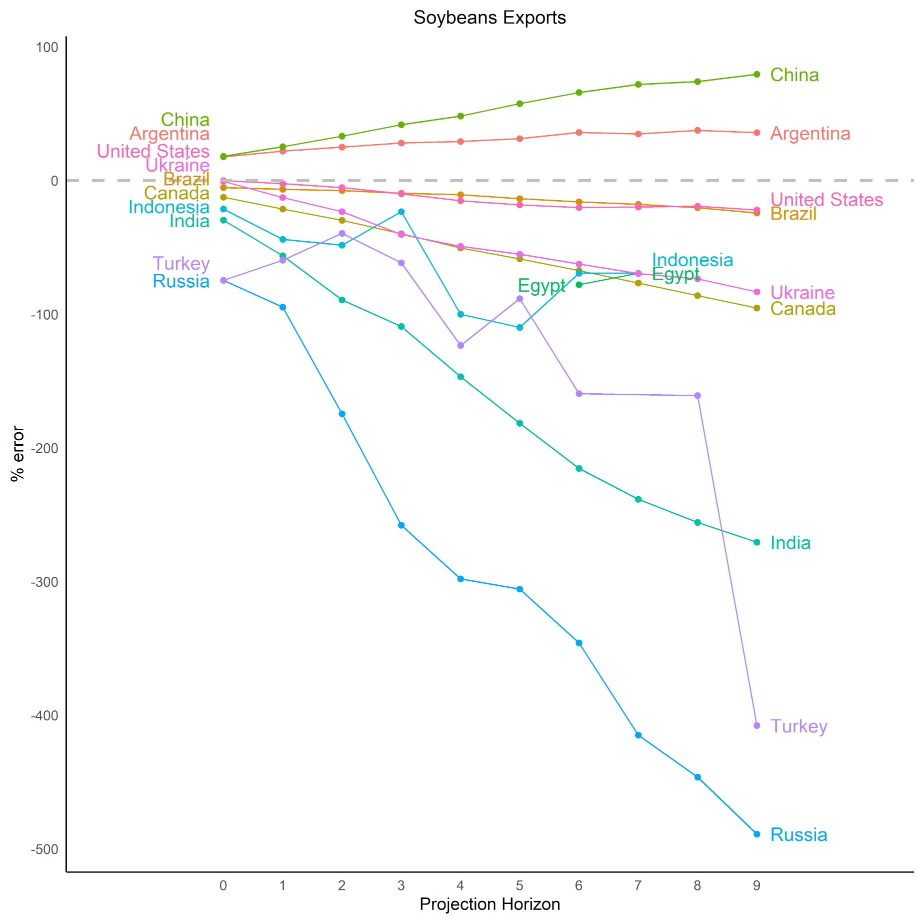

Figure 13: Wheat Yield—Correlation Estimates and Error Calculations

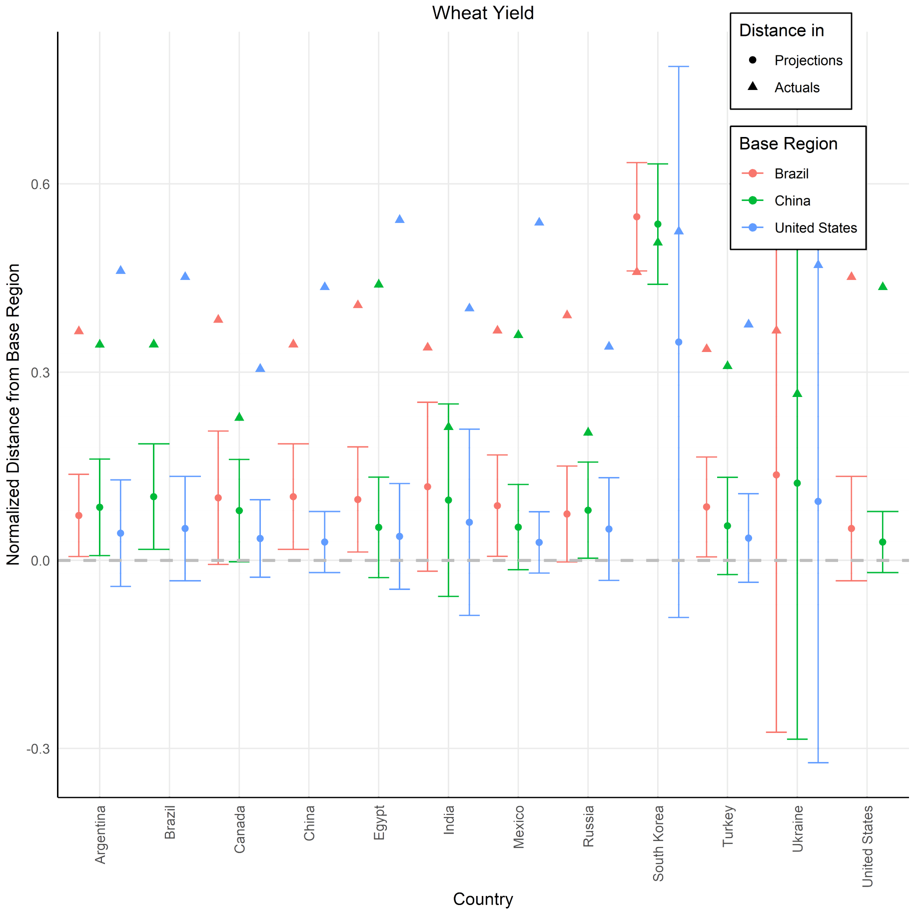

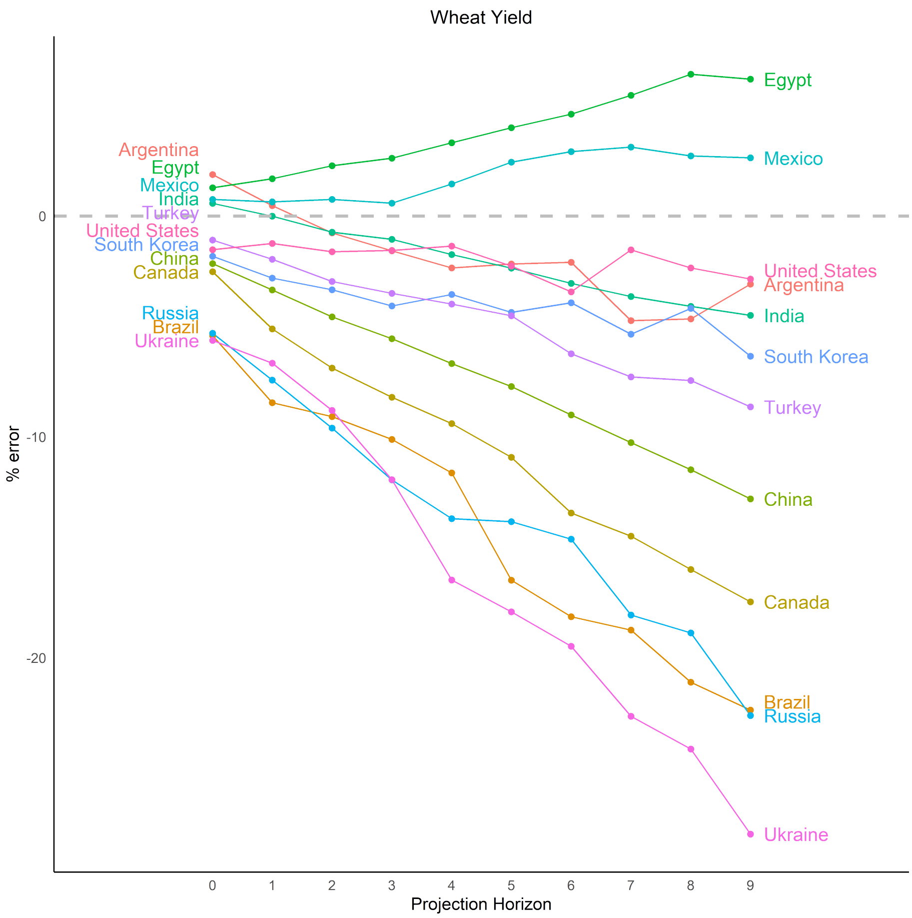

Figure 14: Wheat Area Harvested—Correlation and Error Calculations

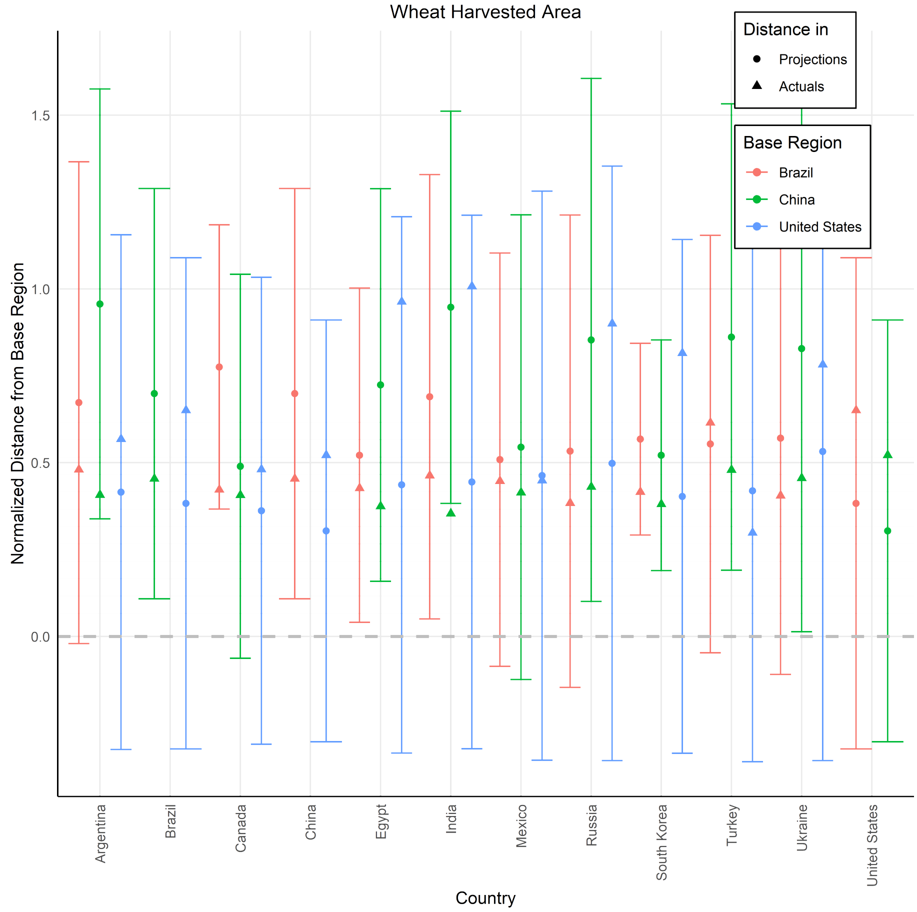

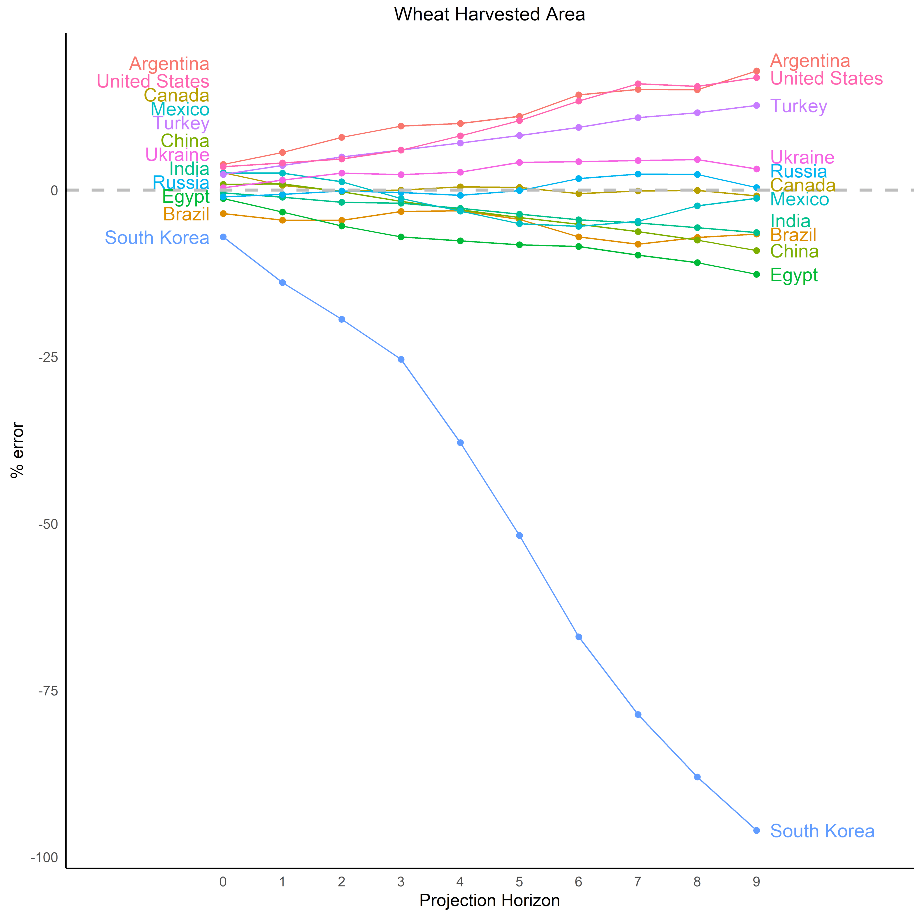

Figure 15: Wheat Total Consumption—Correlation and Error Calculations

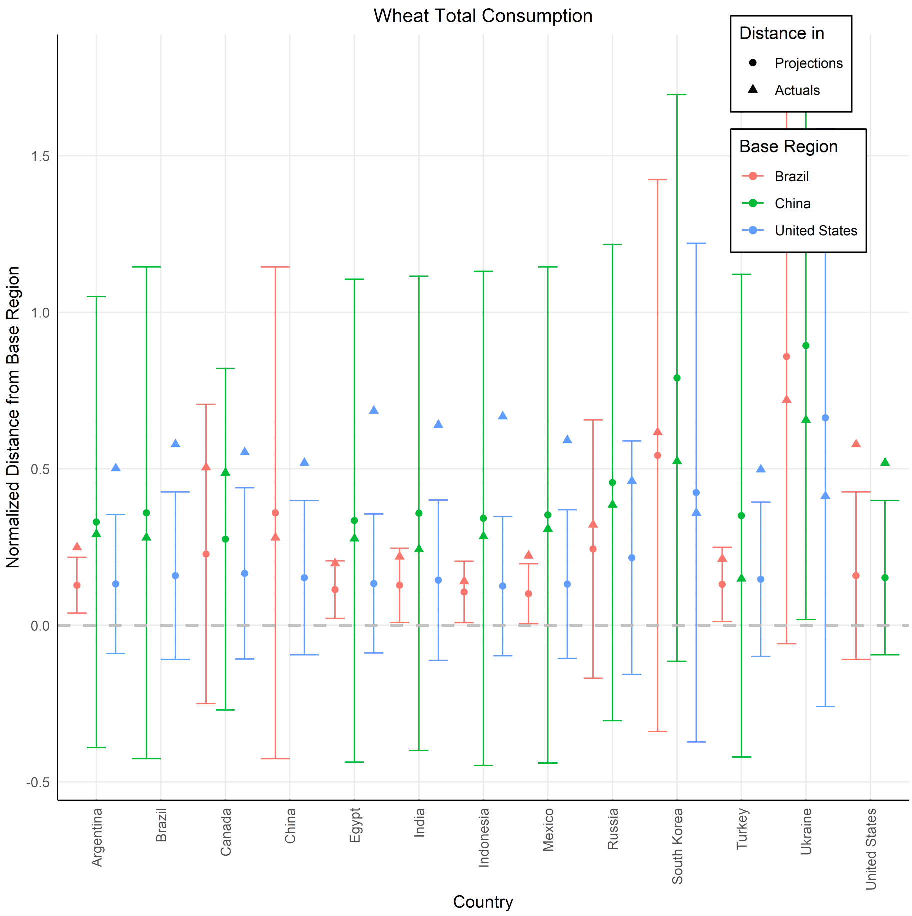

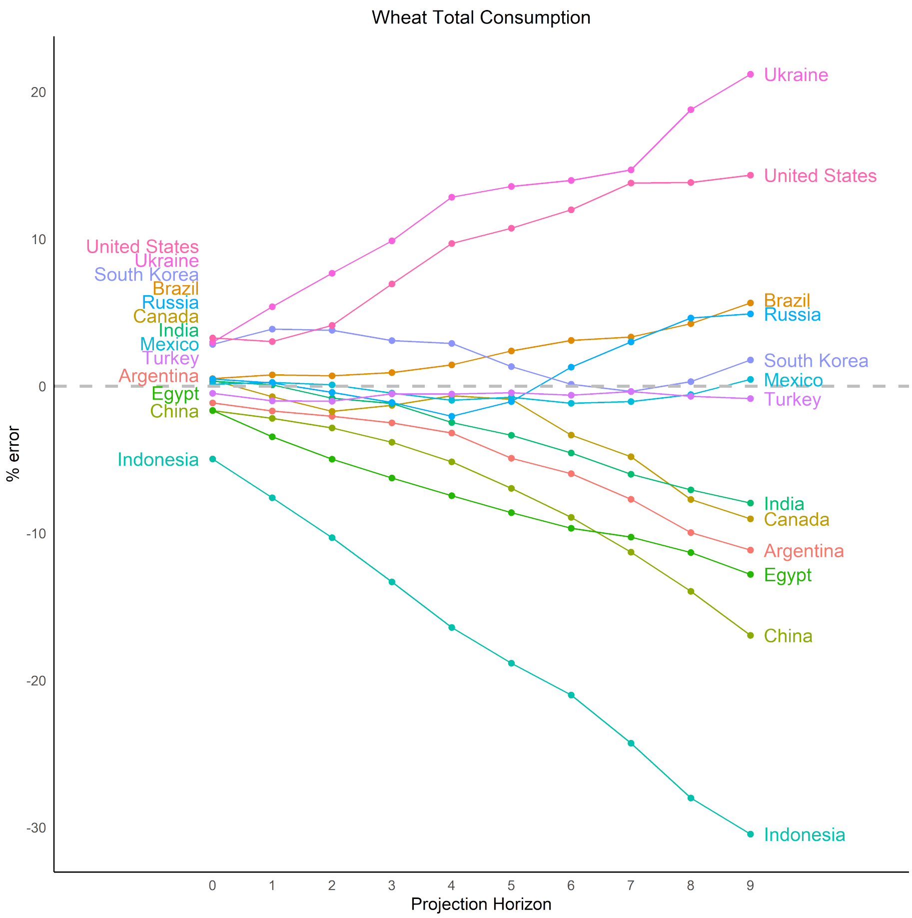

Figure 16: Wheat Ending Stocks—Correlation and Error Calculations

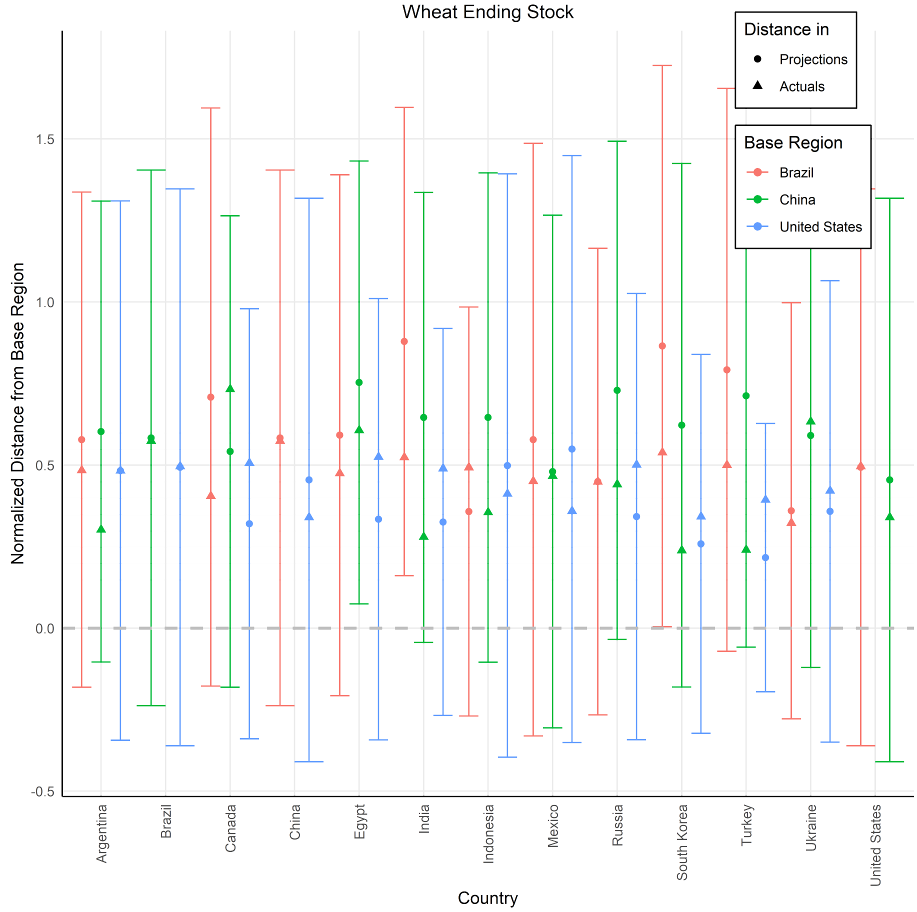

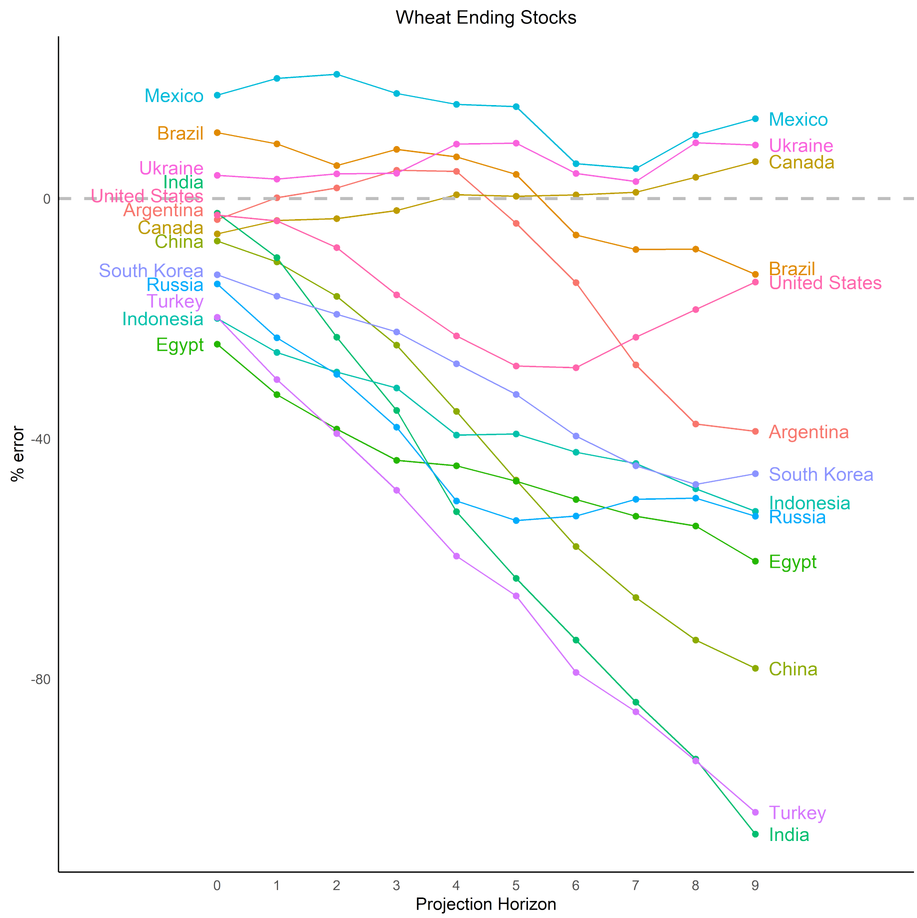

Figure 17: Wheat Imports—Correlation and Error Calculations

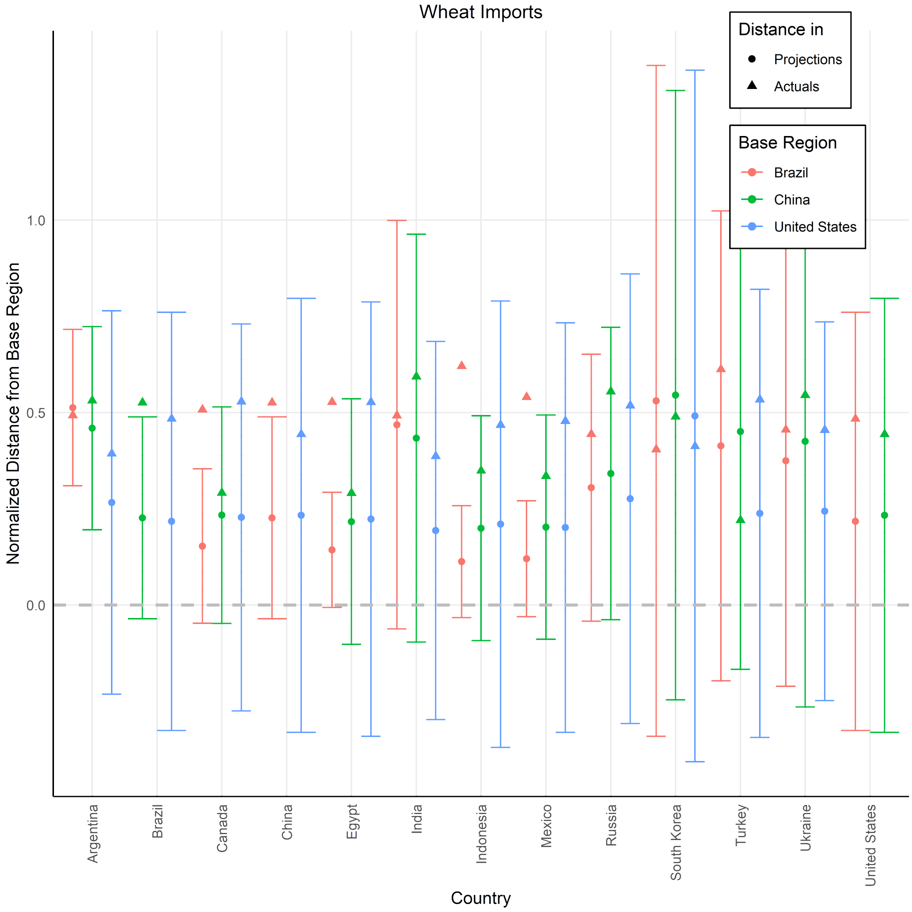

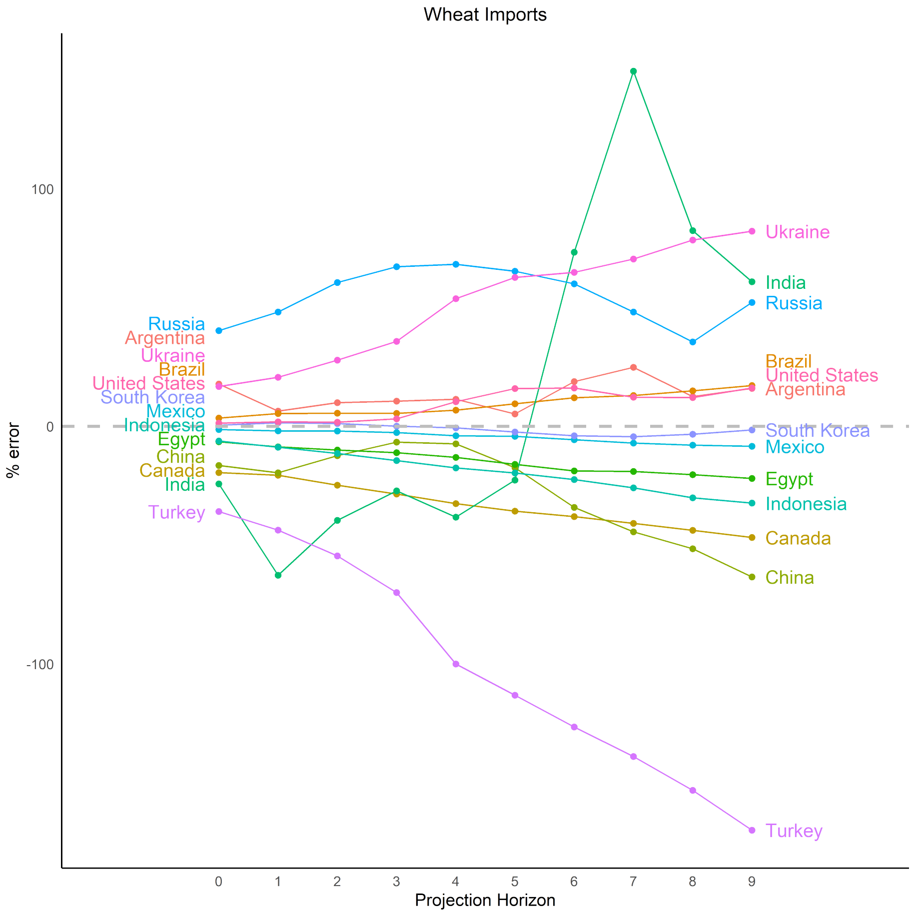

Figure 18: Wheat Exports—Correlation and Error Calculations

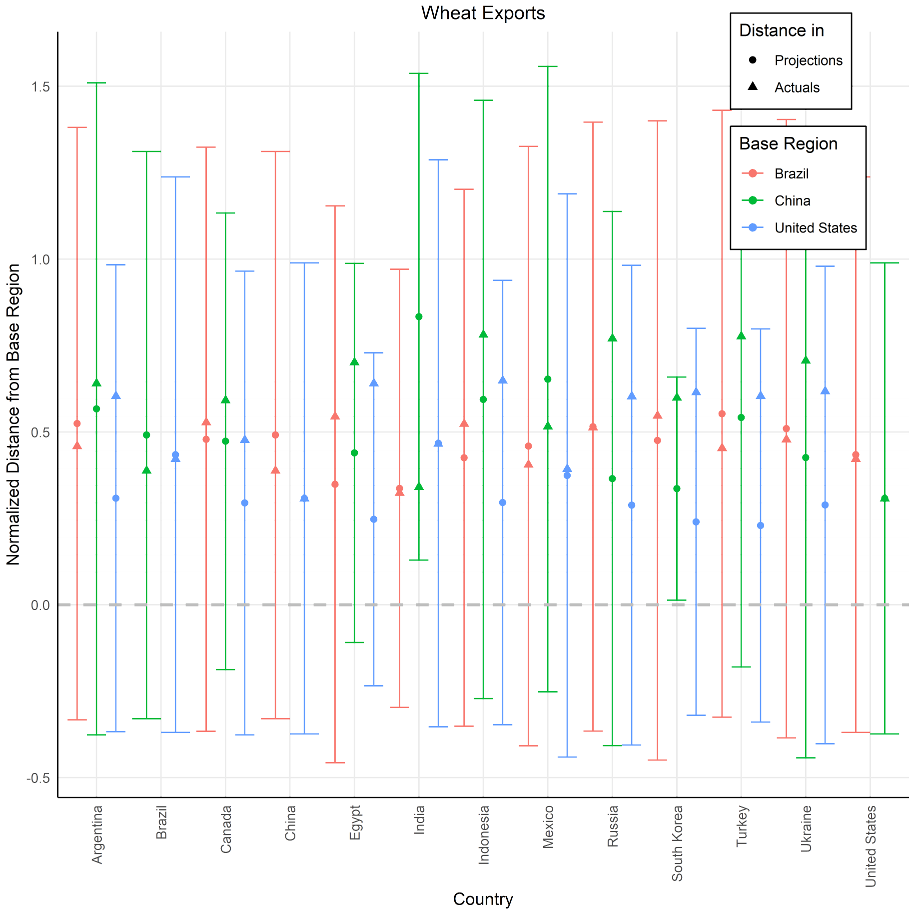

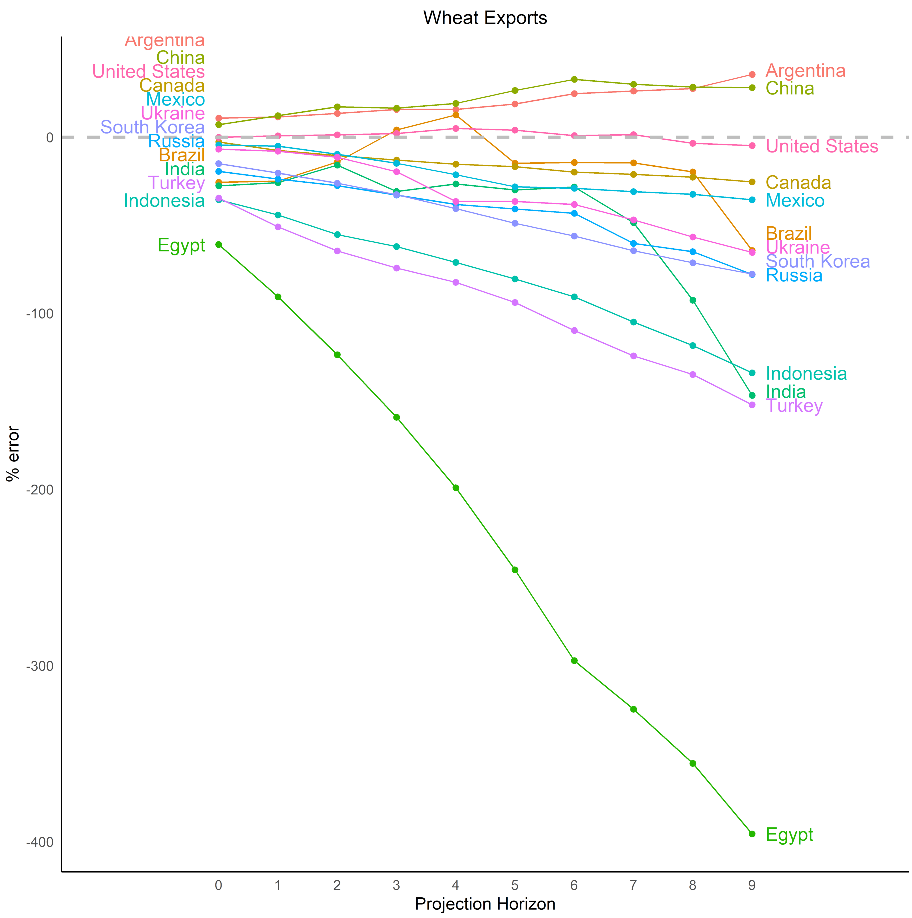

7 Tables

Table 1: Distance and Accuracy - Base country is United States

Variable	Yield	Area Harvested	Imports	Exports	Total Consumption	Ending Stocks
Corn						
Distance from Base$_r$ (logged)	0.0453***	-0.0617	-0.0231	0.0525	-0.0206*	-0.375***
	(0.0139)	(0.0551)	(0.0434)	(0.0788)	(0.0123)	(0.0793)
Soybeans						
Distance from Base$_r$ (logged)	-0.0226***	-0.0947	-0.8817***	0.257*	-0.2369***	-0.3801**
	(0.0052)	(0.1141)	(0.1442)	(0.1314)	(0.0864)	(0.1652)
Wheat						
Distance from Base$_r$ (logged)	-0.0207***	-0.0656***	0.0843***	-0.2666***	-0.0286***	-0.1793***
	(0.0062)	(0.0132)	(0.0188)	(0.0994)	(0.0073)	(0.0341)

This table shows the estimation results for equation $LoggedError_{rjh} = \beta_0 + \beta_1 log(distance(\hat{Y}_{r_{US}}, \hat{Y}_{rj})) + \epsilon_{rjh}$ for each crop and variable. Each column and panel shows the results for a separate regression for the crop-variable labeled in the table. Parentheses contain robust standard errors. *** p<0.01, ** p<0.05, * p<0.1

Table 2: Distance and Accuracy - Base country is China

Variable	Yield	Area Harvested	Imports	Exports	Total Consumption	Ending Stocks
Corn						
Distance from Base$_r$ (logged)	0.0569***	0.0213	0.0832**	-0.552***	-0.0382***	-0.2394***
	(0.0181)	(0.0164)	(0.0395)	(0.1698)	(0.0146)	(0.0788)
Soybeans						
Distance from Base$_r$ (logged)	-0.0364***	0.0391	-0.0834	-0.3317**	-0.2759***	-0.2515***
	(0.0092)	(0.0343)	(0.0724)	(0.1628)	(0.0898)	(0.0837)
Wheat						
Distance from Base$_r$ (logged)	-0.0161***	-0.1111***	0.2483***	-0.0804	-0.0046	-0.0851**
	(0.0062)	(0.0257)	(0.0474)	(0.0867)	(0.0063)	(0.0423)

This table shows the estimation results for equation $LoggedError_{rjh} = \beta_0 + \beta_1 log(distance(\hat{Y}_{r_{China}}, \hat{Y}_{rj})) + \epsilon_{rjh}$ for each crop and variable. Each column and panel shows the results for a separate regression for the crop-variable labeled in the table. Parentheses contain robust standard errors. *** p<0.01, ** p<0.05, * p<0.1

Table 3: Distance and Accuracy - Base country is Brazil

Variable	Yield	Area Harvested	Imports	Exports	Total Consumption	Ending Stocks
Corn						
Distance from Base$_r$ (logged)	0.0692***	0.0139	-0.1235***	0.2285**	-0.0345**	-0.2707***
	(0.0223)	(0.0088)	(0.0412)	(0.1010)	(0.0162)	(0.0677)
Soybeans						
Distance from Base$_r$ (logged)	-0.0293***	-0.0281	-0.9279***	0.047	-0.3403***	-0.183***
	(0.0093)	(0.0256)	(0.1763)	(0.0992)	(0.1093)	(0.0670)
Wheat						
Distance from Base$_r$ (logged)	-0.0249***	-0.0788***	0.1848***	-0.3417**	-0.0413***	-0.0486
	(0.0075)	(0.0190)	(0.0365)	(0.1518)	(0.0090)	(0.0320)

This table shows the estimation results for equation $LoggedError_{rjh} = \beta_0 + \beta_1 log(distance(\hat{Y}_{r_{Brazil}}, \hat{Y}_{r_j})) + \epsilon_{rjh}$ for each crop and variable. Each column and panel shows the results for a separate regression for the crop-variable labeled in the table. Parentheses contain robust standard errors. *** p<0.01, ** p<0.05, * p<0.1