Forecasting end-of-season peanut yield using mid-season crop condition rating signals

by

Taylor Mock Cunningham

B.S., Abraham Baldwin Agricultural College, 2019

A THESIS

Submitted in partial fulfillment of the requirements

for the degree

MASTER OF AGRIBUSINESS

Department of Agricultural Economics

College of Agriculture

KANSAS STATE UNIVERSITY

Manhattan, Kansas

2026

Approved by:

Major Professor
Terry W. Griffin, PhD

ABSTRACT

Peanut production in the United States is concentrated in three geographically distinct regions (southeast, southwest, and Virginia-Carolinas), positioning it as a specialized commodity relative to other major crops. Due to the peanut industry's relatively small and specialized structure, limited literature exists on yield forecasting for peanuts. Substantial research exists on yield forecasting using crop condition ratings for corn and soybeans; however, a notable gap remains for peanuts. Despite being the seventh most valuable crop in the United States, limited empirical work has focused on developing yield forecasting models for the peanut industry, underscoring the importance of this study (USDA NASS 1999). A key challenge in peanut yield forecasting is the limited availability of publicly accessible data. As a privately traded commodity, peanuts rely primarily on USDA reporting, which is consistently available only at the state level.

Data were obtained from the USDA National Agricultural Statistics Service (NASS) based on survey reports of peanut crop conditions and yields for major producing states, including Alabama, Florida, Georgia, the Carolinas, and Texas. Structural break analysis indicated that observations prior to 2012 were not informative for forecasting purposes; therefore, the study period was restricted to 2012–2025. Forecast performance was evaluated weekly throughout the peanut growing season using multiple modeling approaches, including the Good + Excellent, NCI, NCI 2.0, Linear, and Stress weighting models.

Results indicate that the Good + Excellent Scheme consistently outperformed alternative modeling approaches across the evaluated weeks. By week 40 of the calendar

year (early October), the Good + Excellent Scheme achieved the lowest forecast error (RMSE = 0.16; MAE = 0.13), establishing it as the most accurate model among those tested. Forecasts became reliable, beginning around week 32 and remained statistically significant through the end of the growing season. Correlations between Good + Excellent condition ratings and final yield were likewise significant from week 32 onward based on historical crop-year data.

Yield forecasting plays a critical role in the peanut industry by informing marketing and operational decision-making. Accurate forecasts provide value to marketing teams by enabling early assessment of crop quality and anticipated supply for upcoming sales periods. Forecast information also supports planning for storage and processing capacity at shelling and manufacturing facilities. Together, projections of crop size and quality guide marketing strategies and market allocation decisions for the subsequent production year.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

The author wishes to thank my husband, Colby Cunningham, and my daughter, Murphy, for their unwavering support throughout the process of writing this thesis. Their encouragement and patience made this accomplishment possible. I extend my utmost gratitude to my major professor, Dr. Terry Griffin, for the countless hours of guidance, constructive feedback, and for consistently reaching out with support, even when life presented competing priorities. I also acknowledge the USDA National Agricultural Statistics Service (NASS) for collecting and publicly providing the data that made this research possible. Completing this work has been a true privilege. Thank you to Dr. Yeager and Dr. Coffey for agreeing to be on my committee.

CHAPTER I: INTRODUCTION

Commodity growing seasons are characterized by substantial uncertainty, particularly prior to harvest. Crop forecasting serves as a tool to reduce this uncertainty by providing early information on expected production outcomes. Peanuts are a privately traded commodity, limiting the availability of industry-level information. Consequently, publicly available data from the U.S. Department of Agriculture (USDA), largely reported at the state level, plays a critical role in informing market participants. Forecasting the peanut crop therefore provides valuable insights to a wide range of stakeholders across the industry. In-season crop forecasting is important for peanut processing firms for several operational and strategic reasons. Forecasts allow processing facilities to plan storage capacity, processing schedules, and labor requirements in advance of harvest. Forecasting also provides marketing teams with early information regarding expected crop quality and available volumes, enabling more informed decisions about product allocation, forward sales, and whether additional markets can be served in the upcoming year. Within the production and procurement segment of the industry, in-season forecasts assist buying points in forming market expectations and planning purchases of peanuts that have not yet been contracted with producers.

This study evaluates the use of weekly USDA National Agricultural Statistics Service (NASS) crop condition ratings during the growing season to forecast final peanut yields within a structured forecasting framework. The USDA releases the Crop Progress and Condition Report on a weekly basis throughout the growing season (USDA NASS 2018). These reports are based on surveys completed by county extension agents, Farm Service Agency personnel, conservation employees, and farmers, and are published from

March through November. As an example, the most recently available weekly crop condition ratings for Georgia are visually presented in Figure 1.1. Five conditions are reported and defined by USDA as follows:

- Excellent: ideal growing conditions, yield projected to be above normal, little to no stress
- Good: normal growing conditions, yield projected to be normal, some disease and damage, little stress
- Fair: less than normal growing conditions, yield projected to be less than average, more prevalent stress
- Poor: dismal growing conditions, drought, heavy disease and insect pressure, substantial stress to crop
- Very Poor: growing conditions project extreme yield loss, extreme disease and insect damage, highest level of crop stress

Figure 1.1: Example Crop Conditions, 2025 Georgia Peanuts

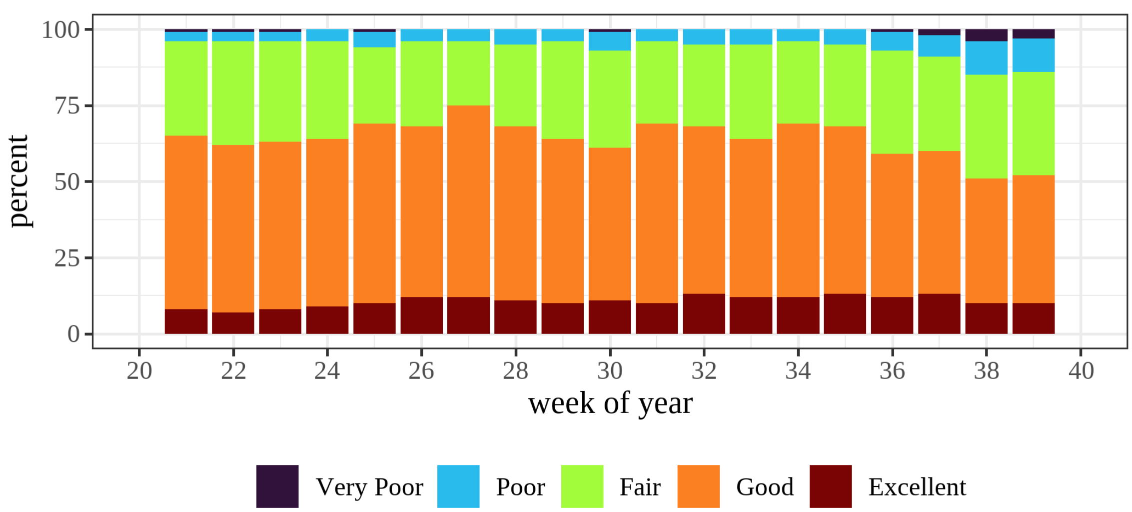

data source: USDA NASS 2026

The five categories are rated into percentages that must sum up to 100. Six states that account for the majority of U.S. peanut planted acreage are Georgia, Alabama, Florida, South Carolina, North Carolina, and Texas.

Despite the widespread availability of weekly crop condition ratings, their effectiveness as in-season forecasting tools for peanuts is not publicly established. Limited research has examined when during the growing season these ratings become informative for predicting final peanut yields or how alternative aggregation approaches influence forecast performance across major peanut-producing states.

While crop condition ratings provide timely information on crop health, they must be transformed into a quantitative framework to generate yield forecasts. This study incorporates weekly crop condition ratings into a forecasting model by aggregating condition categories and evaluating their relationship with final peanut yields throughout the growing season. Multiple aggregation approaches are examined to assess how different representations of crop conditions influence forecast performance over time. The literature review builds on existing research from other commodity groups that examine the use of crop condition ratings for yield forecasting.

1.1 Literature Review

This section reviews the literature related to yield forecasting using crop condition reports and their application in agricultural analysis. The section is organized into three subsections. The first subsection is how crop condition reports are used as information. The second subsection examines the use of crop condition ratings as yield predictors. The final subsection reviews how other commodity groups forecast yields using condition ratings and discusses their relevance to the research question.

1.1.1 Crop Condition Reports as Market Information

To use crop condition ratings in yield forecasting, it is important to understand how these ratings are generated and interpreted. USDA NASS emphasizes that crop condition reports are intended to provide timely information on crop status, rather than direct estimates of final yield (USDA NASS 1999). Condition information is published weekly and made publicly available throughout the growing season. Respondents evaluate crops relative to what is considered normal for the stage of development in their local area. Although condition ratings are inherently subjective, they provide a consistent snapshot of crop health and production expectations as the season progresses.

Bain and Fortenbery (2017) examine the informational role of crop condition reports in winter wheat markets. The authors construct a crop condition index based on the distribution of planted acreage across condition categories and find that the informational content of crop condition reports increases as the growing season progresses. Although the study finds limited price responses to changes in reported conditions, the results indicate that crop condition reports are closely monitored by market participants. Overall, the findings suggest that while crop condition reports provide useful information about production prospects, they do not consistently generate new information sufficient to influence wheat prices.

1.1.2 Crop Condition Ratings as Yield Predictors

The relationship between crop condition ratings and end-of-season yields for soybeans and corn in Iowa is examined by Tang (2016). The study evaluates whether the USDA NASS Good + Excellent (G+E) condition ratings can be used as predictors of yield and identifies the point during the growing season at which these ratings become most informative for forecasting purposes. Historical week-specific crop condition and yield data obtained from the USDA NASS QuickStats database were used to assess the evolution of forecast performance over time (USDA NASS 2026). Results indicate that early-season yield forecasts are highly variable and exhibit limited predictive power; however, forecast accuracy improves as the growing season progresses. Model performance increases notably during reproductive growth stages, when yield formation begins. In the final reporting week, the soybean model explained more than 80% of yield variability, while the corn model explained over 60% of yield variability. Overall, the findings demonstrate a statistically meaningful relationship between crop condition ratings and commodity yields, with predictive strength increasing later in the growing season.

Stokes (2024) proposed a Markov chain framework for modeling the week-to-week transition behavior of USDA crop condition ratings, finding that forecasted conditions generated from the model provided useful input for estimating final yields prior to harvest. This work demonstrates that the temporal structure of crop condition reports contains additional information beyond any single week's rating.

Irwin and Hubbs (2018) evaluate the forecast accuracy of yield models based on Good + Excellent crop condition ratings and assess the presence of forecast bias and error. Using corn and soybean crop condition data, the analysis finds that Good + Excellent ratings exhibit limited predictive value early in the growing season, with forecast accuracy improving as the season progresses and additional information becomes available. However, the study notes that systematic bias can occur in years characterized by adverse growing conditions. Overall, the results suggest that while crop condition ratings contain useful forecasting information, Good + Excellent ratings alone are insufficient, and incorporating additional variables is recommended to improve forecast performance.

Li et al. (2026) compared crop condition forecasting models for corn and soybeans from 1986 to 2022 and reported that no model demonstrated superior forecast accuracy. In addition to their own models (Irwin and Good 2017a, 2017b), they evaluated more complex Schemes proposed by Beguería and Maneta (2020) and Bain and Fortenbery (2017). Simpler models performed as well as more complex, computationally demanding alternatives, suggesting that added complexity does not improve forecast reliability for these commodities.

1.1.3 Yield Forecasting for Other Commodities

Yield forecasting using crop condition ratings has been examined extensively for several major commodities, particularly corn and soybeans. In contrast, peanuts are a privately traded commodity, and relatively little published research has evaluated the use of crop condition ratings for peanut yield forecasting.

Irwin and Good (2017a) extend prior research on the use of USDA crop condition ratings for corn and soybean forecasting by examining how within-season ratings can be adjusted to improve their usefulness in forming expectations about final condition ratings and yield. The study finds that early-season Good + Excellent ratings systematically overestimate the final end-of-season Good + Excellent percentages. The magnitude and seasonal pattern of this bias are analyzed, revealing that Good + Excellent ratings for both commodities exhibited a declining trend from 1986 to 2016. In addition, a structural shift toward increased upward bias in early-season ratings is identified beginning in 1999, although no clear fundamental explanation for this change is provided.

Irwin and Good (2017b) examine when within-season Good + Excellent crop condition ratings begin to provide meaningful information for yield forecasting. The results indicate that early-season crop conditions and final yields are weakly correlated. However, as the growing season progresses, the correlation between condition ratings and yield strengthens. For corn, correlations increase sharply in mid-July, while soybean correlations rise more gradually beginning in mid-August. This pattern reflects the critical reproductive growth stages of each crop. The findings provide market participants with guidance on when crop condition ratings begin to contain reliable forecasting information. This thesis extends this literature to peanuts, where the timing and strength of condition–yield relationships may differ due to crop biology.

CHAPTER II: DATA

The data were retrieved from the USDA National Agricultural Statistics Service via Quick Stats (USDA NASS 2026) using an application programming interface (Dinterman and Eyer 2019). Quick Stats is an online database populated with sample surveys covering all United States Agriculture and with the Census of Agriculture that is performed every five years. The data used in this analysis consist of surveys conducted on field crops and peanuts, at the state level from 2013 to 2025.

Table 2.1 presents summary statistics of state-level peanut yields obtained from the USDA NASS QuickStats database. The table reports yield data for crop years 2012–2025 for the six largest peanut-producing states. Yield is measured in tons per hectare (t/ha). As reported in Table 2.1, Georgia exhibits the highest peanut yields among the states included in the analysis.

Table 2.1: Summary Statistics, Yield (tons per hectare)

State	minimum	Q1	mean	Q3	maximum	standard deviation
Alabama	3.09	3.64	3.75	3.95	4.09	0.27
Florida	3.86	3.92	4.18	4.43	4.48	0.25
Georgia	4.32	4.57	4.69	4.85	4.99	0.22
North Carolina	3.90	4.37	4.55	4.88	4.93	0.36
South Carolina	3.59	3.92	4.17	4.48	4.71	0.37
Texas	2.75	3.12	3.46	3.72	4.06	0.44

2.1 Yield Trends

All available yield data were initially evaluated for suitability in this analysis. A significant shift in the varietal landscape occurred in 2010 when the Georgia-06G cultivar (Branch 2007) achieved widespread commercial adoption. This transition coincided with a period of high productivity, notably in 2012, which saw record-breaking acreage and yields (Smith and Dowdy 2013). Developed through the hybridization of Georgia Green and C-

99R, Georgia-06G was bred to integrate the superior traits of its progenitors into an environmentally adaptive phenotype. For over 13 years, this variety has outperformed competitors due to its high resistance to Tomato Spotted Wilt Virus (TSWV), exceptional yield stability, and robust yield potential (Brown 2023). Consistent with a priori expectations regarding the variety's genetic resilience, a structural adjustment in peanut yields and a subsequent reduction in yield variability became visually discernible after 2012 (Figure 2.1). In this context, a structural break signifies a statistically significant shift in the underlying data-generating process; identifying such breaks is critical, as maintaining parameter stability in the yield–condition relationship is a prerequisite for long-term forecast accuracy.

Figure 2.1: Peanut Yield Trends, 1990 to 2025

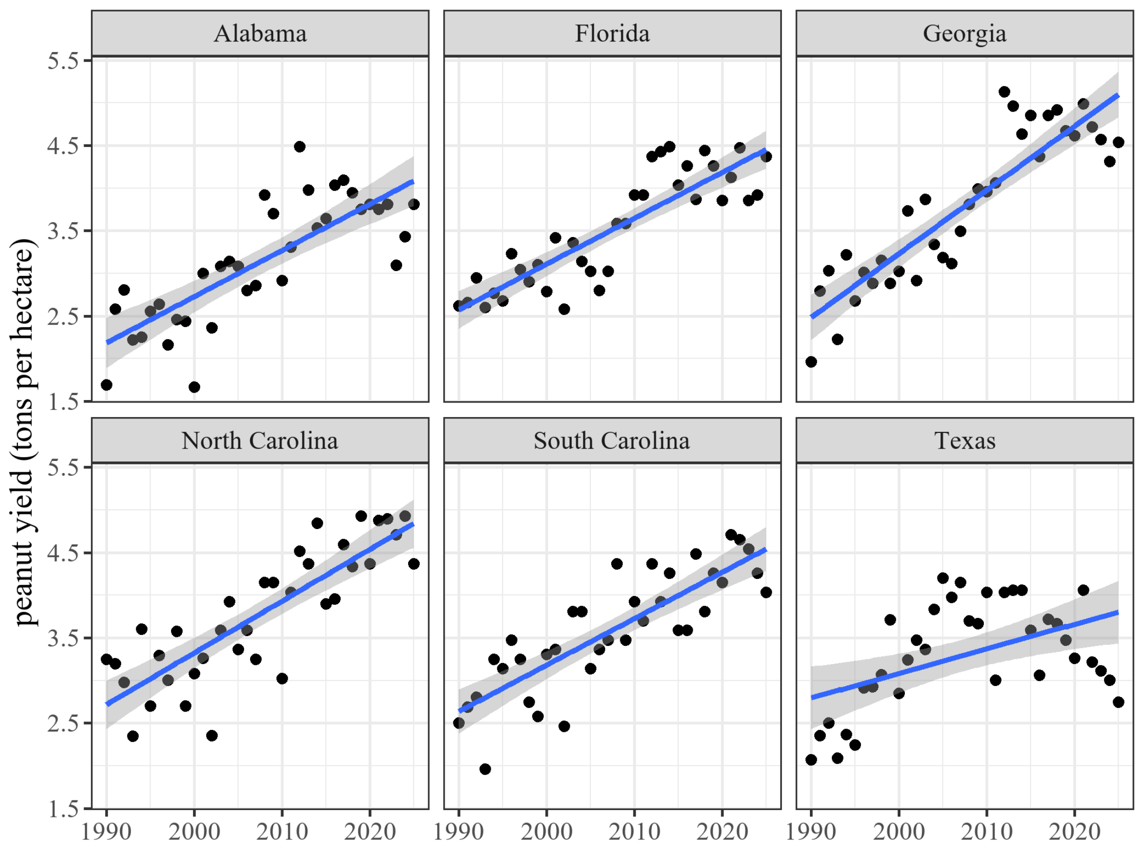

Therefore, quantitative tests were conducted to determine whether yield trends exhibited one or more structural changes between 1990 and 2025 to evaluate the need for yield detrending. A Chow test is a test of structural change. A flat line indicates the absence of a structural change, implying that the slope and intercept of the relationship remain stable over time (Zeileis et al. 2003). The Chow test was performed using the entire time series from 1990 forward (Zeileis et al. 2002). A Chow test revealed there was a structural change in the year 2012 (Table 2.2). The structural change is revealed by having a high F-statistic and low p-value as depicted in Table 2. The F-statistic and p-value establishes the significance of the change is in 2012. A p-value of <0.001 is indicative of a very significant change. The significant change signified a need to look at data from 2013 to 2025. Limiting the analysis to data after 2012 ensures that model parameters are estimated under stable yield conditions, maximizing predictive performance.

Table 2.2: Chow Test for Structural Breaks

breakpoint (year)	F-statistic	p-value
2012	17.36	<0.001

Filtering the yield data to omit observations prior to 2012, positive slope trendlines were replaced by near-zero slopes and slightly negative slopes for some states (Figure 2.2). Yield trends by state for the peanut producing areas of the southeast United States from 2013 to 2025 are represented in Figure 2.2. Texas exhibits the largest decrease in yield of the six states with a drop from 4 tons per hectare to 3 tons per hectare. Alabama's trendline indicates a slight decrease, while Florida and Georgia remain stable over the 14-year period. Only North Carolina and South Carolina have an increase in yield over the 14-year period. North Carolina has a more pronounced increase in yield. During the period, the yield increases from 4 tons per hectare to 5 tons per hectare. Georgia's flat trendline presents no need to detrend the data.

Figure 2.2: Peanut Yield Trends, 2012 to 2025

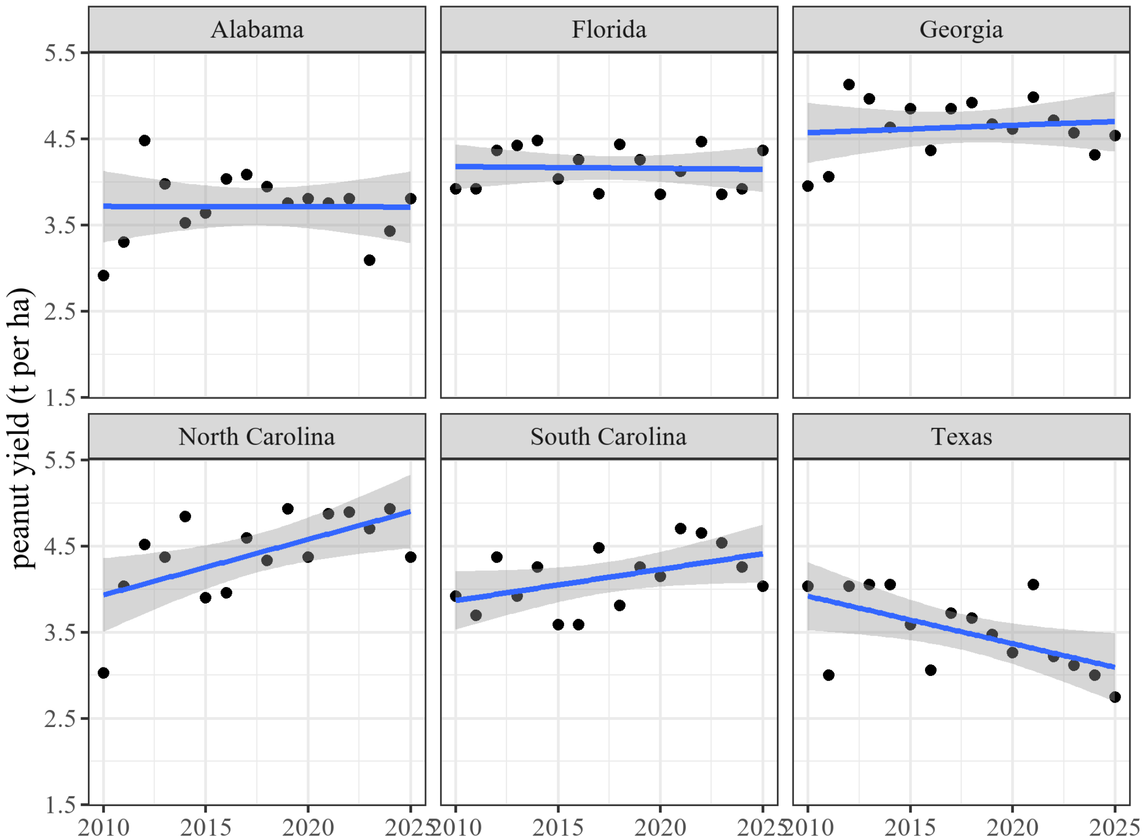

CHAPTER III: METHODS

Text of chapter three begins here. It is important to be consistent in formatting. Spacing should remain at 2.0. You should have a pretty good idea of how to style your subheadings, tables and figures by now. Let us know if you have any questions.

3.1 General Model

Crop yield is modeled as a linear function of a crop condition index (CCI) constructed from USDA weekly crop condition ratings. The general regression model is:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$$

For a single CCI Scheme, the scalar form estimated for state s and crop year t is:

$$\mathrm{Yield}_{s,t} = \beta_{0,w} + \beta_{1,w}\,\mathrm{CCI}_{s,t,w} + \varepsilon_{s,t} \qquad (1)$$

where $\mathrm{CCI}_{s,t,w}$ is one of five index Schemes defined below with parameters estimated independently for each week w as ordinary least squares (OLS).

3.2 Crop Condition Index Schemes

Five CCI Schemes were evaluated. Several modeling Schemes were compared to ensure robustness, account for data uncertainty, and identify the optimal forecasting specification. The benchmark (Good + Excellent) remains the industry standard. Simpler indexes such as NCI and the linear weighted score are more common in older literature, while the Stress index focuses on downside risk and yield floors.

Good + Excellent $(G + E)$ is presented in Equation 2 as

$$\mathrm{CCI1}_{s,y,w} = \%G + \%E \qquad (2)$$

The standard approach is the sum of percentage acreage rated Good or Excellent condition.

Net Condition Index (NCI) (Equation 3):

$$\mathrm{CCI2}_{s,y,w} = (\%G + \%E) - (\%VP + \%P) \qquad (3)$$

NCI subtracts the percentage rated Very Poor and Poor from the percentage rated Good and Excellent, netting favorable against unfavorable conditions.

NCI 2.0 (NCI2) (Equation 4):

The impact of Very Poor (Excellent) conditions is reasonably assumed to be greater than Poor (Good). Rather than an equally weighted index, placing additional weight on the extremes allows the tails of the distribution to exert greater influence than middle ratings:

$$\mathrm{CCI3}_{s,y,w} = (2 \times \%E + 1 \times \%G) - (1 \times \%P + 2 \times \%VP) \qquad (4)$$

Weighted Condition Score (Linear) (Equation 5):

The linear WCS assigns an ordinal weight to each crop condition category, ranging from 1 (Very Poor) to 5 (Excellent). Weekly crop condition percentages are multiplied by their respective weights and summed. Higher WCS values indicate a greater proportion of acres in favorable condition. The linear weighted condition score proposed by Bain and Fortenbery (2017) and evaluated by Li et al. (2026) as the Bain–Fortenbery CCIndex is:

$$\mathrm{CCI4}_{s,y,w} = 1 \times \%VP + 2 \times \%P + 3 \times \%F + 4 \times \%G + 5 \times \%E \qquad (5)$$

where $\%F$ is the percent rated Fair and others as previously defined.

Stress Index (Stress) (Equation 6):

The sum of percentage acreage rated Very Poor and Poor provides an indicator of unfavorable growing conditions:

$$\mathrm{CCI5}_{s,y,w} = \%VP + \%P \qquad (6)$$

3.3 Out-of-Sample Forecasting

Point forecasts from the estimated model take the form:

$$\hat{y}_{s,t} = \hat{\beta}_{0,w} + \hat{\beta}_{1,w}\,\mathrm{CCI}_{s,t,w} \qquad (7)$$

where $\hat{\beta}_w$ is estimated on the training set that excludes year t.

Out-of-sample (OOS) testing evaluates whether the forecasting model generalizes beyond the data used during estimation. Two complementary approaches are employed. Leave-one-year-out (LOYO) cross-validation withholds one crop year at a time (Tashman 2000), using all remaining years for estimation, which provides an unbiased assessment of average forecast accuracy across the full sample. Out-of-time (OOT) validation instead trains the model only on years prior to the forecast year, more closely mimicking the operational environment where future data are unavailable. Together these approaches assess both average model accuracy and real-world forecast reliability.

The accuracy of the OOS is evaluated by Root Mean Square Error (RMSE) and MAE (Mean Absolute Error). The RMSE (Equation 8) measures the extent of the error in each Scheme and is the primary metric for determining forecast accuracy, and is presented as

$$\text{RMSE}_w = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(Y_i - \hat{Y}_{i,w}\right)^2} \qquad (8)$$

where Y is actual end-of-season yield for observation i, Y(i,w) is forecast of yield generated using week w conditions, N is number of test observations (years). Lower RMSE indicates better forecast accuracy. The secondary metric, MAE, is the average absolute forecast error, treating under- and over-predictions symmetrically (Equation 9).

$$\text{MAE}_w = \frac{1}{N}\sum_{i=1}^{N}\left|Y_i - \hat{Y}_{i,w}\right| \qquad (9)$$

CHAPTER IV: RESULTS

Correlations of yield to crop condition ratings are presented in Figure 4.1. The chart illustrates a correlation with each crop condition as the growing season progresses. Three of the conditions, Good, Good + Excellent, and Excellent, are positively correlated. Higher correlations with the three condition variables are observed later in the growing season. Better growing conditions contribute to higher yields. The conditions, Fair, Poor, and Very Poor are negatively correlated. As the growing season progresses, the correlation becomes greater in the negative direction. The chart indicates that there is a relationship between negative growing conditions and reduced yield.

Figure 4.1: Peanut Yield to Crop Condition Rating Correlation

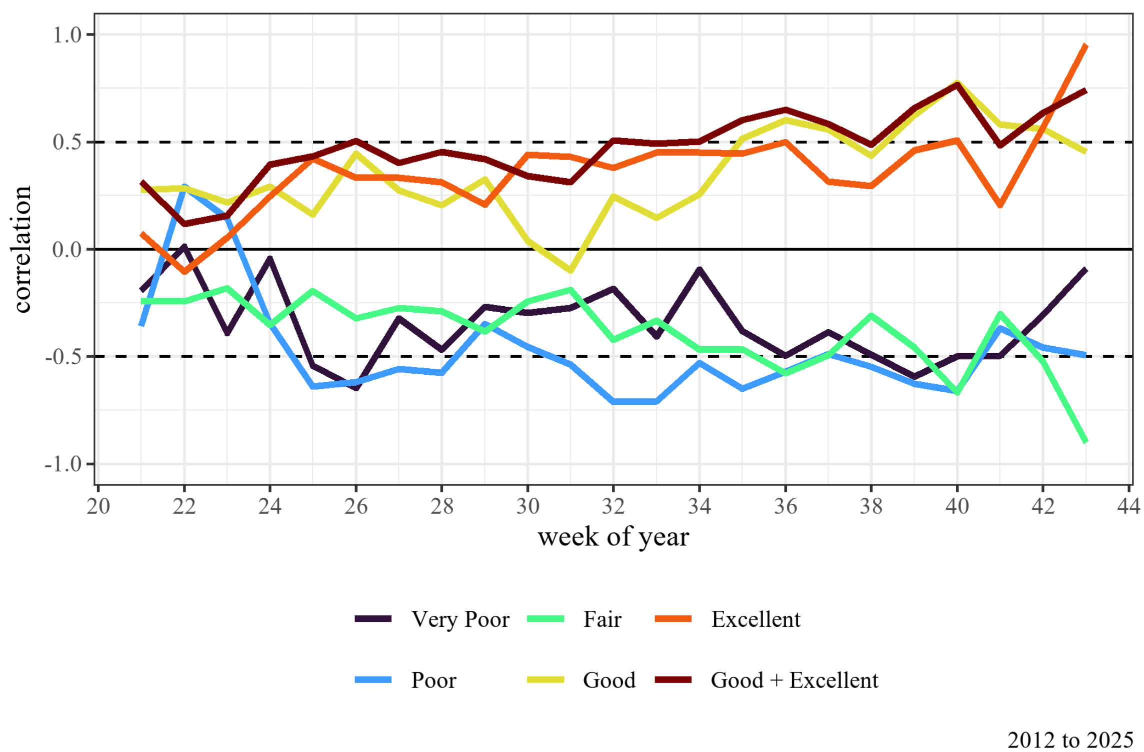

RMSE and MAE are computed for each aggregation Scheme across all forecast weeks (Equation 8) to address two questions: (1) which Scheme provides the most accurate forecasts, and (2) when during the growing season forecasts become sufficiently accurate for operational use. Figure 4.2 compares RMSE trajectories across all five Schemes from week 30 through week 42.

RMSE assigns larger errors more heavily than small errors due to the squared term, making it optimally suitable for assessing forecast reliability. RMSE is expressed in the same units as yield (metric tons/ha). Given the interpretation that a RMSE error of 0.2 tons per hectare means the typical forecast error is approximately 0.2 tons per hectare. The lower the RMSE the better that the forecasting Scheme performs. Week 40 is at the late September early October mark of the year. The crop has completed growing at this point and the final yield has already been established. The best performing Scheme each week is presented in Table 4. By week 40, the Scheme with the best performing forecast was Good + Excellent with an RMSE of 0.16 tons per hectare. Given that Good + Excellent dominated the other candidate Schemes, additional complexity did not benefit modeling for peanuts.

Figure 4.2: Weighting Schemes across Forecast Weeks

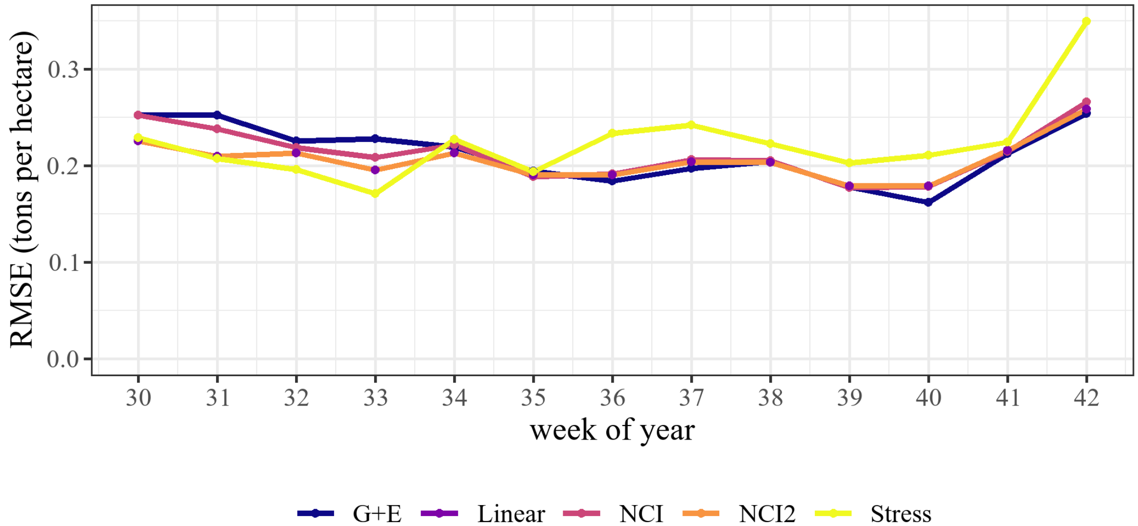

Evaluation metrics for all five Schemes are reported in Table 4.1 and Table 4.2. Final week forecast performance confirms Good + Excellent superiority (Table 4.1). At week 40 (early October), Good + Excellent achieves RMSE of 0.16 tons per hectare, outperforming NCI 2.0 (RMSE = 0.18), Linear (RMSE = 0.18), NCI (RMSE = 0.18), and Stress (RMSE = 0.21). The 12 to 31% higher error rates for complex Schemes underscore the value of simplicity for operational forecasting. The Good + Excellent Scheme also yielded the lowest Mean Absolute Error (MAE = 0.13 tons per hectare; Equation 9), further confirming its superior predictive accuracy.

Table 4.1: Forecast Performance by Scheme (Week 40)

Scheme	RMSE	MAE
G+E	**0.16**	**0.13**
NCI	0.18	0.15
Linear	0.18	0.15
NCI2	0.18	0.15
Stress	0.21	0.19

The Schemes described in Table 4.2 model performance using RMSE and MAE across weeks 30 through 42 and were assessed for forecasting reliability. The top-performing Scheme in each week exhibited marginally lower RMSE and slightly lower MAE, indicating improved model fit for that period. During the later half of the season, the NCI model achieved the highest performance ranking during weeks 35 and 39; however, it ranked as the second-best model in in weeks 40 and 41. The NCI2 Scheme achieved the best performance in three weeks and ranked second-best in four weeks. Overall, the standard Good + Excellent Scheme consistently outperformed the alternative model specifications with ranking as the best model for five of latest seven weeks analyzed. Consequently, subsequent analyses focused exclusively on the Good + Excellent Scheme.

Table 4.2: Best and Second-Best Scheme, Weeks 30 to 42

	Best			Second-Best		
Week	Scheme	RMSE	MAE	Scheme	RMSE	MAE
30	NCI2	0.23	0.16	Linear	0.23	0.16
31	Stress	0.21	0.17	NCI2	0.21	0.18
32	Stress	0.20	0.13	NCI2	0.21	0.16
33	Stress	0.17	0.13	Linear	0.20	0.15
34	NCI2	0.21	0.15	Linear	0.21	0.15
35	NCI	0.19	0.14	Linear	0.19	0.15
36	G+E	0.18	0.14	NCI2	0.19	0.15
37	G+E	0.20	0.15	NCI2	0.20	0.16
38	NCI2	0.2.	0.16	Linear	0.20	0.16
39	NCI	0.18	0.15	G+E	0.18	0.14
40	G+E	0.16	0.13	NCI	0.18	0.15
41	G+E	0.21	0.17	NCI	0.21	0.17
42	G+E	0.25	0.21	Linear	0.26	0.23

4.1 Leave-One-Year-Out (LOYO) Validation

When creating the forecasting model a leave-one-year-out (LOYO) method was used for testing. Each week the model was run with one year left out of the equation while the rest of the years were left in for model estimation. This validation design mirrors the out-of-sample testing approach in Li et al. (2026).

$$\begin{aligned} \text{Training Set}_t &= \{(Y_s, \text{Index}_s) : s \in \mathcal{S},\ s \neq t\} \\ \text{Test Set}_t &= \{(Y_t, \text{Index}_t)\} \end{aligned} \tag{10}$$

where $\mathcal{S}$ is set of all years in sample {2012,2013,...,2025}, t is test year. Training set contains all years except t while the test set contains only year t. Leave-one-year-out (LOYO) cross-validation (Equation 10) provides unbiased assessment of out-of-sample forecast performance. For each year t in the dataset:

- Estimate Equation 10 using all years except t to train the model
- Generate forecast for year t using estimated parameters to test the model
- Calculate forecast error for year t to evaluate the model
- Iterate across all years t∈{2012,...,2025}

The resulting forecasts are truly out-of-sample because year t was never used to estimate the parameters applied to forecast year t. This mimics the operational forecasting scenario where future year outcomes are unknown when models are estimated.

Figure 4.3: Leave-One-Year-Out: Yield Forecasts, Good + Excellent

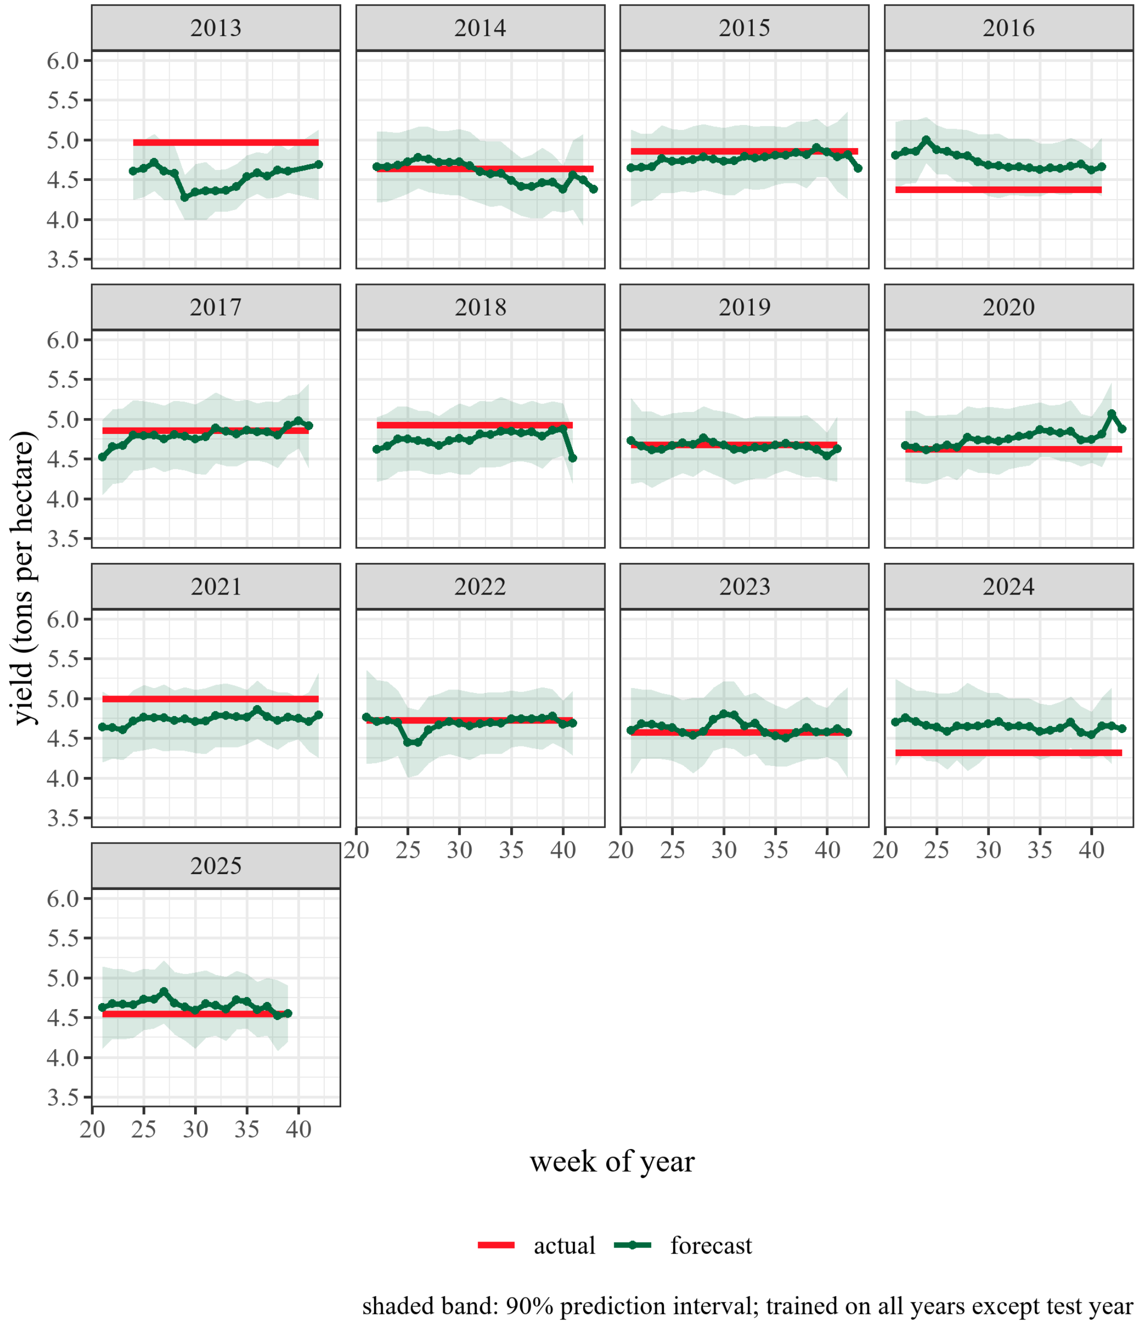

A series of yield forecasts demonstrating the leave-one-year-out (LOYO) approach is presented in Figure 4.3. The model is trained on years 2013 to 2025, with one year withheld at a time to forecast the yield. The green line represents the model's weekly forecast using data available up to that point, while the red line represents the actual observed produced yield of that year. The forecast for 2013 underestimated the actual yield by approximately 0.5 tons per hectare. In contrast, the 2016 forecast over predicted the realized yield. The figure indicates that forecast accuracy improves as the growing season progresses, with predicted yields converging toward observed end-of-season yields. This pattern is expected as additional information becomes available, providing a more comprehensive perspective on final yield outcomes.

Forecast performance varies substantially across years. The model performs exceptionally well in 2022, exhibiting minimal bias throughout the growing season and producing a final forecast within 0.1 tons per hectare of the observed yield. The years 2017, 2019, and 2023 also demonstrate strong performance, with forecast bias converging toward zero by approximately week 40. In contrast, model performance is weaker for 2013, 2016, and 2024. For 2013, the model underestimates yield by approximately 0.5 tons per hectare during the early season, although forecast bias diminishes by harvest. The years 2016 and 2024 exhibit persistent overestimation ranging from 0.3 to 0.5 tons per hectare throughout the season. The 2024 growing season was atypical for peanut production in the Southeastern United States. Excessive rainfall during May disrupted planting operations and necessitated replanting in several areas. This was followed by abnormally hot and dry conditions in June and July, which imposed stress during early vegetative growth. Growing conditions improved in August, with crop condition ratings indicating a more favorable yield outlook. However, rainfall during the harvest period delayed digging and extended harvest into late November and early December, ultimately contributing to reduced crop quality.

This interannual variability suggests that while the model effectively captures typical yield–condition relationships, its predictive performance declines in years characterized by atypical production environments, such as extreme weather events, disease pressure, or yield outcomes that deviate substantially from historical patterns. This behavior is consistent with findings reported by Irwin and Good (2017a) for corn and soybean yield forecasting.

Unlike RMSE, which aggregates all errors regardless of direction, bias distinguishes between under- and over-prediction, forecast bias (Equation 11) measures systematic forecast errors. A well-calibrated model exhibits near-zero bias, indicating that positive and negative errors cancel out over time.

$$\text{bias}_w = \frac{1}{N}\sum_{i=1}^{N}\left(Y_i - \hat{Y}_{i,w}\right) \qquad (11)$$

where positive bias indicates systematic underestimation (forecasts too low), negative bias indicates systematic overestimation (forecasts too high), and zero bias indicates unbiased forecasts on average.

Bias is particularly important for processors and marketers who need forecasts that are correct “on average” rather than conservative or optimistic. We report bias by year to identify whether specific years exhibit systematic under- or over-prediction, which may indicate unusual growing conditions not well-captured by the model.

Bias near zero indicates accurate forecasts. Strong positive bias is observed early in the 2013 season (>10%), indicating the model underestimated actual yields (i.e., observed yield exceeded the forecast). In contrast, persistent negative bias in 2016 and 2024 indicates the model overestimated yields (i.e., forecasts exceeded realized outcomes). This asymmetry suggests the model may struggle under atypical production conditions, particularly in years with unusually high yields (2013) or unusually low yields (2016, 2024), warranting further investigation.

Percent bias from the forecast with a year dropped from the training set is presented in Figure 4.4. The years that are the closest to zero represent no bias, meaning, the model performed accurately for those years. The 2013 season extends beyond the 10% threshold around week 30, reflecting a period in which the model underestimated yield and produced a substantial positive bias. In contrast, both 2016 and 2024 exhibit a consistently strong negative bias throughout their respective forecast seasons. This suggests there was an overestimation of yield for those years.

While LOYO validation assesses average out-of-sample performance, out-of-time (OOT) validation more closely mimics the operational forecasting environment where only historical data are available at the time of the forecast.

Figure 4.4: Leave-One-Year-Out, Good + Excellent

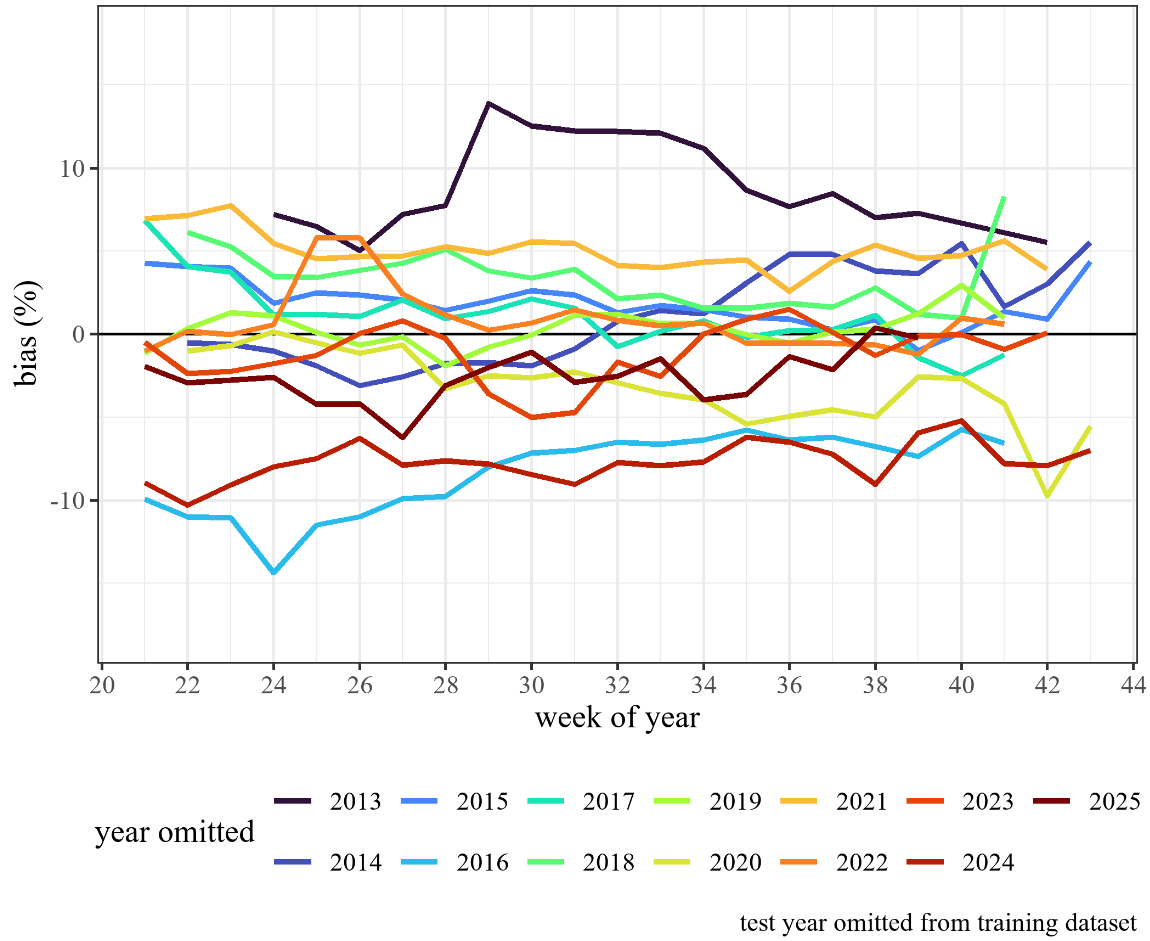

4.2 Out-of-Time (OOT) Validation

Out-of-time (OOT) validation splits the training dataset prior to the testing year in question to avoid data leakage. The OOT validation was employed to mimic real-world forecasting by only using data available at the time of the forecast. The OOT validation procedure is summarized in Figure 4.5, where observations were separated into pre-split training data and post-split test data to evaluate out-of-sample forecasting performance. The line chart represents crop years, beginning with 2015. The initial training set consists of 2012, 2013, and 2014, the first three years of the post-break sample, providing sufficient observations to estimate model parameters for the first OOT forecast year of 2015. Because data are added year-wise to the model, accuracy improves through 2025. OOT validation mimics real world forecasting and offers more stability over time.

Figure 4.5: Out-of-Time Validation, Good + Excellent

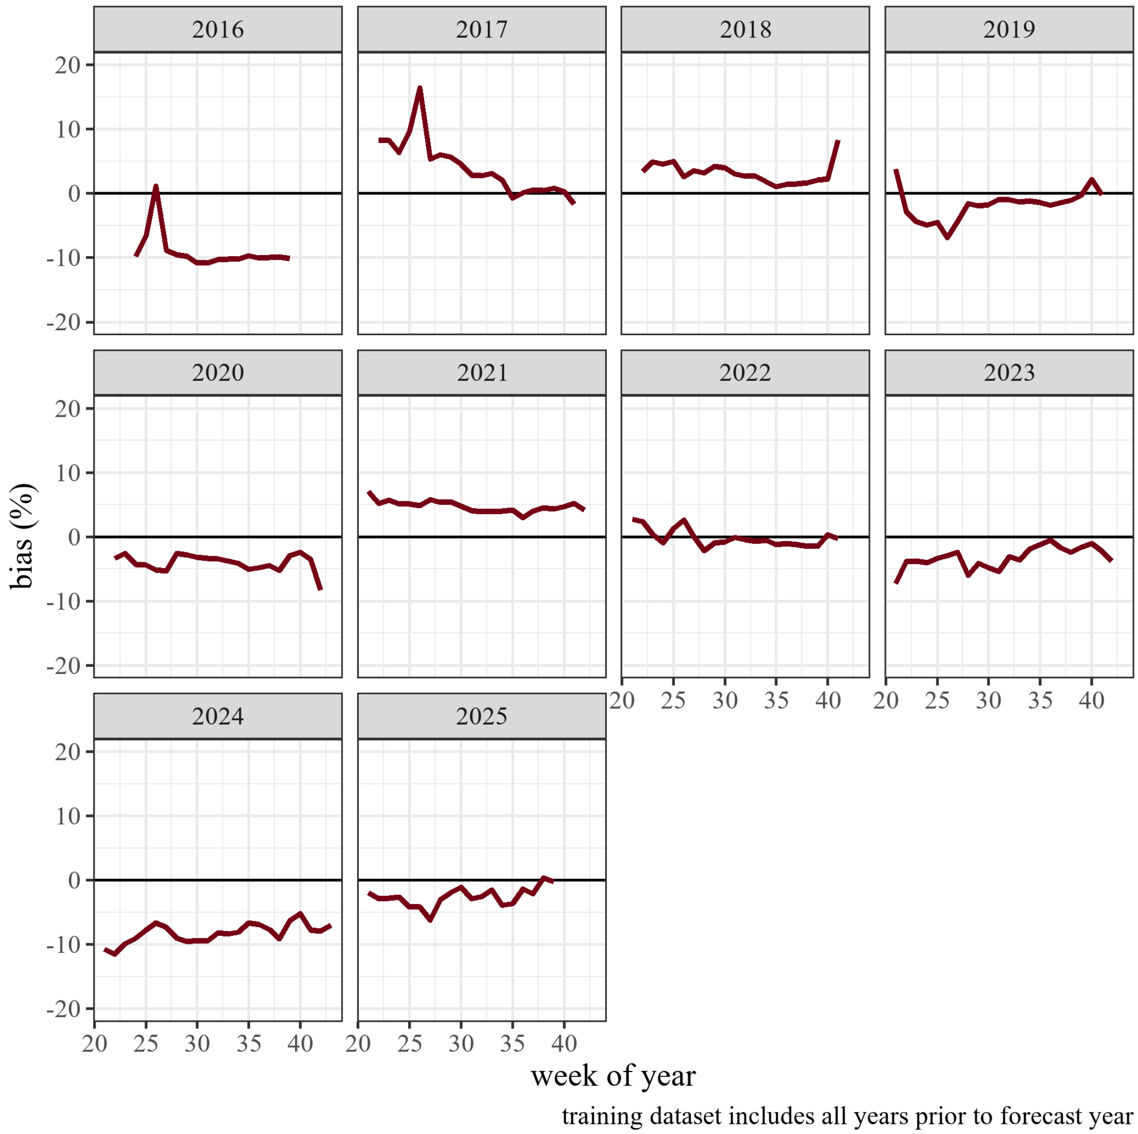

Figure 4.5 presents OOT percent bias by year, where values near zero indicate a well-calibrated model. Bias generally diminishes as the training dataset expands, consistent with the expectation that additional years improve model stability. Exceptions are observed in 2016, 2020, and 2024, where persistent negative bias indicates overestimation of yield.

In contrast, 2022 exhibits negligible bias throughout the forecast period, representing the strongest model performance in the sample. Table 4.3 confirms this pattern: OOT RMSE and MAE are generally lower in later years, though year-specific deviations, most notably 2024, reflect atypical production conditions rather than model degradation. Given the consistent superiority of the Good + Excellent Scheme across both validation frameworks, subsequent analyses are restricted to that Scheme.

Figure 4.6: Out-of-Time: Actual Yield Versus Forecasts, Good + Excellent

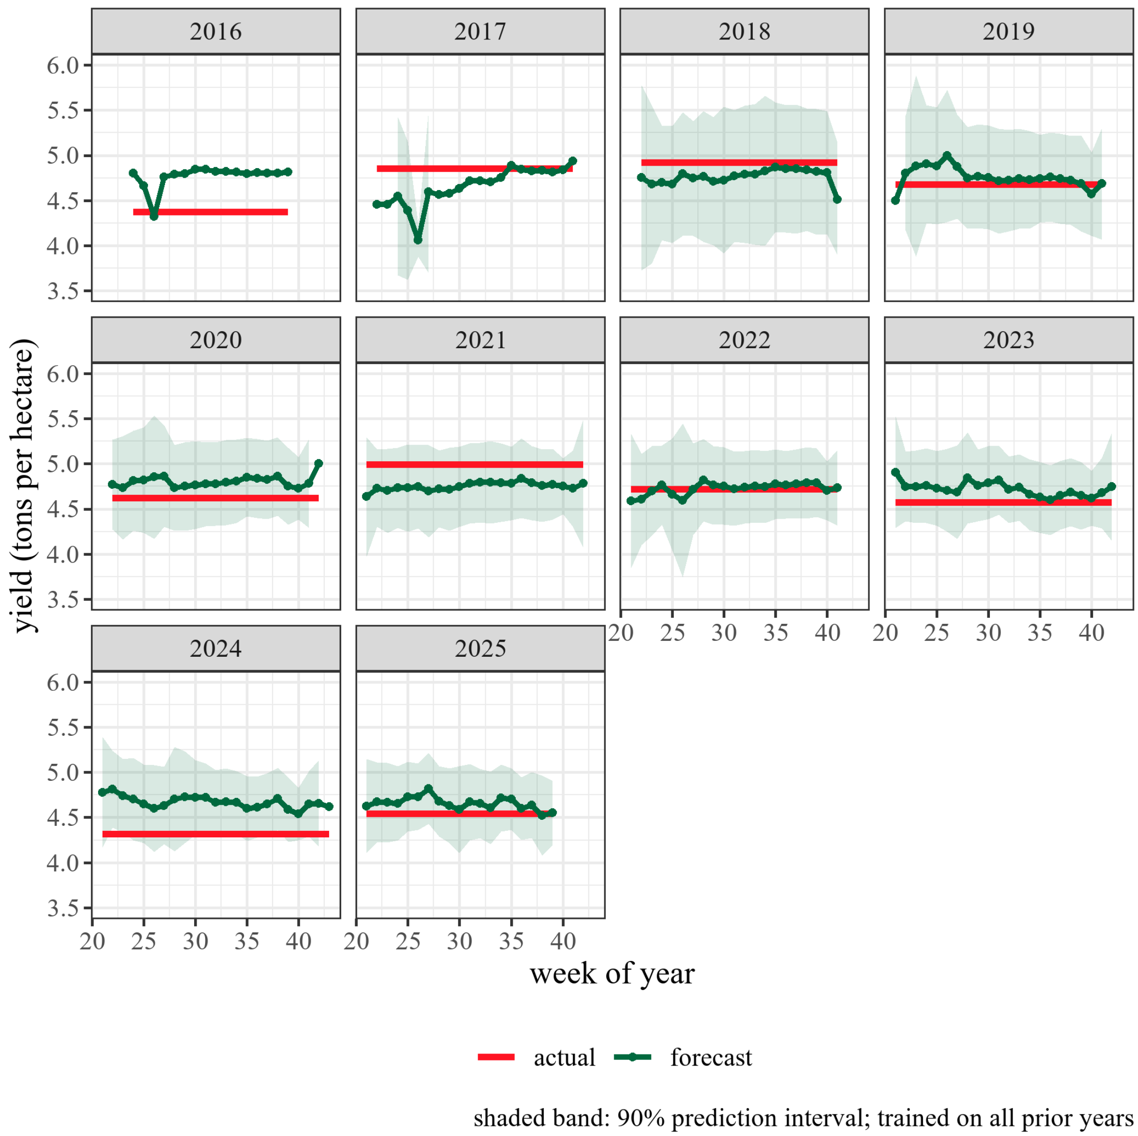

Table 4.3: Forecast Accuracy Metrics, Good + Excellent

	LOYO		OOT (Expanding)	
Year	RMSE	MAE	RMSE	MAE
2013	0.46	0.45		
2014	0.14	0.12		
2015	0.11	0.09		
2016	0.38	0.36	0.42	0.40
2017	0.11	0.08	0.28	0.21
2018	0.19	0.16	0.18	0.16
2019	0.05	0.04	0.14	0.11
2020	0.18	0.15	0.20	0.19
2021	0.26	0.25	0.24	0.24
2022	0.09	0.06	0.06	0.05
2023	0.09	0.07	0.17	0.15
2024	0.34	0.33	0.36	0.36
2025	0.13	0.12	0.13	0.12

4.3 When to Start Paying Attention to Crop Condition Ratings?

Irwin and Good (2017a) reported on when crop condition rates became most useful for predicting yield of corn and soybeans. The analysis demonstrates how model performance evolves over the course of the growing season. The model is evaluated across calendar weeks to determine when it becomes useful within the growing season. The week-by-week framework avoids bias introduced by arbitrarily selecting a single evaluation period and instead identifies the point at which condition data begins to provide meaningful predictive value. The relationship between Good + Excellent crop condition ratings and final yield is presented in Figure 4.7. While Figure 4.8 correlates with the final yield to the Good + Excellent throughout the growing season, Figure 4.7 depicts how correlated Good + Excellent are to the end of year rating. Figure 4.8 provides a weekly assessment of when crop conditions become correlated with final yield for each state. The information becomes useful during different weeks across states. Because Florida's peanut planting season occurs earlier than that of the rest of the southeastern growing region, it was hypothesized that the Good + Excellent yield relationship would emerge earlier in Florida relative to the other states. However, results indicate the opposite pattern: Florida's correlation did not reach the 0.8 threshold until week 38, later than all other states included in the analysis.

Predictive value is limited during the early weeks of the growing season, as correlation coefficients remain below the 0.8 threshold. A correlation of 0.8 indicates a strong model fit. Once the correlation reaches 0.8 or higher, the condition data become operationally meaningful. Beginning in week 33, crop condition ratings become informative for Alabama and South Carolina. As the season progresses, additional states exhibit similarly strong correlations as the yield–condition relationship strengthens.

Figure 4.7: Good + Excellent Crop Condition to End of Year Rating

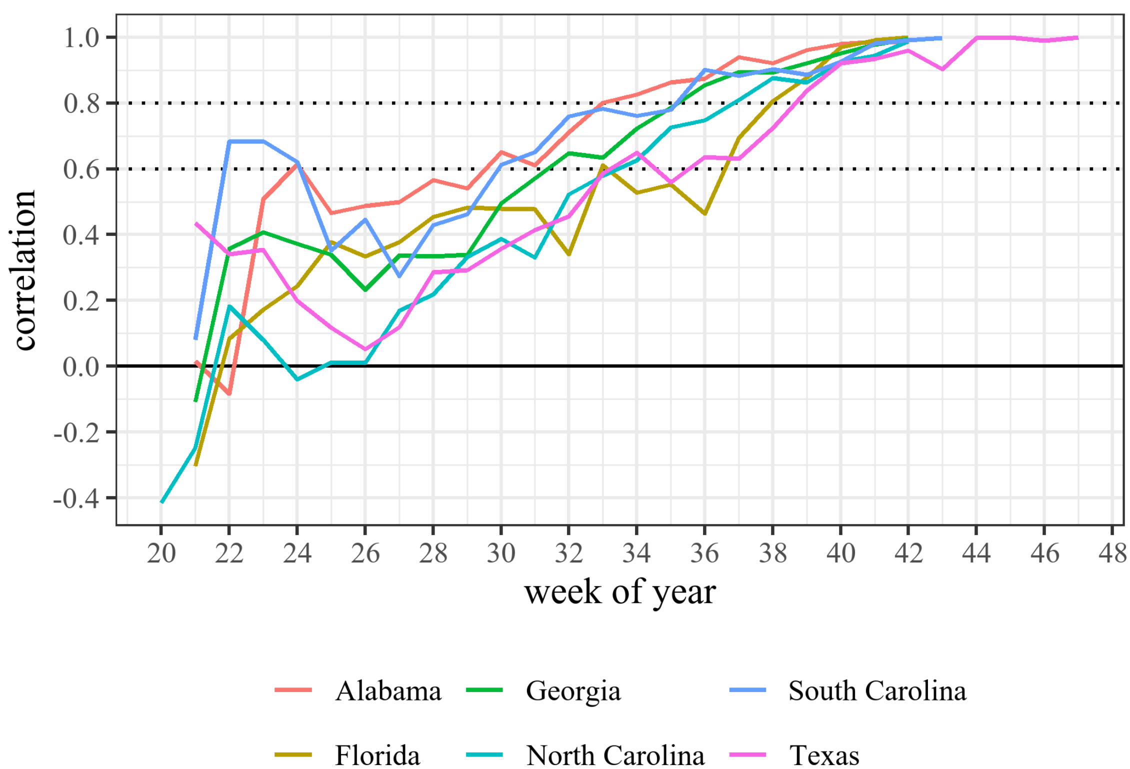

data source: USDA NASS, 2026

Figure 4.8 depicts a correlation between Good + Excellent and final yield. The results suggest that, during the early portion of the year, the relationship between Good + Excellent ratings and yield is not statistically significant. As the growing season progresses, the Good + Excellent rating relationship with the final yield for each state becomes slightly more correlated. Weeks 40 through 42 become the point where the correlation is the highest. Weeks 40 through 42 are late September to early October. During this point a large percentage of the crop is harvested with a known yield.

Figure 4.8: Good + Excellent Crop Condition to Final Yield

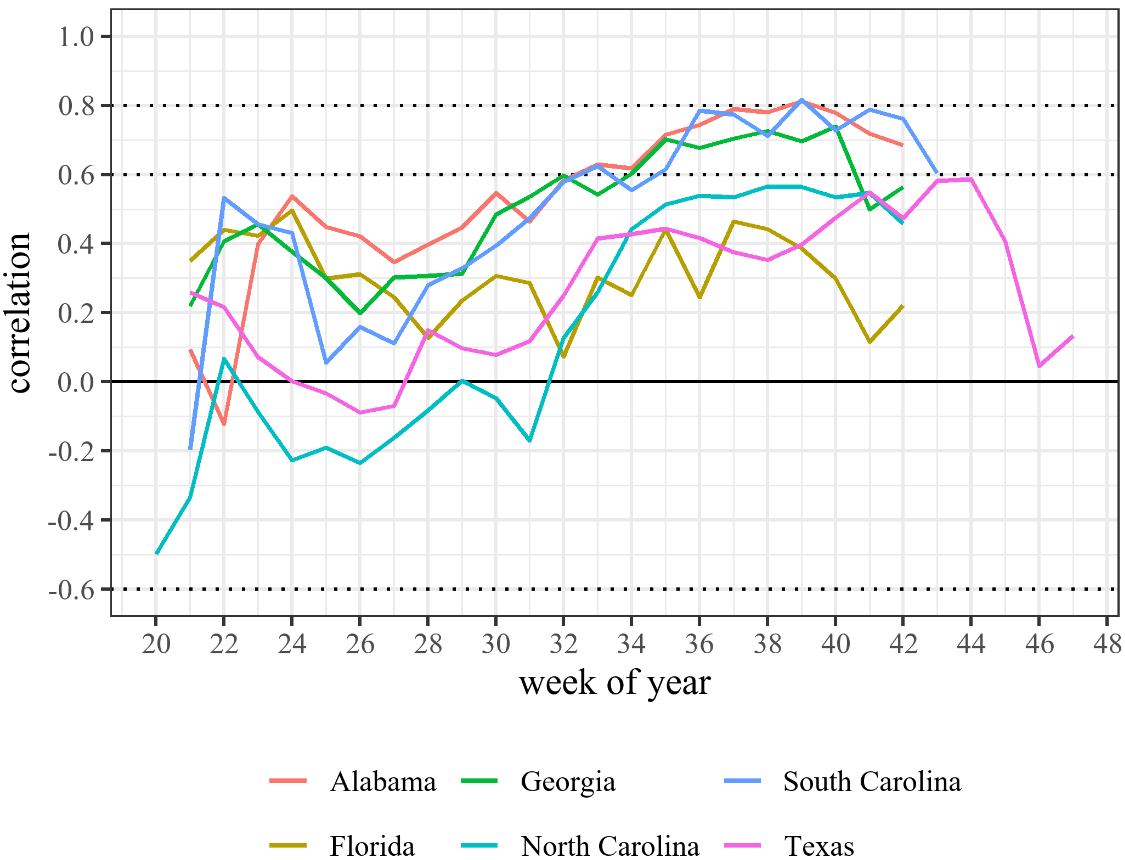

data source: USDA NASS, 2026

Figure 4.9 is a representation of the two previous figures combined. The state of Georgia exhibits exceptional data on when to start paying attention to the forecasting model. The figure demonstrates that both Good + Excellent ratings and yield exhibit limited predictive reliability through the first half of the growing season. As the growing season progresses the reliability increases. Based on the data it is implied around week 32 is when the correlation is gaining strength. At this mark it is time to start paying attention to the Good + Excellent forecast.

Figure 4.9: Good + Excellent Crop Condition to Final Rating and Yield

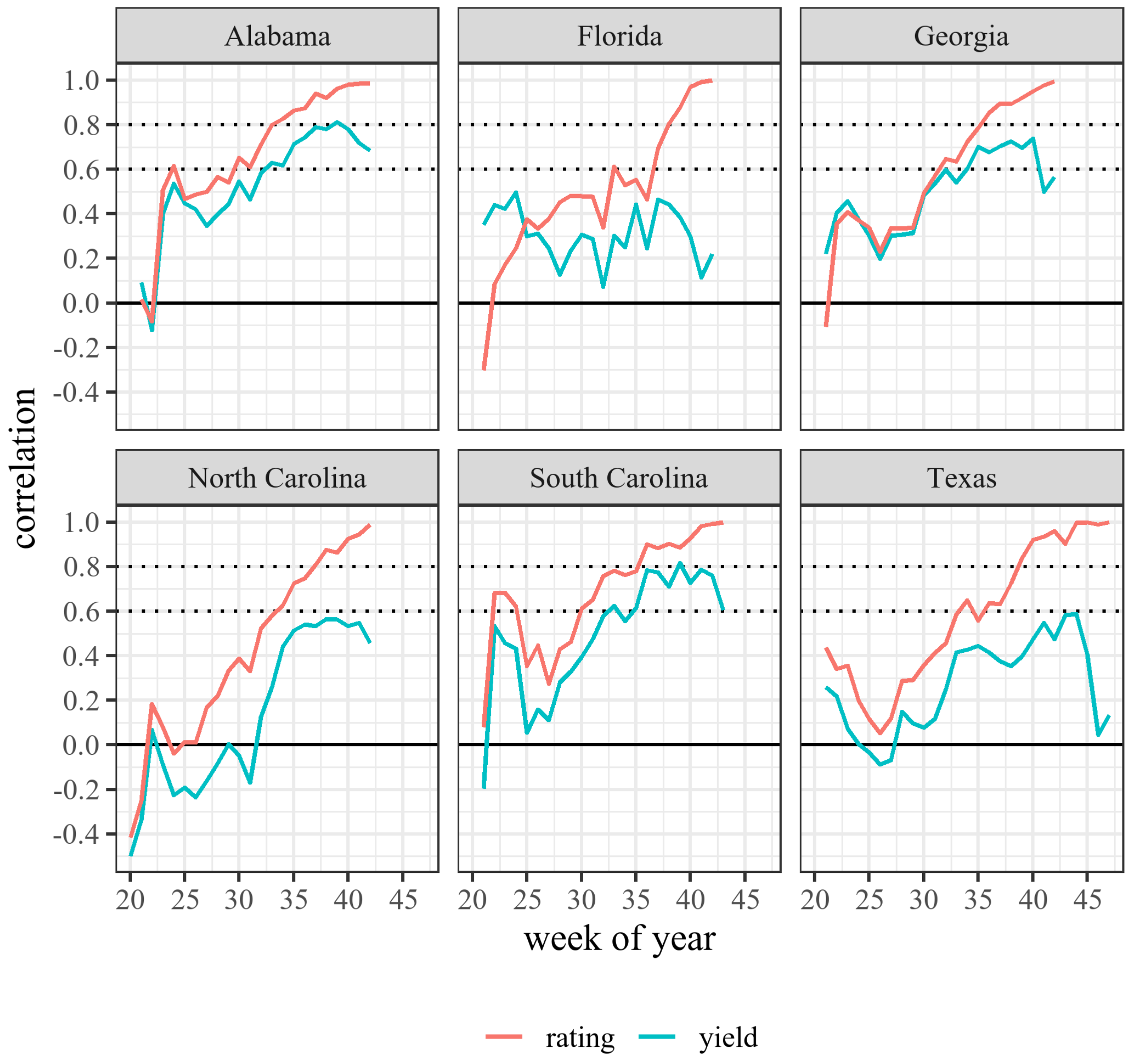

data source: USDA NASS, 2026

The correlation graph follows suit of the life cycle of a peanut. Peanuts are planted in the state of Georgia during week 20. Week 32 would have peanuts averaging around 80 days after planting (DAP). Peanuts start forming seeds inside the pod at the 73 DAP mark (Georgia Peanut Commission 2024). Yield is becoming progressively more important from week 32 forward.

4.4 How Should Crop Condition Ratings Be Used?

Crop condition–based yield forecasts play a critical role in supply chain management by providing early estimates of crop size and quality. These forecasts guide storage and processing decisions, as peanuts may be shelled immediately after harvest or stored in warehouses for up to 12 months prior to shelling depending on market conditions. Forecasts of crop size enable warehousing teams to determine the number of storage facilities required beyond existing capacity, and storage requirements continue through the shelling process, as shelled peanuts may be held under ambient conditions for up to 12 months following shelling (Butts et al. 2017). In addition, crop forecasts inform finished product storage planning.

Beyond operational planning, crop condition–based forecasts support strategic decision-making across the peanut industry. Advance knowledge of expected crop size and quality allows marketing teams to estimate sales volume, assess product condition, and evaluate opportunities for market expansion in subsequent production years. These forecasts also assist farmers in making forward contracting decisions and provide the USDA with information to refine official yield estimates prior to the release of final reports.

CHAPTER V. CONCLUSIONS

This thesis evaluated the use of mid-season USDA crop condition ratings to forecast end-of-season peanut yields. Results indicate that crop condition information becomes increasingly informative as the growing season progresses, with meaningful forecasting value emerging around calendar week 32 and peaking by week 40. From late September to early October, crop condition ratings exhibit their strongest relationship with realized yield, suggesting that late-season conditions provide the most reliable signal for yield expectations.

Forecast accuracy exhibited systematic differences across aggregation Schemes, reflecting variation in bias and stability over the growing season. Five approaches were evaluated, including Good + Excellent, the Net Condition Index (NCI), NCI 2.0, linear, and a stress-based index, Stress. Among these, the Good + Excellent Scheme consistently outperformed more complex alternatives by producing the most accurate and stable yield forecasts. The superior performance of the Good + Excellent approach highlights the value of simplicity, as its reduced structural complexity enhances robustness and limits bias relative to more heavily weighted indices. These results for peanuts are consistent with findings for corn and soybeans (Li et al. 2026), extending that result to a third commodity with distinct biology and market structure. While Li et al. (2026) used 37 years of data and 23 out-of-sample years to reach the “complexity doesn’t pay” conclusion, this peanut analysis reaches the same conclusion with only 14 years post-break.

The informativeness of crop condition ratings depends critically on the timing of observation within the growing season. Crop condition ratings exhibited little to no relationship with final yield across states early in the growing season. As the season progressed, the association between reported conditions and yield strengthened consistently across all six states. This pattern is consistent with findings from the LOYO validation, which indicate that forecast accuracy improves as the growing season progresses and predicted yields converge toward realized end-of-season outcomes. As additional information becomes available later in the season, forecasts become more stable and reflective of final yield potential.

Understanding when crop condition ratings become informative for final yield benefits a wide range of stakeholders. For producers, improved timing of reliable condition signals supports marketing and forward-contracting decisions. Processors and shellers benefit from more accurate expectations of crop size, which inform planning related to storage needs and shelling capacity. The strong performance of the Good + Excellent Scheme further enhances its practical value, as its simplicity facilitates real-time interpretation and clear communication across the supply chain. Overall, these results support the continued use of USDA crop condition ratings as a useful decision-support tool when applied at appropriate stages of the growing season.

Several limitations of this analysis should be acknowledged. The relationship between crop condition ratings and yield is sensitive to weather variability, disease pressure, and harvest timing, and seasons characterized by unforeseen conditions can weaken forecast accuracy. Additionally, this study focused exclusively on yield outcomes and did not account for crop quality, which may also be influenced by late-season conditions. Future research could incorporate additional explanatory variables such as rainfall, temperature, or stress indices, extend the analysis to quality outcomes, and explore regional or sub-state variation to further refine condition-based yield forecasting.

WORKS CITED

Bain, Ryan, and T. Randall Fortenbery. 2017. "Impact of Crop Condition Reports on National and Local Wheat Markets." *Journal of Agricultural and Applied Economics* 49 (1): 97–119. https://doi.org/10.1017/aae.2016.31.

Beguería, Santiago, and Marco P. Maneta. 2020. "Qualitative Crop Condition Survey Reveals Spatiotemporal Production Patterns and Allows Early Yield Prediction." *Proceedings of the National Academy of Sciences* 117 (31): 18317–23. https://doi.org/10.1073/pnas.1917774117.

Branch, W. D. 2007. *Registration of "Georgia-06G" Peanut.* (Madison) 1: 120. https://doi.org/10.3198/jpr2006.12.0812crc.

Brown, Nino. 2023. "Cultivar Spotlight: 'Georgia-06G'." *Institute of Plant Breeding, Genetics & Genomics Newsletter*.

Butts, C. L., M. C. Lamb, R. B. Sorensen, et al. 2017. "Alternative Storage Environments for Shelled Peanuts." *Peanut Science* 44 (2): 111–23. https://doi.org/10.3146/PS17-2.1.

Dinterman, Robert, and Jonathan Eyer. 2019. *usdarnass: USDA NASS Quick Stats API*. https://github.com/rdinter/usdarnass.

Georgia Peanut Commission. 2024. *The Life Cycle of a Peanut: From Planting to Harvest*. https://gapeanuts.com/peanutlifecycle/.

Irwin, Scott, and Darrel Good. 2017a. "How Should We Use Within-Season Crop Condition Ratings for Corn and Soybeans?" *Farmdoc Daily* 7 (101).

Irwin, Scott, and Darrel Good. 2017b. "When Should We Start Paying Attention to Crop Condition Ratings for Corn and Soybeans?" *Farmdoc Daily* 7 (96).

Irwin, Scott, and Todd Hubbs. 2018. "Measuring the Accuracy of Forecasting Corn and Soybean Yield with Good and Excellent Crop Condition Ratings." *Farmdoc Daily* 8 (118).

Li, Jiarui, Scott H. Irwin, and Todd Hubbs. 2026. "Does Complexity Pay? Forecasting Corn and Soybean Yields Using Crop Condition Ratings." *Journal of Agricultural and Resource Economics* 51 (1): 32–52. https://doi.org/10.22004/ag.econ.358998.

Smith, Nathan B., and Sharon Dowdy. 2013. *Georgia Peanut Growers Predicted to Reduce Acreage After 2012 Bumper Crop*. University of Georgia College of Agricultural; Environmental Sciences Field Report; https://fieldreport.caes.uga.edu/news/georgia-peanut-growers-predicted-to-reduce-acreage-after-2012-bumper-crop/.

Stokes, J. R. 2024. "A Markov Chain Model of Crop Conditions and Intrayear Crop Yield Forecasting." *Journal of Forecasting* 43 (3): 583–92. https://doi.org/10.1002/for.3052.

Tang, Jing. 2016. "Utilizing Crop Progress Good and Excellent Condition Ratings for Forecasting Iowa Corn and Soybeans Yield, Part 1." Iowa Farm Bureau.

Tashman, Leonard J. 2000. "Out-of-Sample Tests of Forecasting Accuracy: An Analysis and Review." *International Journal of Forecasting* 16 (4): 437–50. https://doi.org/10.1016/S0169-2070(00)00065-0.

USDA NASS. 1999. *Understanding USDA Crop Forecasts*. Miscellaneous Publication No. 1554. United States Department of Agriculture, National Agricultural Statistics Service and World Agricultural Outlook Board.

USDA NASS. 2018. *National Crop Progress - Terms and Definitions*. May 4.

USDA NASS. 2026. "USDA National Agricultural Statistics Service QuickStats Database." United States Department of Agriculture, National Agricultural Statistics Service. https://quickstats.nass.usda.gov/.

Zeileis, Achim, Christian Kleiber, Walter Krämer, and Kurt Hornik. 2003. "Testing and Dating of Structural Changes in Practice." *Computational Statistics & Data Analysis* 44 (1–2): 109–23. https://doi.org/10.1016/S0167-9473(03)00030-6.

Zeileis, Achim, Friedrich Leisch, Kurt Hornik, and Christian Kleiber. 2002. "strucchange: An R Package for Testing for Structural Change in Linear Regression Models." *Journal of Statistical Software* 7 (2): 1–38. https://doi.org/10.18637/jss.v007.i02.

APPENDIX

Weekly Correlation of Crop Condition Ratings to Peanut Yield

A correlation graph describing the relationship between yield to crop condition over the weeks of the year is presented in Figure A1. The y-axis depicts the correlation or relationship between the variables. No relationship between yield and crop condition if the line for the variable never exceeds the absolute value of 0.5 for the week. The x-axis represents the week of the year. Analysis of the data begins during week 20. Peanuts begin being planted around the first of May. Week 20 depicted on the graph represents anywhere from May 12th to May 18th. The x-axis spans through the entire growing season. The final week is week 40. Week 40 is the first week of October when peanut harvest occurs. Most peanut varieties reach maturity 145 days after planting, with only a few varieties having a longer maturing window. The correlation chart represents combined data from 2013 to 2025 with six different crop conditions listed: Very Poor, Poor, Fair, Good, Good + Excellent, and Excellent. Good + Excellent is the summation of Good and Excellent.

Figure A1: Correlation of Crop Condition Ratings to Peanut Yield

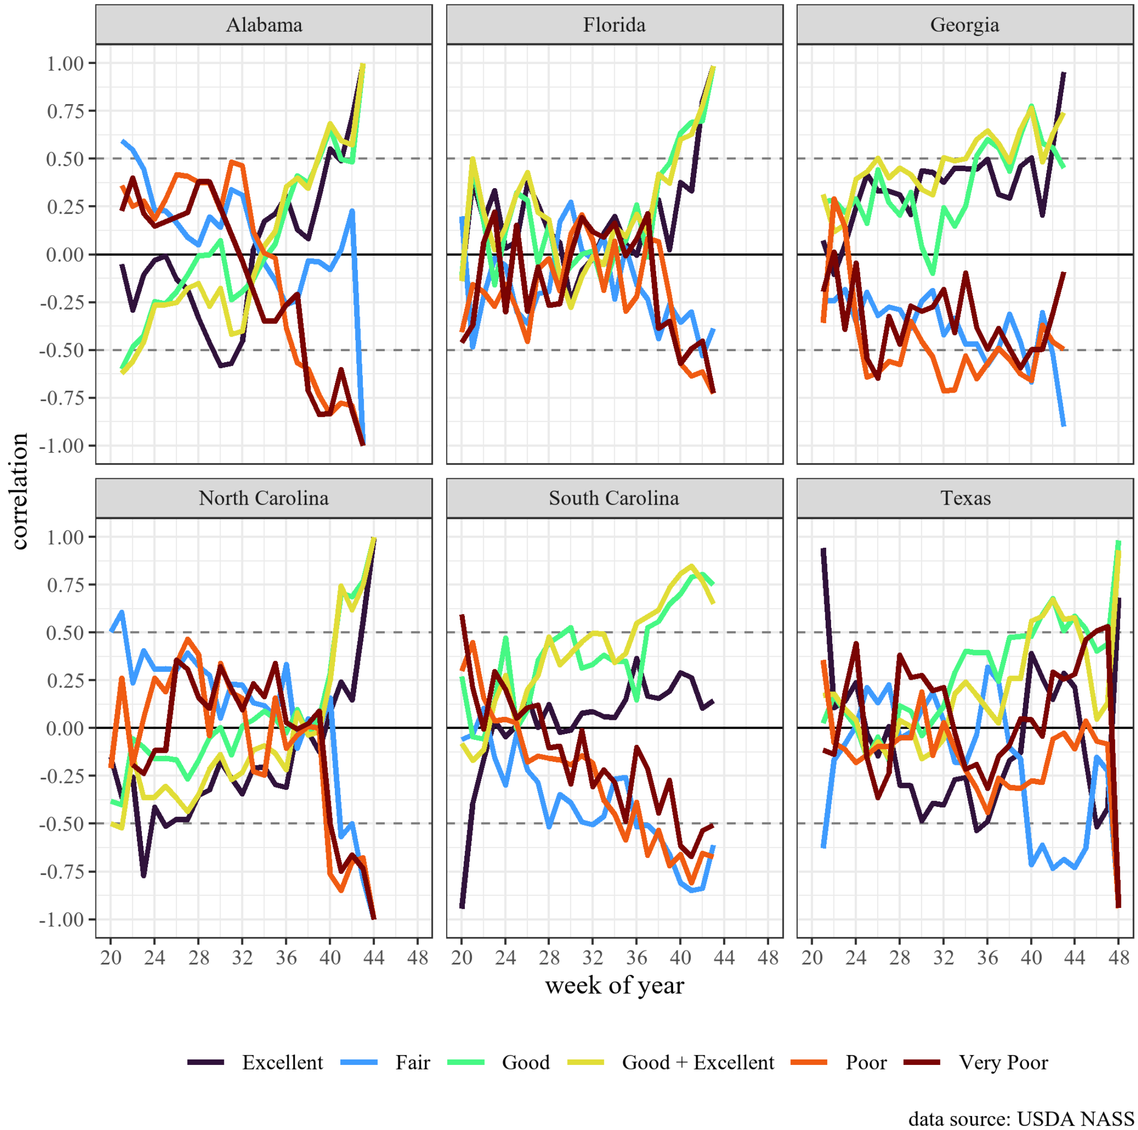

Georgia

Good and Good + Excellent are positively correlated to yield starting at week 32 of the year. Good and Good + Excellent continue to be positively correlated through week 40. Excellent condition trends towards a positive correlation, but is not correlated until week 43 of the year. Week 43 is the end of harvest season. The Fair condition follows the inverse of Excellent. Fair has no correlation until week 43. After week 43 Fair becomes a negative correlation. This means there is no relationship between yield with respect to Fair and Excellent crop conditions during the growing season from 2013 to 2025 in Georgia. Poor condition reaches a negative correlation around the 31-week mark and stays highly correlated in the negative direction through the end of the growing season. Very Poor becomes correlated at week 33 of the year and stays negatively correlated through the rest of the growing season.

Alabama

The Excellent condition variable for the state of Alabama is negatively correlated at the week 22 mark, but that relationship changes immediately and is not correlated again until the 30-week mark. It is negatively correlated at week 30. It finishes out the peanut growing season, becoming positively correlated at week 39 and stays that way to the end of harvest. The Good and Good + Excellent conditions follow the same pattern. They both become positively correlated at the 35-week mark and it stays that way through the end of the growing season. This reveals there is a relationship between Good and Good + Excellent with week 35 forward. Fair condition is marginally correlated for weeks 36 and 37, but is not relevant in any other weeks. The Poor variable becomes relevant or correlated during week 35 of the year. It stays negatively correlated for the rest of the growing season. The Very Poor condition has a vast number of data missing. The data that is present reveals that Very Poor becomes negatively correlated at the 37/38 week and finishes out the growing season correlated.

North Carolina

The Excellent variable is negatively correlated in weeks 22 and 23. It loses correlation from weeks 23 to 41. Good and Good + Excellent follow the same pattern as Excellent. Both are negatively correlated at the 21-week mark. Excellent, Good, and Good + Excellent all become positively correlated at week 40 and stay that way till the end of the growing season. Fair inversely follows the Excellent variable. It is positively correlated at week 22, but loses correlation until week 41. Fair stays negatively correlated till the end of the growing season. Poor and Very Poor are negatively correlated at the 40-week mark.

South Carolina

The Excellent variable for South Carolina begins at the 20-week mark where it is significant in the negative direction. It only stays correlated for 2 weeks and then it never becomes correlated again. The Good and Good + Excellent variables follow the same patterns as for the other states. They are very closely related and are positively correlated at weeks 35 and 36. They stay positively correlated for the remaining weeks. The Fair crop condition becomes negatively correlated early at the 28 weeks mark. Fair has a few ups and downs of being correlated, but at week 35 it remains negatively correlated for the rest of the growing season. Poor and Very Poor follow an inverse relationship to Go and Good + Excellent. Poor and Very Poor both become negatively correlated around the 35-week mark. These conditions stay correlated through the end of the growing season.

Florida

The Excellent variable for Florida becomes positively correlated around week 42 of the year. Week 42 of the year falls into the middle of October. This is the middle of harvest season. Good and Good + Excellent conditions continue to follow suit. They start out in week 20 negatively correlated, but lose correlation up until weeks 41 and 42. In those weeks they are positively correlated and stay that way for the rest of the growing season. The Fair variable never becomes correlated meaning there is no relation between yield to crop conditions through the growing season. Poor Condition hits a negative correlation at the 29 week of the year mark, but does not stay correlated. Poor hits correlation again at the week 40 mark and maintains that relationship till the end of harvest. Very Poor is not correlated until the end of the growing season. Week 39 is when Very Poor becomes negatively correlated.