Fixed-Horizon Commodity Price Forecasts from WASDE Reports

Ivy Mackereth, and Siddhartha Bora

Suggested citation format:

Mackereth, I., and Bora, S. 2026. “Fixed-Horizon Commodity Price Forecasts from WASDE Reports.” Proceedings of the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management, Chicago, Illinois, April 20-21, 2026. [http://www.farmdoc.illinois.edu/nccc134/paperarchive]

Fixed-Horizon Comodity Price Forecasts from WASDE Reports

Ivy Mackereth[*†] and **Siddhartha Bora**[*]

Paper prepared for the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management, 2026.

[*]West Virginia University, Davis College of Agriculture and Natural Resources

[†]We are grateful for the valuable feedback and insightful discussions received from attendees during the Farm Foundation Agricultural Scholars Capstone Presentation at USDA-ERS, Washington, D.C. in December 2025. Corresponding Author: Ivy Mackereth, e-mail: im00031@mix.wvu.edu

Fixed-Horizon Comodity Price Forecasts from WASDE Reports

Abstract

The World Agricultural Supply and Demand Estimates (WASDE) reports play a central role in shaping expectations for major agricultural commodities, yet their fixed-event structure limits their use in forward-looking policy analysis. This study develops 12-month-ahead fixed-horizon price forecasts for corn, soybeans, and wheat by applying an optimal-weighting approach and a standard ad hoc aggregation benchmark to overlapping current- and next-year WASDE releases. We evaluate the constructed forecasts against realized prices and futures market expectations using Diebold-Mariano tests, Mincer-Zarnowitz efficiency regressions, and forecast encompassing tests. Optimal weighting improves point accuracy for all three commodities and forecast efficiency for corn. Accuracy gains are most pronounced for soybeans, moderate for corn, and modest but consistent for wheat. We further show that the fixed-horizon transformation eliminates statistically significant seasonal patterns in soybean forecast errors that persist under both ad hoc and futures-based specifications, suggesting that a meaningful share of the forecast error seasonality documented in prior literature reflects the fixed event design rather than the underlying information environment. These results demonstrate that reframing WASDE forecasts at a constant horizon improves accuracy, efficiency, and interpretability, with direct applications to farm budgeting, forward pricing, and revenue planning.

Keywords: commodity prices, WASDE forecasts, forecast evaluation

1 Introduction

Many economic decisions in agricultural markets require price expectations defined over a constant planning horizon. Yet the most widely used public price forecasts for major commodities are anchored to a marketing-year endpoint that shifts mechanically as the season progresses. The World Agricultural Supply and Demand Estimates (WASDE) report, released monthly by the World Agricultural Outlook Board (WAOB), is one of the most influential public sources of information in agricultural markets, with its forecasts of supply, demand, inventories, and prices shaping expectations and informing decisions by producers, traders, and policymakers (Isengildina Massa, Karali, and Irwin 2024a). WASDE price forecasts are produced within a fixed-event forecasting framework, in which forecasts are updated each month as the terminal event approaches. For most commodities, this is the conclusion of the marketing year. As a result, the effective forecast horizon varies systematically over time. Early in the marketing year, forecasts reflect expectations over a long horizon, while later forecasts pertain to a much shorter remaining period. Forecasts released in different months therefore differ mechanically in horizon length, yet are often treated as comparable. This both limits economic modeling and policy analysis and creates a gap between how forecasts are produced and how they are used, making it difficult to draw clear conclusions about their informational content. Despite this well-recognized limitation, nearly all empirical evaluations of WASDE forecasts continue to work within the fixed-event framework.

This study examines whether translating WASDE price forecasts into a constant 12-month horizon improves their accuracy, efficiency, and interpretability for corn, soybeans, and wheat. We do so by adapting an optimal-weighting framework, based on the covariance structure of monthly price changes, to overlapping current- and next-year WASDE releases. We evaluate the constructed forecasts against realized prices and futures-market benchmarks using Diebold–Mariano (DM) tests of predictive accuracy, Mincer-Zarnowitz (MZ) regressions of forecast efficiency, and pairwise forecast encompassing tests supplemented by ridge-based Granger–Ramanathan (GR) combinations. We find that optimally weighting the information already embedded in WASDE reports yields fixed-horizon price forecasts that are more accurate, more efficient, and more informative than either standard fixed-event releases or futures-derived benchmarks, and that expressing WASDE forecasts at a constant horizon eliminates statistically significant seasonal patterns in soybean forecast errors that persist under both ad hoc and futures-based approaches. The latter finding suggests that the fixed-event horizon design itself, not solely the underlying agricultural information environment, is a source of the forecast error seasonality documented in prior literature.

A large body of literature evaluates the accuracy and informational content of agricultural forecasts. In the context of WASDE, Hoffman et al. (2015) document the performance of USDA price projections for U.S. corn and highlight how forecast accuracy evolves as the marketing year progresses. Other work documents systematic seasonal patterns in WASDE forecast errors and revisions that reflect the interaction between the marketing calendar and the arrival of new crop information (Isengildina Massa, Irwin, and Good 2006; Xiao, Hart, and Lence 2017). More broadly, the evaluation of forecast performance often relies on principles of forecast efficiency and unbiasedness, which provide benchmarks for assessing whether forecast errors systematically incorporate available information (Nordhaus 1987). At the same time, a long-standing literature shows that combining information from multiple forecasts can improve predictive performance relative to relying on a single source (Bates and Granger 1969). These insights suggest that the way forecasts are structured and combined can materially affect both their interpretation and their usefulness for decision-making.

Many applications require forecasts to be defined over a constant time period, such as the next 12 months, rather than tied to a crop production cycle-derived endpoint. When horizons vary across months, forecast accuracy and interpretation can change mechanically over time, making it difficult to evaluate performance or to use the information consistently in empirical models (Xiao, Hart, and Lence 2017; Ganics, Rossi, and Sekhposyan 2024). The fixed-event structure can also generate persistent revision patterns and correlated errors, a phenomenon that is well documented in both agricultural and macroeconomic forecasting literature (Coibion and Gorodnichenko 2015; Ganics, Rossi, and Sekhposyan 2024; Xiao, Hart, and Lence 2017). Methods for translating fixed-event forecasts into fixed-horizon projections have been developed in the literature, including weighted combinations of sequential forecasts (Dovern, Fritsche, and Slacalek 2012; Knüppel and Vladu 2025). However, such methods have rarely been applied in agricultural forecasting contexts, and no prior study has translated WASDE forecasts into a fixed-horizon framework. This gap limits both the interpretability of WASDE forecasts and their usefulness for the many economic decisions that require a price outlook defined at a consistent forward horizon.

Our analysis yields several findings with both methodological and practical implications. Our approach does not replace the existing fixed-event framework. Instead, it provides a complementary way to translate WASDE forecasts into a format better aligned with forward-looking economic analysis. First, fixed-horizon forecasts constructed from WASDE data achieve higher accuracy than those derived from the futures forward curve alone. Second, we show that expressing WASDE forecasts at a fixed horizon eliminates seasonal patterns in soybean forecast errors. Formal tests that account for serial correlation fail to reject the absence of seasonality for all optimal weighting specifications, including the estimated persistence parameter, while strongly rejecting the null under ad hoc and futures approaches. This provides direct evidence that a meaningful share of the forecast error seasonality documented in prior WASDE literature reflects the structure of the fixed-event horizon rather than the agricultural information environment itself. Third, we find evidence that forecast efficiency improves when forecasts are evaluated at a constant forward-looking horizon, suggesting that some of the information frictions documented in prior work may be alleviated by this transformation. Finally, we demonstrate practical applications for policy and planning, including farm budgeting and forward pricing decisions. These applications show that fixed-horizon forecasts provide actionable guidance beyond what is offered in standard WASDE releases.

The paper proceeds as follows. Section 2 reviews the related literature, highlighting prior work on WASDE forecast accuracy, seasonality, and information rigidity, as well as macroeconomic methods for translating fixed-event forecasts to fixed horizons. Building on this review, Section 3 presents a conceptual framework that formalizes the influence of structured information arrival, horizon design, and forecast combination on forecast errors. Section 4 describes the data, including how realized 12-month weighted-average prices are constructed and how WASDE releases are aligned with futures contracts. Section 5 outlines the two weighting methods, ad hoc and optimal, and the evaluation framework used to assess their accuracy and efficiency, covering DM, MZ, and forecast-encompassing (FE) tests, as well as moving-block bootstrap (MBB) inference. Section 6 reports the results, including comparisons with the futures forward curve and analyses of seasonal and efficiency patterns, and illustrates practical applications of fixed-horizon forecasts to forward pricing, crop insurance elections, and farm bill program enrollment. Section 7 concludes.

2 Related Literature

A substantial literature evaluates the performance and informational role of United States Department of Agriculture (USDA) forecasts, particularly those published in WASDE reports. A comprehensive review of the literature is provided by Isengildina Massa, Karali, and Irwin (2024a), who examine more than one hundred studies assessing USDA forecast performance. A central finding in this body of work is that forecast errors for fixed-event USDA projections decline as the marketing year progresses and additional information becomes available. This systematic decline in forecast errors reflects the structured process through which WASDE forecasts incorporate new information about production, demand, and market conditions. Rather than indicating weaknesses in the forecasts, the systematic decline in errors reinforces their credibility as timely assessments of evolving agricultural markets.

The literature further emphasizes the broader economic role of these forecasts. Because WASDE reports are widely followed by producers, traders, and policymakers, they function as a central public signal in agricultural markets. For example, Isengildina Massa, Karali, and Irwin (2024b) show that USDA projections improve market efficiency and reduce information asymmetries among market participants. Empirical studies also document an immediate market effect and show that futures prices respond systematically to the information contained in WASDE reports, indicating that releases convey new information to the market. Similarly, Isengildina-Massa et al. (2008) find that implied volatility in corn and soybean options increases around WASDE release dates, reflecting heightened uncertainty and information arrival during these announcements. In this sense, WASDE forecasts not only provide predictions but also actively shape expectations and price discovery in agricultural markets.

This importance has also motivated comparisons between USDA forecasts and private forecasts, particularly those derived from futures markets. Hoffman et al. (2015) find that WASDE corn price forecasts outperform several futures-based benchmarks across a majority of forecasting periods between 1980 and 2012. At the same time, they show that combining USDA forecasts with futures prices produces even lower errors, suggesting that public and private information sets are complementary rather than competing. Related work by Etienne et al. (2023) further documents how information from WASDE releases interacts with futures markets and is incorporated into prices and public reports. Together, these studies establish that WASDE forecasts contain distinct and economically valuable information.

While the literature largely affirms the informational value of USDA forecasts, it also documents systematic patterns in their revisions. Several studies identify seasonal dynamics in forecast errors and updates that reflect the interaction between the marketing calendar and the arrival of new crop information. For example, Xiao, Hart, and Lence (2017) show that revisions to USDA forecasts follow predictable seasonal patterns, while earlier work by Isengildina Massa, Irwin, and Good (2006) documents similar dynamics in production forecasts for major crops. These patterns arise partly from the institutional structure of the forecasts themselves. Because WASDE projections are expressed relative to a fixed marketing-year endpoint, successive revisions repeatedly update the same target outcome as new information arrives. As a result, some of the persistence and seasonality observed in forecast errors and revisions reflect the design of the forecasting horizon rather than solely the underlying information environment.

This observation raises the natural question of how the dynamics of USDA forecasts might appear if they were expressed relative to a constant time-ahead horizon instead of a fixed marketing-year endpoint. Insights from the macroeconomic forecasting literature provide a useful methodological foundation for addressing this question. A large body of literature shows that combining forecasts from different models or information sets can improve predictive performance. Early work by Bates and Granger (1969) introduced the formal framework for forecast combination, and subsequent studies show that relatively simple combinations often perform as well as or better than individual forecasts in practice (Stock and Watson 2004). A comprehensive overview of this literature is provided by Wang et al. (2023).

Building on this, later work shows that weighted combinations can also be used to reorganize forecasts across horizons. Dovern, Fritsche, and Slacalek (2012) show how fixed-event forecasts can be translated into fixed-horizon projections by combining sequential forecasts. Knüppel and Vladu (2025) extend this framework by formalizing alternative weighting schemes and developing a flexible methodology for constructing fixed-horizon forecasts while preserving the informational content of the original projections. In this sense, fixed-horizon forecasts can be viewed as structured combinations of forecasts made at different points in time. These approaches show that forecasts produced for a fixed target date can be reorganized to align with decision-relevant horizons without altering the underlying information set.

Related insights emerge from the literature on information rigidity and forecast smoothing. Early theoretical work by Mankiw and Reis (2002) proposes a sticky-information model in which agents update their information sets only gradually as new information becomes available. Building on this idea, Coibion and Gorodnichenko (2015) develop an empirical approach showing that forecasters often incorporate new information gradually, leading to persistent patterns in forecast revisions. Similar behavior has been documented in agricultural forecasting contexts. Evidence of correlated smoothing has been observed in USDA production forecasts for corn and soybeans (Isengildina Massa, Irwin, and Good 2006), while persistent revision dynamics have also been identified in ending stocks forecasts (Xiao, Hart, and Lence 2017). When forecasts repeatedly revise the same marketing-year outcome, some persistence in revisions can arise mechanically. Fixed-horizon forecasts, by contrast, target a different outcome each period. This distinction reduces the influence of older information and may allow forecasts to respond more directly to new data. Expressing forecasts at a constant horizon therefore provides an alternative way to examine whether observed persistence reflects slow information updating or the structure of the forecasting target itself

Overall, the literature establishes that WASDE forecasts are accurate, influential, and central to agricultural markets. At the same time, it shows that several commonly observed patterns in forecast errors and revisions are closely tied to the fixed-event structure of marketing-year projections. Despite this insight, most empirical evaluations continue to analyze USDA forecasts within the traditional fixed-event framework.

Our study builds directly on this foundation by reconsidering the horizon through which WASDE forecasts are expressed. Using existing USDA price projections, we construct fixed-horizon forecasts that maintain the underlying information content while expressing predictions at a constant time-ahead horizon. This approach provides a complementary perspective on USDA forecasts, allowing us to examine revision dynamics, seasonality, and responsiveness to new information in a framework that aligns more closely with forward-looking decision horizons. By reframing how these widely used forecasts are evaluated, our analysis extends the literature while reinforcing the central informational role of USDA projections in agricultural markets.

3 Conceptual Framework

While WASDE forecasts are widely regarded as accurate and timely, the structure of the reporting process means that forecasts formed at different points in the marketing year may inherently reflect different levels of precision, as highlighted in prior literature. In this section, we develop a conceptual framework linking structured information arrival, forecast horizon definition, and forecast combination to the empirical comparisons examined in our work. By connecting forecast timing and weighting decisions to the underlying dynamics of agricultural information, this framework motivates the construction and evaluation of fixed-horizon average price forecasts for key commodities.

3.1 Information Arrival and the Fixed-Event Structure

Let P_{t+k} denote the realized farm price in month $t+k$, where t is the forecast origin and $k = 1, 2, \ldots H$ indexes months ahead, with H representing the forecast horizon. Where the information set available at time t is denoted as $\mathcal{I}_t$, the optimal forecast formed at time t is

$$\hat{P}_{t+k|t} = E\left[P_{t+k} \mid \mathcal{I}_t\right]. \tag{1}$$

Agricultural prices evolve as new information arrives each month, which can be represented as

$$P_{t+k} = P_{t+k-1} + \nu_{t+k}, \tag{2}$$

where ν_{t+k} represents new information shocks with seasonal variance

$$Var(\nu_{t+k}) = \sigma_s^2, \tag{3}$$

with s indexing the stage of the crop production cycle.

Prior studies document that fixed-event forecast errors decline as the marketing year progresses (Isengildina Massa, Karali, and Irwin 2024a), and that revisions display seasonal patterns and smoothing behavior (Isengildina Massa, Irwin, and Good 2006; Xiao, Hart, and Lence 2017). Models of information rigidity suggest that forecasters gradually incorporate new information over time (Coibion and Gorodnichenko 2015). While we do not explicitly test for seasonal patterns in fixed-event forecast errors, these findings motivate expressing forecasts over a consistent, fixed horizon, which aligns the timing of information with the prediction target and allows for clearer assessment of forecast accuracy.

3.2 Fixed Horizon Representation

Under the fixed-event design, early-season forecasts are formed under greater uncertainty than later-season forecasts, creating variation in effective forecast horizons across report dates (Isengildina Massa, Karali, and Irwin 2024a). Conceptually, the fixed-horizon forecast at month t can be written as

$$\hat{P_t^H} = \sum_{k=1}^{H} w_{t+k} P_{t+k}, \tag{4}$$

where w_{t+k} are weights that reflect the relative contribution of each month. Expressing the forecast at a consistent horizon aligns the information gap between forecast and realization, so that the variance of forecast errors reflects uncertainty at the fixed horizon rather than calendar timing (Knüppel and Vladu 2025). By formalizing the fixed-horizon forecast in this way, we can directly evaluate whether aligning the information gap improves predictive performance, testing our first hypothesis:

Hypothesis 1. Expressing USDA forecasts at a fixed horizon should improve predictive performance relative to fixed-event forecasts.

The Methods section details how we operationalize this idea for a 12-month horizon using overlapping current- and next-year WASDE projections with specific weighting schemes. This framework naturally extends to evaluating forecast combinations, which motivates our second empirical hypothesis regarding optimal weighting strategies.

3.3 Forecast Combination and Weighting

Forecast performance also depends on how multiple information sources are combined. Suppose the fixed-horizon forecast is constructed from current- and next-year WASDE projections $P_{t+k|t}$ and $P_{t+k+H|t}$. A linear combination is

$$\hat{P_t^H} = \sum_{k=1}^{H} \left[w_k P_{t+k|t} + (1 - w_k) P_{t+k+H|t} \right], \tag{5}$$

where w_k is the weight assigned to the current-year forecast for month k. The expected squared error for this combination is

$$MSE(w_k) = E\left[\left(P_{t+k}^H - \hat{P}_{t+k}^H \right)^2 \right]. \tag{6}$$

Unless the two forecasts have identical variance and perfectly correlated errors, equal weights, such as those derived for our ad hoc aggregations, are generally suboptimal. Optimal weighting, which minimizes MSE with respect to w_k, exploits differences in forecast precision and error correlation. This framework naturally extends to evaluating how multiple forecast sources are combined, motivating our second empirical hypothesis on optimal weighting strategies:

Hypothesis 2. Forecast combinations using optimal, data-driven weights should weakly outperform ad hoc weighting schemes whenever forecast errors differ in variance or are imperfectly correlated.

The Methods section describes how these weights are computed, including adjustments for the marketing-year basis of WASDE reports and the autocorrelation structure of monthly prices, enabling a consistent 12-month fixed-horizon forecast that incorporates both sources of information.

4 Data

Our analysis draws on three primary sources of information: average farm price forecasts from WASDE reports, realized price data, and futures market prices for the commodities of interest. All data are observed at a monthly frequency and aligned to the timing of WASDE report releases. Together, these sources provide a comprehensive view of both market expectations and realized outcomes for corn, soybeans, and wheat.

Realized prices for commodities are constructed using historical monthly cash price data published by the USDA National Agricultural Statistics Service (NASS) and obtained from the NASS QuickStats Database (United States Department of Agriculture, National Agricultural Statistics Service 2024). To align with the economic concept underlying WASDE price forecasts, ex post realized prices are defined as weighted averages of monthly cash prices over the subsequent 12 months. The weights reflect commodity-specific marketing patterns and actual marketed volumes. Total marketed volumes are drawn from the final marketing-year values reported in WASDE (United States Department of Agriculture 2026), and monthly marketing percentages, which capture the timing of physical sales within the marketing year, are obtained from the USDA Economic Research Service (ERS) (U.S. Department of Agriculture, Economic Research Service 2025). Following the ERS methodology for constructing season-average prices, monthly marketing percentages are averaged over the previous five years to obtain stable seasonal shares, which are then combined with total marketed volumes to construct the final monthly weights (United States Department of Agriculture, Economic Research Service 2025).

Futures market data are used both to supplement WASDE information during periods when overlapping marketing-year forecasts are unavailable and to provide an external benchmark for forecast evaluation. We obtain CBOT futures for each commodity using contract prices observed around the date of each WASDE release. For corn and wheat, there are five primary contract months corresponding to March, May, July, September, and December delivery. To construct a continuous series that aligns with the marketing year, we fill in the remaining months using prices from the nearest available contracts, following a procedure similar to the USDA's method for calculating the season-average price (United States Department of Agriculture, Economic Research Service 2025). Under this scheme, the December contract is applied to September through November; the March contract to December through February; the May contract to March and April; the July contract to May and June; and the September contract to July and August. This procedure assigns each marketing-year month to a relevant, actively traded contract, producing a smooth and consistent series for analysis. For soybeans, seven primary futures contracts are available, corresponding to January, March, May, July, August, September, and November. Following the same framework used for corn and wheat, contract prices are assigned to calendar months based on delivery proximity, using only futures contracts that are actively traded and observable in each month. Under this approach, each assigned price reflects information observable at the time. Futures prices are adjusted for basis using USDA ERS data (U.S. Department of Agriculture, Economic Research Service 2025) to align them with cash prices.

Average farm price forecasts are obtained from the monthly WASDE reports published by the USDA. For each release, WASDE reports marketing-year average farm price forecasts for both the current and, when available, the subsequent marketing year. These forecasts are fixed-event projections, as they are tied to the conclusion of the marketing year rather than a constant horizon. Thus, each sequential forecast is an update of the previous forecast describing the same terminal event. For corn and soybeans, the marketing year runs from September 1 through August 31, implying that the effective forecast horizon declines mechanically over the course of the year and resets each September. For wheat, the marketing year runs from June 1 through May 31. The availability of overlapping marketing-year forecasts varies within the year. For corn and soybeans, forecasts for both the current and the upcoming marketing year are available from May through November, allowing for direct construction of 12-month-ahead implied price forecasts during this period. For wheat, overlapping forecasts are available from May through September (Isengildina Massa, Karali, and Irwin 2013). From December through April for corn and soybeans, and October through April for wheat, forecasts for the next marketing year are not yet released, creating gaps in the data needed to construct fixed-horizon expectations. To address this, we supplement the available WASDE forecasts with futures market prices, using the closing value of the nearest relevant futures contract to fill months without overlapping WASDE forecasts. This produces a continuous series of fixed-horizon forecasts across the entire year, ensuring that 12-month-ahead expectations are consistently defined regardless of the timing of WASDE releases.

Together, these data sources provide a comprehensive view of market expectations and realized outcomes for corn, soybeans, and wheat. WASDE forecasts, futures prices, and realized cash prices are carefully aligned in time to prevent look-ahead bias. Forecasts are treated as observed at the time of release, futures are drawn from trading days immediately surrounding each release, and realized prices are based only on subsequent observations. Futures contracts are mapped to months using a nearest-contract approach, producing continuous marketing-year series. By combining USDA-reported information with market-based data, the dataset captures both official projections and real-time market signals, providing a robust foundation for evaluating forecast performance.

5 Empirical Methods

This study constructs 12-month-ahead fixed-horizon forecasts of corn, soybean, and wheat prices using the information embedded in WASDE reports. While WASDE provides forecasts on a fixed-event basis for the current and subsequent marketing years, market participants and policymakers often require a consistent forward-looking measure over a fixed horizon. To address this, we translate overlapping current- and next-year WASDE forecasts into a 12-month forecast. For each forecast month, the realized ex post 12-month price is calculated as a weighted average of the next twelve months' observed farm prices

$$P_t^{12} = \sum_{k=1}^{12} w_{t+k}\, P_{t+k}, \tag{7}$$

In this equation, t denotes the current month, which serves as the forecast origin, and $k = 1, 2, \ldots, 12$ indexes the subsequent 12 months. P_{t+k} represents the observed farm price in month $t + k$, while w_{t+k} is the corresponding weight for that month. The corresponding weight is based on the monthly share of marketed volume. This weighting captures the seasonal distribution of sales and provides a realistic measure of the realized price over the forecast horizon. To construct forecasts for the same 12-month horizon, we implement two complementary approaches: an ad hoc aggregation method and a more sophisticated optimal-weight approximation, presented in order of increasing complexity.

5.1 Ad hoc Aggregation Method

The ad hoc aggregation method provides a straightforward approach to approximate a 12-month fixed-horizon forecast using WASDE's current- and next-year fixed-event forecasts (Dovern, Fritsche, and Slacalek 2012). Let $P_{t+k|t}$ denote the k-month-ahead forecast based on information available at month t. If the report provides forecasts for both the current and next marketing year, the 12-month-ahead forecast can be expressed as

$$\hat{P}_t^{12} \approx \sum_{k=1}^{12} \left[\frac{12-k}{12}\, P_{t+k|t} + \frac{k}{12}\, P_{t+k+12|t} \right] \tag{8}$$

As before, t denotes the current month, which serves as the forecast origin, and $k = 1, 2, \ldots, 12$ indexes the subsequent 12 months. Forecasts are then made for future months, where $P_{t+k|t}$ represents the current-year forecast and $P_{t+k+12|t}$ is the corresponding next-year forecast for the same month in the following marketing year. Here, the weights shift linearly over the next horizon. Specifically, early months are weighted more heavily toward the current-year forecast, and later months receive a larger contribution from the next-year forecast. By construction, the weights for each month sum to one, ensuring that the fixed-horizon forecast is a proper convex combination of the two marketing-year forecasts.

5.2 Optimal-Weight Approximation Method

While the ad hoc method provides a simple and transparent benchmark for constructing fixed-horizon average price estimates using WASDE data, it does not account for the covariance structure of monthly price changes or differences in forecast precision across marketing years. To address this limitation, we implement an optimal-weight approximation following the framework developed by Knüppel and Vladu (2025). This combines current- and next-year WASDE forecasts with weights chosen to minimize the expected squared error over the 12-month horizon.

Formally, let $P_{t+k|t}$ and $P_{t+k+12|t}$ denote the current- and next-year forecasts, respectively. The 12-month fixed-horizon forecast is defined as

$$\hat{P}_t^{12} = \sum_{k=1}^{12} \left[w_k P_{t+k|t} + (1 - w_k) P_{t+k+12|t} \right], \tag{9}$$

where t denotes the current month, $k = 1, \ldots, 12$ indexes months in the forecast horizon, and w_k is the optimal weight in month k. The weight is chosen to minimize the expected squared approximation error:

$$w_k^* = \arg\min_{w_k} E\left[\left(P_{t+k}^{12} - \hat{P}_{t+k}^{12} \right)^2 \right]. \tag{10}$$

We assume that monthly prices follow an AR(1) process:

$$P_{t+k} = \rho P_{t+k-1} + \epsilon_{t+k}, \quad \epsilon_{t+k} \sim N(0, \sigma^2), \tag{11}$$

so that the covariance between any two months t and $t+k$ is

$$\mathrm{Cov}(P_t, P_{t+k}) = \frac{\sigma^2}{1-\rho^2}\rho^{|k|}. \tag{12}$$

For each commodity i, the AR(1) persistence parameter ρ_i is estimated by regressing the 12-month-ahead forecast error $e_{i,t}$ on its 12-month lag:

$$e_{i,t} = \rho_i \, e_{i,t-12} + u_{i,t}, \tag{13}$$

where $u_{i,t}$ is an error term. The resulting estimated coefficient $\hat{\rho}$ is used to construct the optimal-weight forecasts, allowing the weighting scheme to adapt to the observed persistence in forecast errors. To assess sensitivity, forecasts are also constructed using four alternative fixed AR(1) coefficients $\rho \in \{0.0, 0.5, 0.8, 0.99\}$. The choice of these alternate coefficients is motivated by prior literature (Knüppel and Vladu 2025).

To explicitly define the optimal weights, let $\mathcal{M}_{\text{cur}}$ and $\mathcal{M}_{\text{next}}$ denote the sets of months corresponding to the current- and next-marketing-year averages, respectively. Let $k = 1, \ldots, 12$ index months in the forecast horizon and $\ell \in \mathcal{M}_{\text{cur}}$ or $\ell \in \mathcal{M}_{\text{next}}$ index months in the marketing-year averages. Then the expected covariance between the forecast horizon and each marketing-year average is:

$$\mathrm{Cov}\big(P_t^{12}, P_{\text{cur}}\big) = \frac{1}{12\,|\mathcal{M}_{\text{cur}}|}\sum_{k=1}^{12}\sum_{\ell\in\mathcal{M}_{\text{cur}}}\mathrm{Cov}\big(P_{t+k}, P_{t+\ell}\big), \tag{14}$$

$$\mathrm{Cov}\big(P_t^{12}, P_{\text{next}}\big) = \frac{1}{12\,|\mathcal{M}_{\text{next}}|}\sum_{k=1}^{12}\sum_{\ell\in\mathcal{M}_{\text{next}}}\mathrm{Cov}\big(P_{t+k}, P_{t+\ell}\big). \tag{15}$$

The variance of the difference between current- and next-year averages is

$$\mathrm{Var}(P_{\text{cur}} - P_{\text{next}}) = \mathrm{Var}(P_{\text{cur}}) + \mathrm{Var}(P_{\text{next}}) - 2\,\mathrm{Cov}(P_{\text{cur}}, P_{\text{next}}), \tag{16}$$

and the resulting closed-form expression for the optimal weight on the current-year forecast is

$$w_k^* = \frac{\mathrm{Cov}(P_t^{12}, P_{\text{cur}}) - \mathrm{Cov}(P_t^{12}, P_{\text{next}})}{\mathrm{Var}(P_{\text{cur}} - P_{\text{next}})}. \tag{17}$$

In constructing the optimal-weight forecasts, the monthly weights are adjusted to account for the fact that WASDE forecasts are reported on a marketing-year basis, rather than a calendar-year basis. This ensures that the weighted combination of current- and next-year forecasts aligns with the true timing of marketed volumes, producing a consistent 12-month fixed-horizon forecast. The constructed weights, including both ad hoc and optimal weights, are depicted in Figure 1.

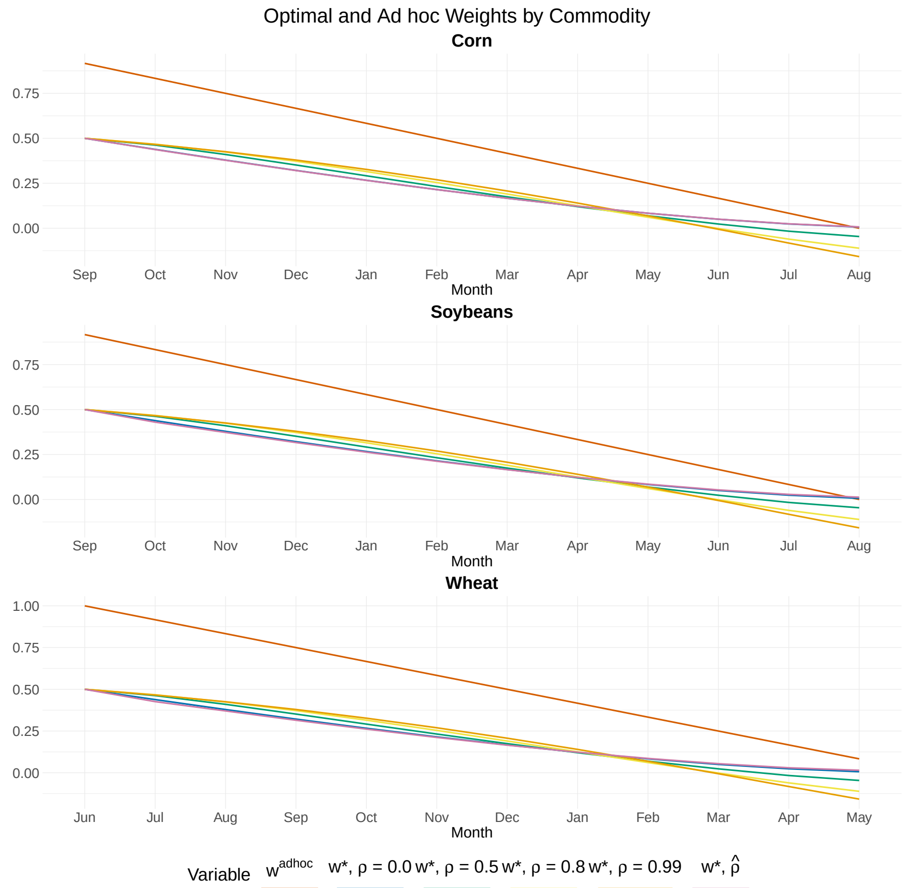

Figure 1: Visual demonstration of the weighting schemes used to construct 12-month-ahead forecasts, presented in marketing-year order. The ad hoc aggregation (red) transitions linearly from current-year to next-year forecasts, with weights adjusted to align with each commodity's marketing year. The optimal-weight schemes include the estimated AR(1) coefficient $\hat{\rho}$ as well as four fixed ρ values $(0, 0.5, 0.8, 0.99)$. Higher ρ values place greater emphasis on prior-month trends. All weights are further adjusted to reflect WASDE reporting on a marketing-year basis, ensuring the 12-month combination aligns with the timing of marketed volumes.

5.3 Forecast Evaluation and Inference

To assess and evaluate the accuracy and robustness of our constructed 12-month-ahead forecasts, we evaluate forecast errors for each weighting method using a combination of standard loss metrics, formal statistical tests, and bootstrap inference. Together, these procedures allow us to quantify the performance of each weighting scheme and assess forecast stability under realistic, out-of-sample conditions. The analysis uses monthly WASDE forecasts and realized farm prices spanning from September 2002 through September 2024 for corn and soybeans and June 2002 through June 2024 for wheat, effectively providing 264 months of observations for each commodity. This time span allows for robust evaluation of 12-month-ahead forecasts.

Forecast performance is measured using the error between the predicted 12-month average price and the realized ex-post 12-month price. For each forecast method i and forecast origin t, the forecast error is defined as

$$e_{i,t} = \hat{P}_{i,t}^{12} - P_t^{12}, \tag{18}$$

where $\hat{P}_{i,t}^{12}$ is the 12-month ahead forecast constructed using either ad hoc or optimal weight aggregations, and P_t^{12} is the realized 12-month average price. These errors form the basis for all accuracy metrics, statistical tests, and inference procedures described.

5.3.1 Loss-Based Accuracy Metrics

We summarize forecast performance using four complementary metrics: (1) mean absolute error (MAE), (2) mean absolute percentage error (MAPE), (3) root mean squared error (RMSE), and (4) root mean squared percentage error (RMSPE). For a forecast method i with errors $e_{i,t}$ over T periods, loss-based accuracy measures are defined as:

$$MAE_i = \frac{1}{T}\sum_{t=1}^{T} |e_{i,t}| \,. \tag{19}$$

$$MAPE_i = \frac{100}{T}\sum_{t=1}^{T} \left|\frac{e_{i,t}}{P_t^{12}}\right| \tag{20}$$

$$RMSE_i = \sqrt{\frac{1}{T}\sum_{t=1}^{T} e_{i,t}^2} \tag{21}$$

$$RMSPE_i = 100\sqrt{\frac{1}{T}\sum_{t=1}^{T} \left(\frac{e_{i,t}}{P_t^{12}}\right)^2} \tag{22}$$

MAE provides a straightforward measure of typical forecast accuracy by averaging absolute deviations, treating all errors equally. RMSE instead squares forecast errors before averaging, placing greater weight on large deviations and making it more sensitive to outliers. MAPE and RMSPE scale errors relative to the realized price, facilitating comparison across periods and commodities with different price levels. Similar to their level-based counterparts, RMSPE penalizes large proportional errors more heavily than MAPE. For each of these measures, lower values indicate superior forecast performance.

5.3.2 Predictive Accuracy

To formally evaluate differences in predictive accuracy across forecasting methods, we implement the DM test (Diebold and Mariano 1995; D. Harvey, Leybourne, and Newbold 1997). The DM test assesses whether the mean loss differential between two competing forecasts is statistically different from zero while accounting for serial correlation that may arise when evaluating multi-step-ahead forecasts.

Pairwise comparisons are conducted between the futures-based forecast, the ad hoc weighted aggregation, and the four optimal-weight approximations. The optimal-weight forecasts differ according to the assumed AR(1) autocorrelation parameter of monthly prices ($\rho = 0, 0.5, 0.8, 0.99$). Forecast errors are computed monthly for each method over the sample period.

For each pairwise comparison, the loss differential is defined using squared forecast errors,

$$d_t = e_{1,t}^2 - e_{2,t}^2, \tag{23}$$

where $e_{1,t}$ and $e_{2,t}$ denote the forecast errors from the two competing methods at time t. The DM statistic is calculated as

$$DM = \frac{\bar{d}}{\sqrt{\widehat{Var}(d_t)/T}}, \tag{24}$$

where $\bar{d}$ represents the sample mean of the loss differential and $\widehat{Var}(d_t)$ is a consistent estimate of its variance, allowing for serial correlation. T denotes the number of observations.

Because the forecasts evaluated in this study correspond to fixed 12-month marketing-year horizons, the DM test is implemented using a 12-period forecast horizon to account for serial correlation in the loss differentials that arises from overlapping multi-step forecast errors. Following D. Harvey, Leybourne, and Newbold (1997), the small-sample correction to the DM statistic is applied to improve finite-sample inference.

The procedure is repeated for all pairwise combinations of forecast methods within each commodity, producing a matrix of DM statistics and associated p-values that summarize relative predictive accuracy across forecasting approaches.

5.3.3 Forecast Efficiency

Following our evaluation of forecast significance, we assess whether forecasts are unbiased or efficient. To do this, we implement MZ forecast efficiency tests (Mincer and Zarnowitz 1969). That is, for each forecast constructed using the weighting method i, we estimate the MZ regression

$$P_t^{12} = \alpha_i + \beta_i \hat{P}_{i,t}^{12} + u_{i,t}, \tag{25}$$

where P_t^{12} is the realized ex post 12-month average price and $\hat{P}_{i,t}^{12}$ is the corresponding forecast constructed using method i.

Under the null hypothesis of forecast efficiency, the intercept and slope satisfy

$$H_0 : \alpha_i = 0 \quad \text{and} \quad \beta_i = 1. \tag{26}$$

We estimate α_i and β_i along with heteroskedasticity- and autocorrelation-consistent (HAC) standard errors with a lag-length of 12 to account for serial correlation induced by overlapping 12-month forecast horizons. We then use Wald tests to evaluate the joint null hypothesis for each forecast method. Failure to reject the joint null hypothesis indicates that the forecast is consistent with unbiasedness and efficiency. Deviations suggest systematic bias ($\alpha_i \neq 0$) or inefficiency ($\beta_i \neq 1$).

5.3.4 Forecast Encompassing

While the DM and MZ tests assess relative predictive accuracy and forecast efficiency, they do not indicate whether one forecast contains the information of another. To evaluate this, we implement pairwise FE tests (Clements and D. I. Harvey 2010).

For each pair of forecasts (i, j), we estimate the encompassing regression:

$$P_t^{12} = \alpha + \beta_i \hat{P}_{i,t}^{12} + \beta_j \hat{P}_{j,t}^{12} + u_t, \tag{27}$$

where P_t^{12} is the realized 12-month average price, and $\hat{P}_{i,t}^{12}$ and $\hat{P}_{j,t}^{12}$ are the forecasts from methods i and j.

The null hypotheses can be interpreted as follows:

- β_i significant and β_j not significant: method i *encompasses* method j.
- β_j significant and β_i not significant: method j *encompasses* method i.
- Both β_i and β_j are significant: each forecast contains incremental information beyond the other.
- Neither coefficient is significant: neither forecast adds meaningful information beyond the other.

We estimate β_i and β_j using ordinary least squares with HAC standard errors to account for serial correlation in the overlapping 12-month forecast horizons. These tests are conducted for all pairwise combinations of alternative weighting methods and the benchmark futures forecast for each commodity.

To formally exploit the complementary information identified by the FE tests, we construct Granger-Ramanathan (GR) combined forecasts using ridge regression (Granger and Ramanathan 1984). The combined forecast is

$$\hat{P}_t^{GR} = \sum_{i=1}^{N} w_i \hat{P}_{i,t}, \quad \mathbf{w} = \arg\min_{w} \left\{ \sum_t \left(P_t^{12} - \sum_i w_i \hat{P}_{i,t} \right)^2 + \lambda \sum_i w_i^2 \right\}, \tag{28}$$

where $\hat{P}_{i,t}$ are the individual forecasts, w_i are the GR weights, and λ is the ridge penalty. The ridge term shrinks the weights toward zero, which helps stabilize the regression when forecasts are highly correlated and prevents any single forecast from dominating the combination. Because our dataset includes forecasts for multiple persistence levels ($\rho = 0.0, 0.5, 0.8, 0.99, \hat{\rho}$), the GR regression simultaneously determines the relative contribution of each forecast across these specifications, allowing the model to capture complementary information and produce more accurate 12-month-ahead predictions.

5.4 Tests for Robustness

To ensure that our results are not sensitive to specific assumptions or sample choices, we implement several robustness checks. First, we use an MBB to assess the sampling uncertainty of forecast differences. Forecast errors are resampled in consecutive blocks of length $L = 12$, with replacement, $B = 1000$ times. For each replicate, we recompute forecast metrics and differences relative to the futures-based benchmark. By preserving time dependence in the errors, this approach captures shocks common to all forecasts and yields empirical distributions of MAE differences. From these distributions, we calculate median values, 95% confidence intervals, and the probability that an alternative weighting scheme outperforms the benchmark. The MBB provides a distribution-free check of both statistical and economic significance, complementing the DM tests.

Next, we examine whether forecast accuracy exhibits systematic seasonality. As an initial diagnostic, we aggregate monthly forecast errors over the full sample and conduct one-way ANOVA tests to assess whether MAEs differ across calendar months. The null hypothesis is equality of MAEs across months. Because fixed-horizon forecast errors are overlapping and therefore serially correlated, ANOVA does not provide reliable inference and is interpreted as descriptive. Our primary test instead estimates a regression of absolute 12-month-ahead forecast errors on calendar-month dummy variables and conducts a HAC-robust joint Wald test of their significance. To account for time-series dependence, we implement the HAC estimator of Newey and West (1987) with automatic bandwidth selection following Newey and West (1994). Under the null hypothesis, all monthly coefficients are jointly zero, implying no seasonal pattern in forecast accuracy. This framework provides statistically valid inference and serves as our main test of seasonality.

Together, these robustness checks reinforce the reliability of our conclusions and provide a comprehensive evaluation of both statistical and economic performance of the weighting schemes.

6 Results and Discussion

To provide context for evaluating our applied weighting schemes in constructing 12-month-ahead average price estimates from WASDE forecasts, we first examine the patterns of the estimated forecasts for corn, soybeans, and wheat. Figure 2 presents monthly prices for all three commodities, highlighting trends over the sample period.

A key input in the construction of our optimal weighting forecasts is the AR(1) coefficient ρ, which captures the persistence of forecast errors and informs the degree to which past errors should influence current weights. We estimate commodity-specific AR(1) processes for 12-month-ahead futures forecast errors to assess this persistence. For corn, the estimated persistence is $\hat{\rho} = -0.017$ (p = 0.793), indicating no evidence of serial correlation. Soybeans exhibit weak negative autocorrelation, with $\hat{\rho} = -0.107$ (p = 0.087), while wheat shows statistically significant negative autocorrelation, with $\hat{\rho} = -0.149$ (p = 0.018). Overall, the coefficients are small in magnitude, suggesting limited persistence in forecast errors across commodities, with mild mean reversion for soybeans and wheat.

To explore sensitivity, we also construct forecasts using four alternative ρ values, 0, 0.5, 0.8, and 0.99, in addition to the estimated $\hat{\rho}$. This allows us to compare the performance of the optimal weighting forecasts against conventional, fixed AR(1) assumptions and to visualize how varying ρ affects forecast responsiveness.

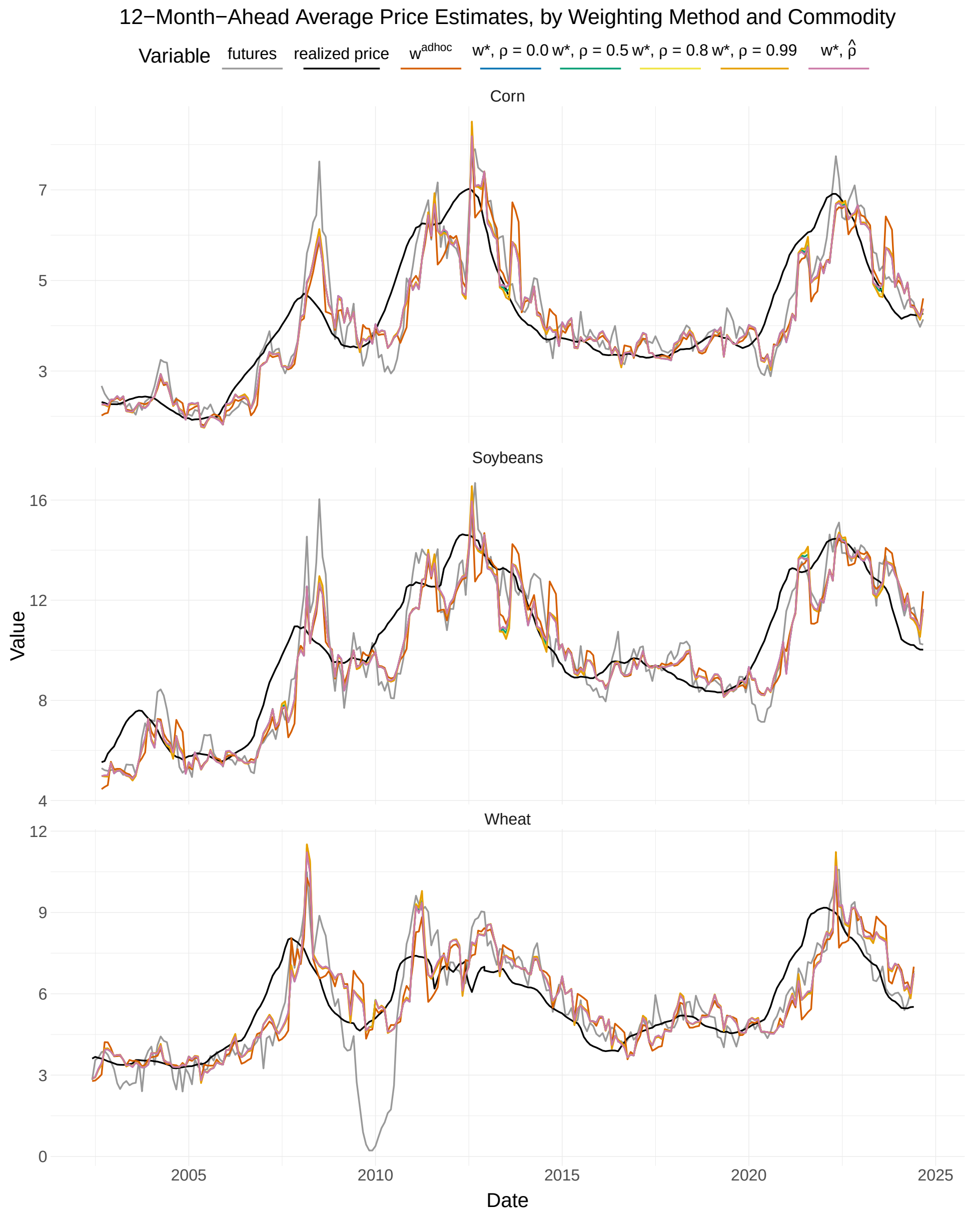

Figure 2: Marketing-year price forecasts constructed under alternative weighting schemes are plotted alongside the realized marketing-year prices. Forecasts include both the estimated $\hat{\rho}$ and four alternative ρ values (0, 0.5, 0.8, 0.99). Panels are faceted by commodity.

As shown in Figure 2, price forecasts generally follow realized prices for each commodity. Differences across AR(1) coefficients (ρ) are generally small. All optimal weighting schemes display similar responses over time. Forecasts based on futures data capture the overall trend but tend to overshoot during rapid price changes, while WASDE-informed weights often show slightly smaller jumps. Ad hoc forecasts are less volatile than futures but show larger fluctuations than optimal weighting schemes. Overall, the figure provides a visual summary of these patterns, highlighting price movements across forecasts before we conduct a more formal evaluation.

To assess the robustness of these gains, we examine forecast errors quantitatively. Rolling 36-month MAE plots (Figure 3) indicate that improvements from optimal weighting are generally persistent over time, with occasional spikes corresponding to extreme price movements.

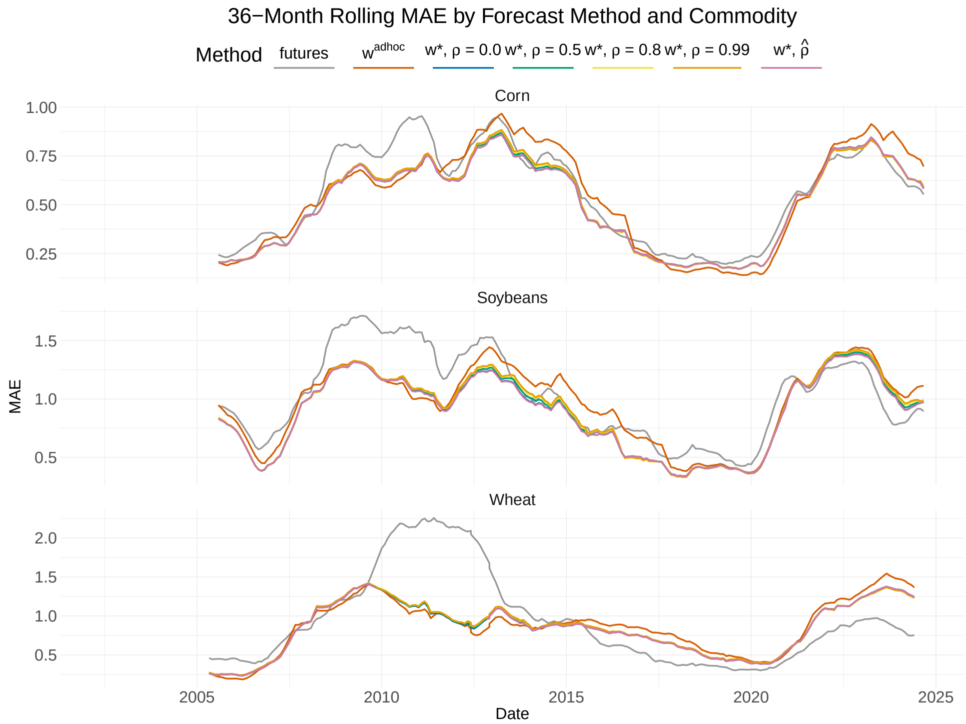

Figure 3: Thirty-six-month rolling mean absolute errors (MAE) for corn, soybeans, and wheat are shown for each forecast method. Each panel corresponds to a commodity, and colored lines represent alternative weighting schemes, including optimal weighting under estimated $\hat{\rho}$ and fixed ρ values, as well as ad hoc and futures-based forecasts.

As shown in Table 1, subsample MAE and RMSE by commodity across three periods (Before 2012, 2012–2019, After 2019) show that optimal weighting consistently reduces errors relative to both ad hoc and futures-based forecasts. The improvements are not driven by a single period or extreme spike, suggesting that gains from optimal weighting are stable and broadly applicable across different market conditions. Together, these results provide both visual and quantitative evidence that our weighting schemes improve predictive accuracy for corn, soybeans, and wheat.

Table 1: Subsample Forecast Accuracy by Commodity, Period, and Metric

commodity	period	Metric	Futures	Ad hoc	$\rho = 0.0$	$\rho = 0.5$	$\rho = 0.8$	$\rho = 0.99$	$\rho = \hat{\rho}$
Corn	**2012–2019**	**MAE**	0.41	0.40	0.37	0.38	0.38	0.38	0.37
Corn	**2012–2019**	**RMSE**	0.56	0.64	0.57	0.58	0.58	0.59	0.57
Corn	**After 2019**	**MAE**	0.67	0.77	0.69	0.69	0.69	0.69	0.69
Corn	**After 2019**	**RMSE**	0.77	0.95	0.86	0.87	0.87	0.87	0.86
Corn	**Before 2012**	**MAE**	0.56	0.50	0.47	0.47	0.48	0.48	0.47
Corn	**Before 2012**	**RMSE**	0.80	0.66	0.62	0.62	0.63	0.63	0.62
Soybeans	**2012–2019**	**MAE**	0.67	0.72	0.60	0.61	0.61	0.62	0.60
Soybeans	**2012–2019**	**RMSE**	0.90	0.98	0.84	0.86	0.88	0.89	0.84
Soybeans	**After 2019**	**MAE**	1.19	1.33	1.24	1.25	1.26	1.27	1.24
Soybeans	**After 2019**	**RMSE**	1.50	1.63	1.53	1.53	1.53	1.53	1.53
Soybeans	**Before 2012**	**MAE**	1.24	1.05	0.99	1.00	1.00	1.00	0.99
Soybeans	**Before 2012**	**RMSE**	1.61	1.38	1.28	1.29	1.30	1.30	1.28
Wheat	**2012–2019**	**MAE**	0.61	0.73	0.68	0.69	0.69	0.69	0.68
Wheat	**2012–2019**	**RMSE**	0.83	0.85	0.80	0.81	0.81	0.81	0.80
Wheat	**After 2019**	**MAE**	0.76	1.27	1.18	1.18	1.17	1.17	1.18
Wheat	**After 2019**	**RMSE**	0.94	1.49	1.33	1.34	1.34	1.34	1.33
Wheat	**Before 2012**	**MAE**	1.21	0.76	0.78	0.79	0.79	0.80	0.79
Wheat	**Before 2012**	**RMSE**	1.73	1.08	1.08	1.09	1.11	1.11	1.08

6.1 Forecast Evaluation and Inference

Building on observed price patterns, we evaluate the accuracy of forecasts generated using futures-based, ad hoc, and optimal weighting methods. Table 2 reports key metrics (MAE, RMSE, MAPE, RMSPE) for each method over the full sample period.

For corn, the estimated AR(1) coefficient is effectively zero. Using this coefficient in the optimal weighting scheme yields the lowest MAE, identical to that obtained with $\rho = 0.0$. Higher ρ values (0.5, 0.8, 0.99) produce slightly larger errors, but all optimal weighting forecasts remain more accurate than ad hoc (0.523) and futures-based (0.529) forecasts. The MAE improvement from ad hoc to the lowest-error optimal weighting translates to roughly \$0.04 per bushel, with a similar improvement over futures-based weighting.

For soybeans, the estimated AR(1) coefficient reflects weak negative autocorrelation. Optimal weighting using this coefficient reduces MAE and other error metrics relative to ad hoc (0.991) and futures-based (1.023) benchmarks, while higher ρ values produce slightly larger errors. The MAE reduction amounts to approximately \$0.09 per bushel relative to ad hoc weighting, and roughly \$0.12 relative to futures-based forecasts.

For wheat, the estimated AR(1) coefficient indicates statistically significant negative autocorrelation. Using this coefficient improves forecast accuracy relative to ad hoc (0.854) and futures-based (0.894) forecasts, but MAE and other error metrics are slightly higher than the $\rho = 0.0$ specification. This suggests that, while statistically significant, the negative autocorrelation has limited practical importance. Higher ρ values (0.5, 0.8, 0.99) further increase errors, reflecting over-responsiveness to recent deviations. The MAE improvement compared with ad hoc weighting is approximately \$0.03 per bushel, and nearly \$0.07 per bushel compared with futures-based weighting.

Across all commodities, incorporating the estimated AR(1) coefficients provides a balance of responsiveness and stability, with forecast improvements consistent across absolute, overall, and percentage error measures. Low or near-zero ρ values perform similarly when forecast errors show little persistence, while higher ρ values increase sensitivity but may slightly reduce precision. These results demonstrate that the optimal weighting approach remains robust across alternative AR(1) specifications, confirming the value of data-driven persistence estimates. While the absolute reductions in forecast errors translate to relatively small dollar differences per bushel, they are meaningful in commodity markets, where even modest improvements in price forecasts can influence hedging, marketing, and revenue decisions over a marketing year. Applied at the production scale, these reductions suggest that fixed-horizon price forecasts using WASDE data and optimal weighting provide more consistent and economically relevant gains than either ad hoc or futures-based approaches.

To formally assess whether these differences in forecast accuracy are statistically significant, we conduct DM tests. By accounting for temporal correlation in forecast errors, the DM test allows us to determine whether the gains from optimal weighting are statistically significant relative to ad hoc and futures-based approaches. This provides a clear framework for understanding how the estimated persistence affects forecast performance and highlights where optimal weighting delivers the most consistent improvements across

Table 2: Forecast Errors by Commodity

Method	MAE	RMSE	MAPE	RMSPE	N
Corn					
Ad hoc	0.523	0.723	11.925	15.414	264
Futures	0.529	0.716	12.390	16.249	264
Optimal ($\rho = 0.0$)	0.484	0.663	11.137	14.152	264
Optimal ($\rho = 0.5$)	0.486	0.667	11.188	14.238	264
Optimal ($\rho = 0.8$)	0.487	0.671	11.213	14.289	264
Optimal ($\rho = 0.99$)	0.488	0.672	11.220	14.294	264
Optimal ($\rho = \hat{\rho}$)	0.484	0.662	11.136	14.149	264
Soybeans					
Ad hoc	0.991	1.311	9.868	12.909	264
Futures	1.023	1.368	10.382	13.788	264
Optimal ($\rho = 0.0$)	0.904	1.207	8.988	11.800	264
Optimal ($\rho = 0.5$)	0.911	1.217	9.043	11.886	264
Optimal ($\rho = 0.8$)	0.918	1.226	9.093	11.953	264
Optimal ($\rho = 0.99$)	0.921	1.230	9.117	11.980	264
Optimal ($\rho = \hat{\rho}$)	0.903	1.205	8.975	11.784	264
Wheat					
Ad hoc	0.854	1.100	14.579	17.767	269
Futures	0.894	1.306	15.761	23.212	269
Optimal ($\rho = 0.0$)	0.824	1.044	13.933	16.608	269
Optimal ($\rho = 0.5$)	0.828	1.052	14.006	16.729	269
Optimal ($\rho = 0.8$)	0.831	1.060	14.063	16.826	269
Optimal ($\rho = 0.99$)	0.831	1.062	14.067	16.851	269
Optimal ($\rho = \hat{\rho}$)	0.825	1.044	13.950	16.610	268

commodities.

As shown in Figure 4, optimal weighting improves forecast accuracy for soybeans relative to both futures-based and ad hoc methods, with predominantly negative DM statistics and significant p-values. Ad hoc forecasts are frequently outperformed, whereas futures-based forecasts yield smaller yet consistent gains. For corn, optimal weighting improves forecast accuracy over ad hoc methods, as shown by negative and significant DM statistics. However, despite negative DM statistics, the tests do not yield significant improvement over the futures baseline. Forecast improvements for wheat follow similar patterns, but are more modest, reflecting the difficulty of predicting 12-month-ahead prices. However, low-persistence optimal weights consistently offer an advantage.

While the DM tests identify which methods perform best in terms of predictive ac-

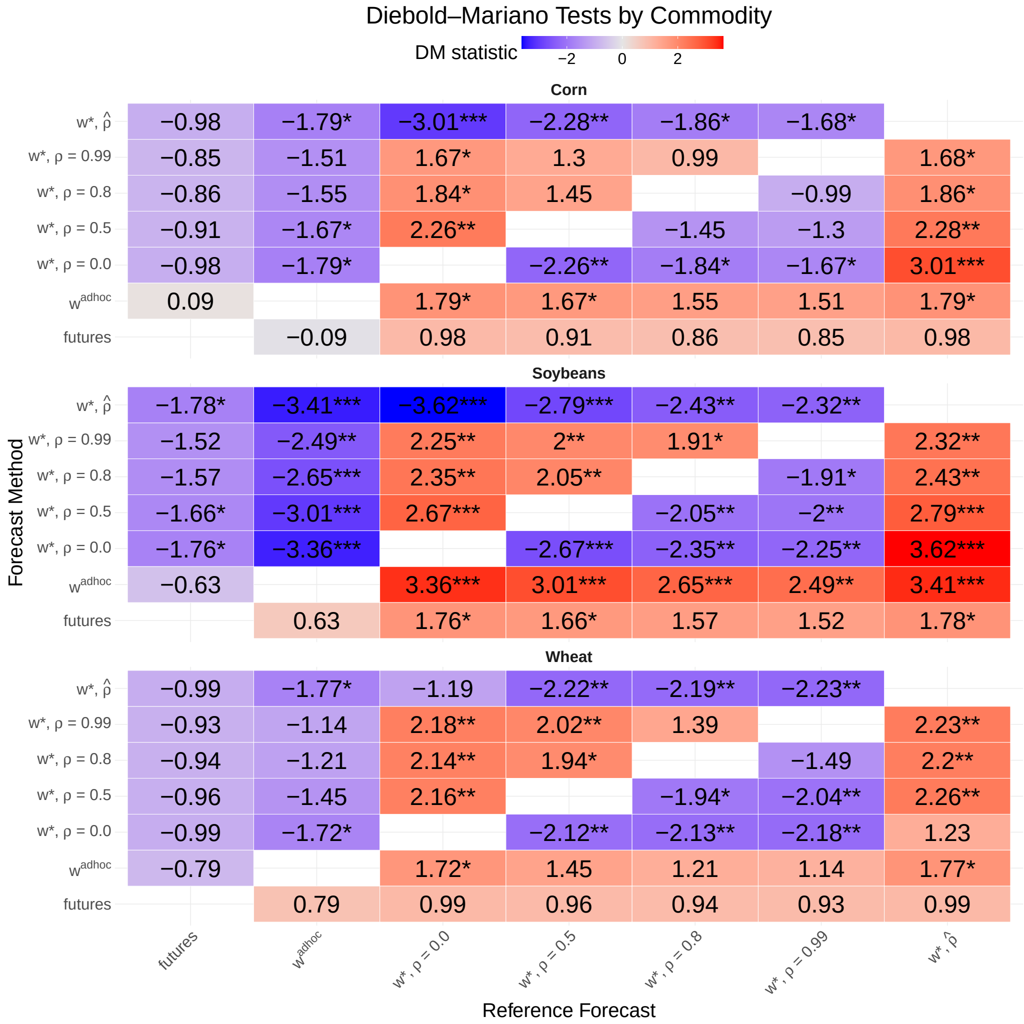

Figure 4: Heatmap of DM test statistics for all pairwise comparisons of forecasting methods under squared-error loss. Each cell shows the DM statistic for the row method relative to the column method. Negative values (blue) indicate the row method outperforms the reference (lower forecast errors), while positive values (red) indicate worse performance. Shading reflects the magnitude of the statistic, and stars indicate significance (*** p<0.01, ** p<0.05, * p<0.10). Panels are shown separately by commodity.

curacy, they do not assess whether forecasts are unbiased or efficient. To examine these properties, we implement MZ regressions with HAC standard errors to account for serial correlation, testing the joint null hypothesis for each method using the Wald test. These regressions provide a complementary evaluation of forecast quality, and the results are presented in Table 3.

For corn, most weighting schemes, including ad hoc and optimal weights across the range of ρ values, fail to reject the null hypothesis of $\alpha = 0$ and $\beta = 1$, as reflected by their large joint p-values. Intercepts (α) are small and slopes (β) are near one, indicating

Table 3: Mincer–Zarnowitz Forecast Efficiency Tests by Commodity

Method	alpha	beta	Joint p	Sig.
Corn				
Ad hoc	0.53	0.89	0.127	
Futures	0.68	0.83	0.062	*
Optimal ($\rho = 0.0$)	0.39	0.92	0.204	
Optimal ($\rho = 0.5$)	0.41	0.92	0.185	
Optimal ($\rho = 0.8$)	0.43	0.91	0.162	
Optimal ($\rho = 0.99$)	0.44	0.91	0.149	
Optimal ($\rho = \hat{\rho}$)	0.39	0.92	0.205	
Soybeans				
Ad hoc	1.77	0.86	0.011	**
Futures	2.20	0.80	0.006	***
Optimal ($\rho = 0.0$)	1.55	0.88	0.009	***
Optimal ($\rho = 0.5$)	1.59	0.88	0.008	***
Optimal ($\rho = 0.8$)	1.63	0.87	0.006	***
Optimal ($\rho = 0.99$)	1.65	0.87	0.006	***
Optimal ($\rho = \hat{\rho}$)	1.55	0.88	0.009	***
Wheat				
Ad hoc	1.59	0.71	<0.001	***
Futures	2.41	0.58	<0.001	***
Optimal ($\rho = 0.0$)	1.48	0.72	<0.001	***
Optimal ($\rho = 0.5$)	1.51	0.71	<0.001	***
Optimal ($\rho = 0.8$)	1.53	0.71	<0.001	***
Optimal ($\rho = 0.99$)	1.54	0.71	<0.001	***
Optimal ($\rho = \hat{\rho}$)	1.48	0.72	<0.001	***

Joint null hypothesis: $\alpha = 0$ and $\beta = 1$, indicating unbiased and efficient forecasts. Significance levels are indicated by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

forecasts are generally unbiased and efficient. Futures-based forecasts show slight inefficiency, with a larger α and a joint p-value below 0.10. Optimal weighting, including the estimated $\hat{\rho}$, reduces α and brings β closer to one, improving calibration. Across all tested ρ values, low-persistence weighting, particularly $\rho = 0.0$ and $\hat{\rho}$, delivers the most reliable forecasts, consistent with prior research showing that corn forecasts are typically well-calibrated but can benefit from weighting schemes that mitigate small systematic errors (Isengildina Massa, Irwin, and Good 2006; Xiao, Hart, and Lence 2017).

For soybeans and wheat, all forecasts exhibit positive α and $\beta < 1$, with joint p-values often significant at the 0.05 or 0.01 level, reflecting systematic underreaction in

WASDE forecasts. This pattern aligns with literature documenting slow incorporation of new information and persistent forecast biases, particularly for crops with more variable production (Coibion and Gorodnichenko 2015; Isengildina Massa, Irwin, and Good 2006; Xiao, Hart, and Lence 2017). Optimal weighting reduces the magnitude of α and increases β relative to ad hoc and futures forecasts, suggesting that combining information from multiple forecasts partially corrects the observed inefficiency. The estimated $\hat{\rho}$ performs comparably to the lowest-persistence weighting ($\rho = 0.0$), indicating that modest persistence is sufficient to improve calibration for these more volatile crops.

Taken together, these results demonstrate that optimal weighting meaningfully enhances predictive accuracy and efficiency. For corn, low-persistence weighting, including $\hat{\rho}$, largely eliminates inefficiency. For soybeans and wheat, the optimal weights framework provides measurable forecast improvements despite persistent underreaction. Expressing forecasts at a fixed horizon further aligns predictions with decision-relevant timeframes, allowing a more direct assessment of responsiveness to new information and seasonal revisions.

Building on the efficiency gains observed in MZ regressions, we conduct FE tests to assess whether optimal weighting also captures predictive information contained in alternative forecasts. Specifically, these tests examine whether forecast errors from competing methods provide additional explanatory power once the optimal forecast is included. Figure 5 presents pairwise FE results for corn, soybeans, and wheat. Each tile reports the outcome of testing whether the row method encompasses the column method using Newey-West adjusted regressions. Tile color reflects the direction and strength of encompassing, and stars denote standard statistical significance levels. Specifically, blue tiles indicate that the row method (Method 1) encompasses the column method (Method 2), and red tiles show the opposite. Values near white indicate weak or no dominance of either method.

The pairwise forecast encompassing tests show nuanced patterns across commodities, highlighting the benefits of optimal weighting while leaving room for forecasts to provide additional information. For corn, both futures and ad hoc forecasts are largely encom-

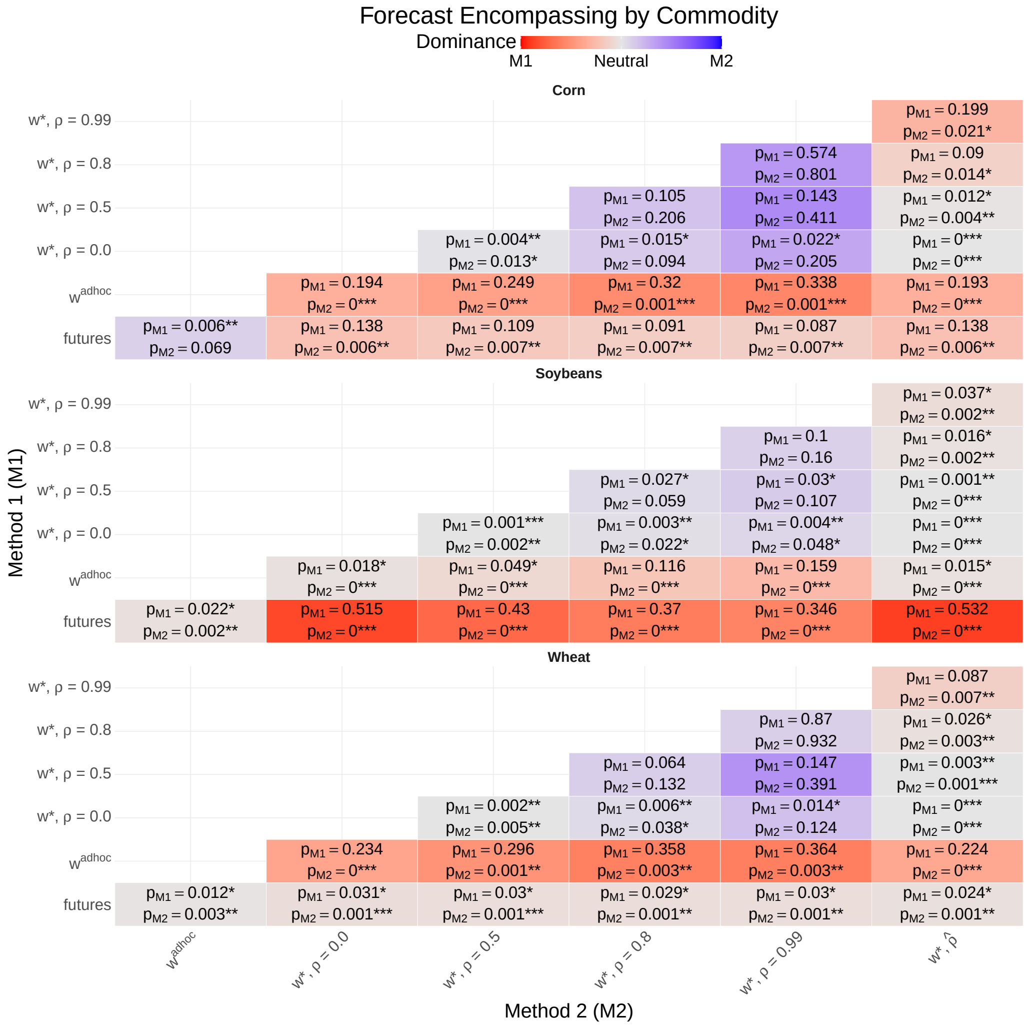

Figure 5: Pairwise forecast encompassing by commodity. Each cell compares Method 1 (row) to Method 2 (column). Colors indicate direction and strength: blue denotes Method 1 encompasses Method 2, red denotes Method 2 encompasses Method 1, and values near white indicate weak or no dominance. Cell labels report p-values for Method 1 (top) and Method 2 (bottom), and stars indicate significance (*** p<0.01, ** p<0.05, * p<0.10). Panels are shown separately by commodity.

passed by optimal forecasts, especially at low persistence ($\rho = 0.0$) and the estimated $\hat{\rho}$, though some comparisons indicate incremental contributions. In soybeans, optimal forecasts frequently encompass futures prices, but ad hoc comparisons are less consistent, often indicating that both forecasts provide useful predictive information. Wheat shows more cases where forecasts contribute complementary information, with futures and optimal forecasts both providing predictive information, while ad hoc forecasts are generally encompassed. Overall, optimal weighting schemes capture additional predictive information across commodities, often encompassing ad hoc and futures forecasts while

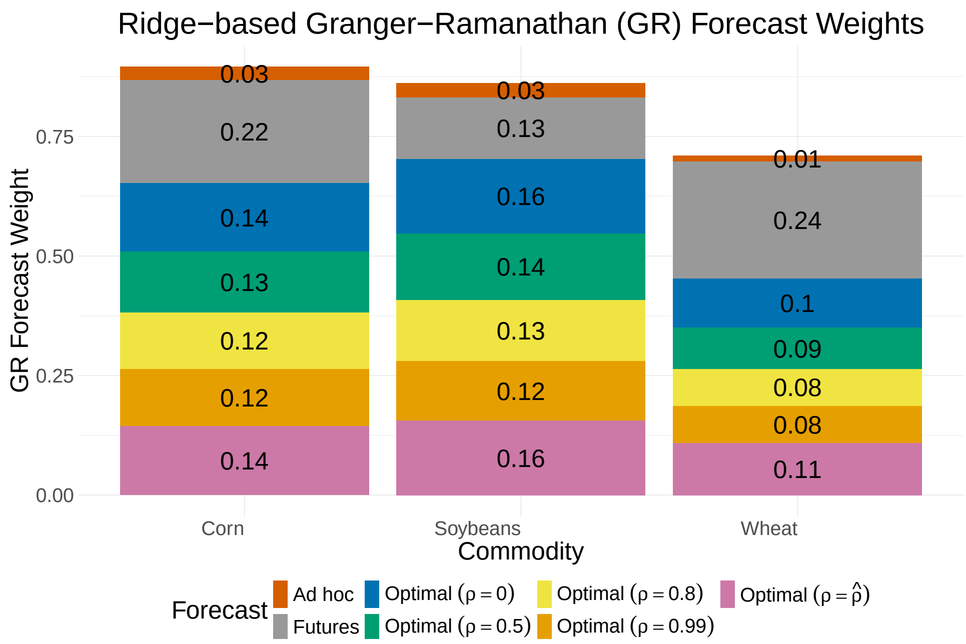

Figure 6: Ridge-regularized Granger-Ramanathan (GR) forecast weights for each commodity. Weights are estimated using ridge regression to combine individual forecasts of 12-month-ahead prices, including futures, ad hoc forecasts, and model-based optimal forecasts at multiple persistence levels ($\rho = 0.0, 0.5, 0.8, 0.99, \hat{\rho}$). The ridge penalty stabilizes the weights when forecasts are highly correlated, ensuring a balanced contribution of each method to the combined forecast.

also providing incremental insights where encompassing is partial.

GR combined forecasts using ridge regression (Granger and Ramanathan 1984) confirm that integrating futures and optimal weighting improves performance, as shown in Figure 6. Corn and wheat assign the largest weights to futures, with additional contributions from multiple optimal weighting schemes, while soybeans draw most of their predictive power from optimally aggregated WASDE-based forecasts. Across all commodities, weights are spread primarily across low- and estimated-persistence ($\rho = 0.0$ and $\hat{\rho}$) schemes, and ad hoc forecasts contribute minimally, demonstrating that combining forecasts in a structured way captures complementary information and provides a more complete and informative picture of expected price movements.

Overall, these results indicate that optimal weighting contains incremental predictive information beyond either futures prices or simple ad hoc combinations. Across all commodities, low-persistence schemes consistently provide the most robust encompassing, reinforcing that modest persistence in the weighting scheme is sufficient to improve both predictive accuracy and informational efficiency in 12-month-ahead forecasts. GR combined forecasts confirm these findings, showing that corn and wheat assign substantial weight to futures alongside multiple optimal weighting schemes, while soybeans rely primarily on optimally aggregated WASDE-based forecasts.

6.2 Robustness of the Results

To assess the stability of the accuracy gains and to test whether the fixed-horizon transformation affects the seasonal structure of forecast errors, we conduct two complementary analyses. The first uses MBB inference to confirm that accuracy improvements are not driven by isolated periods or temporal dependence patterns. The second directly tests whether expressing WASDE forecasts at a fixed horizon eliminates the seasonal patterns in forecast errors documented in prior literature. This is a question with implications beyond robustness, addressing whether observed seasonality in WASDE errors reflects the information environment or the design of the forecasting horizon itself.

As our first robustness check, we implement MBB tests. This bootstrapping approach accounts for potential serial correlation in the data by resampling blocks of consecutive observations, allowing us to evaluate whether the observed improvements from optimal weighting remain statistically significant when temporal dependence is considered. Specifically, we resample forecast errors in consecutive blocks of $L = 12$, with replacement, $B = 1000$ times. For each replicate, we recompute forecast metrics and differences relative to the futures-based benchmark. By comparing bootstrapped distributions of forecast errors across weighting schemes, we can determine the robustness of our earlier findings and ensure that gains are not driven by isolated periods or patterns in the data. The results of our MBB tests are depicted in Figure 7. In the figure, median and 95% confidence intervals (pointranges) are overlaid on violin plots of the bootstrapped MAE differences, visualizing both the distribution and stability of the results.

For corn, the optimal-weight forecasts consistently show slightly lower MAE than futures, with median improvements of roughly 0.04–0.05. The ad hoc weighting scheme performs nearly equivalently to futures, with a median difference close to zero and confi-

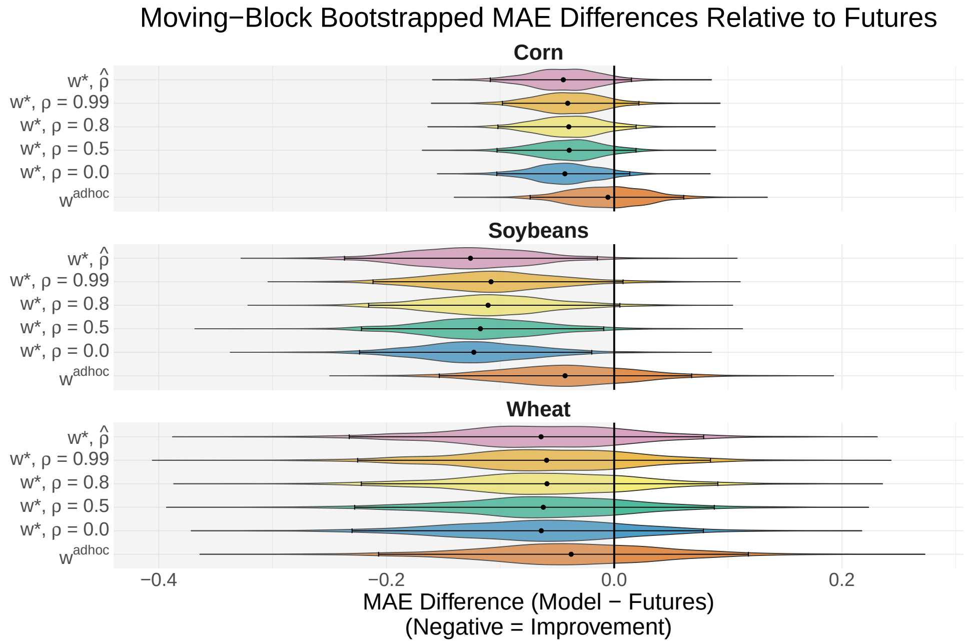

Figure 7: Moving-block bootstrapped MAE differences relative to the futures benchmark for various weighting schemes, faceted by commodity. Violin shapes show the full distribution across 1,000 bootstrap replicates, while points and horizontal bars indicate the median and 95% confidence intervals. Negative values indicate that the weighting method outperforms the futures benchmark. Consecutive monthly blocks are resampled to account for temporal dependence in the forecast errors.

dence intervals spanning both positive and negative values. For soybeans, the gains are more pronounced. Optimal-weight forecasts reduce MAE by approximately 0.11–0.12 on average, with lower ρ values yielding the largest improvements. In these cases, the bootstrap confidence intervals remain entirely below zero, indicating stable improvements in forecast accuracy relative to futures. For wheat, the results are more mixed. Median MAE differences are modestly negative for both the ad hoc and optimal-weight forecasts, suggesting some improvement relative to futures. However, the bootstrap confidence intervals are wide and span both negative and positive values, indicating greater variability and less consistent performance across resampled blocks. Overall, the MBB results suggest that applying optimal weights to WASDE-based forecasts improves predictive accuracy relative to the futures benchmark, especially for soybeans. The violin plots illustrate both the magnitude and distribution of these improvements, highlighting stronger and more stable gains for soybeans, moderate improvements for corn, and more variable outcomes for wheat.

To evaluate whether forecast performance varies systematically over the calendar year, we formally test for seasonality in monthly forecast errors. We begin with ANOVA as a descriptive benchmark, but because fixed-horizon errors are overlapping and serially correlated, our primary inference relies on a HAC-robust joint Wald test of calendar-month indicators. This approach allows us to assess whether any seasonal variation remains after accounting for time-series dependence and whether the gains from optimal weighting are consistent throughout the year.

Table 4: ANOVA and HAC-robust tests for seasonality in monthly forecast errors across commodities and weighting schemes

Commodity	Weights	ANOVA p	Sig. (ANOVA)	HAC p	Sig. (HAC)
Corn	Ad hoc	0.498		0.087	*
Corn	Futures	0.392		0.390	
Corn	Optimal ($\rho = 0.0$)	0.941		0.370	
Corn	Optimal ($\rho = 0.5$)	0.951		0.259	
Corn	Optimal ($\rho = 0.8$)	0.964		0.103	
Corn	Optimal ($\rho = 0.99$)	0.974		0.054	*
Corn	Optimal ($\rho = \hat{\rho}$)	0.941		0.372	
Soybeans	Ad hoc	0.199		0.000	***
Soybeans	Futures	0.361		0.014	**
Soybeans	Optimal ($\rho = 0.0$)	1.000		0.405	
Soybeans	Optimal ($\rho = 0.5$)	1.000		0.536	
Soybeans	Optimal ($\rho = 0.8$)	1.000		0.338	
Soybeans	Optimal ($\rho = 0.99$)	0.999		0.221	
Soybeans	Optimal ($\rho = \hat{\rho}$)	1.000		0.372	
Wheat	Ad hoc	0.268		0.258	
Wheat	Futures	1.000		0.999	
Wheat	Optimal ($\rho = 0.0$)	0.862		0.095	*
Wheat	Optimal ($\rho = 0.5$)	0.839		0.103	
Wheat	Optimal ($\rho = 0.8$)	0.823		0.124	
Wheat	Optimal ($\rho = 0.99$)	0.820		0.145	
Wheat	Optimal ($\rho = \hat{\rho}$)	0.854		0.067	*

As shown in Table 4, ANOVA results show no evidence of seasonal differences in mean absolute errors for any commodity, with consistently large p-values across weighting schemes, including the estimated optimal specification ($\rho = \hat{\rho}$). However, once serial correlation is accounted for using the HAC-robust joint test, a different pattern emerges.

For corn, there is modest evidence of seasonality under the ad hoc weights ($p = 0.087$)

and borderline significance under the highly persistent optimal specification ($\rho = 0.99$, p = 0.054), while seasonality remains insignificant for futures and for other optimal weighting schemes. For soybeans, the HAC test indicates strong seasonality in the ad hoc (p < 0.001) and futures (p = 0.014) specifications, but no evidence of seasonality under any optimal weighting scheme. Wheat shows little evidence of seasonal variation in forecast errors overall. The HAC results are insignificant for the ad hoc and futures specifications, with only weak, borderline evidence under $\rho = 0$ (p = 0.095) and $\rho = \hat{\rho}$ (p = 0.067).

Taken together, the HAC results suggest that ignoring serial dependence can obscure economically meaningful seasonal patterns. For soybeans, statistically significant seasonal patterns under both ad hoc and futures-based specifications are entirely eliminated under optimal weighting across all persistence levels, including the estimated persistence parameter. For corn, seasonality under ad hoc weighting is reduced to borderline significance, while the estimated persistence specification shows no evidence of seasonal variation. This finding supports the interpretation that a substantial share of the seasonal dynamics documented in prior WASDE literature is an artifact of the fixed-event horizon design rather than a feature of how agricultural information arrives

Together, these robustness checks provide additional support for the improvements observed when applying optimal weights to fixed-horizon WASDE forecasts. The MBB results confirm that the gains remain stable when accounting for temporal dependence, indicating that the improvements are not driven by isolated periods or short-run patterns in the data. The seasonality tests indicate that forecast performance does not exhibit strong systematic variation across calendar months under the optimal weighting framework, and this holds across the full range of ρ values. Although some seasonal patterns appear in the ad hoc and futures specifications, particularly for soybeans, these patterns largely disappear when optimal weights are applied.

Overall, the results indicate that applying optimal weights to WASDE forecasts enhances 12-month-ahead price predictions for major agricultural commodities. Forecast accuracy improves consistently across multiple metrics (MAE, RMSE, MAPE, RMSPE), with the largest gains observed for soybeans and moderate improvements for corn. For wheat, gains are more variable, reflecting the inherent difficulty of forecasting this commodity due to its heightened sensitivity to unpredictable weather, global supply shocks, and policy changes. Beyond numerical improvements, optimal weighting produces forecasts that are broadly unbiased, more efficient, and contain incremental information relative to ad hoc or futures-based approaches. Economically, the reductions in forecast errors translate to meaningful differences at the bushel level, with potential implications for hedging, marketing, and revenue planning. Robustness checks demonstrate that these benefits persist across different temporal windows, seasonal patterns, and assumptions about price persistence. Taken together, the evidence suggests that incorporating optimal weighting into WASDE-based forecasting frameworks delivers more accurate, reliable, and informative price predictions for agricultural markets.

6.3 Applications of Fixed Horizon Forecasts

A distinguishing feature of the fixed-horizon forecasts constructed in this study is their generality with respect to the planning calendar. Because the forecast is defined as a rolling 12-month average originating from any calendar month, it is not anchored to the agricultural marketing year and can be applied to decision contexts organized around the calendar year, the fiscal year, or any other planning interval. This flexibility is a direct consequence of the fixed-horizon design and represents a practical advantage over both standard WASDE releases, which are tied to marketing-year endpoints, and futures-based forecasts, which must be mapped to a specific contract sequence that may not align with the user's planning horizon. The applications described below show how this property, combined with documented accuracy gains, fills a concrete gap in the price information available to producers, program administrators, and agribusinesses.

The fixed horizon forecasts directly support forward-pricing decisions, hedge ratio adjustments, and storage choice at the farm level, which arise continuously throughout the year and are not adequately served by a forecast whose effective horizon contracts as the marketing year progresses. Because the fixed-horizon forecast is reconstructed each month using the most recent WASDE release, producers have access to a consistent 12-month price expectation regardless of when their pricing decision arises. The accuracy gains documented in this study translate directly to the farm level: for a representative corn operation producing 100,000 bushels annually, the RMSE advantage of optimal weighting over the futures benchmark 5.4 cents/bu implies a reduction in expected squared price error of approximately $5,400 per year, with larger gains for soybeans (16.3 cents per bushel, roughly $8,150 annually for a 50,000-bushel operation) and modest but meaningful improvements for wheat.

Beyond forward pricing, crop revenue insurance elections present a related but distinct application. Federal crop revenue insurance elections require producers to commit to a coverage tier before the policy year begins, relying on forward price expectations to assess the relative value of higher versus lower coverage levels. The rolling fixed-horizon forecast provides a more accurate 12-month price signal at the time of this election than the futures alone, supporting both the choice of coverage tier and the selection between Revenue Protection and Revenue Protection with Harvest Price Exclusion.

The availability of rolling 12-month forecasts is informative for Farm Bill program enrollments and administration. The choice between Agriculture Risk Coverage County (ARC-CO) and Price Loss Coverage (PLC) hinges on where the season-average price is expected to settle relative to statutory reference prices and benchmark revenue levels, making the rolling fixed-horizon forecast a natural input to annual program enrollment decisions. A producer who elected PLC based on prevailing futures prices can monitor each subsequent month whether the optimal-weight forecast continues to support that election or whether revised fundamentals have altered the relative attractiveness of the two programs. For program administrators and budget analysts working on a federal fiscal-year cycle, the calendar-year flexibility of the rolling forecast is similarly valuable, permitting consistent forward projections of program expenditure obligations without requiring adjustment for the offset between marketing-year and fiscal-year endpoints.

Beyond farm-level decisions, grain elevators, ethanol producers, flour millers, and other agricultural processors require a price outlook defined at a consistent forward interval for capital budgeting, procurement, and hedging program design. For agribusinesses reporting on a calendar-year basis, standard WASDE releases introduce a structural misalignment between the forecast object and the firm's planning cycle that the rolling fixed-horizon forecast resolves directly. The forecast is recomputed monthly at a constant horizon, delivering an updated, horizon-stable price signal each month as new WASDE information becomes available.

7 Conclusions

WASDE average farm price forecasts play a central role in agricultural markets, shaping expectations, guiding decisions, and informing research and policy. Our analysis shows that these forecasts can be made even more useful by translating them into a constant 12-month horizon and combining them using optimal weights. This approach allows consistent comparisons across months, clearer interpretation of predictive performance, and alignment with the timing of real economic decisions. Applying optimal weights reduces forecast errors for corn, soybeans, and wheat, improves bias and efficiency, and adds incremental informational content, with these improvements holding broadly across seasons, time periods, and forecast assumptions. Gains are largest for soybeans and corn, while wheat benefits more modestly, likely due to its greater sensitivity to weather, global supply shocks, and policy changes. Beyond accuracy gains, the seasonal analysis reveals that the statistically significant forecast error seasonality documented in prior WASDE literature for soybeans disappears entirely under optimal weighting, pointing to the fixed-event horizon design, not the agricultural information environment, as its source. Even small per-bushel improvements can translate into meaningful differences when applied at the national level.

Despite these gains, some limitations remain. Translating forecasts into a fixed horizon relies on assumptions about price persistence that may not always hold, and extreme market shocks or policy changes can still generate forecast errors. Systematic under-reaction in WASDE forecasts persists for more volatile crops like soybeans and wheat, meaning that optimal weighting improves but does not fully correct inherent inefficiencies. For months when overlapping WASDE forecasts are unavailable, we approximate 12-month-ahead prices using the closing value of the nearest futures contract, which provides a practical solution but may not fully capture short-term price dynamics. Wheat, in particular, remains the most challenging commodity to predict.

These findings have broad practical relevance. By providing a consistent, forward-looking price signal, fixed-horizon forecasts can support a wide range of decisions for producers, agribusinesses, and program administrators. They improve the information available for forward pricing, risk management, and revenue planning, while also offering a horizon-aligned input for policy programs and financial projections. The flexibility of a rolling 12-month forecast means it can be applied to different planning cycles, including calendar year, fiscal year, or marketing-year considerations, making it a useful tool across diverse decision contexts. While the improvements are modest for individual bushels, they accumulate to meaningful economic differences when scaled to production or program levels.

Looking ahead, there are clear avenues for further work. Incorporating additional market information or applying the framework to other commodities could broaden its relevance. Testing the forecasts under different regional or extreme market conditions could assess their generalizability. More broadly, integrating these methods into operational decision-support tools, such as farm-level revenue dashboards, procurement planning platforms, and program budget simulators, could help translate improved forecast design into actionable insights, enhancing the practical value of WASDE reports for researchers, policymakers, and market participants alike.

References

Bates, J. M. and C. W. J. Granger (1969). "The Combination of Forecasts". In: *Journal of the Operational Research Society* 20.4, pp. 451–468. DOI: https://doi.org/10.2307/3008764.

Clements, Michael P. and David I. Harvey (2010). "Forecast encompassing tests and probability forecasts". In: *Journal of Applied Econometrics* 25.6, pp. 1028–1062. DOI: 10.1002/jae.1097.

Coibion, Olivier and Yuriy Gorodnichenko (2015). "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts". In: *American Economic Review* 105.8, pp. 2644–2678. DOI: 10.1257/aer.20110306.

Diebold, Francis X. and Roberto S. Mariano (1995). "Comparing Predictive Accuracy". In: *Journal of Business & Economic Statistics* 13.3, pp. 253–263. ISSN: 07350015. URL: http://www.jstor.org/stable/1392185 (visited on 11/21/2025).

Dovern, Jonas, Ulrich Fritsche, and Jiri Slacalek (Nov. 2012). "Disagreement among Forecasters in G7 Countries". In: *The Review of Economics and Statistics* 94.4, pp. 1081–1096. ISSN: 0034-6535. DOI: 10.1162/REST_a_00207.

Etienne, Xiaoli L. et al. (2023). "Forecasting the U.S. season-average farm price of corn: Derivation of an alternative futures-based forecasting model". In: *Journal of Commodity Markets* 30, p. 100333. ISSN: 2405-8513. DOI: https://doi.org/10.1016/j.jcomm.2023.100333. URL: https://www.sciencedirect.com/science/article/pii/S2405851323000235.

Ganics, Gergely, Barbara Rossi, and Tatevik Sekhposyan (2024). "From Fixed-Event to Fixed-Horizon Density Forecasts: Obtaining Measures of Multihorizon Uncertainty from Survey Density Forecasts". In: *Journal of Money, Credit and Banking*. DOI: https://doi.org/10.1111/jmcb.13105.

Granger, C. W. J. and R. Ramanathan (1984). "Improved Methods of Combining Forecasts". In: *Journal of Forecasting* 3.2, pp. 197–204. DOI: 10.1002/for.3980030207.

Harvey, David, Stephen Leybourne, and Paul Newbold (1997). "Testing the equality of prediction mean squared errors". In: *International Journal of Forecasting* 13.2, pp. 281–291.

Hoffman, Linwood A. et al. (2015). "Forecast performance of WASDE price projections for U.S. corn". In: *Agricultural Economics* 46.S1, pp. 157–171. DOI: 10.1111/agec.12204.

Isengildina Massa, Olga, Scott H. Irwin, and Darrel L. Good (2006). "Are revisions to USDA crop production forecasts smoothed?" In: *American Journal of Agricultural Economics* 88.4, pp. 1091–1104. DOI: 10.1111/j.1467-8276.2006.00918.x.

Isengildina Massa, Olga, Berna Karali, and Scott H. Irwin (2013). "When do the USDA forecasters make mistakes?" In: *Applied Economics* 45.36, pp. 5086–5103. DOI: 10.1080/00036846.2013.818213. URL: https://doi.org/10.1080/00036846.2013.818213.

— (2024a). "Are USDA Forecasts Optimal? A Systematic Review". In: *Journal of Agricultural and Applied Economics* 56.3, pp. 496–527.

— (2024b). "What do we know about the value and market impact of the US Department of Agriculture reports?" In: *Applied Economic Perspectives and Policy* 46.2, pp. 698–736. DOI: 10.1002/aepp.13409.

Isengildina-Massa, Olga et al. (2008). "Impact of WASDE Reports on Implied Volatility in Corn and Soybean Markets". In: *Agribusiness* 24.4, pp. 473–490. DOI: 10.1002/agr.20169.

Knüppel, Malte and Andreea L. Vladu (2025). "Approximating Fixed-Horizon Forecasts Using Fixed-Event Forecasts". In: *Journal of Applied Econometrics* 40.4, pp. 359–379. DOI: https://doi.org/10.1002/jae.3114.

Mankiw, N. Gregory and Ricardo Reis (2002). "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve". In: *Quarterly Journal of Economics* 117.4, pp. 1295–1328. DOI: 10.1162/003355302320935034.

Mincer, Jacob and Victor Zarnowitz (1969). "The Evaluation of Economic Forecasts". In: *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Per-*

formance. Ed. by Jacob Mincer. New York: National Bureau of Economic Research, pp. 1–46.

Newey, Whitney K. and Kenneth D. West (1987). "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". In: *Econometrica* 55.3, pp. 703–708. DOI: 10.2307/1913610.

— (1994). "Automatic Lag Selection in Covariance Matrix Estimation". In: *Review of Economic Studies* 61.4, pp. 631–653. DOI: 10.2307/2297912.

Nordhaus, William D. (1987). "Forecasting Efficiency: Concepts and Applications". In: *Review of Economics and Statistics* 69.4, pp. 667–674.

Stock, James H. and Mark W. Watson (2004). "Combination Forecasts of Output Growth in a Seven-Country Data Set". In: *Journal of Forecasting* 23.6, pp. 405–430. DOI: 10.1002/for.928.

U.S. Department of Agriculture, Economic Research Service (2025). *Season-Average Price Forecasts*. https://www.ers.usda.gov/data-products/season-average-price-forecasts. Season-average price forecast data downloaded; accessed 2026-02-09.

United States Department of Agriculture (2026). *Historical WASDE Report Data*. https://www.usda.gov/historical-wasde-report-data-3. Accessed: 2026-02-09.

United States Department of Agriculture, Economic Research Service (2025). *Season-Average Price Forecasts: Documentation*. https://www.ers.usda.gov/data-products/season-average-price-forecasts/documentation. Accessed: 2026-01-27.

United States Department of Agriculture, National Agricultural Statistics Service (2024). *QuickStats Database*. https://quickstats.nass.usda.gov/. Accessed: 2026-02-09.

Wang, Xiaoqian et al. (2023). "Forecast combinations: An over 50-year review". In: *International Journal of Forecasting* 39.4, pp. 1518–1547. ISSN: 0169-2070. DOI: https://doi.org/10.1016/j.ijforecast.2022.11.005.

Xiao, Jinzhi, Chad E. Hart, and Sergio H. Lence (2017). "USDA Forecasts of Crop Ending Stocks: How Well Have they Performed?" In: *Applied Economic Perspectives and Policy*.