Measuring Physical Risk and the Price of Protection: Evidence from Agricultural Derivatives Markets

Michael K. Adjemian[†] Michel A. Robe[‡] Valentina G. Bruno[§]

August 2026

Abstract

Option-implied volatility is too often treated as a direct measure of agricultural price risk. It is not: it combines the variance that market participants expect with the price they pay to insure against it, the variance risk premium (VRP). We develop a forecast-free decomposition that separates the two using matched option-implied and subsequent realized variance. The approach attributes the effect of any observed market state or identified news shock on implied variance to expected physical risk *vs.* the cost of transferring risk. Consequently, option-implied volatility can remain stable even as expected crop-price risk and the cost of insurance move in opposite directions. Using daily fixed-tenor, at-the-money corn and soybean options from 2014–2026, we find a positive average protection premium for corn; the soybean premium is indistinguishable from zero, a difference we attribute to market structure. Commodity scarcity and approaching USDA crop reports raise expected physical variance, and for soybeans so do NOAA-CPC hot-forecast innovations during the critical crop window. Financial-market uncertainty and commercial hedging pressure raise the corn protection premium, and poor agricultural-supply news raises it immediately. Neither component trended upward over our sample, so we find no evidence that climate change has yet raised either the physical risk crop traders expect or the price of bearing it. These findings are consistent with a stylized partial-equilibrium model that

[*] Adjemian is an academic consultant with the Commodity Futures Trading Commission (CFTC). The analyses and conclusions expressed in this paper are those of the authors and do not necessarily reflect the views of other Commission staff, the Commission's Economic Research Section, or the Commission.

[†] Department of Agricultural and Applied Economics, University of Georgia, Athens, GA. Email: michael.adjemian@uga.edu. Corresponding author.

[‡] Robins School of Business, University of Richmond, Richmond, VA.

[§] Kogod School of Business, American University, Washington, DC, and the National Bureau of Economic Research.

we introduce. To illustrate the first-order importance of our results, we observe that the USDA's crop-insurance formula uses option-market volatility to measure expected corn price risk—but it does not remove the VRP when assessing crop insurance rates. After accounting for that extra protection charge, we find that the average corn expected volatility measure is about 7% lower—a difference worth more than $5.8 billion in additional corn revenue-insurance premiums over crop years 2015–2024, according to our back-of-the envelope calculation, including $3.7 billion in taxpayer-paid federal premium subsidies and $2.1 billion in producer-paid premium.

Keywords: implied volatility; price risk measurement; variance risk premium; commodity options; crop insurance; RMA. **JEL:** G13, G22, Q11, Q14.

1 Introduction

Quantifying agricultural price uncertainty and its drivers is critical to managing risk and, ultimately, to the stability of the food supply. For major U.S. row crops, derivatives markets provide high-frequency, forward-looking measures that are widely used by researchers and market participants to summarize this risk.

Option-implied variance is often treated as a variance forecast for the underlying commodity, and researchers have produced a body of literature comparing its accuracy to many candidate competitors, including the realized variance of the immediate past (e.g., Egelkraut et al. (2007)).However, it is not a one-dimensional measure of expected crop-price variability alone. Instead, option-implied variance combines the expected *physical variance* (the variance that traders expect the underlying futures price to realize) with the *variance risk premium* (VRP, the compensation that sellers of protection require for bearing that variance risk). Conflating the option-implied measure with physical uncertainty can therefore confound a change in the anticipated distribution of crop prices with a change in the market cost of transferring crop-price risk.[1]

Decomposing the implied variance between uncertainty and its price permits a sharper interpretation of movements in option-implied volatility. For example, weather, inventories, and scheduled crop information could alter the expected level and volatility of future crop prices, while hedging demand and financial risk-bearing conditions may alter the cost of insuring a given distribution. Separating the two channels makes it possible to compare risk across crops and over time without conflating changes in agricultural fundamentals with changes in the organization and capacity of risk-transfer markets.

[1] Technically, this statement applies to model-free implied variance Neuberger (1994); Britten-Jones and Neuberger (2000); Demeterfi et al. (1999). Our empirical measure departs from it by a wedge that depends on the shape of the volatility smile. We derive and explain the implications of that empirical choice.

This distinction is important for understanding how key factors affect commodity markets. For example, the United States Department of Agriculture's (USDA) Risk Management Agency (RMA) relies on option-implied variances to price federal crop insurance worth over $200 billion in annual risk protection (USDA 2025). Measuring the extent to which option prices also reflect factors other than physical risk is therefore possibly worth billions of dollars to producers or to U.S. taxpayers. More broadly, a large body of literature is focused on measuring the effects of climate change on agricultural production; another is concerned with the financialization of commodities—including the massive growth in derivatives trading by a non-traditional class of index traders who have no commercial interest in the underlying commodity itself. Both developments have the potential to affect derivatives pricing. Without decomposing the option-implied variance, however, it may be that stable agricultural option prices over time mask the rising impact of climate effects insofar as new financial traders add sufficient liquidity to markets for the cost of protection from physical crop risk to decrease. In that case, a rising physical component and a falling premium would offset one another.

This paper develops a measurement framework that separates these two objects without the need to first estimate a model of expected variance. We exploit the fact that, under rational expectations, subsequent realized variance in a market equals expected physical variance plus a forecast error that is orthogonal to the information available when the variance forecast is made. For any factor known at that date, estimating its relationship with the subsequent realized variance therefore recovers its relationship with expected physical variance. And estimating its relationship with the difference between the implied and subsequent realized variance recovers that factor's relationship with the VRP. Given the adding-up identity, those two effects sum exactly to the coefficient produced by regressing the option-implied variance on the same factor.[2] This logic also applies to the estimation of dynamic effects. Neither exercise requires a forecast of expected physical variance, which would otherwise enter as a generated regressor and pass on its own estimation error into the coefficients (Pagan, 1984).[3]

Applied empirically to daily corn and soybean options in 2014–2026, the framework shows that identical-looking changes in option-implied volatility arise through different channels.

[2] Our research design builds on the conditional-expectation decompositions of Fama (1984) and Fama and French (1987) and the variance-swap approach of Carr and Wu (2009). Fama and Fama–French place the forward discount or commodity basis—the object being decomposed—on the right-hand side, while Carr–Wu average the matched variance payoff and regress the realized variance on the implied variance. We instead place external economic states and identified innovations on the right-hand side, and estimate how their relationship with implied variance is split between expected physical variance and the VRP.

[3] We apply a forecast only to produce the historical decomposition, and report it separately. Because the VRP path is constructed as implied variance less the fitted physical path, forecast error shifts level between the two components one-for-one rather than attenuating both. The regression estimates are free of it; the decomposition is not, and we acknowledge that reality.

Figure 1 previews our headline results. Splitting each row by component rather than by crop makes two things visible at once. When a factor has a significant relationship with total implied variance—in the left panel—the sign is what economics would lead one to expect. And when it does not, the reason is usually that expected physical variance and the price of protection are moving in *opposite* directions rather than that nothing is happening. Anticipated crop information and market tightness operate principally through expected physical variance, while financial uncertainty and hedging pressure are more closely associated with the price of protection.

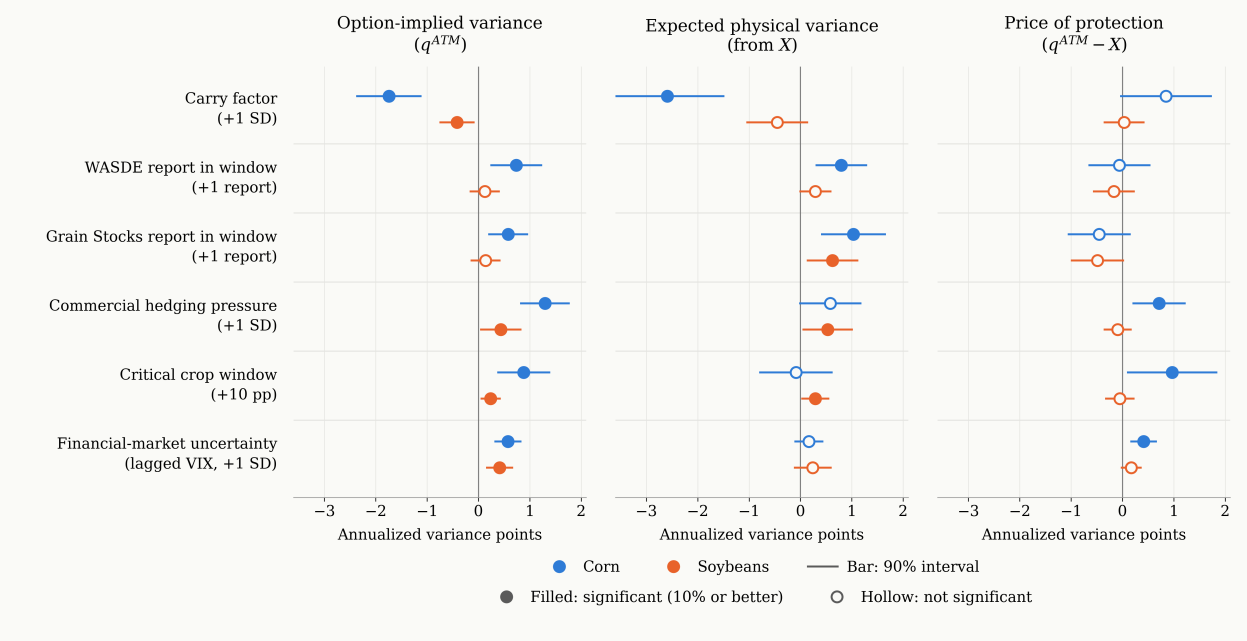


Figure 1: Decomposing option-implied variance, by crop

Notes: Coefficients estimated from the 60-day tenor specification, expressed in annualized variance points. Each row reports the same regressor estimated on option-implied variance, subsequent realized variance, and their difference, using identical observations and controls, so the left panel values equal the sum of the middle and right panels. Filled markers are significant at ten percent level or better; bars are 90 percent intervals from 1000 paired 65-session calendar-block wild-bootstrap draws. Table 3 reports the full specification.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities, Databento CBOT futures contracts, CFTC Commitments of Traders, USDA release calendars, and Cboe VIX.

Formally, in order to understand what drives option-implied variance changes in daily corn and soybean markets, we apply level and dynamic regressions to decompose factor effects into the two channels, using fixed-tenor, at-the-money call and put implied volatilities from 2014 through 2026.[4] We square the two sides and average them to construct two-sided

[4] Tenor is the number of days remaining until the option expires; a fixed-tenor series holds that horizon constant as contracts roll. An at-the-money option is one whose strike price equates to the current futures price.

implied-variance measures, and attach them to a single futures contract that expires no earlier than 30 days after the end of the variance window.[5] Using three tenors allows us to gauge how the same factor affects the implied-variance curve. We report all estimates in annualized variance points, defined as annualized variance multiplied by 100. An option-implied volatility of 30% corresponds to $0.30^2 \times 100 = 9$ variance points. A corresponding premium level of 0.63 variance points implies an expected physical variance of $9 - 0.63 = 8.37$ points, or 0.0837 in decimal terms. Its square root is $\sqrt{0.0837} = 0.289$, so in this example expected physical volatility is 28.9% against an implied volatility of 30%. The premium would therefore account for roughly one point of implied volatility. In our sample, on average, the corn market charges a positive VRP—0.50 annualized variance points at a 30-day tenor, 0.63 at 60, and 0.70 at 90—while the premium for soybeans is statistically indistinguishable from zero at any tenor. That difference may emerge for several reasons, the most likely being structural. In our model the VRP is positive when end users and financial firms are net purchasers of variance protection. An ethanol plant buys corn and sells fuel priced relative to gasoline. Because it has to absorb corn price increases it buys options as a hedge. A soybean crusher, on the other hand, buys beans and sells meal and oil—but those prices track the bean price, so its margin is little affected and it has less reason to hedge with options.

We introduce a partial-equilibrium model to organize the factors likely to affect each component. Expected physical variance depends on (a) uncertainty about supply—theorized to rise over time as a consequence of climate change—and (b) the sensitivity of commodity prices to that quantity uncertainty. Empirically, however, we find that neither variance expectations nor VRPs exhibit a sustained upward trend (for either corn or soybeans) between 2014 and 2026. Thus, we find no evidence that climate effects produced a sustained increase in row-crop market participants' expectations of physical price risk. It is worth stressing that we cannot distinguish if climate effects are simply not evident or if producers and/or markets have managed to adjust to them. In the theoretical model, low levels of stocks (implicit in the commodity futures term structure) make a given production surprise harder to buffer. As the model predicts, we find in the baseline empirical specification that corn variance is higher when the futures curve signals scarcity of the commodity;[6] because the VRP does not respond to it, that relationship falls entirely in the physical channel.

In our model, the VRP is the price that clears a market between participants demanding

[5] Doing so avoids mechanical changes in the realize variance introduced by rolling from one futures contract to the next.

[6] For soybeans the sign of the relationship is identical, but the finding is not statistically significant. It may be that soybean variance is less sensitive to storage since the South American crop follows just six months afterwards.

variance protection and dealers or other traders who have limited capacity to supply it. Increased commercial hedging pressure and tighter financial conditions can therefore raise the premium, as can the growing season itself—insofar as weather risk concentrated in a few weeks is difficult to hedge or offset against other clients. We find that, at a 60-day tenor, the corn VRP behaves as predicted: it is higher when the VIX is elevated and commercial hedging pressure is high; it also tends to be higher when the measurement window spans the critical growing period for the crop. These relationships are weaker for soybeans. According to the model, financial participation can affect either the demand for protection or the capacity to bear its risk, so its predicted effect is not signed with certainty. Our empirical results fail to establish that index-trader participation (measured using public CFTC trader position data) has a negative relationship with the premium. Together with the absence of VRP trend, our findings therefore do not support the possibility that rising physical risk from climate change was concealed by a falling premium as these markets became financialized.

We control for scheduled government crop reports, since prior work shows theoretically and empirically that their release ought to resolve uncertainty and therefore lower option-implied variance. Our level regressions indicate that this effect operates through the physical channel: for both crops, approaching USDA Grain Stocks releases are associated with higher expected physical variance at the shorter tenors. World Agricultural Supply and Demand Estimates (WASDE) reports produce the same baseline relationship for corn. Our dynamic results confirm that expected physical variance for both crops decreases sharply on both types of USDA publication days—most strongly for Grain Stocks reports.

Other dynamic evidence offers further insight. During soybeans' critical crop window, a hot forecast innovation from the U.S. National Oceanic and Atmospheric Administration (NOAA) Climate Prediction Center is associated with higher subsequent physical variance (corn shows the same mean effect, but it is statistically insignificant). Poor agricultural supply news (identified with Adjemian and Jo (2025)'s instrument) raises corn's implied variance at impact, by increasing the VRP. Soybean implied variance also rises and so does its premium, though the latter is statistically insignificant. Oil supply news work through the physical channel. A Känzig (2021)-identified price increase has no precise effect for over a month, then raises subsequent realized variance in both crops by more than it raises implied variance, meaning that the price of protection falls. That delay is consistent with the biofuel channel: higher crude prices improve blending profits and raise the derived demand for row crops, but crush margins and decisions adjust over weeks rather than days, so the variance arrives after the shock. Monetary tightening reverses that pattern in corn, lowering expected physical variance while raising the price of protection for several weeks—the risk-bearing mechanism that our model describes. Climate news attention, measured with media

coverage indices and unsigned by sentiment, splits the same way. Transition-risk attention raises expected physical variance——clearly in soybeans, less precisely in corn——without a lasting change in the VRP. Physical-risk attention works the other way, lowering expected variance in both crops. Since the physical risk index measures media coverage rather than damage, a spike in physical-risk news can resolve uncertainty rather than create it.

Finally, our decomposition has a direct monetary implication. The Risk Management Agency uses option-implied volatility to construct the volatility factor it uses to rate federal crop insurance. We reconstruct the procedure that RMA uses to determine how much different the premium would be if the VRP component were excluded, and we apply it to an illustrative range of RMA products. For corn, our results indicate that including the price of protection raises the total premium rate by about 7% at the median across products. The resulting additional premium is shared between producer-paid premiums and taxpayer federal premium subsidies; our back-of-the-envelope calculation pegs the ten-year sum at $5.8 billion. Matching our level regression findings, soybean products' corresponding premium change is not distinguishable from zero.

2 Background

Financial derivatives are forward-looking instruments: their prices combine forecasts with compensation for bearing risk. Fama (1984) studies how the forward discount apportions between expected depreciation and a time-varying premium. French (1986) and Fama and French (1987) apply the same question to commodity futures, focusing on whether the futures–spot spread forecasts a subsequent spot-price change, reflects a futures risk premium, or both. Later work focuses on how commodity futures premia respond to inventory levels and physical scarcity (Gorton et al., 2013), to net hedging demand, and to limits on intermediary risk-bearing capacity (Acharya et al., 2013). A similar logic applies to options. When an option-implied variance changes, the price shift alone does not reveal whether traders expect the underlying futures price to vary more over the option horizon or whether protecting against that variance has become more expensive (or both). Conducting a decomposition is needed to determine how implied variance evolves over time, and why.

For equities, Carr and Wu (2009) show how option prices convey risk-neutral expected variance information about the variance-swap, Bollerslev et al. (2009) use the implied–realized variance difference to study expected stock returns, and Bekaert et al. (2013) and Bekaert and Hoerova (2014) separate expected physical variance from the variance premium and examine how uncertainty, risk aversion, and monetary policy shape their evolution. Trolle and Schwartz (2010) and Prokopczuk et al. (2017) establish the existence of variance

risk premia in energy and other commodity markets.

Up to now, the agricultural economics literature has been concerned mainly with testing the forecasting ability of option-implied variance.[7] Egelkraut et al. (2007), for example, show that corn option-implied forward volatility predicts subsequent realized volatility. Wang et al. (2012) estimate the corn VRP and compare the predictive content of alternative forward-variance measures. More recently, the literature has focused on the implications of its forecasting errors. Xi et al. (2019) document average implied–realized gaps for corn and soybeans in 2010-2016. They find that implied volatility exceeds expected volatility by about 15% for both crops, and consider the crop-insurance implications.[8] He et al. (2023) decompose the agricultural commodity VRP into jump and diffusion components, and locate much of premium in jump-tail risk—providing evidence that it varies seasonally and around USDA reports. He and Bian (2025) document compensation for report-day risk, while Cao and Robe (2022) and Yang and McKenzie (2025) study how scheduled USDA releases impact the total option-implied uncertainty for the main U.S. row crops.

Although the above papers investigate its forecast accuracy, the existence of average premia, or jump and announcement effects, the literature has not yet offered a common regressor decomposition—for the same observations and over the same horizons—that explains movements in the components of the implied variance itself. Doing so can establish whether its mean, trend, and response to economic states and news arise from expected physical variance or from the price of protection.

Our crop-insurance application follows Bulut and Fortenbery (2022), who find useful forecast content in RMA's volatility factor—particularly for corn—while Xi et al. (2019) emphasize the overstatement implications associated with the implied–realized volatility gap that they document. An option-derived factor can forecast physical variance while also embedding a time-varying price of variance protection. We combine RMA's volatility factor construction procedure with the crop-specific variance gaps we estimate, and also the premium-rate sensitivities illustrated as an example by Sherrick (2015) to measure how much of that price is carried into an administrative insurance rate intended to reflect physical price risk.

[7] An exception is the work of Adjemian et al. (2017), who use time series methods to decompose the path of overall implied volatility for agricultural commodities.

[8] They compare 22-trading-day realized variance with a 30-day, across-strike variance-swap rate from 2010–2016. We instead use fixed-tenor two-sided ATM implied variance and same-contract subsequent variance using data from 2014–2026. Their 15% is an implied-volatility statistic, whereas our roughly 7% corn result is a median illustrative insurance-rate effect. The sample difference may be enough to explain the different results—certainly our historical decomposition confirms that the VRP is time varying. But option horizon, strike coverage, or contract matching may also factor in; we intend to expand our sample backward in time to further explore the discrepancy.

3 Decomposing Implied Variance

Option-implied variance at time t is the sum of expected physical variance—uncertainty about the price of the underlying asset over the period the option represents—and the price of bearing that risk, the VRP. Realized variance is itself the sum of expected variance and of a conditionally mean-zero surprise, so differencing implied and realized variance removes expected variance exactly. The average of that difference therefore identifies the average premium, and its projection onto any factor known at t recovers the relationship of that factor with the premium, without the need to estimate expected variance at all. Likewise, projecting realized variance onto the same factor recovers its relationship with expected variance. The same logic applies to the estimation of dynamic effects in response to an innovation at t. Although the exact decomposition applies to a model-free variance, our at-the-money (ATM) proxy adds a measurement wedge whose implications we describe below.

For commodity i, information date t, and horizon h, futures contract $c(i,t,h)$ trades at settlement price $F_{i,c,\tau}$; it expires no earlier than $h+30$ calendar days after t. We follow that same contract through the end of the window, so the realized variance describes one price series and never a roll between contracts.[9] On trading day τ, where $\tau-1$ is the preceding settlement day for the same contract, the annualized realized variance is calculated as the normalized sum of squared changes in log futures prices

$$X_{i,h,t} = \frac{365}{h} \sum_{\tau \in \mathcal{T}_i(t,t+h)} \left[\log F_{i,c,\tau} - \log F_{i,c,\tau-1}\right]^2, \tag{1}$$

where $\mathcal{T}_i(t,t+h)$ contains the product's exchange trading days strictly after t and through $t+h$.[10]

At t, the expected physical variance is

$$P_{i,h,t} = \mathbb{E}^{\mathbb{P}}[X_{i,h,t} \mid \mathcal{I}_t], \tag{2}$$

where $\mathbb{P}$ is the physical probability measure and $\mathcal{I}_t$ contains the information available at t. That is, evaluated at the data-generating probabilities, taking the expectation in (2) produces the market's forecast of physical variance if expectations are rational. But taking the same expectation evaluated over the probabilities implicit in option prices—which may

[9] Calculating returns with a continuation series risks recording the price difference between two different contracts when the series rolls. Following a named contract instead prevents that difference from contaminating the realized variance.

[10] All components in the decomposition below use the same annualized variance units. The RMA application converts annualized variance to its crop-specific February-to-October horizon only at the final translation stage.

overweigh less desirable outcome states—produces

$$q_{i,h,t} = \mathbb{E}^{\mathbb{Q}}[X_{i,h,t} \mid \mathcal{I}_t], \tag{3}$$

where $\mathbb{Q}$ is the pricing measure. This is the fair variance-swap rate, quoted in the same annualized variance units as $X_{i,h,t}$.[11] The VRP is the difference between the two measures,

$$\mathrm{RP}_{i,h,t} = q_{i,h,t} - P_{i,h,t}, \qquad \text{so that} \qquad q_{i,h,t} = P_{i,h,t} + \mathrm{RP}_{i,h,t}. \tag{4}$$

Our sign convention makes $\mathrm{RP} > 0$ when variance protection is expensive relative to expected physical variance.[12] The realization surprise

$$\varepsilon_{i,h,t} = X_{i,h,t} - P_{i,h,t}$$

has conditional mean zero by construction. Thus, for every integrable state Z_t known at t,

$$\mathbb{E}^{\mathbb{P}}[\varepsilon_{i,h,t} \mid \mathcal{I}_t] = 0, \qquad \mathbb{E}^{\mathbb{P}}[\varepsilon_{i,h,t} Z_t] = 0. \tag{5}$$

Let $\beta_Y(Z)$ denote the population coefficient on Z_t in a linear regression of Y on Z_t and a common set of controls known at t, and let $\widehat{\beta}_Y(Z)$ denote its sample analogue. Equation (5) and the identity $q = X + (q - X)$ imply that

$$\beta_X(Z) = \beta_P(Z), \qquad \beta_{q-X}(Z) = \beta_{\mathrm{RP}}(Z), \qquad \widehat{\beta}_q(Z) = \widehat{\beta}_X(Z) + \widehat{\beta}_{q-X}(Z). \tag{6}$$

The final equality in (6) holds exactly when the three projections use the same observations, right-hand-side variables, weights, outcome scaling, and linear estimator. It shows that the relationship between the implied variance and any state known at t is exactly allocated between the expected physical variance and the market price of variance protection. The same argument applies dynamically. At a post-event local-projection horizon, the innovation is in the information set when the later realized-variance window opens and is therefore orthogonal to the corresponding realization surprise. For our agricultural-supply, oil-supply, and monetary-policy designs, the $t = 0$ option quote incorporates the shock-session innovation while the realized-variance window begins after that session. With identical observations and specifications for the three outcomes at each horizon, their estimated responses therefore add

[11] The fair rate is the fixed leg that makes the contract worth zero at inception, and it can be replicated from a strip of options across strikes (Neuberger, 1994; Britten-Jones and Neuberger, 2000; Demeterfi et al., 1999).

[12] Some studies, such as He et al. (2023), instead report $P - q$, or its ex-post counterpart $X - q$. Their sign is the reverse of ours.

up exactly and admit the same component interpretation.

This projection logic follows Fama (1984) and Fama and French (1987), who regress complementary realizations on the forward discount or commodity basis. Because their right-hand-side variable is the price difference being decomposed, their coefficients allocate its own variation between expectations and the premium and sum up to one. Carr and Wu (2009) provide the variance-market counterpart by averaging the matched payoff $X - q$ and regressing X on q. Because we want to know how market developments operate through the two channels to affect option-implied variance, we use the same matched payoff approach but place external economic states and identified innovations Z_t on the right-hand side, so the component coefficients sum to $\beta_q(Z)$.[13]

The ideal q exploits information from options across many *strikes*, the exercise prices at which option holders may buy or sell the underlying. *Call* options gain value as the price of the underlying rises, while *put* options rise in value as the underlying price falls; both types of options are listed over many strikes. Our options data, sourced from LSEG, instead offer fixed-tenor at-the-money implied volatility, drawn from options spanning a constant h-day window, with a strike price just at the current futures price. Squaring an annualized implied volatility yields an annualized implied variance; averaging the call and put sides generates the two-sided measure

$$q^{\mathrm{ATM}}_{i,h,t} = \frac{1}{2}\left[\left(\frac{IV^{\mathrm{call}}_{i,h,t}}{100}\right)^2 + \left(\frac{IV^{\mathrm{put}}_{i,h,t}}{100}\right)^2\right] = q_{i,h,t} + \nu_{i,h,t} = P_{i,h,t} + \mathrm{RP}_{i,h,t} + \nu_{i,h,t}. \tag{7}$$

The wedge ν collects the gaps between the ATM measure and the exact decomposition in (4), due to the use of a single strike rather than a strip of strike prices, the fixed-tenor interpolation, liquidity, quote timing, the fact that LSEG does not document the precise futures contract underlying each tenor, and any other mismatches. Its sign is not known: an ATM quote can lie above or below the corresponding strip-based variance rate.

Replacing q with q^{ATM} preserves the exact estimated adding-up relation in (6), with q^{ATM}

[13] In work related to agricultural commodities, our closest antecedent is Triantafyllou et al. (2015); they are principally concerned with comparing forecasts of future realized variance, and include model-free implied variance itself among the predictors. They also regress $X - q$ and implied variance q on economic variables, but in separate exercises with nonidentical right-hand sides. As a result, their coefficients do not allocate a common relationship between an economic state and implied variance into expected-variance and premium components. Diersen and Wang (2025) likewise estimate complementary realized-volatility and implied-minus-realized-volatility equations on implied volatility, a report indicator, and their interaction. Working with volatility units, they are concerned with whether weekly options contain distinct information. Neither paper places all three outcomes in a single variance-unit projection system, so that every state and news response adds back to the observed implied-variance response across tenors and dynamic horizons.

substituted for q, but changes the population interpretation to

$$\beta_{q^{\text{ATM}}-X}(Z) = \beta_{\text{RP}}(Z) + \beta_{\nu}(Z). \tag{8}$$

Accordingly, $\mathbb{E}[q^{\text{ATM}} - X] = \mathbb{E}[\text{RP} + \nu]$, and the forecast-free mean and projections combine the model-free VRP with the ATM wedge. A dated forecast-based measure additionally contains the forecast discrepancy because

$$q^{\text{ATM}}_{i,h,t} - \widehat{P}_{i,h,t} = \text{RP}_{i,h,t} + \nu_{i,h,t} + \left(P_{i,h,t} - \widehat{P}_{i,h,t}\right). \tag{9}$$

Unlike the realization surprise, this last term need not be orthogonal to Z_t. We therefore use the matched $q^{\text{ATM}} - X$ measure for our headline means, state relationships, and news responses, and reserve $\widehat{P}$ for the historical decompositions and the RMA application. For simplicity, throughout the paper we call both quantities risk-premium measures while retaining these qualifications.

4 A model of commodity implied volatility

We link the implied variance decomposition to a theoretical partial-equilibrium model to formalize the VRP as a price that clears a market, and motivate the factors that we associate with the risk premium or with the expected physical variance. Deriving comparative statics indicates what signs to expect from empirical work.

4.1 Expected variance: storage and climate

Two elements determine P: how much surprise is attached to the commodity supply side, and how much of that surprise affects prices. The first is nature, which supplies exogenous shocks: weather over the window generates a (bullish) production surprise w with conditional variance $h(c_t)$, where c_t is the weather-uncertainty state—how unresolved the season's outcomes remain at t—and $h_c > 0$. Climate change, if it is observable, shifts h upward, so a given c carries more production variance than it once did.

Storage is the second. Storers respond endogenously to weather shocks. Intertemporal arbitrage permits the carrying of inventory toward periods of scarcity, cushioning the price response to a shortfall when stocks are positive; the same shortfall produces a larger price change as stocks approach zero, because nothing can be carried out of an empty silo. Write the price change over the window as $g(s_t, w)$, where s_t is a stocks-to-use proxy,[14] with $g_w > 0$,

[14] Stocks-to-use is the ratio of ending inventories to total use over the period, a standard measure of scarcity

$g_{ww} > 0$, and $g_{ws} < 0$: a shortfall raises price, by more the larger the shortfall, and by less the fuller the silo. Linearizing g in w around $w = 0$ yields:

$$g(s_t, w) \approx g(s_t, 0) + g_w(s_t, 0)\, w.$$

Let $\varphi(s_t) \equiv g_w(s_t, 0)$, the price impact of a unit shortfall at stocks s_t. Then $\varphi > 0$ and $\varphi_s < 0$ follow from $g_w > 0$ and $g_{ws} < 0$. The first term is known at t and drops out of the variance, so the price change varies only through $\varphi(s_t)\, w$, and

$$P_t = \varphi(s_t)^2\, h(c_t) + \sigma^2_{u,t}, \tag{10}$$

where σ^2_u collects variance from demand, macro, and policy shocks, taken to be conditionally uncorrelated with w. Three comparative statics follow:

$$\begin{aligned} &\frac{\partial P}{\partial c} = \varphi(s_t)^2\, h_c(c_t) > 0, \\ &\frac{\partial P}{\partial s} = 2\varphi(s_t)\, \varphi_s(s_t)\, h(c_t) < 0, \qquad \frac{\partial^2 P}{\partial s\, \partial c} = 2\varphi(s_t)\, \varphi_s(s_t)\, h_c(c_t) < 0. \end{aligned} \tag{11}$$

Climate uncertainty raises expected physical variance, and how much it raises it depends on the silo: the same increase in c_t raises P more when the stocks-to-use ratio is low.

Of course, not all unfavorable news raise expected variance. Bullish news about the crop raise its price, but in (10) its *variance* changes only when news affects how much uncertainty is left, c_t, or the degree to which a given shock can be cushioned through s_t. A fully anticipated crop loss does neither, so although it may raise prices it does not update their expected variance P. A report that resolves uncertainty lowers c_t and therefore lowers P, even when its content is bullish. One more implication of (10) is that common weather draws can make w, and therefore the realization surprises ε, correlated across crops.[15]

4.2 The market for variance protection

Three kinds of agents use the market. Commercial firms handle the physical commodity—farmers, merchants, and processors whose revenues and costs depend on its price—and generally buy options to hedge.[16] Financial firms hold commodity positions as investments: index funds, exchange-traded funds, and managed money traders, whose positions change

in commodity markets. Because it is only measured infrequently, in our empirical work we proxy for s_t using the slope of the futures price term structure.

[15] The physical link between corn and soybeans is why our inference resamples the two crops together.

[16] In practice, commercial market participants sometimes sell options if they have a directional view about prices. Our model permits them to.

as money flows into and out of their products. "Dealers" take the other side of both.[17] A dealer can hedge an option's directional exposure, but not the risk that variance itself comes in high or low. That residual stays on the book, and the dealer charges a market-clearing price to carry it: the VRP.

Let d_t denote customers' net demand for variance exposure—commercial and financial firms together—in variance-equivalent units. It can take either sign. Producers creating price floors, end users buying protection against price spikes, and managed money buying straddles ahead of reports push d up; programs that sell options against physical or financial positions push it down. Only the variance component of these trades enters d; directional exposure can be hedged with futures, and skew is handled by options distant from the current futures price. What remains is the variance risk the dealers carry.

A customer who buys variance protection pays q and receives X, so the dealer on the other side earns $q - X = \mathrm{RP} - \varepsilon$:[18] the premium for certain, less the realization surprise. Expected profit per unit is RP with risk $\mathrm{Var}_t(\varepsilon)$. A dealer choosing position θ solves

$$\max_{\theta_t} \Big\{ \theta_t \,\mathrm{RP}_t - \tfrac{\kappa_t}{2}\,\omega_t^2\,\theta_t^2 \Big\}, \qquad \omega_t^2 = \mathrm{Var}_t(\varepsilon_t), \tag{12}$$

where κ_t is risk-bearing tightness—the inverse of capacity, high when dealer capital is scarce—and ω_t^2 is the conditional variance of the variance payoff, in squared-variance units. The first-order condition determines the supply schedule $\theta_t = \mathrm{RP}_t/(\kappa_t\omega_t^2)$: dealers carry more variance risk only at a higher price. Clearing at $\theta_t = d_t$ gives

$$\mathrm{RP}_t = \kappa_t\,\omega_t^2\,d_t. \tag{13}$$

If demand is price-sensitive and downward sloping, $d_t = d_t^0 - b\,\mathrm{RP}_t$, the premium becomes $\kappa_t\omega_t^2 d_t^0/(1+\kappa_t\omega_t^2 b)$, attenuated relative to (13) but with every sign preserved. The displayed comparative statics below use the inelastic-demand benchmark in (13); elastic demand attenuates their magnitudes.

$$\begin{aligned} &\frac{\partial \mathrm{RP}}{\partial d} = \kappa_t\omega_t^2 > 0, \\ &\frac{\partial \mathrm{RP}}{\partial \kappa} = \omega_t^2 d_t \gtrless 0, \qquad \frac{\partial \mathrm{RP}}{\partial \omega^2} = \kappa_t d_t \gtrless 0, \qquad \text{both signed by } d_t. \end{aligned} \tag{14}$$

[17] Grain option markets have no designated market maker of the kind that equity markets rely on. We use the term "dealer" for the marginal supplier of variance exposure, whatever its institutional form: exchange market makers, proprietary trading firms, swap dealers hedging over-the-counter grain business, or at times commercial or managed-money programs that sell premium. The model requires only that this capacity be limited and vary over time.

[18] Recall equation (4).

More net demand for protection raises the premium. Tighter dealer capacity or a riskier variance payoff makes the premium larger in absolute value, since whoever carries the risk must be paid more to carry it, but the direction of the change follows d_t. When customers are net buyers, dealers sell protection and charge for it. When customers are net sellers, dealers take the other side and do so only at a discount, so the premium is negative. But whether d_t is empirically positive or negative is down to market structure rather than a parameter that the model sets.

To link the price of variance protection to the physical states, we assume that greater weather uncertainty raises the conditional variance of the variance payoff and that fuller stocks reduce it:

$$\frac{\partial \omega^2}{\partial c} > 0 \qquad \text{and} \qquad \frac{\partial \omega^2}{\partial s} < 0.$$

Because $q = P + \text{RP}$, differentiating with respect to c affects q through both channels. It enters P through h, and enters RP through ω^2. Holding demand, dealer tightness, stocks, and the other variance sources fixed, we have:

$$\frac{\partial q}{\partial c} = \underbrace{\varphi(s_t)^2 \, h_c(c_t)}_{\text{expected variance}} + \underbrace{\kappa_t d_t \, \frac{\partial \omega^2}{\partial c}}_{\text{risk premium}} . \tag{15}$$

The first term in (15) is positive. In the second, c enters only through ω^2, and $\partial \omega^2 / \partial c > 0$, so more weather uncertainty grows the multiplier on d_t in (13). That raises the premium in absolute value, whatever the sign of d. Stocks work through the same two channels,

$$\frac{\partial q}{\partial s} = \underbrace{2\varphi(s_t) \, \varphi_s(s_t) \, h(c_t)}_{\text{expected variance}} + \underbrace{\kappa_t d_t \, \frac{\partial \omega^2}{\partial s}}_{\text{risk premium}} , \tag{16}$$

where the first term is negative. But notice that in the second term of (16), s enters only through ω^2, and $\partial \omega^2 / \partial s < 0$, so a fuller silo shrinks the multiplier on d_t in (13). That attenuates the premium no matter the sign on d.

Financialization can affect the premium in two ways. Let f represent the extent of participation in the market by financial traders. We assume it raises client demand for variance protection, $\partial d / \partial f > 0$ because managed money buys options ahead of scheduled information, and that it adds risk-bearing capital $\partial \kappa / \partial f < 0$, because proprietary trading firms and hedge funds are among the marginal suppliers of variance exposure. Commodity index traders sit outside both channels: they hold long-only futures positions and neither buy nor supply variance, so they reach the premium only if deeper futures markets lower the

cost of delta-hedging. Given (14),

$$\frac{\partial \text{RP}}{\partial f} = \underbrace{\kappa_t \omega_t^2 \frac{\partial d}{\partial f}}_{>0} + \underbrace{\omega_t^2 d_t \frac{\partial \kappa}{\partial f}}_{<0 \text{ when } d_t>0}, \tag{17}$$

so when customers are net buyers ($d_t > 0$) the two terms oppose each other and financialization's net effect on the VRP is an empirical question. On the other hand, if customers are net sellers, then $d_t < 0$ and both terms in the equation above are positive—financialization raises the premium.

4.3 Extending Across Markets

Physical and financial ties link markets in ways that challenge the scalar representation that we have used for the model up to now. Some crops are linked, physically: for example, corn and soybeans compete for acreage, and government programs that direct corn (soybean) conversion into ethanol (biodiesel) tie it to energy prices (Avileis, 2025; Raman and Robe, 2026). On the financial side, dealers manage books across markets. A dealer who is short variance in both corn and soybeans holds two positions that may lose together, if the same weather surprises for both commodities. The scalar ω^2 cannot easily represent that reality, so we use vector notation to stack markets, indexing variables by tenor where needed, and let net demand and supply for variance exposure be

$$\mathbf{d}_t(\mathbf{RP}_t) = \mathbf{d}_t^0 - B_t \mathbf{RP}_t, \qquad \boldsymbol{\theta}_t = \frac{1}{\kappa_t} \Omega_t^{-1} \mathbf{RP}_t,$$

so that market clearing gives

$$\mathbf{RP}_t = \left[B_t + \frac{1}{\kappa_t} \Omega_t^{-1} \right]^{-1} \mathbf{d}_t^0, \tag{18}$$

where Ω_t is positive definite (so higher premia invite more dealer supply), and B_t is positive semidefinite (so customer demand does not rise with the price of protection). Together, these properties make the bracketed matrix in (18) invertible.

Notice that equation (13) is the single-market special case of (18), with $B_t = 0$; the price-elastic scalar case above has $B_t = b$. Writing the model in vector notation offers additional intuition. Physical links across crops enter P. The scalar model located those shocks in σ_u^2.

Here, expected variance in crop i is

$$P_{i,t} = \varphi_i(s_{i,t})^2\, h_i(c_{i,t}) + [W_t\, \mathbf{g}_t]_i + \sigma^2_{u,i,t}, \qquad \frac{\partial P_{i,t}}{\partial g_{j,t}} = W_{ij,t}, \tag{19}$$

where $\mathbf{g}_t$ stacks shocks originating in each market and W_t specifies the physical links that carry them across crops: acreage competition, and the ethanol (resp. biodiesel) margin that ties corn (resp. soybeans) to energy prices. Carry and crop states are informative about P through the first term, and a shock in one market reaches the other through $W_{ij,t}$.

The composition of traders and their market presence may affect the demand for variance protection $\mathbf{d}^0$, while κ is still the common tightness of risk-bearing capacity—which is why trader positions and financial conditions are informative about RP. Outside-market news can change both the physical and the VRP component, operating through W_t or through dealer balance sheets; an oil shock, for instance, can change corn's expected variance through the ethanol margin, or raise the premium by tightening dealer capacity.

The matrix Ω_t records how the realization surprises $\varepsilon_{i,h,t}$ co-move across markets. A drought that slashes both corn and soybean yields makes both crops realize more variance than expected at the same time, so the off-diagonal entry for corn i and soybeans j, $\Omega_{i,j,t}$, is positive. Its diagonal entry $\Omega_{i,i,t}$ is the scalar ω_t^2 from (12) for crop i.

For clarity, assume that customers buy a fixed amount of protection no matter what it costs, so that $B_t = 0$—the same inelastic benchmark used in the single-market case in (13). Then (18) becomes $\mathbf{RP}_t = \kappa_t \Omega_t \mathbf{d}^0_t$, and the premium on crop i is:

$$\mathrm{RP}_{i,t} = \kappa_t \left(\Omega_{i,i,t}\, d^0_{i,t} + \sum_{j \neq i} \Omega_{i,j,t}\, d^0_{j,t} \right), \qquad \frac{\partial \mathrm{RP}_{i,t}}{\partial \Omega_{i,j,t}} = \kappa_t\, d^0_{j,t}. \tag{20}$$

A dealer cannot simply lay off corn variance risk by taking a soybean position, because the two tend to lose at the same time. So when customers are net buyers of soybean variance, that correlation raises the price of corn protection even if nothing about corn *per se* has changed. The two cross-market channels are $W_{ij,t}$ and $\Omega_{ij,t}$: a shock arriving through the first shows up in X, and one arriving through the second shows up in $q^{\mathrm{ATM}} - X$.

Our regressions cannot recover Ω_t, B_t, κ_t, or $\mathbf{d}^0_t$ individually, since only their product appears in the premium. And when B_t is not diagonal, the inverse in (18) can flip the sign of an individual entry. As a result, we treat the comparative statics derived in the model as suggestive benchmarks rather than structural estimates.

5 Empirical design

Most of our empirical work exploits the exact adding-up nature of option-implied variance, and does not depend on a forecasting model, as documented in appendix table A1. Only the historical decomposition does. This section runs in the order of that table, before describing our policy application and the method we use to draw inference on our findings.

5.1 Connecting the model to our empirical work

Table 1 collects the model's most useful comparative statics for our empirical specifications. Weather uncertainty raises expected variance directly, and (weather events that produce) low stocks amplify it further by removing the storage buffer between production shocks and price effects; both can affect the VRP through ω^2. Trader behavior enters only the premium, and the two channels run through different participants. Commercial hedging pressure and managed-money option buying raise d, while proprietary trading firms and hedge funds supply variance and so lower κ. Index traders affect the options market indirectly: they hold long-only futures positions and neither buy nor supply variance, but affect the premium only if deeper futures markets lower the cost of delta-hedging. That combination forms an offsetting-trends hypothesis: physical variance can rise while expanding dealer capacity compresses the VRP, leaving option-implied variance stable. The fifth row in Table 1 concerns a scheduled release that resolves uncertainty, holding its content fixed.

5.2 The average variance risk premium

Subtracting realized variance from the measured ATM value yields

$$q_t^{\mathrm{ATM}} - X_t = \mathrm{RP}_t + \nu_t - \varepsilon_t. \tag{21}$$

Because $\mathbb{E}[\varepsilon_t] = 0$,

$$\mathbb{E}[q_t^{\mathrm{ATM}} - X_t] = \mathbb{E}[\mathrm{RP}_t] + \mathbb{E}[\nu_t].$$

As a result, the sample mean identifies the average premium plus the average ATM measurement wedge without the need to forecast the expected future physical variance. It identifies the average economic VRP itself if $\mathbb{E}[\nu_t] = 0$, or if we estimate the mean wedge from option chains across strikes and subtract it.[19] We refer to the observed gap in (21) as the ATM-

[19] For robustness we intend to follow this approach in the next draft, since constructing a model-free strip variance directly would remove ν rather than correct for it. Serial and standard expiries bracket our horizons on most days, so constant-maturity series are constructible at all three tenors, though strikes far from the money are thinly traded and we would need to be judicious about the cutoff point as far as the appropriate

Change	**Effect on** P	**Effect on** RP	**Where we test it**
Climate uncertainty $c \uparrow$	$+$, by (11)	Away from **zero** through ω^2, the variance of the variance payoff; $+$ if net demand for protection $d > 0$	Climate-state relationships and weather-forecast and climate-news response specifications for X and $q^{\mathrm{ATM}} - X$
Stocks $s \uparrow$	$-$, by (11)	Toward **zero** if $\partial\omega^2/\partial s < 0$; $-$ if $d > 0$	State regressions using vintaged stocks and carry
Net variance demand $d \uparrow$	No direct effect	$+$, by (14)	State regressions of $q^{\mathrm{ATM}} - X$ on position proxies
Dealer capacity $\uparrow$ ($\kappa \downarrow$)	No direct effect	Toward **zero**	Associations with financial conditions; trends in $q^{\mathrm{ATM}} - X$
Approach and resolution of scheduled uncertainty	$+$ before release and $-$ afterward if the conditional distribution first widens and then narrows	Away from **zero** before release and toward **zero** afterward if event demand or ω^2 first rises and then falls	WASDE and Grain Stocks dummy impact paths; report-window state regressions

Table 1: Model implications and where the paper tests them.

Notes: The table lists what our model implies about each component of option-implied variance and where the paper examines it. P is expected physical variance, the variance the market expects the futures price to realize over the window; RP is the economic VRP. Neither is observed directly. We instead observe q^{ATM}, at-the-money implied variance, and X, the variance the futures market subsequently realizes. Because $X = P + \varepsilon$, regressing X on states known at t recovers their relationships with P. Because $q^{\mathrm{ATM}} - X = \mathrm{RP} + \nu - \varepsilon$, where ν is the ATM measurement wedge, regressing $q^{\mathrm{ATM}} - X$ on those states recovers their relationships with $\mathrm{RP} + \nu$, and with RP when the wedge is constant or orthogonal to the regressors. The WASDE and Grain Stocks dummy impact paths are descriptive timing responses. Their negative horizons measure anticipation of publicly scheduled releases, not leakage of the eventual report content. *Position proxies* are commercial hedging pressure and the commodity-index-trader long share from the CFTC's weekly Commitments of Traders. *Signed report content* refers to the directional price impact of the subsequently observed USDA surprise.

based average VRP, subject to this qualification. Calculating it still requires measurement choices: constructing q^{ATM}, selecting the named futures contract for X, and matching their horizons.

5.3 Average State Relationships

For crop i and tenor h, we estimate the same linear specification on the same observations for the option-implied variance, the expected physical variance, and the VRP:

$$Y^j_{i,h,t} = \alpha^j_{i,h} + Z'_{i,t}\beta^j_{i,h} + \delta^j_{i,h}\widetilde{t} + \sum_{r=1}^{2}\left[a^j_{r,i,h}\sin(2\pi r d_t/365.25) + b^j_{r,i,h}\cos(2\pi r d_t/365.25)\right] + u^j_{i,h,t}, \tag{22}$$

where j indexes all three outcomes, $Y^j_{i,h,t} \in \{q^{\text{ATM}}_{i,h,t},\, X_{i,h,t},\, q^{\text{ATM}}_{i,h,t} - X_{i,h,t}\}$. The intercept is $\alpha^j_{i,h}$ and $u^j_{i,h,t}$ is the regression error. The vector $Z_{i,t}$ contains all ten state variables of the preferred specification, entered jointly, with coefficients $\beta^j_{i,h}$. The centered linear time trend $\widetilde{t}$ carries coefficient $\delta^j_{i,h}$. The last term is two annual harmonics: d_t is the day of the year, so the $r = 1$ sine and cosine pair traces one cycle per year and the $r = 2$ pair traces two, which together allow an asymmetric seasonal shape without absorbing the state variation that we want to measure. Every coefficient is estimated separately by crop and implied variance tenor.

With the same observations, regressors, weights, and linear estimator in all three equations, and with the outcomes entered in common variance units, exact adding-up under linearity gives:

$$\widehat{\beta}^{\,q^{\text{ATM}}} = \widehat{\beta}^{\,X} + \widehat{\beta}^{\,q^{\text{ATM}}-X}. \tag{23}$$

Since $X_{i,h,t} = P_{i,h,t} + \varepsilon_{i,h,t}$ and $\varepsilon_{i,h,t}$ is orthogonal to every regressor, the coefficient in a regression of X on $Z_{i,t}$ equals the projection coefficient for expected variance (see appendix table A1). Meanwhile $q^{\text{ATM}} - X = \text{RP} + \nu - \varepsilon$, so the regression of $q^{\text{ATM}} - X$ on $Z_{i,t}$ recovers the projection of $\text{RP} + \nu$, and the risk-premium projection itself when the ATM wedge projects to zero on the included states. (This requires that we not drop observations based on how the variance turned out, which our contract-selection rules respect.)

The trend term $\widetilde{t}$ works the same way, because calendar time is known at t. Its coefficient in the X equation estimates the trend in physical variance, its coefficient in the $q^{\text{ATM}} - X$ equation estimates the trend in $\text{RP}+\nu$, and the two estimates sum up to the trend in option-implied variance. This is the fourth row of Table 1: dealer capacity compresses the premium; if that happens alongside a rise in expected variance, option-implied variance will appear

set to include.

stable even though its components are not. We estimate the trend both inside (22) and in a reduced version that replaces $Z_{i,t}$ with the trend alone, which runs on the full matched q^{ATM}/X sample through April 2026 rather than on the shorter state sample.

These regressions decompose conditional relationships, but they do not guarantee causality. Stocks, trader positions, and financial conditions can all respond to omitted economic forces. Droughts do not, but they can arrive alongside them, so its coefficient can pick up their effects. We therefore describe these coefficients as state relationships and reserve causal language for separately identified innovations. Overlapping realized-variance windows make the estimates noisy, and persistent state variables retain the usual small-sample caveat; the bootstrap addresses dependence from overlapping outcomes, not persistent-regressor bias.

5.4 Impulse responses

We estimate impulse response functions using local projections (Jordà, 2005) to understand the dynamic relationships between innovations of interest and the two components of option-implied variance. Throughout this subsection, we suppress the crop and tenor subscripts. Every projection below is estimated separately by crop i and tenor h. Let e index events with identified innovations S_e. For the agricultural-supply, oil-supply, and monetary-policy designs, $t(e,k)$ is the kth exchange trading day beginning with the shock session, so $t(e,0)$ is the session on the event's reaction date. The q^{ATM} quote at $t(e,0)$ incorporates the shock, while $X_{t(e,0)}$ is realized over the forward window beginning after that session. Our weather and news-attention designs continue to begin with the first crop session strictly after the source date. For $k = 0, \ldots, 60$ we estimate the same local projection for

$$Y^j_{e,k} \in \left\{ q^{\mathrm{ATM}}_{t(e,k)},\ X_{t(e,k)},\ q^{\mathrm{ATM}}_{t(e,k)} - X_{t(e,k)} \right\} :$$

$$Y^j_{e,k} = \alpha^j_k + \beta^j_k S_e + C'_{e,k}\Gamma^j_k + u^j_{e,k}, \tag{24}$$

where j indexes the three outcomes, α^j_k is the intercept, $u^j_{e,k}$ the regression error, $C_{e,k}$ the vector of controls with coefficients Γ^j_k, and β^j_k—the coefficient of interest—the response at horizon k. Note that $X_{t(e,k)}$ covers the forward window beginning after the origin session; it does not represent cumulative variance from the event through k. The sample, shock scaling, regressors, weights, and linear estimator are identical across the three outcomes, which enter in common variance units, so the point estimates add exactly at every horizon:

$$\widehat{\beta}^{q^{\mathrm{ATM}}}_k = \widehat{\beta}^{X}_k + \widehat{\beta}^{q^{\mathrm{ATM}}-X}_k. \tag{25}$$

Scheduled-report timing is handled with a complementary dummy impact design. WASDE

and Grain Stocks dummy variables and their joint-release interaction enter together. For these impulse responses, event time zero is the official crop session containing the release; the option quote incorporates report-day information, while the subsequent-realized-variance window begins after that session. Our impulse responses use identified innovations, while stocks-to-use and other state levels enter the matched state regressions. For price-instrumented shocks, the shock-day option quote is paired with subsequent realized variance measured after that session. We report changes in financial pressure as dynamic associations. We estimate the report-publication-indicator coefficients on trading days leading up to the report, as well as post-report days, because publication dates are public before their content is observed. Pre-report trading sessions therefore measure anticipation of the scheduled news rather than leakage of the eventual signed surprise.

The controls in $C_{e,k}$ are the matching-tenor pre-event log q^{ATM}; the two pre-event log ratios $\log q^{\mathrm{ATM}}_{i,30} - \log q^{\mathrm{ATM}}_{i,60}$ and $\log q^{\mathrm{ATM}}_{i,90} - \log q^{\mathrm{ATM}}_{i,60}$, which summarize the front and back of the implied-variance curve; lagged 22-session realized variance; CarryPC1; the commercial hedging pressure; the prior own-crop futures return; the outcome window's critical-period overlap; WASDE and Grain Stocks overlap indicators; and the annual harmonics. For weather, we use a crop-area aggregated NOAA-CPC hot innovation.

Two of our innovations arrive as futures-price movements rather than survey surprises, so a price change could reflect the news or something else happening that day. We therefore instrument for them writing Δp_e for the event return and z_e for the instrument; the first stage is:

$$\Delta p_e = \pi z_e + \Lambda' C_e + v_e, \tag{26}$$

where π measures how much of the event return the instrument explains and v_e is the first-stage error. For agricultural supply news, Δp_e is the roll-safe own-crop event return and z_e is the Adjemian and Jo (2025) instrument; for oil supply news, Δp_e is the source-consistent front WTI return and z_e is the Känzig (2021) instrument. We estimate the reduced form on the same sample,

$$Y^j_{e,k} = \rho^j_k z_e + C_e' \Gamma^j_k + u^j_{e,k}, \qquad \beta^{j,IV}_k = \frac{\rho^j_k}{\pi}, \tag{27}$$

so $\beta^{j,IV}_k$ replaces β^j_k from (24) as the response at horizon k. Because all three outcomes share the same first stage and the same regressors, the responses still add exactly as in (25). The result is an implied variance component response to the identified news shock normalized by its price effect.

To interpret these responses as dynamic effects, the innovation and controls must be known by time t; conditional mean-zero realization surprises then make the residualized innovation orthogonal to the surprise, which lets the response of X stand in for the response

of the expected variance and lets the response of $q^{\mathrm{ATM}} - X$ stand in for the response of the VRP plus wedge. A causal interpretation requires more: the residualized innovation must be unrelated to disturbances that would otherwise change the outcome. The surprise construction, timing, sign convention, scaling, and controls provide that argument, innovation by innovation.[20]

The appendix repeats equation (24) with $\widehat{P}$ and $q^{\mathrm{ATM}} - \widehat{P}$ in place of X and $q^{\mathrm{ATM}} - X$. Comparing the two sets of responses shows how much of the split comes from the forecast rather than the data. This matters because q^{ATM} is a predictor in the forecast equation, and q^{ATM} contains the premium, so $\widehat{P}$ can absorb premium variation that belongs in the residual.

5.5 Historical decomposition

While none of the empirical specifications above requires a forecast, recovering the historical path of the components that make up implied variance does, because a particular date supplies just one realization. For example, an unexpectedly calm period of 60 days can make $q_t^{\mathrm{ATM}} - X_t$ large even *when the true VRP never changed*, so the ex-post difference cannot represent the premium with accuracy. Consequently, we construct

$$\widehat{\mathrm{RP}}_{i,60,t} = q_{i,60,t}^{\mathrm{ATM}} - \widehat{P}_{i,60,t} = \mathrm{RP}_{i,60,t} + \nu_{i,60,t} + \left(P_{i,60,t} - \widehat{P}_{i,60,t}\right). \tag{28}$$

The residual in (28) contains the ATM wedge and forecast error as well as the economic premium.

A single prespecified equation, estimated separately by crop and tenor, produces the daily forecast. Its predictors, collected in $\boldsymbol{\zeta}_{i,h,m}$ for each month m, are all market data. Building on prior work, they are: sine and cosine harmonics terms for the annual seasonal cycle; own realized variance, in annualized logs, over the trailing week, month, and quarter (5, 22, and 66 trading days); futures carry for both crops, which storage ties to expected variance in Section 4.1; the other crop's trailing-month realized variance; and the log matching-tenor ATM implied variances of the own crop, the other crop, and an equal-weighted average of matching-tenor Chicago wheat and soybean-oil implied variance. Realized-variance and carry terms are lagged one trading day, so that every input is available when the forecast is made. The forecast equation is

$$\log X_{i,h,m} = \boldsymbol{\zeta}_{i,h,m}'\theta_{i,h} + u_{i,h,m}. \tag{29}$$

[20] State levels such as stocks-to-use never enter as impulses, a surprise partly constructed from futures-price reactions is not paired with a same-trading day futures-derived outcome, and changes in financial pressure without a separate source of identification are reported as dynamic associations.

Estimation is recursive: on date t, the equation is fit only to months whose outcomes are already resolved (so the forecast uses nothing that was unknown at t): $\mathcal{M}_{i,h}(t) = \{m : e_i(m) < t,\ X_{i,h,m} > 0,\ \boldsymbol{\zeta}_{i,h,m}$ complete$\}$, where $e_i(m)$ is the final date of month m's window:

$$\widehat{\theta}_{i,h}(t) = \arg\min_{\theta} \sum_{m \in \mathcal{M}_{i,h}(t)} \left[\log X_{i,h,m} - \boldsymbol{\zeta}'_{i,h,m}\theta\right]^2, \qquad \widehat{u}_{i,h,m}(t) = \log X_{i,h,m} - \boldsymbol{\zeta}'_{i,h,m}\widehat{\theta}_{i,h}(t). \tag{30}$$

Exponentiating a forecast from a log equation understates the variance level, so we use the smearing correction of Duan (1983) to repair the bias and recover the daily forecast

$$\widehat{P}_{i,h,t} = \exp\left(\boldsymbol{\zeta}'_{i,h,t}\widehat{\theta}_{i,h}(t)\right) \left[\frac{1}{|\mathcal{M}_{i,h}(t)|} \sum_{m \in \mathcal{M}_{i,h}(t)} \exp\{\widehat{u}_{i,h,m}(t)\}\right]. \tag{31}$$

Coefficients update only when months resolve, but forecasts are issued every trading day; this avoids the problem of overfitting associated with overlapping daily windows. A forecast is issued once $|\mathcal{M}_{i,h}(t)| \geq \max\{48,\ 4p_{i,h}\}$, where $p_{i,h} = \dim(\theta_{i,h})$.

5.6 Policy application: the RMA volatility factor and its counterfactual

The USDA RMA offers revenue insurance to domestic producers in order to protect them from revenue shortfalls due to price or yield. For most corn and soybean policies, it does so using two futures-market inputs: (i) it sets the projected price following futures prices during a February discovery window, and sets the harvest price from a second window in October; and (ii) it draws a volatility factor from near-the-money option-implied volatility on the final five trading days of the February window. The RMA volatility factor, then, contains both the expected physical risk and the VRP (Woodard and Sproul, 2016).[21] Our objective is to understand how the latter affects the cost of revenue insurance to producers.

For this step, we rely on Annual Commodity Exchange Price Provisions that designate each crop year's contracts and discovery dates, and RMA's published methodology 2015; 2026. Let $\mathcal{D}_y$ contain the final five trading days of the February discovery period for crop year y, and let $\sigma^{\text{RMA,ann}}_{i,d,y}$ denote the annualized implied volatility underlying RMA's published time-adjusted observation for crop i. We define the fraction of the year remaining in the

[21] Because RMA builds its factor from near-the-money implied volatility rather than a model-free strip, its input carries a wedge similar to ours—which is why subtracting our ATM-based premium estimate from the agency's ATM-based factor is the right step.

insurance horizon as

$$a_{d,y} = \frac{N_{d,y}}{L_y}, \tag{32}$$

where $N_{d,y}$ is the number of calendar days from discovery date d to October 16, the midpoint of the harvest-price discovery period, and L_y is 365 or 366. The daily contribution to the volatility factor is $\sigma_{i,d,y}^{\text{RMA,ann}}\sqrt{a_{d,y}}$. The unrounded and reported factors are therefore:

$$\widetilde{VF}_{i,y}^{\text{RMA}} = \frac{1}{5}\sum_{d\in\mathcal{D}_y}\sigma_{i,d,y}^{\text{RMA,ann}}\sqrt{a_{d,y}}, \qquad VF_{i,y}^{\text{RMA}} = \mathcal{R}_{0.01}\left(\widetilde{VF}_{i,y}^{\text{RMA}}\right), \tag{33}$$

where $\mathcal{R}_{0.01}(\cdot)$ rounds to the nearest hundredth.

Our forecasts are produced using fixed-tenor windows and can use a different named contract from the harvest contract RMA uses to produce its volatility factor. We therefore do not use the historical forecast to produce the policy result. Instead, we report an explicitly illustrative fixed-tenor translation that determines how the typical corn and soybean factor would change if the market input carried no average premium.

Let the forecast-free average ATM variance gap at tenor h be

$$\overline{g}_{i,h} = \frac{1}{T_{i,h}}\sum_{t=1}^{T_{i,h}}\left(q_{i,h,t}^{\text{ATM}} - X_{i,h,t}\right). \tag{34}$$

This estimates the average VRP (plus the average ATM measurement wedge). We use $h = 60$ for the headline calculation and $h = 30, 90$ for sensitivity analysis. Because the RMA input is expressed as volatility, we first recover its annualized variance, $q_{i,d,y}^{\text{RMA,ann}} = (\sigma_{i,d,y}^{\text{RMA,ann}})^2$, subtract the estimated gap in variance units, and only then convert the result back to the volatility units used by RMA. The calibrated physical-only factor is:

$$\widetilde{VF}_{i,y,h}^{P,\text{cal}} = \frac{1}{5}\sum_{d\in\mathcal{D}_y}\sqrt{a_{d,y}\max\left\{q_{i,d,y}^{\text{RMA,ann}} - \overline{g}_{i,h}, 0\right\}}. \tag{35}$$

Thus, following RMA's procedure, we take the square root separately for each discovery day before averaging the five observations, and run 1000 paired 65-trading day block-bootstrap draws of $\overline{g}_{i,h}$ through the same nonlinear calculation. The resulting intervals reflect uncertainty in the estimated average gap conditional on the observed RMA inputs.

To illustrate how the corn and soybean factor results can affect insurance premiums, we use the case-farm calculations from Sherrick (2015). Sherrick considers a McLean County, Illinois, corn farm with a trend-adjusted APH of 180 bushels, a projected price of $4.15, and a volatility factor rising by 10 percent, from 0.20 to 0.22. His Table 3 reports the resulting

percentage premium change m_j for each combination j of RP or RP-HPE, coverage level, and optional, basic, or enterprise units. We convert each finite difference into a log arc elasticity,

$$\eta_j = \frac{\log(1 + m_j/100)}{\log(1.10)}, \tag{36}$$

and apply that elasticity to the relative difference between the mean observed and calibrated physical-only factor for each crop, calculating the percentage difference in the policy premium due to including the estimated volatility risk premium:

$$\Delta\text{Premium}_{i,j} = 100 \times \left[\left(\frac{\overline{VF}_i^{\text{RMA}}}{\overline{VF}_{i,60}^{P,\text{cal}}} \right)^{\eta_j} - 1 \right], \qquad i \in \{\text{corn}, \text{soybeans}\}. \tag{37}$$

For corn, this is an illustrative transfer of the premium sensitivity from Sherrick's case farm. For soybeans, the factor change is estimated separately from the soybean data, but its premium effect is calculated using Sherrick's corn sensitivity schedule (because he does not report corresponding numerical calculations for soybeans). The soybean calculation is therefore a more approximate illustration.

5.7 Using fixed tenor implied volatility and its implications

Our measure q^{ATM} is built from a single strike. The ideal q in (3) exploits the whole options strip, so the two differ by the wedge ν. For the average premium, the mean of $q^{\text{ATM}} - X$ is $\mathbb{E}[\text{RP}] + \mathbb{E}[\nu]$, so we recover the economic premium only if the wedge averages to zero. For a state coefficient or trend,

$$\beta_{q^{\text{ATM}} - X} = \beta_{\text{RP}} + \beta_\nu - \beta_\varepsilon, \tag{38}$$

where β_ν and β_ε come from the same projection applied to the wedge and to the realization surprise. Equation (5) sets $\beta_\varepsilon = 0$ whenever every time-varying regressor is known at t. A constant, therefore, is soaked up by the intercept and means but leaves slopes unaffected; a time-varying wedge contaminates a coefficient only if it correlates with that regressor.

Its actual effect is an empirical question about the implied volatility smile—the pattern of implied volatility across strikes. If the smile steepens when markets are stressed, q^{ATM} understates the strip precisely in those states, and our VIX and hedging-pressure coefficients would blend the premium with the wedge. Constructing model-free implied variance from full option chains would remove ν. In future work, we intend to do so.[22]

[22] A similar wedge affects the RMA factor, since the administrative input is itself a near-the-money implied volatility. Subtracting our ATM-based premium estimate from the agency's ATM-based factor is therefore the right comparison.

5.8 Inference

Option-implied variance windows overlap from one date to the next, and corn and soybeans respond to common shocks. All inference therefore uses a wild block bootstrap (Cameron et al., 2008).[23] We use 1,000 draws, 65-session blocks for state regressions, and 125-session blocks for local projections. The outer shaded local-projection bands are 95 percent pointwise percentile confidence intervals, formed from the empirical 2.5th and 97.5th percentiles of the bootstrap coefficient draws; the inner shaded bands are the corresponding approximately 68 percent pointwise intervals.

For the historical decomposition, each replication rebuilds the eligible set of monthly windows, re-estimates $\widehat{\theta}_{i,h}$ and the smearing factor, and only then produces the daily path; a single full-sample refit would not preserve the recursive design. For the RMA calibration, each paired 65-trading day block-bootstrap draw of the sample-average gap is passed through the daily square root, RMA time adjustment, and five-day average, holding the observed RMA inputs fixed.

6 Data

Our data provider, LSEG, reports daily call and put ATM implied volatility at fixed 30-, 60-, and 90-day tenors for corn and soybeans. We construct annualized ATM implied variance using equation (7). The raw series begins in June 2010. We reserve observations through December 2013 for initial forecast estimation and variable construction and report results from January 2014 onward. We summarize the variance curve using $q^{\mathrm{ATM}}_{i,60,t}$, $q^{\mathrm{ATM}}_{i,30,t} - q^{\mathrm{ATM}}_{i,60,t}$, and $q^{\mathrm{ATM}}_{i,90,t} - q^{\mathrm{ATM}}_{i,60,t}$. Because the available LSEG metadata do not establish exact nesting across tenors, we interpret the differences as curve slopes rather than exact forward-variance buckets.

We construct realized variances from named-contract futures settlements. At each date and horizon, we select one contract using the 30-day expiry-buffer rule in Section 3 and follow that contract throughout the realized-variance window. We require a valid settlement price for every return in the window, so a single missing or non-positive price drops that observation. We do not interpolate missing settlements, splice contracts, or use continuation-series

[23] Specifically, each replication rescales the estimated residuals rather than drawing new observations. Two adjustments make that rescaling respect the data. First, the HC2 correction inflates residuals from observations distant from the center of the regressors, since those observations pull the fit toward themselves and produce residuals that understate the true dispersion (MacKinnon and White, 1985). Second, the sample is sorted into calendar blocks, and each replication draws one random multiplier per block and applies it to every residual in that block. Both crops and all tenors within a block share the multiplier, so observations that move together in the observed sample move together in the artificial one, and their dependence survives resampling.

returns across a roll. Each crop follows its own exchange calendar, while joint specifications use dates common to both crops. Recursive forecasts use only outcomes whose windows had closed before the forecast date.

Table 2 reports the resulting sample sizes. The forecast-free analysis uses 2,934 common dates with complete outcomes for both crops at all three tenors. The joint state-variable analysis is limited by the shorter state-data panel and complete availability of the lagged CFTC and climate-news variables, leaving 2,198 common dates per crop. The recursive forecast sample is composed differently because it requires sufficient resolved training observations at each forecast date.

Table 2: Construction and sample accounting

Sample or restriction	Corn	Soybeans
Matched 60-day outcomes		
Daily observations with call and put ATM IV	3,846	3,845
Also matched to a valid named contract	3,804	3,803
Complete matched ATM and realized variance	3,723	3,722
State-variable sample		
Common all-crop, all-tenor sample from January 2014	2,934	2,934
Within state-data coverage through May 12, 2023	2,209	2,209
Lagged CFTC variables available	2,200	2,200
Complete PRI and TRI news innovations	2,198	2,198
Recursive forecast sample		
Scored 60-day forecasts with resolved outcomes	2,214	2,193

Notes: Counts are observations per crop. The common matched sample spans January 2, 2014–April 24, 2026. The state sample begins January 15, 2014 after conservative CFTC timing and ends May 12, 2023; two missing PRI/TRI dates leave 2,198 observations. Scored forecasts use only outcomes resolved before each forecast date.

6.1 State variables and news innovations

In addition to a linear trend and annual and semi-annual sine and cosine terms, the headline joint specification includes the following state variables. We construct a carry factor from the nearest four futures contracts for each commodity. For each adjacent pair, we calculate the annualized log-price slope $s_{j,j+1,t} = \log(F_{i,j+1,t}/F_{i,j,t})/[(T_{i,j+1} - T_{i,j})/365.25]$, for $j = 1, 2, 3$. We standardize these three slopes, extract their first principal component separately by crop using only 2010–2013 observations, and then freeze the centering, scaling, and loadings. The component is signed so that larger values indicate greater contango. The regression uses its value from the prior exchange trading day.

The report variables count WASDE and Grain Stocks publication dates strictly after the option origin and on or before the end of the outcome window, $(t, t+h]$. USDA release dates are published ahead of time and known to market participants; these variables measure exposure to scheduled information releases, not report surprises. Critical-period overlap is the share of calendar days in the outcome window falling between July 10 and July 25 for corn, or July 8 and August 13 for soybeans. We fix these windows using 2010–2013 USDA NASS Crop Progress data: the corn window runs from 25 to 75 percent silking, and the soybean window from 25 percent blooming to 75 percent setting pods. The Oceanic Niño Index is NOAA's three-month Niño 3.4 sea-surface-temperature anomaly. We use the latest available value, enter it on the first day of the second month after the corresponding three-month season ends, and carry it forward until the next value becomes available.

Trader position variables are drawn from the CFTC's weekly Supplemental Commitments of Traders report, which combines futures and options and identifies commodity-index traders separately. Commercial hedging pressure is (commercial short positions excluding CIT − commercial long positions excluding CIT)/total open interest; a larger value therefore indicates greater net short commercial pressure. The commodity-index-trader variable is CIT long positions divided by the same total open-interest denominator. The COT reports Tuesday positions the following Friday, but the CFTC does not archive actual release dates and holidays shift the schedule. To be sure that positions data are known, we therefore wait seven days after the Tuesday reference date before letting either variable enter a regression. Each value is then carried forward daily until the next release becomes available.

The VIX variable is the CBOE closing VIX value from the latest trading day strictly before the option origin. The physical- and transition-risk news variables are innovations in the daily Reuters-based Physical Risk Index and Transition Risk Index of Bua et al. (2024). We construct them as residuals from separate AR(1) models for physical- and transition-risk news concern, assign each complete news day to the first crop trading day strictly afterward, and sum innovations when multiple news days map to the same trading day.

CarryPC1, the commercial hedging pressure, and the CIT long share are standardized separately within crop over the 2014–2023 state panel. The ONI and VIX are standardized using the pooled state-panel moments. PRI and TRI innovations are standardized using their 2005–2013 daily standard deviations before being mapped to crop trading days. Report counts remain in units of releases, while critical-period overlap is divided by 0.10, so its coefficient represents a ten-percentage-point increase. The common complete-case sample contains 2,198 observations per crop and tenor from January 15, 2014, through May 12, 2023.

6.1.1 Dynamic effects data

We construct WASDE and Grain Stocks publication indicators from the official USDA release calendars and include their joint-release interaction. Because release dates are publicized well in advance, we study both the anticipation of a scheduled report as well as its residual effect. Event-time zero is the official crop session containing the release. At event time zero, the option quote incorporates the report information, while the subsequent-realized-variance window begins after that session.[24]

USDA news innovations difference the pre-report Bloomberg analyst-level survey median against USDA actuals. For WASDE, the reported quantity is domestic ending stocks; for Grain Stocks, it is quarterly stocks. A "positive surprise" means "tighter-than-expected stocks". We divide the surprise by its fixed pre-2014 crop-and-report standard deviation. The report regressions also control for the survey's interquartile range and number of contributors, other-crop news, and coincident USDA releases; the Grain Stocks specification includes Prospective Plantings or Acreage surprises on joint release dates. Because these specifications use the subsequently observed signed surprise, their plotted responses begin strictly after release and contain no pre-release coefficients.

The Adjemian and Jo (2025) agricultural-supply instrument is the production-weighted principal component of event-day futures returns for corn, soybeans, wheat, rice, and oats around USDA supply announcements. We normalize it by its fixed pre-2014 standard deviation and use it to instrument the event-day return on the nearby named futures contract for the crop being studied. The Känzig (2021)'s instrument is the principal component of WTI futures returns across maturities on OPEC announcement dates, again normalized by its fixed pre-2014 standard deviation. Its endogenous regressor is the front-contract WTI return.[25]

For monetary policy, we use the Bauer–Swanson surprise around scheduled FOMC announcements (Bauer and Swanson, 2023). It is the first principal component of 30-minute changes in money-market futures covering the next four quarters. We express the surprise in ten-basis-point units and include Bauer and Swanson's six pre-announcement information variables, thereby performing their orthogonalization within the local projection.

Our weather innovation is sourced from the NOAA Climate Prediction Center's (CPC)

[24] USDA moved most major releases from 8:30 a.m. to noon Eastern on January 11, 2013, so an 8:30 report published before the trading session opened while a noon report is released mid-session. Our analysis sample begins in January 2014 and is therefore entirely under the current convention; the change affects only the 2010–2013 period used to fix pre-sample standard deviations.

[25] The exclusion restriction in each LP-IV specification is that the instrument affects subsequent implied and realized variance through the instrumented event return rather than through other news arriving in the event-day window.

6–10 and 8–14-day probabilistic outlooks. We aggregate station-level probabilities using crop-area weights and define the hot index as the probability of above-normal minus below-normal temperature, and remove the component predicted by a model fixed on pre-2014 outlooks. The resulting residual is scaled by its fixed pre-2014 standard deviation.

The agricultural-supply, oil-supply, and monetary-policy responses include the shock session. Weather and news-attention innovations are assigned to the first crop trading day strictly after their source date, so their initial reactions are excluded from the response outcomes. In every case, the innovation is in the information set when the subsequent-realized-variance window opens, satisfying forecast-error orthogonality.

7 Results

7.1 Level regressions: average premia, trends, and factor relationships

Table 3 reports the results for the 60-day ATM tenor; matching results for the 30- and 90-day tenors appear in Appendix Tables B1 and B2. We feature the 60-day results as our headline findings. Because of our exact decomposition, the option-implied coefficient equals the sum of the corresponding coefficients from the subsequent-realized-variance and price-of-protection equations, by commodity. Adjusted R^2 values do not add, so we report them only as equation-specific measures of fit. Our modeled variables explain a large share of option-implied variance (after accounting for degrees of freedom), rising from 65% at the 30-day tenor to 76% at the 90-day. Virtually all of the eighteen models in the three tables strongly reject the null hypothesis of a unit root in the residual according to the Phillips–Perron test; the only exception is 90-day tenor soybean option-implied variance, although this is rejected at the 8% level. Coefficient values in the tables are reported in variance points—annualized variance multiplied by 100. Over the sample period, corn implied variance averages 5.54 variance points while realized same-contract variance averages 4.91, leaving a positive average premium of 0.63 points. Over the sample, corn's VRP represents about 11 percent of its implied variance. It is positive and statistically distinguishable from zero at five percent at every tenor: 0.50 points at 30 days, 0.63 at 60 days, and 0.70 at 90 days. Comparable soybean estimates are -0.17, -0.14, and -0.08, and all are statistically indistinguishable from zero.

That corn has a positive price-of-protection component in our data, on average, while soybeans do not could be for several reasons. As the largest domestic crop by acreage and insured liability, corn may attract demand for protection from both sides of the market:

producers from falling prices, and feeders and ethanol plants from rising ones. The ethanol plant's position is the sharpest case. It buys corn and sells fuel priced off gasoline, so a corn rally comes out of its margin with nothing in its own output to offset it, while a soybean crusher sells meal and oil at prices that move with the bean price and has far less left to hedge. The dealers and speculators who take the other side of both trades hold variance that they cannot diversify away, and so they charge for holding it. In addition, corn resolves most of its yield uncertainty in the days around pollination in July, generating significant revisions to supply expectations. Physical variance that arrives so abruptly can be expensive to hedge. Soybean yields develop over a longer stretch in August, and much of their physical variance comes from export and trade news (with China as the single dominant buyer in the market) rather than from weather. That arrival of that demand-side news is unscheduled and not as predictably seasonal, so the market may not price it into a particular contract ahead of time. When it lands, realized variance can exceed what was charged, which would match our average premium findings.[26]

None of the linear trend coefficients in Table 3 is statistically distinguishable from zero, and all of the point estimates are negative. The 30- and 90-day estimates lead to the same conclusion. We interpret these results to indicate that the sample we consider does not support a sustained, possibly climate-driven, rise in expected physical variance or the price to guard against it. The regression-residual Phillips–Perron tests reject a unit root at ten percent for all 18 crop–tenor–outcome regressions and at five percent for 17 of them; these tests are reported beneath the regression R^2 values in tables 3, B1, and B2.

Contango markets—where later expiring futures contracts are priced higher than nearer deliveries—are associated with lower implied and expected physical variance. For corn, a one-standard-deviation increase in the carry factor (the first principal component of the futures constellation slope) implies 1.74 and 2.59 fewer points of implied and physical variance, respectively. The positive 0.85-point VRP makes up the difference but isn't statistically significant. For soybeans, implied variance falls by 0.42 points and expected physical variance by 0.45 points, though the latter is not significant. Carry is closely related to storage adequacy, so these results are consistent with our model's storage-buffer feature.

Approaching government reports tend to raise option-implied variance, and they operate through expected physical variance. (a) A single Grain Stocks report in the forward 60-day window is associated with an additional 0.58 implied variance points for corn, and 1.03 points of higher expected physical variance. For soybeans, the expected physical variance rises by 0.63 points. Corresponding VRP coefficients are imprecisely estimated in both cases. (b) WASDE reports have a similar impact on corn: they raise implied variance by 0.74 points

[26] We intend to test these explanations empirically in future work.

Table 3: Option-implied variance, realized variance, and the premium: 60-day tenor

	Corn			Soybeans		
	q^{ATM}	X	$q^{ATM} - X$	q^{ATM}	X	$q^{ATM} - X$
Panel A: Full-sample moments						
Mean	5.537***	4.912***	0.625**	3.516***	3.653***	-0.137
	(0.409)	(0.464)	(0.251)	(0.184)	(0.270)	(0.161)
Linear trend, per year	-0.021	-0.005	-0.016	-0.065	-0.043	-0.022
	(0.085)	(0.095)	(0.060)	(0.044)	(0.067)	(0.042)
Full-sample observations	2,934	2,934	2,934	2,934	2,934	2,934
Panel B: Joint state specification						
Carry factor (one SD)	-1.742***	-2.591***	0.849	-0.416**	-0.450	0.033
	(0.387)	(0.676)	(0.543)	(0.209)	(0.366)	(0.241)
WASDE reports in target window	0.737**	0.796***	-0.059	0.124	0.292	-0.167
	(0.306)	(0.306)	(0.367)	(0.179)	(0.191)	(0.247)
Grain Stocks reports in target window	0.580**	1.032***	-0.452	0.138	0.625**	-0.487
	(0.237)	(0.384)	(0.374)	(0.177)	(0.305)	(0.316)
Critical-window overlap (10 pp)	0.882***	-0.086	0.968*	0.239**	0.288*	-0.050
	(0.314)	(0.435)	(0.534)	(0.121)	(0.166)	(0.175)
ENSO state (one SD)	0.202	0.301	-0.098	-0.106	0.084	-0.191
	(0.181)	(0.256)	(0.203)	(0.218)	(0.362)	(0.234)
Commercial hedging pressure (one SD)	1.296***	0.583	0.713**	0.437*	0.530*	-0.092
	(0.294)	(0.369)	(0.316)	(0.245)	(0.299)	(0.167)
CIT long share (one SD)	-0.257	-0.663*	0.406	-0.475*	-0.368	-0.107
	(0.270)	(0.392)	(0.315)	(0.278)	(0.362)	(0.213)
Lagged VIX (one SD)	0.576***	0.164	0.412***	0.413**	0.241	0.172
	(0.160)	(0.172)	(0.158)	(0.161)	(0.225)	(0.125)
Physical-risk news (one SD)	0.050	0.025	0.026	-0.003	0.004	-0.007
	(0.051)	(0.045)	(0.066)	(0.039)	(0.044)	(0.036)
Transition-risk news (one SD)	-0.016	-0.005	-0.011	0.038	0.082	-0.045
	(0.048)	(0.058)	(0.062)	(0.034)	(0.064)	(0.050)
Fourier seasonality and linear trend	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.718	0.622	0.178	0.435	0.444	0.196
Regression residual Phillips–Perron p-value	0.002	0.006	0.002	0.030	0.009	0.001
Joint-sample observations	2,198	2,198	2,198	2,198	2,198	2,198

Notes: Outcomes are annualized variance multiplied by 100. Panel A uses 2,934 common dates from January 2, 2014–April 24, 2026; Panel B uses 2,198 dates from January 15, 2014–May 12, 2023. All three Panel B equations use identical rows and regressors, so coefficients add exactly up to rounding. Standard errors use 1000 paired 65-session calendar-block wild-bootstrap draws. The regression-residual Phillips–Perron row tests the null that the fitted residual from the exact Panel B equation has a unit root. Its auxiliary residual AR has no deterministic terms because Panel B already includes the intercept, linear trend, Fourier terms, and states. The p-value is calibrated by 1000 null-imposed 65-session calendar-block wild-bootstrap paths: centered residual increments are sign-flipped within blocks, cumulated separately across contiguous official-session segments, and the full Panel B equation is re-estimated in every draw. The calculation never links observations across a session gap and uses overlap-matched Newey–West bandwidths of 21, 43, and 62 sessions at 30, 60, and 90 days. Small Phillips–Perron p-values reject the null of a unit root in the residual of the Panel B level regression. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

and expected physical variance by 0.80 points. These findings are consistent with the model's information-arrival feature: an approaching, known physical signal inside the option horizon raises expected physical variance, not the price of protection.

In contrast, hedging pressure and aggregate financial stress tend to affect implied variance through the price of protection—statistically significantly so for corn. A one-standard-deviation increase in commercial hedging pressure raises the commodity's implied variance by 1.30 points. Of that increase, 0.71 points appear in the VRP. And a one-standard-deviation increase in the VIX raises implied variance by 0.58 points, with 0.41 points appearing in the premium. In both cases, the residual physical-variance increase is imprecise. The VIX result is robust across all three tenors and all reviewed specifications. These empirical results are consistent with the risk-bearing side of the theoretical model. More commercial hedging is associated with a larger quantity of corn risk the market must absorb, while an elevated VIX and overall market uncertainty indicates that risk is becoming more costly to bear. For soybeans, VIX increases also raise option-implied variance, although the premium coefficient is not statistically distinguishable from zero. Moreover, while commercial hedging pressure likewise increases that crop's option variance, it appears to do so through the expected physical channel (both coefficients are significant at the 10% level). The crush asymmetry discussed earlier is a potential explanation here as well. A soybean crusher carries little residual exposure to bean-price variance but still knows a great deal about the crop, so its positions signal the conditions that set c and s in (10) without adding to net demand for protection. In contrast, corn hedging raises d and, through (13), the premium.

Similar to hedging pressure, the critical-growing-window overlap does not produce the same results across commodities. A 10-percentage-point increase in overlap raises corn implied variance by 0.88 points, but the physical coefficient is near zero and the 0.97-point premium coefficient is only marginally significant. For soybeans, the implied coefficient is 0.24 but its physical variance coefficient is marginally significant at 0.29. These split results are likely down to how a calendar overlap is defined: it marks a relatively long window in which weather can still generate significant surprise, USDA reports arrive, and producers want protection; so the state regression cannot separate them. Soybeans allocate the response to expected variance, corn allocates them to premium. Because a single calendar variable cannot fairly tell whether the crops face different biology or different demand for protection, our dynamic designs below test the physical channel directly. CIT positions, ENSO, and the physical- and transition-risk news states in table 3 do not produce comparably stable premium effects.

7.2 Dynamic regressions

7.2.1 Scheduled government reports: timing and content

Our event-time results in figures 2 and 3 distinguish anticipation of a known USDA publication from its resolution on the publication day.[27] Implied variance rises as WASDE and Grain Stocks approach and falls on the publication day. The decline is predominantly allocated to subsequent physical variance, especially for Grain Stocks. The Grain Stocks point estimates also show a higher price-of-protection residual that offsets part of the physical-variance decline in the option quote.

7.2.2 Identified news

Figure 4 reports estimates for the effect of a one-percentage-point news-induced own-crop futures return on the components of implied variance for both commodities, estimated by instrumental-variables local projection. Poor agricultural-supply news raises option variance in both commodities on the shock day. For corn, the immediate response is explained by a higher price-of-protection before a change in subsequent physical variance is visible. Soybeans show the same positive implied-variance response, but the premium allocation is less precise. The Adjemian and Jo (2025) instrument is strong in both crops: Table C2 reports first-stage F statistics of 1,685.40 for corn and 92.80 for soybeans on 122 events.

A regression with total implied variance as its only outcome would miss the key part of this result. In our model, supply news can change the price of bearing the remaining risk before the amount of physical variance expected after the shock session changes measurably. The shock-day corn response is consistent with that distinction. Unexpectedly hot CPC forecasts during the critical crop window work the other way; see Figure 5. The first plotted session is strictly after the forecast issue date. For soybeans, the innovation is associated with higher subsequent physical variance rather than a change in the price of protection. Corn shows the same mean allocation, but it is imprecisely estimated. Hot conditions stress crops and can lead to yield losses. The premium response is not distinguishable from zero. A hot growing-season forecast therefore changes the expected distribution of future crop prices rather than the compensation required to bear it.

Oil-supply and monetary-policy news appear to operate through different channels. (i) Figure 6 shows that a positive WTI return induced by the Känzig (2021) oil-supply-news instrument has no precise shock-day effect on grain and oilseed variance. Beginning roughly 30 trading sessions later, subsequent physical variance rises for both crops, most clearly for

[27] To estimate these dynamic effects, we use local projections with outcome variance points calculated using the 60-day options tenor.

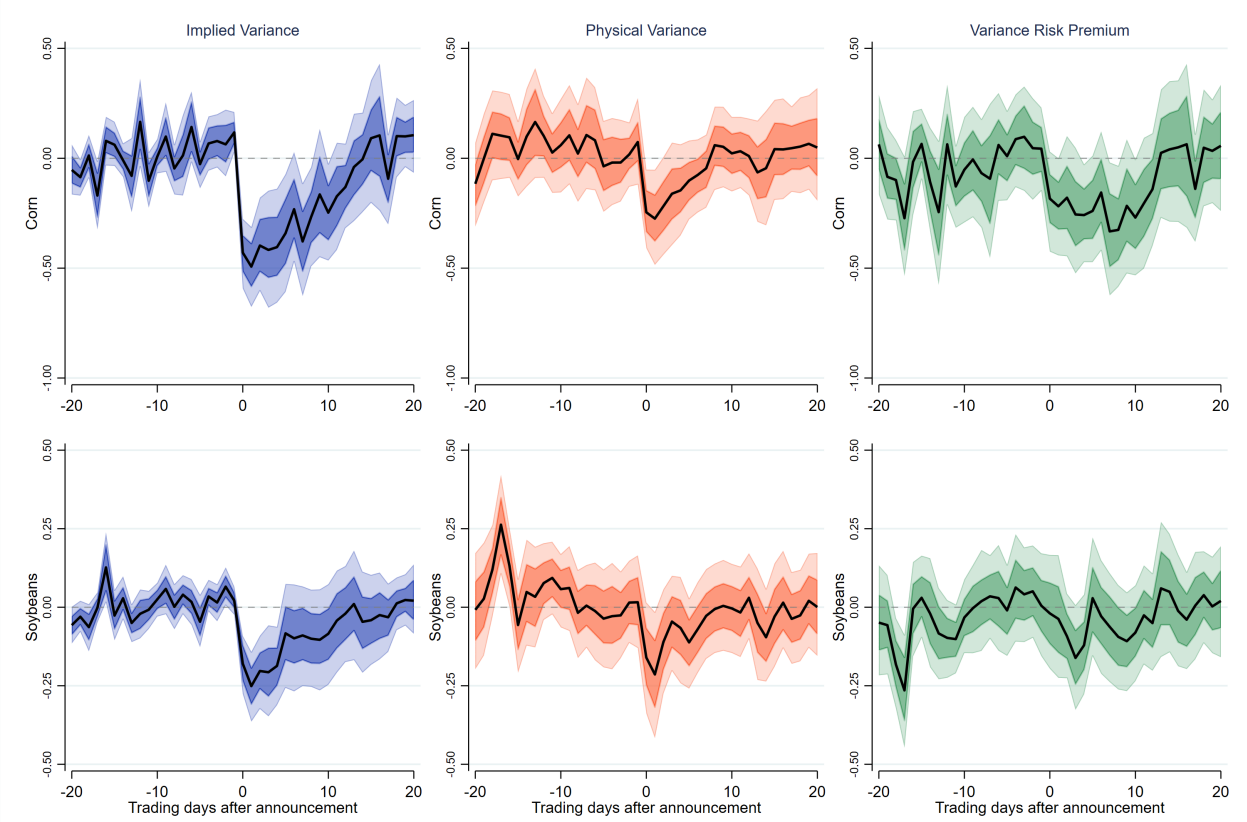


Figure 2: WASDE dummy impact response around the report date at the 60-day tenor

Notes: Rows report corn and soybeans; columns report ATM implied variance, subsequent realized variance, and their difference, the variance-premium measure. WASDE and Grain Stocks dummy variables and their joint-release interaction enter together; the displayed path is the standalone WASDE response. Negative event times describe the pre-release path associated with the known scheduled occurrence. Event time zero is the official crop session containing the release. The option quote incorporates the report information, while the subsequent-realized-variance window begins after that session. Shaded bands are 95 and approximately 68 percent pointwise percentile confidence intervals from 1,000 wild-block bootstrap draws using 125-session blocks. These are descriptive timing responses, not responses to signed report content.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities, Databento CBOT named-futures settlements, and USDA WASDE and Grain Stocks release calendars.

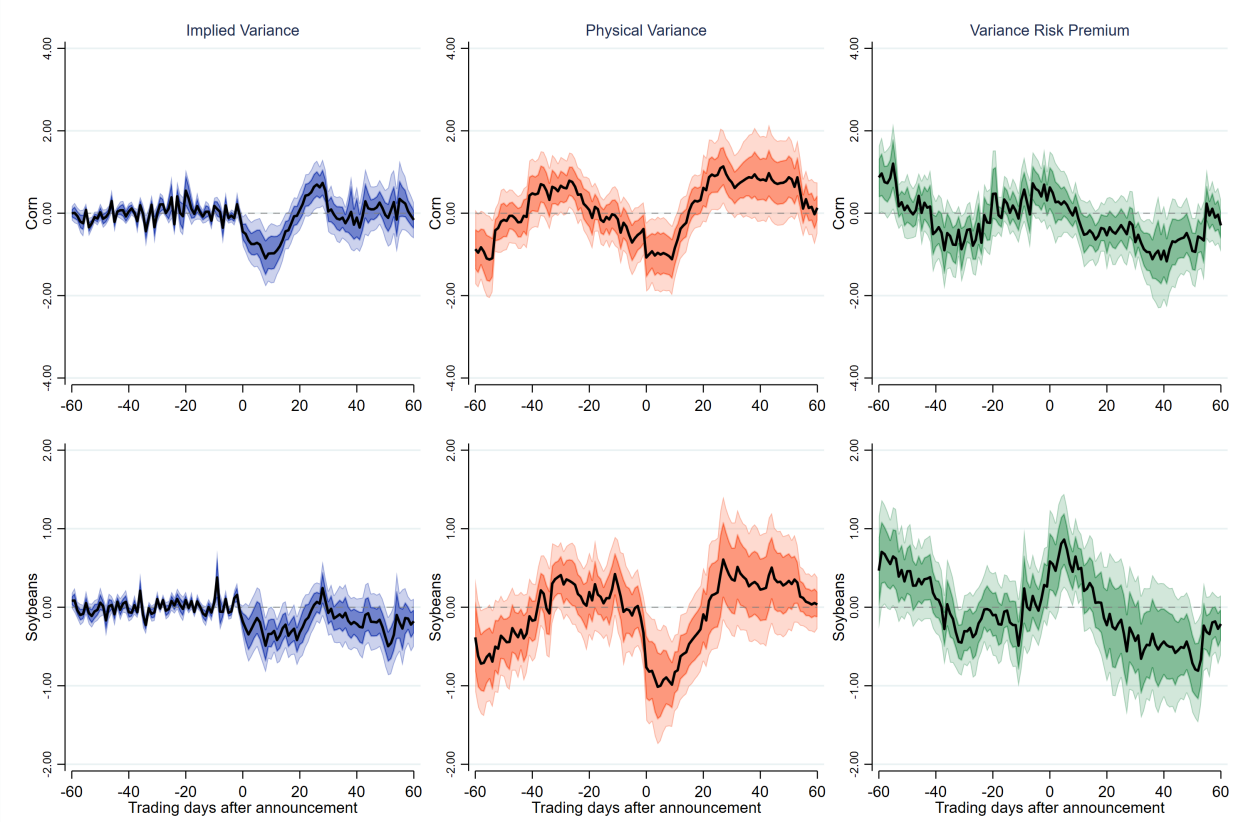


Figure 3: Grain Stocks dummy impact response around the report date at the 60-day tenor

Notes: Rows report corn and soybeans; columns report ATM implied variance, subsequent realized variance, and their difference, the variance-premium measure. WASDE and Grain Stocks dummy variables and their joint-release interaction enter together; the displayed path is the standalone Grain Stocks response. Negative event times describe the pre-release path associated with the known scheduled occurrence. Event time zero is the official crop session containing the release. The option quote incorporates the report information, while the subsequent-realized-variance window begins after that session. Shaded bands are 95 and approximately 68 percent pointwise percentile confidence intervals from 1000 wild-block bootstrap draws using 125-session blocks. These are descriptive timing responses, not responses to signed report content.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities, Databento CBOT named-futures settlements, and USDA WASDE and Grain Stocks release calendars.

corn. Implied variance rises by less, so the price-of-protection residual falls at the point estimates. Higher energy prices can change production, transportation, and biofuel profitability, making a delayed response in the crop-price distribution plausible. (ii) Monetary tightening produces nearly the opposite breakdown in corn. Figure 7 shows that, for an unexpected monetary tightening on the FOMC announcement day, expected physical variance falls while the price of protection rises. The sum of two movements leave the implied variance much less responsive than either of its two components, individually. The pattern persists for several weeks in the point estimates. At least for corn, this result fits the model's risk-bearing mechanism: tighter financial conditions raise the compensation required to supply protection even when the amount of expected physical variance falls. Soybean responses do not exhibit as clear a pattern.

To study the dynamic effects of climate-related news, we include Bua et al. (2024)'s news-attention measures in the dynamic analysis. Those authors build two global indices from Reuters articles, a *physical risk index* based on physical hazards (like extreme weather, drought, sea levels, desertification, and ocean temperature changes), and a *transition risk index* oriented to regulatory, technology and government policy. Each index scores how closely a day's news resembles a reference document assembled from scientific and authoritative climate texts. It is a similarity measure, not a count or a sentiment score, so a higher value means the day's coverage looked more like the relevant portion of climate vocabulary—not necessarily that conditions worsened. Figure 8 shows that an unexpected increase in physical-risk attention is followed by lower implied variance in both crops. For corn, the decline channels are the expected physical variance through the first thirty sessions and the price of protection thereafter; for soybeans, the channel is expected physical variance throughout, with the premium effect flat and not distinguishable from zero. Because physical-risk attention is not physical damage, more attention can resolve uncertainty as well as signal worsening conditions, so the negative response that we estimate does not contradict our model. Transition-risk attention works in the other direction, empirically. Figure 9 displays higher implied variance for both crops following a TRI innovation, although the shape of the crop variance response differs.

Soybean expected physical variance is higher on impact, while its implied variance takes about twenty-five sessions to catch up—so the measured premium is negative over the first month and returns to zero thereafter. Corn reverses that order: implied variance rises first, and expected physical variance follows only at longer horizons. In neither crop does the price of protection show a persistent increase, so transition-policy discussion appears to change the expected distribution of future crop prices rather than the compensation required to bear it. Because the index measures news attention rather than policy stringency, these remain

conditional associations rather than identified policy effects.

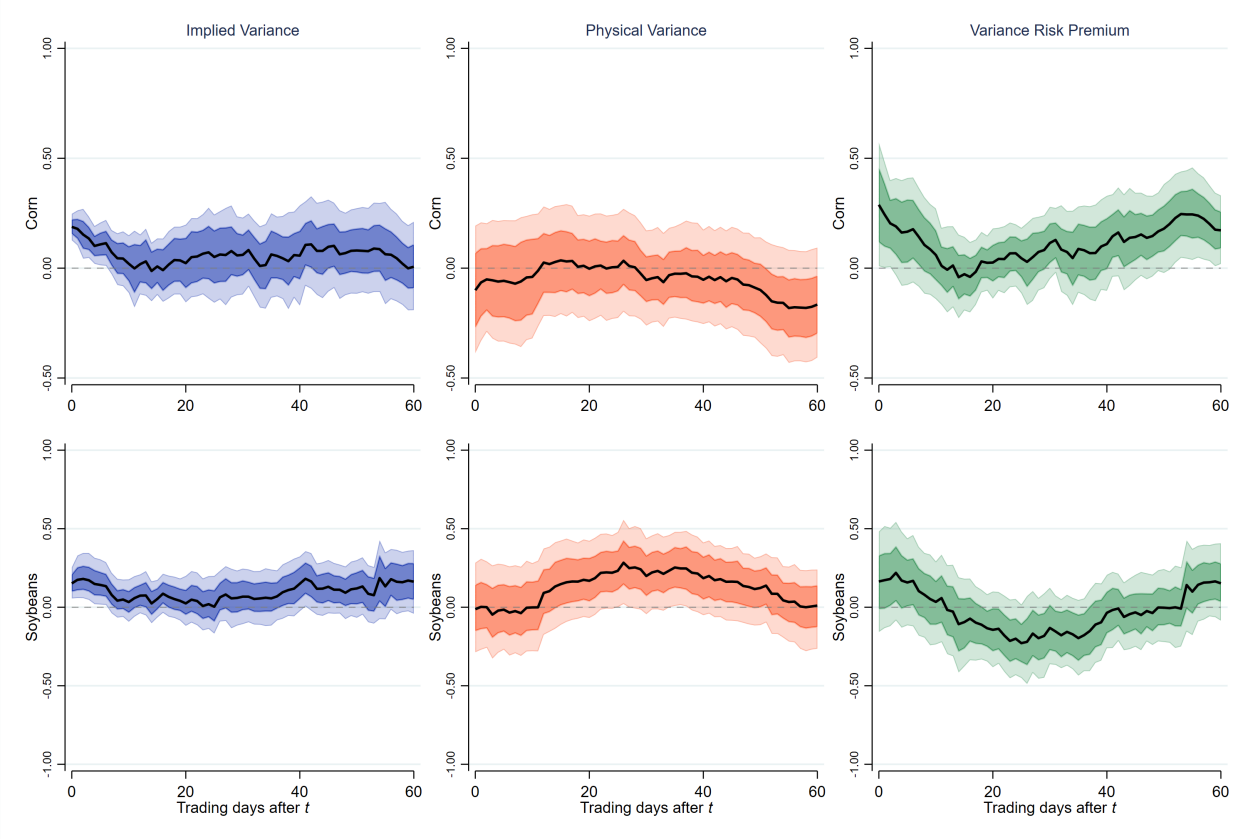


Figure 4: Agricultural-supply-news LP-IV at the 60-day tenor

Notes:

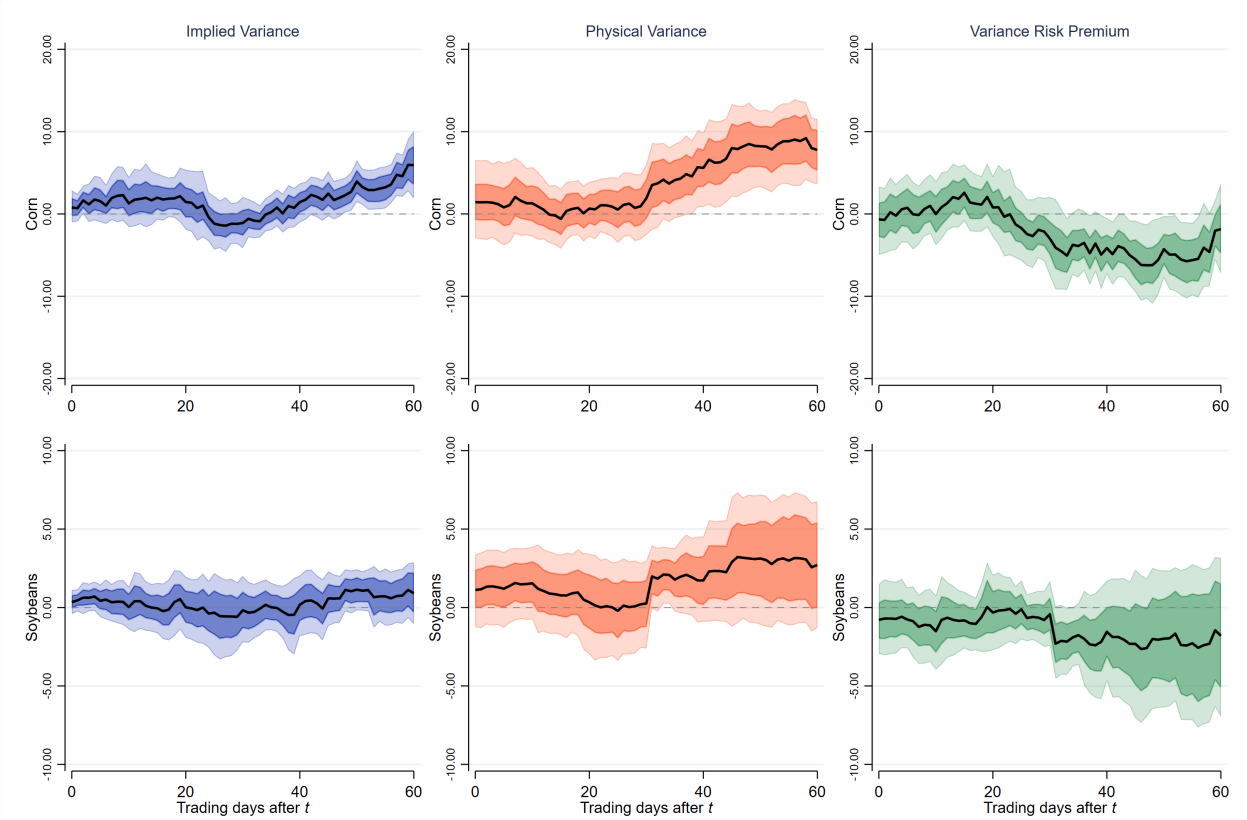


Figure 6: Oil-supply-news LP-IV response at the 60-day tenor

Notes: Rows report corn and soybeans; columns report ATM implied variance, subsequent realized variance, and their difference, the variance-premium measure. Responses are annualized variance points per 10-percent news-induced front-month WTI futures return. The exercise contains 37 events and is treated as appendix evidence. Event time $t = 0$ is the official crop session on the announcement date; the option quote incorporates the shock-day reaction, while the subsequent-realized-variance window begins after that session. Shaded bands are 95 and approximately 68 percent pointwise percentile confidence intervals from 1000 wild-block bootstrap draws using 125-session blocks.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities, Databento CBOT futures contracts, and the Känzig (2021) oil-supply-news data and WTI event returns.

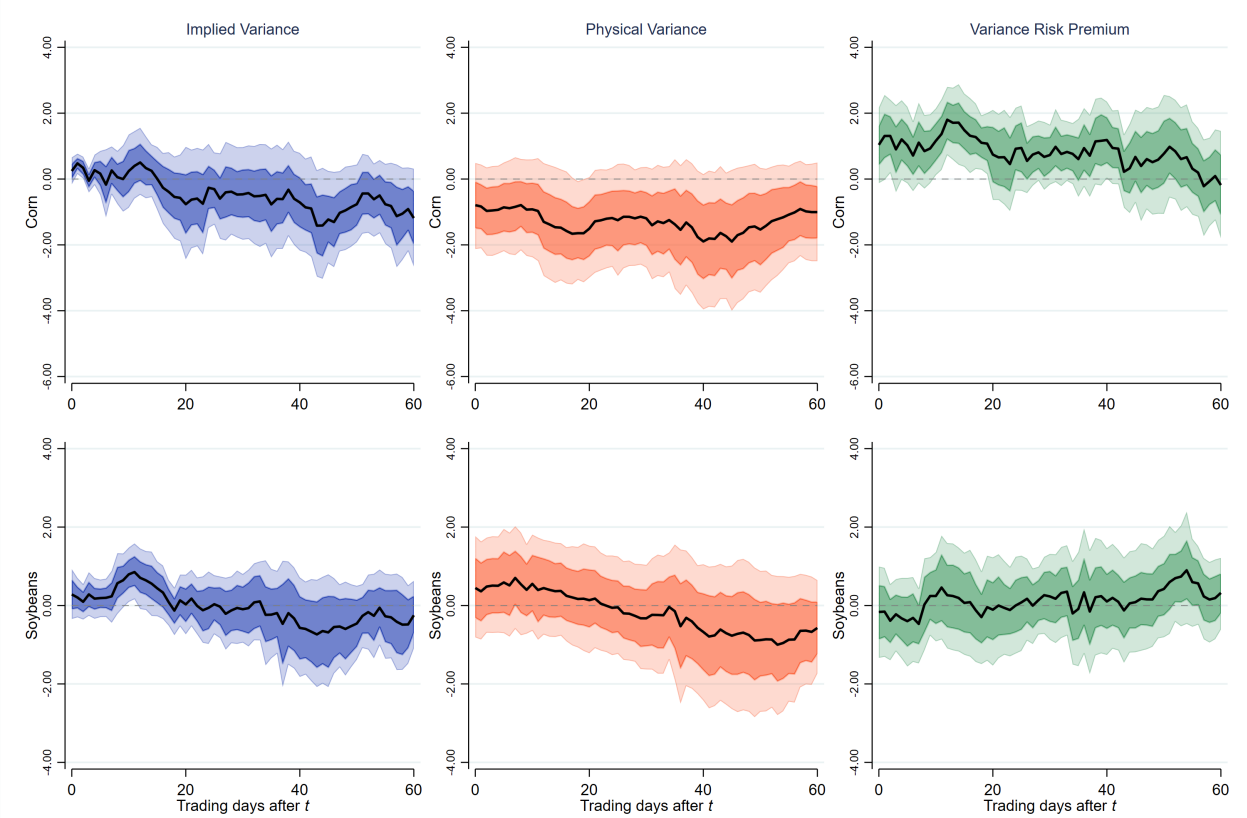


Figure 7: Monetary-policy-news response at the 60-day tenor

Notes: Rows report corn and soybeans; columns report ATM implied variance, subsequent realized variance, and their difference, the variance-premium measure. The shock is a 10-basis-point Bauer–Swanson orthogonalized monetary-policy surprise. Regressions include their six pre-announcement information controls. Event time $t = 0$ is the official crop session containing the FOMC announcement; the option quote incorporates the announcement-day reaction, while the subsequent-realized-variance window begins after that session. Shaded bands are 95 and approximately 68 percent pointwise percentile confidence intervals from 1000 wild-block bootstrap draws using 125-session blocks.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities, Databento CBOT futures contracts, and the Bauer and Swanson (2023) monetary-policy-surprise data.

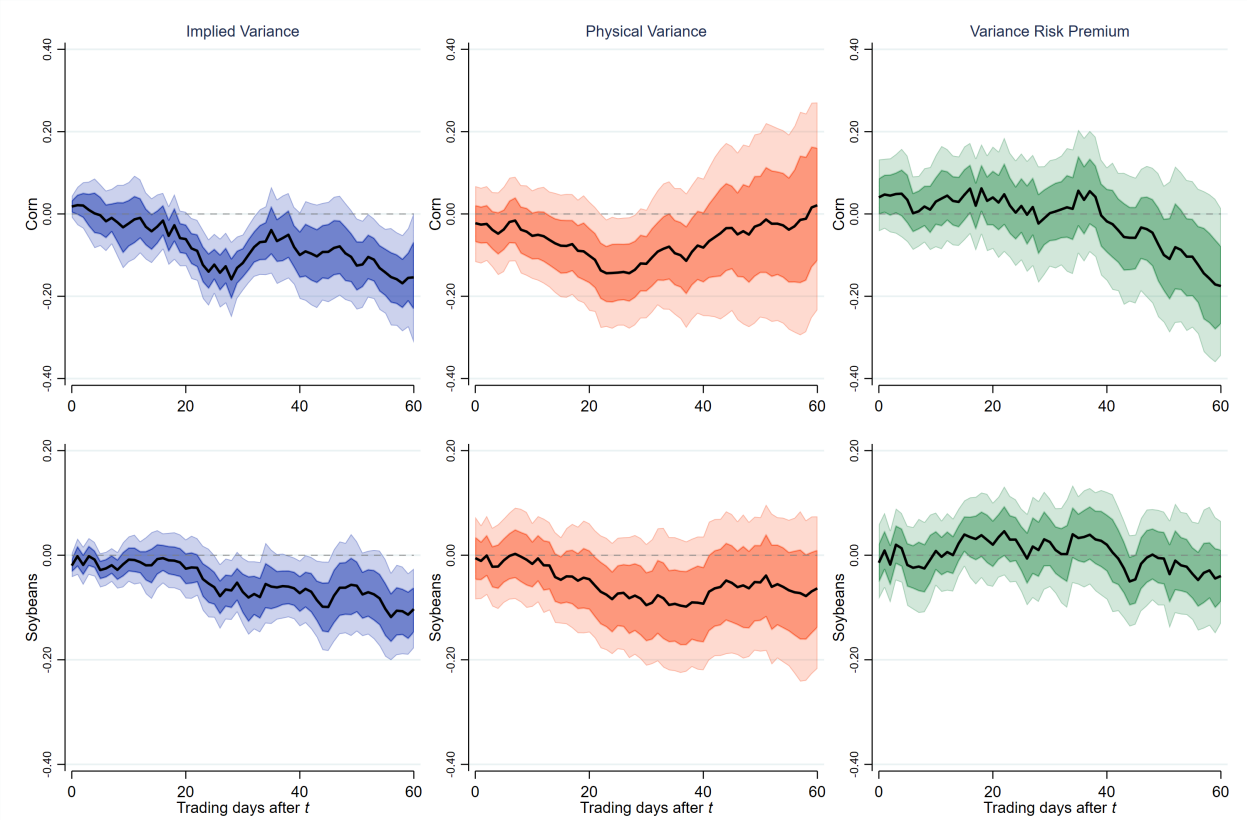


Figure 8: Physical-risk news-attention response at the 60-day tenor

Notes: Rows report corn and soybeans; columns report ATM implied variance, subsequent realized variance, and their difference, the variance-premium measure. The shock is a one-fixed-pre-2014-standard-deviation unexpected daily increase in Reuters physical-risk attention; physical- and transition-risk innovations enter jointly. Each complete news day is assigned to the first crop trading day strictly afterward. Shaded bands are 95 and approximately 68 percent pointwise percentile confidence intervals from 1000 wild-block bootstrap draws using 125-session blocks. The innovation measures news attention, not physical damage or an objective weather shock.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities, Databento CBOT futures contracts, and the Reuters-based Physical Risk Index of Bua et al. (2024).

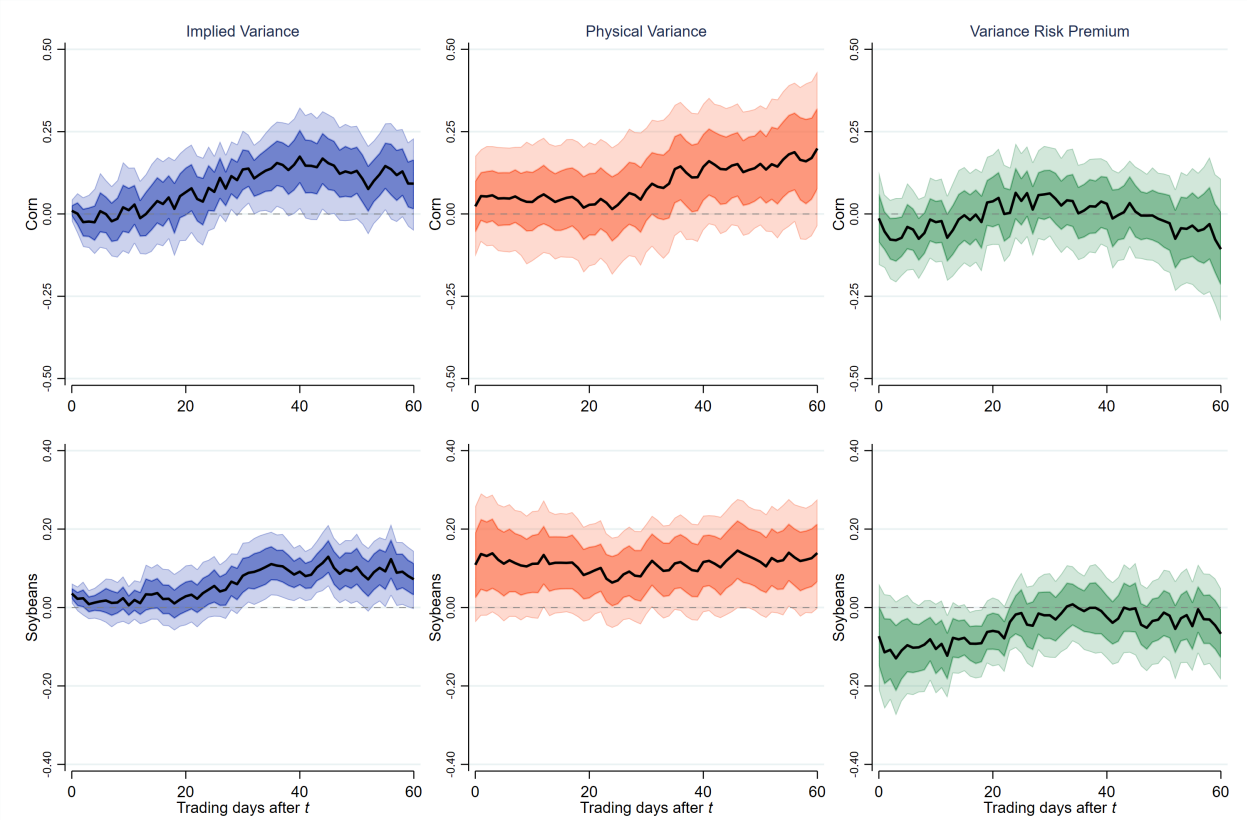


Figure 9: Transition-risk news-attention response at the 60-day tenor

Notes: Rows report corn and soybeans; columns report ATM implied variance, subsequent realized variance, and their difference, the variance-premium measure. The shock is a one-fixed-pre-2014-standard-deviation unexpected daily increase in Reuters transition-risk attention; physical- and transition-risk innovations enter jointly. Each complete news day is assigned to the first crop trading day strictly afterward. Shaded bands are 95 and approximately 68 percent pointwise percentile confidence intervals from 1000 wild-block bootstrap draws using 125-session blocks. The innovation measures news attention, not policy stringency or a structural policy shock.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities, Databento CBOT futures contracts, and the Reuters-based Transition Risk Index of Bua et al. (2024).

7.3 Historical decomposition

Figure 10 plots monthly averages of our estimates of each crop price's implied variance and of its components over time.[28] Appendix figures B1 and B2 report the other tenors. In the figures, broad movements in implied variance follow our estimate of expected physical variance quite closely, while the VRP is generally smaller and changes sign.[29]

The contrast in the figure between the 2021 and 2022 crops is illustrative and useful for dating, rather than identifying, the mechanisms at work in the evolution of implied volatility. The large 2021 increase in both crops' implied variance occurs almost entirely because (our forecast of) their expected physical variance rises, which is what the model predicts for that year: tight stocks after heavy Chinese buying meant that the March acreage surprise and the Northern Plains drought occurred in a market with little inventory to cushion them, so $\varphi(s_t)$ in (10) was large. The two crops part ways the following year. Corn's implied variance stays high even as its expected physical variance recedes, because the price-of-protection residual turns positive; soybean implied variance tracks lower than its expected physical variance because its premium turns negative. Russia's invasion of Ukraine removed a major corn exporter from world trade while leaving the global soybean trade largely intact. Processors, feedlots and ethanol plants may have bid for corn protection specifically, raising d_t in (13) and the corn premium with it, while soybean net demand stayed relatively constant. Monetary tightening and the crude oil spike raised κ_t and ω_t^2 for both crops, which in the model scales each premium away from zero rather than moving them together—possibly amplifying the corn premium and pushing the soybean premium further negative.

7.4 RMA crop-insurance premium counterfactual

RMA uses an option-implied volatility factor to rate the price-risk component of federal crop insurance, but the agency observes the (total) implied variance rather than its expected

[28] Recall that, unlike our results above, the historical decomposition depends on a forecasting model that introduces error. Table C1 reports the diagnostics of the forecast we selected: *the matching-tenor baseline model*; it has the lowest quasi-likelihood loss for both crops. That loss compares realized to forecast variance as a ratio rather than a difference, so it is scale-free and comparable across crops; mean squared error is not, since soybean variance is lower in level. All six specifications overpredict, and the corn mean error ranges from -0.084 to -0.594 across variants—a spread that transfers one for one into the path the premium traces out, against a forecast-free corn mean of 0.625. The Mincer–Zarnowitz coefficients indicate where the overprediction falls: slopes below one paired with positive intercepts mean the forecast is too dispersed, running high exactly when it predicts high variance. We conclude that the premium path therefore understates the premium most in its most volatile episodes.

[29] Our theoretical model accounts for that possibility: in (13) the premium is $\kappa_t \omega_t^2 d_t$ and, since dealer tightness and payoff risk are both positive, only net demand for protection d_t can turn negative. That implies customers were net sellers—e.g. farmers selling covered calls against stored grain, with no offsetting demand for protection large enough to outweigh them—so dealers bought variance at a discount rather than charging for it. This said, we stress that the historical decomposition is subject to forecast errors.

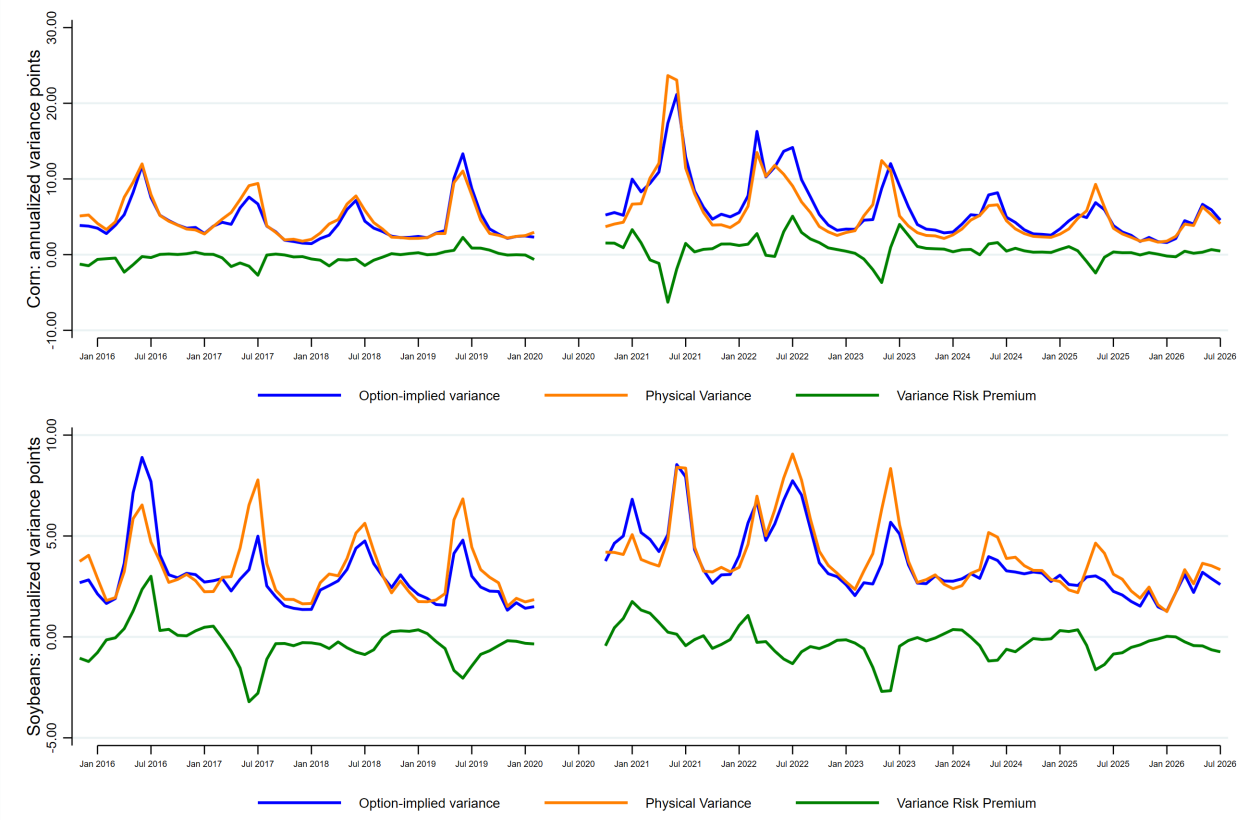


Figure 10: Historical decomposition, 60-day tenor, using the recursive forecast

Notes: Lines are monthly means of annualized ATM implied variance $100 \times q^{ATM}$, forecast physical variance $100 \times \widehat{P}$, and the residual variance-premium measure $100 \times (q^{ATM} - \widehat{P})$. The three series add exactly in the displayed variance units. The residual contains the ATM measurement wedge and physical-variance forecast error as well as the economic premium. Missing months are shown as breaks.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities and Databento CBOT futures contracts.

physical variance component. Its market input therefore contains both the amount of price variance traders expect and the price that the option market charges to bear it. Here we consider what RMA's factor would look like under a physical-risk-only benchmark. To do so, we remove each crop's forecast-free average implied–realized variance gap from RMA's daily variance inputs and then apply the agency's square-root time adjustment and five-day averaging procedure unchanged. Table 4 displays our estimates of that policy choice. At the headline 60-day tenor, removing the average corn premium lowers the mean unrounded factor from 18.44 percent to 17.30 percent. The 1.14-factor-point difference has a 95 percent interval of $[0.15, 2.06]$ and equates to about 6.61 percent of the physical-only factor. For soybeans, the difference is -0.30 factor points and is not statistically distinguishable from zero; its interval is $[-1.04, 0.39]$.

Our back-of-the-envelope calculation of the effects of including the VRP in the policy price is based on the average-premium results in table 3, so is focused on corn (whose options carry a positive and significant average premium in that table). For clarity, here "policy price" means the total dollar charge for the insurance contract—what RMA records as *Total Premium*—whereas the VRP remains the option market's price of protection.

The median result in table 4 indicates that including the estimated VRP raises the policy price by 6.8 percent relative to the price that would result under the physical-risk-only benchmark. Let C^{IV} denote the observed aggregate policy price and C^{P} its physical-risk-only counterfactual. Then $C^{IV} = (1 + 0.068)C^{P}$, so the portion of the observed policy price associated with retaining the estimated VRP is

$$\Delta C = C^{IV} - C^{P} = \frac{0.068}{1.068} C^{IV} = 0.0637 C^{IV}. \tag{39}$$

RMA's Summary of Business records $90.46 billion in total policy prices for nationwide corn Revenue Protection and Revenue Protection with Harvest Price Exclusion policies over crop years 2015–2024.[30] Applying equation (39) produces

$$0.0637 \times \$90.46 \text{ billion} = \$5.76 \text{ billion}.$$

Therefore, we estimate that including the average corn VRP in RMA's volatility factor input raised the price of these policies by about $5.8 billion over the decade, compared to the counterfactual.

RMA reports that federal premium subsidies financed $57.94 billion, or 64.0 percent, of this policy book. Holding policy choices and subsidy shares fixed, approximately $3.69 billion of the $5.76 billion is therefore federal premium-subsidy outlay—for which taxpayers are responsible—while approximately $2.06 billion is paid by producers.

[30] Authors' calculations from the USDA Risk Management Agency Summary of Business Report Generator, nationwide corn, insurance-plan codes 02 (Revenue Protection) and 03 (Revenue Protection with Harvest Price Exclusion), crop years 2015–2024.

Table 4: Illustrative premium effect of RMA's inclusion of the variance risk premium

	Revenue Protection (RP)			RP–Harvest Price Excl.		
Coverage	Opt.	Basic	Enterp.	Opt.	Basic	Enterp.
Panel A: Corn						
50%	2.67%	2.14%	2.93%	1.34%	1.80%	1.94%
	[0.33, 4.99]	[0.27, 3.98]	[0.37, 5.49]	[0.17, 2.49]	[0.23, 3.36]	[0.24, 3.61]
55%	3.06%	3.13%	4.19%	2.93%	3.79%	6.61%
	[0.38, 5.74]	[0.39, 5.86]	[0.52, 7.87]	[0.37, 5.49]	[0.47, 7.12]	[0.81, 12.55]
60%	3.46%	3.99%	5.10%	3.72%	5.10%	7.51%
	[0.43, 6.49]	[0.50, 7.50]	[0.63, 9.64]	[0.46, 6.99]	[0.63, 9.64]	[0.92, 14.33]
65%	4.25%	4.71%	6.28%	4.19%	6.21%	10.67%
	[0.53, 8.00]	[0.58, 8.88]	[0.77, 11.92]	[0.52, 7.87]	[0.76, 11.79]	[1.29, 20.60]
70%	6.15%	7.13%	8.74%	9.06%	9.96%	18.14%
	[0.76, 11.66]	[0.87, 13.57]	[1.06, 16.75]	[1.10, 17.39]	[1.21, 19.19]	[2.13, 36.09]
75%	6.67%	7.32%	8.81%	9.96%	11.62%	18.08%
	[0.82, 12.68]	[0.90, 13.95]	[1.07, 16.88]	[1.21, 19.19]	[1.40, 22.53]	[2.12, 35.96]
80%	7.00%	7.25%	8.35%	10.67%	11.56%	15.66%
	[0.86, 13.31]	[0.89, 13.82]	[1.02, 15.99]	[1.29, 20.60]	[1.39, 22.40]	[1.86, 30.85]
85%	6.93%	7.19%	7.97%	10.03%	10.47%	13.46%
	[0.85, 13.19]	[0.88, 13.69]	[0.97, 15.22]	[1.21, 19.32]	[1.27, 20.22]	[1.61, 26.29]
Panel B: Soybeans						
50%	-0.84%	-0.67%	-0.92%	-0.42%	-0.57%	-0.61%
	[-2.78, 1.11]	[-2.24, 0.89]	[-3.05, 1.22]	[-1.41, 0.56]	[-1.90, 0.75]	[-2.04, 0.81]
55%	-0.96%	-0.98%	-1.30%	-0.92%	-1.18%	-2.02%
	[-3.18, 1.27]	[-3.25, 1.30]	[-4.30, 1.73]	[-3.05, 1.22]	[-3.91, 1.57]	[-6.63, 2.71]
60%	-1.08%	-1.24%	-1.58%	-1.16%	-1.58%	-2.29%
	[-3.58, 1.43]	[-4.11, 1.65]	[-5.20, 2.10]	[-3.84, 1.54]	[-5.20, 2.10]	[-7.47, 3.08]
65%	-1.32%	-1.46%	-1.93%	-1.30%	-1.91%	-3.18%
	[-4.36, 1.76]	[-4.82, 1.94]	[-6.32, 2.58]	[-4.30, 1.73]	[-6.26, 2.55]	[-10.29, 4.33]
70%	-1.89%	-2.17%	-2.64%	-2.73%	-2.99%	-5.18%
	[-6.20, 2.53]	[-7.11, 2.92]	[-8.59, 3.57]	[-8.88, 3.70]	[-9.68, 4.05]	[-16.37, 7.22]
75%	-2.04%	-2.23%	-2.66%	-2.99%	-3.45%	-5.17%
	[-6.69, 2.74]	[-7.29, 3.00]	[-8.65, 3.59]	[-9.68, 4.05]	[-11.12, 4.71]	[-16.32, 7.20]
80%	-2.14%	-2.21%	-2.53%	-3.18%	-3.43%	-4.54%
	[-6.99, 2.87]	[-7.23, 2.97]	[-8.24, 3.41]	[-10.29, 4.33]	[-11.06, 4.68]	[-14.44, 6.27]
85%	-2.12%	-2.19%	-2.42%	-3.00%	-3.13%	-3.95%
	[-6.93, 2.84]	[-7.17, 2.95]	[-7.89, 3.26]	[-9.74, 4.08]	[-10.13, 4.25]	[-12.66, 5.42]

Notes: Each entry is the percentage premium increase from retaining RMA's option-implied volatility factor rather than removing the corresponding crop's estimated variance-risk-premium component, calibrated as the full-common-sample mean implied-minus-subsequent variance difference. Bracketed values are 95% paired 65-session moving-block bootstrap intervals, in percentage points. Opt., Basic, and Enterp. refer to RMA's optional, basic, and enterprise unit structures: optional units insure acreage separately by section, basic units combine a producer's acreage in a county under a single share arrangement, and enterprise units combine all of a producer's acreage of the crop in a county. Premium rates rise as units are disaggregated because pooling reduces the insured yield's variability. The common premium-sensitivity template is Sherrick's (2015, Table 3) McLean County, Illinois corn finite-difference schedule; the soybean panel is therefore a Sherrick-style sensitivity calibration, not a county-specific soybean premium rerating. The factor calibration averages 12 complete discovery years from 2014 through 2026; 2020 is excluded because only four of the five factor dates were available in our data.

8 Conclusions

A commodity option's price captures both what the market expects about the variance of the benchmark futures price and what it charges to carry that risk (the variance risk premium or VRP). Each component has different implications for producers, for traders, and for the federal programs that use options prices as inputs. Our approach to separating them is simple. Realized variance differs from expected variance by a surprise that is uncorrelated with anything known when the window opens. Thus, a regression of realized variance on known states recovers their relationship with expected physical variance, and the same regression on implied minus realized recovers their relationship with the VRP. We run these two regressions, along with a regression with option-implied variance on the left-hand side, on the same data with the same right-hand-side specification, so the three sets of coefficients add exactly. Together, they offer an accounting of important option determinants and the channel through which each operates. The averages, the state relationships, and the dynamic responses that we estimate require no variance forecast. Only the historical decomposition does, so it inherits that forecasting model's error.

We establish a positive average VRP for corn; in our sample, the premium for soybeans is not statistically distinguishable from zero. We hypothesize that crush asymmetry to variance exposure may explain the difference: an ethanol plant carries corn-price exposure that it must hedge, while a soybean crusher sells meal and oil at prices that move with the bean price, so has little residual exposure.

Commercial hedging pressure and the VIX affect the VRP. Neither component of the implied variance trends between 2014 and 2026. As a result, we have no confirmation in the data yet of strong climate effects on either market physical crop risk estimates or the price of that risk, possibly as a result of producer adaptation. Our impulse response results offer a range of insightful findings, too. Including that poor agricultural supply news raises the corn premium on impact, before any change in subsequent physical variance is visible. A regression with implied variance as its only outcome would simply show a positive coefficient and could not explain *why* it is positive.

Telling those components apart has a direct application for crop insurance. The federal program prices revenue protection using a volatility factor built from implied volatility—so, in addition to expected physical variance, whatever the market charges for bearing variance enters the USDA rate. For corn, removing the average premium from a calculation applied to a range of illustrative insurance products lowers their median total premium rate by around 7 percent. Applied RMA's reported national corn revenue protection premium volume over crop years 2015–2024, that difference amounts to $5.8 billion in extra total premium, allo-

cating roughly $3.7 billion to taxpayer-paid federal premium subsidies and $2.1 billion to producer-paid premium. Whether a public program should charge producers the private market's price of risk is a policy question, but its implications can only be understood if the two components are measured accurately.

A Categorizing Our Estimates

Empirical Target	Model Expression	Forecast?	Interpretation
Average observed premium	Mean of $q^{\mathrm{ATM}} - X$	No	Mean of RP $+ \nu$, since $\mathbb{E}[\varepsilon] = 0$, where ν is the ATM measurement wedge
Components versus states	X and $q^{\mathrm{ATM}} - X$ on Z_t	No	Orthogonality recovers the P projection; the $q^{\mathrm{ATM}} - X$ projection is for RP $+ \nu$
Trends in components	X and $q^{\mathrm{ATM}} - X$ on time	No	The same orthogonality and adding-up argument
Dynamic responses	q_s^{ATM}, X_s, and $q_s^{\mathrm{ATM}} - X_s$	No	Dummy impact paths describe timing around known scheduled releases; the reported LP-IV responses require separate shock-identification conditions
Historical decomposition and RMA calibration	$\widehat{P}_t$, $q_t^{\mathrm{ATM}} - \widehat{P}_t$, and the RMA premium effects	Historical path: Yes; RMA: No	The historical decomposition inherits recursive-forecast error; the RMA exercise maps the forecast-free average gap through the administrative formulas and Sherrick's example

Table A1: Five empirical questions.

Notes: The table lists the primary empirical objectives of the paper and what each requires. P is expected physical variance and RP is the VRP; neither is observed. We instead observe q^{ATM}, at-the-money option-implied variance, and X, the variance the same named futures contract subsequently realizes. Z_t is any state known at t. Orthogonality refers to equation (5): the realization surprise ε is uncorrelated with information known at t. The first four rows therefore require no forecast of P. Within the fifth row, the dated historical decomposition requires a recursive forecast, whereas the RMA calibration uses the forecast-free sample mean of $q^{\mathrm{ATM}} - X$.

B Additional tenors

Table B1: Option-implied variance, realized variance, and the premium: 30-day tenor

	Corn			Soybeans		
	q^{ATM}	X	$q^{ATM} - X$	q^{ATM}	X	$q^{ATM} - X$
Panel A: Full-sample moments						
Mean	5.455***	4.955***	0.500**	3.513***	3.686***	-0.173
	(0.413)	(0.473)	(0.218)	(0.192)	(0.279)	(0.151)
Linear trend, per year	-0.019	-0.008	-0.012	-0.074	-0.045	-0.029
	(0.086)	(0.098)	(0.051)	(0.047)	(0.069)	(0.037)
Full-sample observations	2,934	2,934	2,934	2,934	2,934	2,934
Panel B: Joint state specification						
Carry factor (one SD)	-1.947***	-2.569***	0.622	-0.447**	-0.321	-0.125
	(0.450)	(0.782)	(0.497)	(0.223)	(0.375)	(0.209)
WASDE reports in target window	0.430*	0.716**	-0.286	0.193	0.199	-0.006
	(0.240)	(0.350)	(0.314)	(0.159)	(0.293)	(0.266)
Grain Stocks reports in target window	1.371***	1.510**	-0.139	0.477***	1.064**	-0.586*
	(0.363)	(0.588)	(0.573)	(0.180)	(0.430)	(0.342)
Critical-window overlap (10 pp)	0.296*	-0.191	0.486	0.195**	0.112	0.083
	(0.157)	(0.336)	(0.320)	(0.085)	(0.140)	(0.157)
ENSO state (one SD)	0.166	0.261	-0.095	-0.177	-0.050	-0.127
	(0.205)	(0.277)	(0.201)	(0.220)	(0.357)	(0.189)
Commercial hedging pressure (one SD)	1.284***	0.927**	0.357	0.371	0.644*	-0.273
	(0.356)	(0.448)	(0.253)	(0.242)	(0.339)	(0.169)
CIT long share (one SD)	-0.448	-1.072**	0.624*	-0.480*	-0.495	0.015
	(0.303)	(0.437)	(0.327)	(0.272)	(0.378)	(0.212)
Lagged VIX (one SD)	0.840***	0.427**	0.414***	0.510***	0.359	0.151
	(0.177)	(0.197)	(0.160)	(0.166)	(0.241)	(0.140)
Physical-risk news (one SD)	0.064	0.101	-0.037	-0.017	0.038	-0.055
	(0.068)	(0.063)	(0.065)	(0.042)	(0.056)	(0.040)
Transition-risk news (one SD)	-0.005	-0.014	0.009	0.057	0.118	-0.061
	(0.056)	(0.083)	(0.075)	(0.043)	(0.084)	(0.066)
Fourier seasonality and linear trend	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.648	0.518	0.115	0.406	0.341	0.146
Regression residual Phillips–Perron p-value	0.001	0.001	0.001	0.024	0.001	0.001
Joint-sample observations	2,198	2,198	2,198	2,198	2,198	2,198

Notes: Outcomes are annualized variance multiplied by 100. Panel A uses 2,934 common dates from January 2, 2014–April 24, 2026; Panel B uses 2,198 dates from January 15, 2014–May 12, 2023. All three Panel B equations use identical rows and regressors, so coefficients add exactly up to rounding. Standard errors use 1000 paired 65-session calendar-block wild-bootstrap draws. The regression-residual Phillips–Perron row tests the null that the fitted residual from the exact Panel B equation has a unit root. Its auxiliary residual AR has no deterministic terms because Panel B already includes the intercept, linear trend, Fourier terms, and states. The p-value is calibrated by 1000 null-imposed 65-session calendar-block wild-bootstrap paths: centered residual increments are sign-flipped within blocks, cumulated separately across contiguous official-session segments, and the full Panel B equation is re-estimated in every draw. The calculation never links observations across a session gap and uses overlap-matched Newey–West bandwidths of 21, 43, and 62 sessions at 30, 60, and 90 days. Small Phillips–Perron p-values reject the null of a unit root in the residual of the Panel B level regression. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table B2: Option-implied variance, realized variance, and the premium: 90-day tenor

	Corn			Soybeans		
	q^{ATM}	X	$q^{ATM} - X$	q^{ATM}	X	$q^{ATM} - X$
Panel A: Full-sample moments						
Mean	5.561***	4.866***	0.695**	3.531***	3.608***	-0.076
	(0.377)	(0.442)	(0.275)	(0.173)	(0.253)	(0.164)
Linear trend, per year	-0.032	-0.008	-0.025	-0.060	-0.045	-0.015
	(0.081)	(0.090)	(0.065)	(0.041)	(0.063)	(0.045)
Full-sample observations	2,934	2,934	2,934	2,934	2,934	2,934
Panel B: Joint state specification						
Carry factor (one SD)	-1.710***	-2.496***	0.786	-0.415**	-0.554*	0.140
	(0.278)	(0.620)	(0.644)	(0.187)	(0.335)	(0.251)
WASDE reports in target window	1.173***	0.555	0.618	0.339*	0.271	0.068
	(0.326)	(0.342)	(0.527)	(0.192)	(0.214)	(0.288)
Grain Stocks reports in target window	0.295	0.345	-0.050	0.124	0.084	0.040
	(0.223)	(0.349)	(0.413)	(0.153)	(0.242)	(0.253)
Critical-window overlap (10 pp)	0.104	-0.381	0.486	0.231	0.242	-0.011
	(0.411)	(0.621)	(0.790)	(0.172)	(0.263)	(0.263)
ENSO state (one SD)	0.123	0.320	-0.197	-0.138	0.185	-0.323
	(0.144)	(0.240)	(0.198)	(0.192)	(0.349)	(0.252)
Commercial hedging pressure (one SD)	1.206***	0.392	0.814**	0.386*	0.401	-0.014
	(0.215)	(0.335)	(0.336)	(0.220)	(0.266)	(0.189)
CIT long share (one SD)	-0.193	-0.482	0.289	-0.460*	-0.319	-0.142
	(0.230)	(0.349)	(0.307)	(0.244)	(0.334)	(0.223)
Lagged VIX (one SD)	0.490***	-0.060	0.551***	0.346**	0.156	0.191
	(0.154)	(0.222)	(0.178)	(0.155)	(0.255)	(0.154)
Physical-risk news (one SD)	0.051	-0.033	0.084	-0.003	-0.022	0.018
	(0.040)	(0.044)	(0.066)	(0.035)	(0.040)	(0.038)
Transition-risk news (one SD)	-0.019	-0.030	0.011	0.033	0.023	0.010
	(0.040)	(0.058)	(0.065)	(0.031)	(0.054)	(0.040)
Fourier seasonality and linear trend	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.761	0.664	0.233	0.441	0.479	0.229
Regression residual Phillips–Perron p-value	0.003	0.027	0.024	0.080	0.032	0.005
Joint-sample observations	2,198	2,198	2,198	2,198	2,198	2,198

Notes: Outcomes are annualized variance multiplied by 100. Panel A uses 2,934 common dates from January 2, 2014–April 24, 2026; Panel B uses 2,198 dates from January 15, 2014–May 12, 2023. All three Panel B equations use identical rows and regressors, so coefficients add exactly up to rounding. Standard errors use 1000 paired 65-session calendar-block wild-bootstrap draws. The regression-residual Phillips–Perron row tests the null that the fitted residual from the exact Panel B equation has a unit root. Its auxiliary residual AR has no deterministic terms because Panel B already includes the intercept, linear trend, Fourier terms, and states. The p-value is calibrated by 1000 null-imposed 65-session calendar-block wild-bootstrap paths: centered residual increments are sign-flipped within blocks, cumulated separately across contiguous official-session segments, and the full Panel B equation is re-estimated in every draw. The calculation never links observations across a session gap and uses overlap-matched Newey–West bandwidths of 21, 43, and 62 sessions at 30, 60, and 90 days. Small Phillips–Perron p-values reject the null of a unit root in the residual of the Panel B level regression. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.



Figure B1: Historical decomposition, 30-day tenor, using the recursive forecast

Notes: Lines are monthly means of annualized ATM implied variance $100 \times q^{ATM}$, forecast physical variance $100 \times \widehat{P}$, and the residual variance-premium measure $100 \times (q^{ATM} - \widehat{P})$. The three series add exactly in the displayed variance units. The residual contains the ATM measurement wedge and physical-variance forecast error as well as the economic premium. Missing months are shown as breaks.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities and Databento CBOT futures contracts.

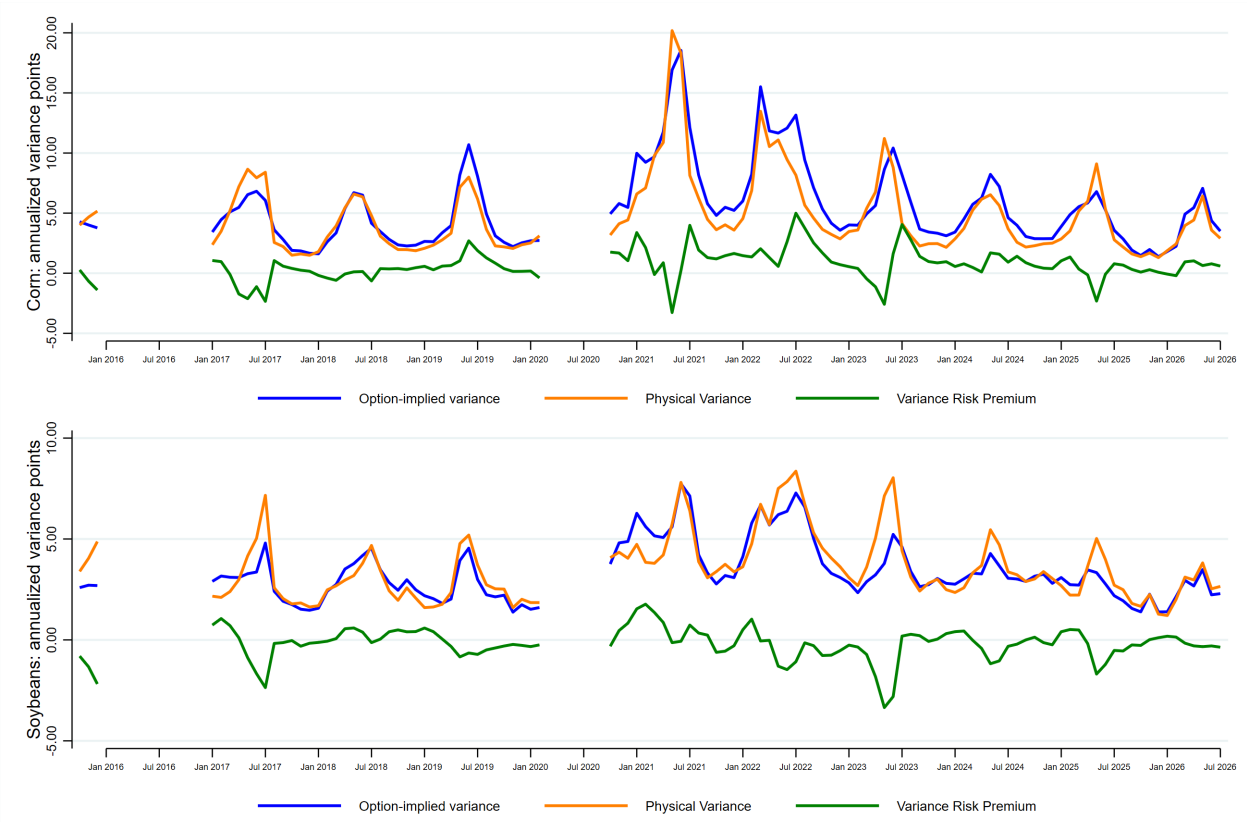


Figure B2: Historical decomposition, 90-day tenor, using the recursive forecast

Notes: Lines are monthly means of annualized ATM implied variance $100 \times q^{ATM}$, forecast physical variance $100 \times \widehat{P}$, and the residual variance-premium measure $100 \times (q^{ATM} - \widehat{P})$. The three series add exactly in the displayed variance units. The residual contains the ATM measurement wedge and physical-variance forecast error as well as the economic premium. Missing months are shown as breaks.

Sources: Authors' calculations using LSEG Workspace fixed-tenor ATM implied volatilities and Databento CBOT futures contracts.

C Forecast and first-stage diagnostics

Table C1: Recursive 60-day physical-variance forecast diagnostics

Crop	Specification	QLIKE	MSE	Mean $X - \widehat{P}$	MZ intercept	MZ slope	N
Corn	Matching-tenor baseline	0.091	0.000711	-0.304	0.387	0.866	2,214
Corn	Own-curve specification	0.097	0.000729	-0.084	0.256	0.931	2,214
Corn	No-own-market-IV specification	0.100	0.000833	-0.594	0.356	0.826	2,214
Soybeans	Matching-tenor baseline	0.089	0.000342	-0.227	0.390	0.832	2,193
Soybeans	Own-curve specification	0.097	0.000353	-0.180	0.393	0.842	2,193
Soybeans	No-own-market-IV specification	0.090	0.000352	-0.465	0.451	0.766	2,193

Notes: Expanding-window forecasts use only resolved monthly outcomes and are scored daily. MSE is in squared annualized-variance units; the mean error and Mincer–Zarnowitz intercept are variance points, so negative $X - \widehat{P}$ denotes overprediction. The matching-tenor baseline uses the matching own-market ATM tenor; the own-curve model adds the two own-market curve slopes; the no-own-market-IV model removes own-market ATM variance.

Table C2: LP-IV first stages at the 60-day tenor

Instrument family	Crop	First-stage coefficient	HAC robust F	Partial R^2	Events
Agricultural supply news	Corn	2.423	1,685.4	0.954	122
Agricultural supply news	Soybeans	1.388	92.8	0.488	122
Oil supply news	Corn	2.658	1,069.2	0.960	37
Oil supply news	Soybeans	2.763	1,652.9	0.968	37

Notes: The agricultural instrument predicts the roll-safe own-crop named-futures event return; the oil instrument predicts the front WTI-futures event return. Both endogenous returns are measured in percentage points. HAC robust F statistics use the implemented 125-session calendar-distance covariance estimator.

References

Viral V. Acharya, Lars A. Lochstoer, and Tarun Ramadorai. Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics*, 109(2):441–465, 2013. doi: 10.1016/j.jfineco.2013.03.003.

Michael K. Adjemian and Jungkeon Jo. The macroeconomic effects of agricultural supply news. SSRN Working Paper 5405106, 2025. August 25, 2025.

Michael K. Adjemian, Valentina G. Bruno, Michel A. Robe, and Jonathan Wallen. What drives volatility expectations in grain and oilseed markets? In *Proceedings of the 2017 Agricultural and Applied Economics Association Annual Meeting*, Chicago, IL, July 2017. URL `https://ageconsearch.umn.edu/record/258452`.

Fernando G. Avileis. Commodity volatility in the biofuel era. Paper presented at the 2025 Meeting of the Agricultural and Applied Economics Association, Denver, CO, August 2025.

Michael D. Bauer and Eric T. Swanson. A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1):87–155, 2023.

Geert Bekaert and Marie Hoerova. The VIX, the variance premium and stock market volatility. *Journal of Econometrics*, 183(2):181–192, 2014. doi: 10.1016/j.jeconom.2014.05.008.

Geert Bekaert, Marie Hoerova, and Marco Lo Duca. Risk, uncertainty and monetary policy. *Journal of Monetary Economics*, 60(7):771–788, 2013. doi: 10.1016/j.jmoneco.2013.06.003.

Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11):4463–4492, 2009. doi: 10.1093/rfs/hhp008.

Mark Britten-Jones and Anthony Neuberger. Option prices, implied price processes, and stochastic volatility. *The Journal of Finance*, 55(2):839–866, 2000. doi: 10.1111/0022-1082.00228.

Giovanna Bua, Daniel Kapp, Federico Ramella, and Lavinia Rognone. Transition versus physical climate risk pricing in european financial markets: A text-based approach. *The European Journal of Finance*, 30(17):2076–2110, 2024. doi: 10.1080/1351847X.2024.2355103.

Harun Bulut and T. Randall Fortenbery. Reevaluating the use of volatility factor in crop insurance premium rating. *Journal of the Agricultural and Applied Economics Association*, 1(2):124–135, 2022. doi: 10.1002/jaa2.17.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427, 08 2008. ISSN 0034-6535. doi: 10.1162/rest.90.3.414. URL `https://doi.org/10.1162/rest.90.3.414`.

An N. Q. Cao and Michel A. Robe. Market uncertainty and sentiment around USDA announcements. *Journal of Futures Markets*, 42(2):250–275, 2022. doi: 10.1002/fut.22283.

Peter Carr and Liuren Wu. Variance risk premiums. *The Review of Financial Studies*, 22 (3):1311–1341, 2009. doi: 10.1093/rfs/hhn038.

Kresimir Demeterfi, Emanuel Derman, Michael Kamal, and Joseph Zou. A guide to volatility and variance swaps. *The Journal of Derivatives*, 6(4):9–32, 1999.

Matthew Diersen and Zhiguang Wang. Weekly options on grain futures. *Journal of Agricultural and Applied Economics*, 57(1):60–85, 2025. doi: 10.1017/aae.2024.38. URL `https://doi.org/10.1017/aae.2024.38`.

Naihua Duan. Smearing estimate: A nonparametric retransformation method. *Journal of the American Statistical Association*, 78(383):605–610, 1983. doi: 10.1080/01621459.1983.10478017.

Thorsten M. Egelkraut, Philip Garcia, and Bruce J. Sherrick. The term structure of implied forward volatility: Recovery and informational content in the corn options market. *American Journal of Agricultural Economics*, 89(1):1–11, 2007. doi: 10.1111/j.1467-8276.2007.00958.x.

Eugene F. Fama. Forward and spot exchange rates. *Journal of Monetary Economics*, 14 (3):319–338, 1984. ISSN 0304-3932. doi: https://doi.org/10.1016/0304-3932(84)90046-1. URL `https://www.sciencedirect.com/science/article/pii/0304393284900461`.

Eugene F. Fama and Kenneth R. French. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *The Journal of Business*, 60(1): 55–73, 1987. doi: 10.1086/296385.

Kenneth R. French. Detecting spot price forecasts in futures prices. *The Journal of Business*, 59(2):S39–S54, 1986. doi: 10.1086/296338.

Gary B. Gorton, Fumio Hayashi, and K. Geert Rouwenhorst. The fundamentals of commodity futures returns. *Review of Finance*, 17(1):35–105, 2013. doi: 10.1093/rof/rfs019.

Xinyue He and Siyu Bian. The announcement volatility risk premium in agricultural markets: Evidence from USDA report releases. *Journal of Agricultural Economics*, 76(3):651–665, 2025. doi: 10.1111/1477-9552.12647.

Xinyue He, Siyu Bian, and Teresa Serra. The pricing of variance risks in agricultural futures markets: Do jumps matter? *European Review of Agricultural Economics*, 50(4):1428–1452, 2023.

Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005. doi: 10.1257/0002828053828518.

Diego R. Känzig. The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125, 2021. doi: 10.1257/aer.20190964.

James G. MacKinnon and Halbert White. Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, 29 (3):305–325, 1985.

Anthony Neuberger. The log contract. *The Journal of Portfolio Management*, 20(2):74–80, 1994. doi: 10.3905/jpm.1994.409478.

Adrian Pagan. Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, 25(1):221–247, 1984.

Marcel Prokopczuk, Lazaros Symeonidis, and Chardin Wese Simen. Variance risk in commodity markets. *Journal of Banking & Finance*, 81:136–149, 2017. doi: 10.1016/j.jbankfin.2017.05.003.

Vikas Raman and Michel A. Robe. The renewable fuel standard and market comovements: Evidence from option-implied volatility correlations. Paper presented at the 2026 Meeting of the World Finance Conference, Dublin, July 2026.

Bruce J. Sherrick. Understanding the implied volatility (IV) factor in crop insurance. farmdoc daily, February 2015. URL `https://farmdocdaily.illinois.edu/2015/02/understanding-implied-volatility-crop-insurance.html`. Department of Agricultural and Consumer Economics, University of Illinois at Urbana–Champaign.

Athanasios Triantafyllou, George Dotsis, and Alexandros H. Sarris. Volatility forecasting and time-varying variance risk premiums in grains commodity markets. *Journal of Agricultural Economics*, 66(2):329–357, 2015. doi: 10.1111/1477-9552.12101. URL `https://doi.org/10.1111/1477-9552.12101`.

Anders Bjerre Trolle and Eduardo S. Schwartz. Variance risk premia in energy commodities. *The Journal of Derivatives*, 17(3):15–32, 2010. doi: 10.3905/jod.2010.17.3.015.

USDA Risk Management Agency. RMA volatility factor calculation methodology. `https://legacy.rma.usda.gov/pubs/2015/volatilitymethodology.pdf`, 2015.

USDA Risk Management Agency. Volatility factor calculation methodology (27-CEPP-VF), 2026. Released June 2026 for 2027 and succeeding crop years.

Zhiguang Wang, Scott W. Fausti, and Bashir A. Qasmi. Variance risk premiums and predictive power of alternative forward variances in the corn market. *Journal of Futures Markets*, 32(6):587–608, 2012. doi: 10.1002/fut.20527.

Joshua D. Woodard and Thomas W. Sproul. Identification of hedging pressure effects in commodity futures-options markets: Implications for price discovery. Working Paper, June 2016.

Wenwen Xi, Dermot Hayes, and Sergio Horacio Lence. Variance risk premia for agricultural commodities. *Agricultural Finance Review*, 79(3):286–303, 2019.

Yao Yang and Andrew M. McKenzie. Do corn options update volatility expectations in the wake of USDA reports? *Journal of Futures Markets*, 45(10):1852–1868, 2025. doi: 10.1002/fut.70020.