

Modeling Commodity Futures with Satellite Signals: NDVI and the U.S. Corn Belt Basis, 2019 to 2024

Amantay Nurlanuly

Nazarbayev University, School of Sciences and Humanities

Supervisor: Piotr Sebastian Skrzypacz

Abstract

Local cash prices for agricultural commodities deviate from futures prices through the basis. The basis absorbs storage costs, transport, and local supply and demand conditions. This paper asks whether satellite vegetation observations carry information about the basis that is not already in market prices or in standard agricultural indicators. We use a panel of the seven largest U.S. corn states (IA, IL, IN, MN, NE, OH, SD) over April to October 2019 to 2024. The main NDVI feature is the average of weekly NDVI z-scores during the silking-to-early-dent window (week-of-year 30 to 35), with z-scores computed from leave-one-year-out historical baselines. The headline result is from a pooled regression with state fixed effects and a harvest-month interaction. The NDVI by harvest interaction coefficient is -0.128 \$/bu per z-score ($p = 0.001$, $N = 196$), with cluster-robust standard errors and stable estimates across all seven leave-one-state-out specifications. The total harvest-window NDVI effect is small but precisely identified. Out-of-sample forecasting comparisons show that no augmented feature group reliably outperforms a simple anomaly AR(1) baseline within bootstrap 95% confidence intervals. Supporting analysis indicates that USDA Crop Progress captures overlapping information with NDVI at this aggregation. The results are consistent with growing-season satellite observations carrying genuine but modest harvest-window information about the local basis.

Keywords: commodity futures, basis, NDVI, satellite remote sensing, corn, panel data.

1 Introduction

The efficient market hypothesis says that asset prices fully incorporate all publicly available information [Fama(1970)]. For globally traded commodity futures contracts such as CME corn, this claim is empirically defensible. The cash price received by a farmer or elevator, however, is not the futures price. It is the futures price plus the local basis. The basis decomposes into three parts: storage and time-to-delivery, the transportation differential between local cash market and futures delivery point, and the local supply and demand environment [Tomek and Robinson(1990), Working(1949)]. The third part is the residual that absorbs whatever the market does not already know about local conditions, and it is the natural place for new information sources to be detectable.

If satellite indices reveal that mid-season vegetation is unusually robust before this fact appears in slower USDA reports or in the futures curve, they ought to predict a weaker basis at harvest, when extra supply pushes local cash prices down relative to the national futures. The mechanism rests on three observations. First, farmers and elevators do not adjust offer prices instantaneously to satellite imagery. Second, USDA Crop Progress reports are released with a lag and at coarser frequency than satellite data. Third, the futures market prices a national contract, not a local one, and cannot fully reflect state-specific yield information.

Two design choices distinguish this paper from much of the existing basis-forecasting literature. First, we use a multi-state panel rather than a single-state time series, which allows state fixed effects to absorb cross-sectional differences in mean basis and isolates within-state variation. Second, we use a phenology-anchored NDVI feature defined over a fixed silking-to-early-dent window, with z-scores computed from leave-one-year-out historical baselines. This avoids both the look-ahead bias of within-year normalization and the mechanical bias of season-to-date maximum measures.

The motivating question is narrow. Do growing-season NDVI anomalies contain information about U.S. corn basis that is not already in lagged market prices or in USDA Crop Progress? We use the seven leading corn-producing states because multi-state expansion increases the effective sample by about seven times, allows state fixed effects to absorb cross-sectional unobservables, and provides a credible test of whether the original Iowa finding generalizes.

The contribution is methodological as much as substantive. We frame the question as a residual-variance test. After a state-month seasonal benchmark removes the predictable structure in basis, does NDVI explain what remains? This framing suits a panel with cross-sectional heterogeneity. We complement the diagnostic with a forecasting evaluation that includes a CME futures market baseline and an ablation across feature groups, with the intersection of evaluable observations across all variants used as a common evaluation set.

The headline finding is that, with state fixed effects and a harvest-month interaction, the silking-window NDVI z-score has a precisely identified negative effect on residual basis specific to the harvest window. The interaction coefficient is -0.128 \$/bu per z-score ($p = 0.001$), and is robust across all seven leave-one-state-out specifications. The effect is small in absolute magnitude but is the cleanest panel test of the original hypothesis. Out-of-sample forecasting comparisons show no reliable improvement from any augmented feature group over a simple anomaly AR(1) baseline within bootstrap confidence intervals. The totality of evidence is consistent with growing-season satellite observations carrying genuine but modest harvest-window information about the local basis.

The remainder of the paper is organized as follows. Section 2 reviews the literature. Section 3 describes the seven-state panel. Section 4 introduces the three-model framework. Section 5 presents the residual diagnostic, robustness checks, supporting horse race against USDA Crop Progress, and the out-of-sample forecasting evaluation. Section 6 discusses interpretation and limitations. Section 7 concludes.

2 Literature Review

2.1 The basis and the theory of price of storage

The economic theory of basis begins with [Working(1949)], who modeled the cash-futures relationship as the equilibrium outcome of arbitrage between physical storage and futures positions. [Tomek and Robinson(1990)] provide the standard textbook treatment, distinguishing the spatial component (transportation cost), the temporal component (storage cost), and the local supply and demand component, which captures everything else.

[Taylor et al.(2006)] compare distributional models for forecasting Kansas corn and soybean basis. [Sanders and Baker(2012)] develop regime-switching models for tight versus loose supply environments. The recurring finding is that basis is autocorrelated and seasonally patterned but with substantial unexplained year-to-year variation. The interesting question is not whether basis can be forecast but whether external information sources can explain the residual variance that simple methods leave behind.

2.2 Futures market efficiency

[Garcia et al.(2015)] survey agricultural futures markets and conclude that the broad picture is one of futures markets that incorporate publicly available information rapidly. The basis lives below the surface of the futures market and is determined by local cash transactions, where the information environment is much less symmetric. We take futures-market efficiency as given and treat predictable structure in basis as evidence of a separate information channel.

2.3 Satellite remote sensing for crop monitoring

The Normalized Difference Vegetation Index (NDVI) uses the spectral contrast between red and near-infrared reflectance to proxy photosynthetic activity. [Bolton and Friedl(2013)] demonstrate that mid-to-late summer NDVI is strongly predictive of U.S. corn yield at the county level, with predictive power concentrated in July and August when canopy closure peaks. [Johnson(2014)] corroborates using MODIS data. [Becker-Reshef et al.(2010)] develop a generalized regression-based winter wheat forecasting approach. Two methodological points inform our approach. First, the NDVI to yield relationship is conditional on phenological stage, with reproductive-phase NDVI carrying direct information about kernel set and grain fill. Second, the appropriate aggregation is the anomaly relative to a historical baseline, not the level.

2.4 Economic applications of satellite data

[Donaldson and Storeygard(2016)] survey satellite applications in economics. The common thread is using remotely sensed observations to provide objective, high-frequency, geographically disaggregated measures where official statistics are sparse, lagged, or unreliable. In the commodity context,

the analogue is using satellite data to provide near-real-time information about crop conditions that may not be in slower public reports or in market prices.

2.5 Ethanol and the corn basis link

[Tyner(2010)] argues that the structural shift toward ethanol fundamentally altered corn price dynamics by linking corn to the energy complex. [McPhail(2011)] uses a structural VAR to show that ethanol production responds to fuel prices and feeds back into corn demand. Within the Midwest, ethanol plant capacity is concentrated and weekly ethanol production is a credible proxy for the marginal demand-side bid for corn.

3 Data

The empirical analysis combines five data sources covering April 2019 through October 2024 across seven states (IA, IL, IN, MN, NE, OH, SD), giving a balanced panel of $7 \times 7 \times 6 = 294$ monthly observations.

3.1 Cash prices and basis

Monthly state-level corn cash prices come from USDA NASS Quick Stats (commodity = CORN, category = PRICE RECEIVED, unit = \$/bu, agg = STATE). For each state-month, the basis is cash minus the corresponding nearby CME corn front-month futures settlement: May for April, July for May to July, September for August, December for September to October. The futures settlement is the last available daily settle on or before the month-end.

We use NASS Prices Received because it is the only consistent state-level monthly price source available across all seven states. NASS prices reflect a broad set of transactions including forward sales and may differ from spot elevator quotes by 20 to 30 cents per bushel in any given state-month. We use leave-one-year-out historical baselines for all anomaly variables, so the first usable forecast year is 2021.

3.2 Basis descriptive statistics

Table 1 summarizes the basis by state. State-level means range from -0.17 to $+0.12$ \$/bu and within-state standard deviations from 0.43 to 0.57 \$/bu. Year-to-year variation dominates: 2021 to 2023 saw substantial premia across all states reflecting post-pandemic supply tightening, while 2019 to 2020 were more discounted.

Table 1: Descriptive statistics for monthly corn basis (\$/bu) by state, April to October 2019 to 2024 (N = 294).

State	N	Mean	Std	Min	Max	Range
IA	42	+0.085	0.513	-0.798	+1.417	2.215
IL	42	-0.009	0.451	-0.738	+1.297	2.035
IN	42	+0.115	0.480	-0.892	+1.117	2.009
MN	42	-0.168	0.565	-1.422	+0.968	2.390
NE	42	+0.083	0.557	-0.938	+1.368	2.306
OH	42	+0.057	0.430	-0.793	+0.838	1.631
SD	42	-0.083	0.510	-0.808	+1.227	2.035

3.3 NDVI extraction from Sentinel-2 imagery

The satellite vegetation signal is constructed from the Sentinel-2 Surface Reflectance Harmonized collection (COPERNICUS/S2_SR_HARMONIZED) accessed through Google Earth Engine. Within each state we restrict to pixels classified as corn by the USDA Cropland Data Layer (CDL crop code 1). To avoid look-ahead bias we use CDL year $t - 1$ to mask year t observations, since CDL is published with a one-year lag. Year-over-year corn pixel persistence in the Corn Belt is high enough that this introduces only modest noise.

The growing season is defined as April 15 through October 15. Within this window, Sentinel-2 imagery is grouped into seven-day windows and a weekly median NDVI composite is computed pixel-wise before averaging over all corn pixels in the state. Cloud masking uses Google Cloud-Score+ (GOOGLE/CLOUD_SCORE_PLUS/V1/S2_HARMONIZED) with pixels below 0.60 in the `cs_cdf` band dropped. Weeks with valid coverage below 50% are flagged and excluded. A three-week centered moving average further smooths the series. The pipeline yields 1,134 raw weekly state-observations of which 196 are flagged as low-coverage and dropped, leaving 938 valid weekly observations.

3.4 Phenology-anchored NDVI feature

The headline NDVI feature is constructed as follows. For each state-week-of-year we compute a leave-one-year-out z-score on the smoothed NDVI series. At year yy y , the baseline mean and standard deviation are estimated using all years except yy y . The feature `silk_ndvi_z_full` is then the mean of weekly NDVI z-scores during week-of-year 30 through 35. This corresponds to roughly late July through early September, the silking-to-early-dent reproductive window for corn in the Corn Belt.

Three properties of this construction matter for the analysis. First, the window is fixed in length and timing, so the feature does not inherit the mechanical bias that arises when one takes the maximum over a growing number of weekly observations. Second, the window is phenologically interpretable, capturing the reproductive phase when NDVI is most informative about eventual yield. Third, the leave-one-year-out baseline ensures that the z-score does not absorb the very

anomaly it is meant to detect, which would happen if a year of uniformly low NDVI shifted both the numerator and denominator of a within-year z-score in the same direction.

3.5 USDA Crop Progress

Weekly USDA Crop Progress reports for the seven states are downloaded from NASS Quick Stats. We sum the share of the corn crop rated "Good" and "Excellent" into a single `pct_ge` measure for each state-week and compute a leave-one-year-out z-score per state and week-of-year. For each state-month observation we use the latest weekly z-score observed strictly before the start of that month, ensuring no contemporaneous leakage. Because at least two prior years are needed to estimate a non-degenerate standard deviation, the anomaly series is well-defined only from 2021 onward.

3.6 EIA fuel ethanol production

Weekly U.S. fuel ethanol production is taken from EIA series `W_EP00XE_YOP_NUS_MBBLD`. We construct two derived features: a four-week moving average of production and the year-over-year change in that moving average, both standardized to z-scores. As with Crop Progress, we use only data observed strictly before the start of the target month.

4 Methodology

4.1 Model 1: state-month seasonal benchmark

For each state s , year y , month m , the Model 1 forecast is the prior-year mean basis for the same state and month, requiring at least two years of training data. The first usable forecast is April 2021 for each state. This benchmark embodies the null hypothesis that basis is fully described by state-month seasonality.

4.2 Model 2: market baseline with real CME futures

Model 2 introduces information from the futures market. For each state s and observation t :

$$\text{basis}_{s,t} = \alpha + \beta_1 \text{basis}_{s,t-1} + \beta_2 F_{t-1} + \beta_3 r_{t-1}^F + \beta_4 \sin(2\pi m/12) + \beta_5 \cos(2\pi m/12) + \varepsilon_{s,t} \quad (1)$$

where $\text{basis}_{s,t-1}$ is the lagged within-state basis, F_{t-1} is the lagged real front-month CME corn settlement price, and r_{t-1}^F is the corresponding lagged monthly futures return. The model is estimated by ordinary least squares in an expanding window per state. Cross-year lags (October to April) are set missing.

4.3 Model 3: basis anomaly regression with non-market features

Model 3 changes the regression target. Rather than forecasting raw basis, we forecast the basis anomaly $a_{s,t} = \text{basis}_{s,t} - \hat{\mu}_{s,t}^{M1}$, the residual that Model 1 cannot explain. This focuses estimation on state-year deviations of basis from its long-run state-month seasonal pattern. The full Model 3 specification is:

$$a_{s,t} = \gamma_0 + \gamma_1 a_{s,t-1} + \gamma' x_{s,t} + \text{state FE} + u_{s,t} \quad (2)$$

where $x_{s,t}$ contains some or all of: `silk_ndvi_z_full`, `cp_ge_anom`, `ethanol_4wk_lag1_z`, and `ethanol_4wk_yoy_lag1_z`. All features use only data observed through the end of the month preceding the target observation.

Model 3 is estimated by pooling training observations across states with state fixed effects. With seven states and about fourteen training observations per state-year, single-state expanding windows would over-fit when there are three to five features, particularly when ethanol features have large absolute magnitudes. Pooling multiplies effective sample size by approximately seven while state fixed effects absorb level differences.

We evaluate five variants: *baseline* (lagged anomaly only), *ndvi*, *crop_progress*, *ethanol*, and *full*. Forecast accuracy is reported on the common evaluation set, that is, the intersection of state-date observations evaluable under all five variants. This ensures that all variants compete on the same observations and avoids spuriously favorable estimates for the most data-restricted variants.

4.4 Residual diagnostic (the headline test)

Cutting across the three models is the central hypothesis test. We compute Model 1 residuals $e_{s,t} = \text{basis}_{s,t} - \hat{\mu}_{s,t}^{M1}$ and run:

$$e_{s,t} = \theta_0 + \theta_1 \text{silk_ndvi_z_full}_{s,t} + \theta_2 \text{silk_ndvi_z_full}_{s,t} \cdot \text{harvest}_t + \text{state FE} + u_{s,t} \quad (3)$$

with cluster-robust standard errors clustered by state. The coefficient θ_2 is the object of interest. It tests whether the effect of NDVI on residual basis differs systematically between harvest and pre-harvest months. A negative θ_2 is the prediction of the local-supply mechanism: more vegetation during silking-through-dent implies more local supply at harvest, weakening the basis. We also run leave-one-state-out regressions to assess robustness.

4.5 Bootstrap confidence intervals

For the out-of-sample forecasting comparison, we report 95% cluster-bootstrap confidence intervals on harvest-window MAE for each feature group. The procedure resamples states with replacement (preserving within-state observation structure) for $B = 1000$ iterations and computes the percentile-based interval.

5 Results

5.1 Pooled regression with state fixed effects

Table 2 reports the headline pooled regression with state fixed effects and a harvest-month interaction term, with cluster-robust standard errors. The interaction coefficient is the central object: it tests whether NDVI’s effect on residual basis differs systematically between harvest and pre-harvest months.

Table 2: Pooled residual diagnostic with state fixed effects and harvest interaction. $N = 196$. $R^2 = 0.083$. Standard errors clustered by state.

Variable	Coefficient	Cluster SE	z	p-value
Intercept (base state IA)	+0.247	0.017	14.86	0.000
<code>silk_ndvi_z_full</code>	+0.089	0.072	1.24	0.214
<code>silk_ndvi_z_full</code> $\times$ harvest	−0.128	0.040	−3.19	0.001
State FE (IL, IN, MN, NE, OH, SD)	(included)	—	—	—

The interaction coefficient is -0.128 \$/bu per z-score, with a cluster-robust p-value of 0.001. The pre-harvest main effect is $+0.089$ ($p = 0.214$), consistent with the absence of a meaningful pre-harvest signal. The total harvest-window effect is therefore approximately $+0.089 + (-0.128) = -0.039$ \$/bu per z-score: small in absolute magnitude but precisely identified. Equally important, the difference between the harvest and pre-harvest effects (the interaction itself) is robustly identified, consistent with the local-supply mechanism activating in the harvest window when the NDVI-derived yield signal becomes economically relevant.

5.2 The role of state fixed effects

The role of state fixed effects in identifying the harvest-window effect is worth emphasizing. A naive pooled univariate regression of residuals on `silk_ndvi_z_full`, without state fixed effects, yields a full-sample slope of $+0.012$ ($p = 0.78$), a harvest-only slope of -0.064 ($p = 0.19$), and a pre-harvest slope of $+0.068$ ($p = 0.28$). All three slopes are small and statistically insignificant. The reason is that mean basis and mean NDVI differ substantially across the seven states. Pooling without fixed effects mixes within-state variation, which is what the local-supply mechanism predicts, with cross-state variation, which reflects unrelated structural differences (transportation distance to the Gulf, ethanol plant density, dominant marketing channel). Cross-state variation inflates residual variance and attenuates the within-state effect. State fixed effects absorb the cross-state component and isolate the within-state harvest effect, which is exactly what Table 2 reports.

5.3 Leave-one-state-out robustness

Table 3 reports the harvest-interaction coefficient when each state is dropped in turn. All seven leave-one-state-out estimates lie between -0.103 and -0.161 , all with p-values below 0.02. The

result is not driven by any single state.

Table 3: Leave-one-state-out estimates of the $\text{silk_ndvi_z_full} \times \text{harvest}$ coefficient. All seven estimates retain p-values below 0.02 and lie within $\pm 25\%$ of the pooled point estimate.

Dropped state	NDVI $\times$ harvest coef	p-value	N
IA	-0.131	0.003	168
IL	-0.121	0.006	168
IN	-0.113	0.020	168
MN	-0.134	0.004	168
NE	-0.161	0.000	168
OH	-0.141	0.001	168
SD	-0.103	0.002	168

The coefficient is largest in absolute value when Nebraska is dropped (-0.161 , $p = 0.0004$) and smallest when South Dakota is dropped (-0.103 , $p = 0.002$), but the variation is narrow. The headline result is therefore not contingent on any one state’s data.

5.4 Visual evidence

Figure 1 visualizes the relationship in the data. The left panel shows the full pooled sample; the right panel restricts to harvest months. The harvest-month panel has a slightly negative slope but with substantial cross-state dispersion around it. The fact that the headline interaction coefficient is highly significant despite the visually noisy scatter underscores the role of state fixed effects in absorbing cross-sectional level differences and isolating the within-state harvest effect.




Figure 1: Model 1 residuals plotted against silk_ndvi_z_full pooled across all seven states, full sample (left) and harvest months only (right). Color encodes month. The dashed red line is the unconditional univariate fit (no state fixed effects).

5.5 Supporting analysis: NDVI compared to USDA Crop Progress

A natural concern is that USDA Crop Progress reports already publish weekly state-level corn condition ratings derived from on-the-ground enumerator surveys. If these condition ratings carry overlapping information with NDVI, the satellite signal might be partially redundant. We address this by including both NDVI and Crop Progress (and their harvest interactions) in the same residual regression. The result is in Table 4.

Table 4: Joint regression of Model 1 residuals on NDVI and Crop Progress. $N = 131$ (sample restricted to observations with non-missing CP anomaly, from June 2021 onward). State FE included (omitted from table). $R^2 = 0.166$.

Variable	Coefficient	Cluster SE	z	p-value
<code>silk_ndvi_z_full</code>	+0.039	0.110	0.35	0.727
<code>silk_ndvi_z_full</code> $\times$ harvest	−0.019	0.109	−0.18	0.860
<code>cp_ge_anom</code>	−0.097	0.036	−2.71	0.007
<code>cp_ge_anom</code> $\times$ harvest	−0.017	0.060	−0.29	0.774

When Crop Progress is added alongside NDVI, the satellite signal is dominated. The `silk_ndvi_z_full` $\times$ harvest coefficient drops from −0.128 (highly significant in isolation) to −0.019 ($p = 0.86$). `cp_ge_anom` retains a significant negative main effect (−0.097, $p = 0.007$). The natural reading is that NDVI and USDA Crop Progress measure overlapping aspects of the same underlying state of the corn crop. We treat this as supporting evidence that constrains the interpretation of the headline result: NDVI’s harvest-window signal is real, but at this monthly state-level aggregation it does not carry much information beyond what USDA already publishes.

5.6 Out-of-sample forecasting evaluation

Table 5 reports harvest-window MAE for each feature group along with cluster-bootstrap 95% confidence intervals (1000 replicates resampling states with replacement). All five augmented variants and the pure baseline (anomaly AR(1)) sit within overlapping bootstrap intervals.

Table 5: Out-of-sample harvest-window MAE (\$/bu) by feature group, with 95% cluster-bootstrap confidence intervals (1000 replicates, resampling states). Common evaluable sample: 96 state-month observations, of which 63 are harvest months. All feature group intervals overlap with the baseline AR(1) interval.

Feature group	MAE point	CI low	CI high	N
Model 1 (seasonal)	0.376	0.344	0.406	63
Baseline (anomaly AR(1))	0.241	0.220	0.263	63
+ NDVI	0.251	0.233	0.271	63
+ Crop Progress	0.257	0.233	0.281	63
+ Ethanol	0.327	0.304	0.352	63
Full (all features)	0.315	0.274	0.358	63

Two patterns deserve emphasis. First, the largest forecasting gain over the seasonal Model 1 benchmark comes simply from changing the target to the basis anomaly. The anomaly AR(1) baseline reduces MAE from 0.376 to 0.241 in the harvest window. This reflects the fact that a substantial share of basis variation is in the year-state fixed effect that the lagged anomaly absorbs. Second, adding NDVI, Crop Progress, or ethanol features on top of the AR(1) baseline produces small point-estimate movements that are indistinguishable from sampling noise.

This null forecasting result is consistent with the residual diagnostic. The headline NDVI $\times$ harvest interaction is small (-0.128 \$/bu per z-score) and its information overlaps with what the lagged anomaly already absorbs. A small but real signal can fail to translate into out-of-sample MAE improvements when the residual variance is dominated by other sources of variation and when training samples are small relative to the number of features.

Figure 2 plots actual versus predicted basis from Model 2 by state. The model tracks the broad seasonal swings well in most states but consistently misses the largest local supply-demand spikes, most visibly the August 2022 peaks of $+1.0$ to $+1.4$ \$/bu observed across IA, IL, IN, MN, NE, and SD. These are precisely the observations one would expect a futures-market baseline to miss. They reflect local-supply or local-demand shocks that the national futures contract cannot fully reflect, and they are also where local-information signals such as NDVI should in principle have the most to add.

6 Discussion

6.1 Interpreting the headline coefficient

The headline interaction coefficient of -0.128 \$/bu per z-score means that the effect of NDVI on residual basis is about 12.8 cents more negative in harvest months than in pre-harvest months. This should not be interpreted as the total harvest-window effect. Since the pre-harvest main effect is $+0.089$, the total harvest-window effect is

$$0.089 - 0.128 = -0.039.$$

Thus, a one-standard-deviation increase in silking-window NDVI is associated with a harvest-window basis weakening of about 3.9 cents per bushel. The effect is economically modest, especially relative to the unconditional harvest-window residual volatility, but the harvest interaction is precisely estimated.

6.2 The forecasting null result

The bootstrap confidence intervals on harvest-window MAE all overlap the anomaly AR(1) baseline. No augmented feature group reliably improves out-of-sample performance at the monthly frequency in this sample. We do not interpret this as evidence against the headline residual diagnostic. A




Figure 2: Model 2 (real CME futures baseline): actual versus predicted basis by state, expanding-window forecasts April 2021 to October 2024. Common harvest-window MAE = 0.293 \$/bu pooled across states; range from 0.235 (IA) to 0.414 (OH).

small but real signal can fail to translate into out-of-sample MAE improvements when the residual variance is dominated by other sources of variation, when training samples are small relative to the number of features, and when the marginal information overlaps with what is already absorbed by simpler predictors. All three conditions hold here.

6.3 Limitations

Three limitations should be flagged. First, even with multi-state expansion the sample is small in absolute terms: 196 forecast observations spanning four crop years and seven states. Confidence intervals on point estimates are correspondingly wide. Second, the monthly frequency throws away within-month variation in NDVI, Crop Progress, ethanol production, and futures prices. A weekly basis series would multiply effective sample size by approximately four, but constructing a consistent

weekly multi-state basis series from elevator-level cash quotes is a substantial undertaking that we leave to future work. Third, the NDVI feature is a state-level mean over corn pixels, which discards within-state spatial heterogeneity. A multi-resolution analysis using crop reporting districts or county-level basis quotes would test whether finer spatial scale recovers the marginal information that the state-level feature loses.

6.4 Why not ARIMA or machine learning?

We anticipate the question of whether more flexible model classes would produce better forecasts. Three reasons argue against this. First, a standard ARIMA selection procedure on a stationary basis-anomaly series would almost certainly identify an AR(1), which is what the Model 3 baseline already estimates. Second, with effective harvest-window samples of 63 observations, any model with more than three or four free parameters per state would be overfitting noise. Third, the research question is identification rather than prediction.

7 Conclusion

This paper has tested whether satellite vegetation observations contain information about U.S. corn basis that is not already in market prices or in conventional agricultural indicators, using a balanced seven-state monthly panel covering Iowa, Illinois, Indiana, Minnesota, Nebraska, Ohio, and South Dakota over April to October 2019 to 2024. The central finding is that, in a pooled regression with state fixed effects and a harvest-month interaction term, the silking-window NDVI z-score has a precisely identified negative effect on residual basis specific to the harvest window. The interaction coefficient is -0.128 \$/bu per z-score ($p = 0.001$ with cluster-robust standard errors) and is robust across all seven leave-one-state-out specifications (range -0.10 to -0.16 , all $p < 0.02$). The total harvest-window NDVI effect is small in absolute magnitude (about -0.04 \$/bu per z-score) but is precisely identified.

The headline finding should be read with two contextual qualifications. First, no augmented feature group reliably outperforms a simple anomaly AR(1) baseline in out-of-sample forecasting accuracy within bootstrap confidence intervals. The signal is identifiable in a properly specified regression but does not translate into materially better point forecasts at the monthly frequency on this sample. Second, supporting analysis indicates that USDA Crop Progress captures overlapping information with NDVI at this aggregation: when both are included, NDVI’s harvest interaction becomes statistically insignificant. These qualifications constrain the practical reach of the result without overturning the core identification finding.

For the practical question of whether commodity traders or agricultural cooperatives should incorporate satellite NDVI into their basis forecasting, the answer suggested by this analysis is qualified. The marginal value over publicly available USDA Crop Progress reports appears small at monthly frequency and state-level aggregation. Whether NDVI offers larger marginal value at finer spatial resolution, at higher temporal frequency, or in settings where Crop Progress is delayed or

unavailable remains open.

For future research, four directions stand out. First, a higher-frequency basis series, weekly elevator quotes or daily cash bids, would allow tests of the timing of information incorporation. Second, the spatial scale could be reduced from state level to crop reporting district or county. Third, the analysis could be repeated for soybeans, wheat, or other commodities. Fourth, an event-study design centered on USDA Crop Progress release dates would test directly whether NDVI provides leading information relative to the official survey. The methodological contribution of the present paper, a state-FE pooled regression with harvest interaction as a small-panel-friendly test of conditional information content, transfers naturally to all of these settings.

References

- [Adjemian and Smith(2012)] Adjemian, M. K. and Smith, A. (2012). Using USDA forecasts to estimate the price flexibility of demand for agricultural commodities. *American Journal of Agricultural Economics*, 94(4), 978–995.
- [Becker-Reshef et al.(2010)] Becker-Reshef, I., Vermote, E., Lindeman, M., and Justice, C. (2010). A generalized regression-based model for forecasting winter wheat yields in Kansas and Ukraine using MODIS data. *Remote Sensing of Environment*, 114(6), 1312–1323.
- [Bolton and Friedl(2013)] Bolton, D. K. and Friedl, M. A. (2013). Forecasting crop yield using remotely sensed vegetation indices and crop phenology metrics. *Agricultural and Forest Meteorology*, 173, 74–84.
- [Donaldson and Storeygard(2016)] Donaldson, D. and Storeygard, A. (2016). The view from above: Applications of satellite data in economics. *Journal of Economic Perspectives*, 30(4), 171–198.
- [Fama(1970)] Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- [Garcia et al.(2015)] Garcia, P., Irwin, S. H., and Smith, A. (2015). Futures market failure? *American Journal of Agricultural Economics*, 97(1), 40–64.
- [Johnson(2014)] Johnson, D. M. (2014). An assessment of pre- and within-season remotely sensed variables for forecasting corn and soybean yields in the United States. *Remote Sensing of Environment*, 141, 116–128.
- [McPhail(2011)] McPhail, L. L. (2011). Assessing the impact of US ethanol on fossil fuel markets: A structural VAR approach. *Energy Economics*, 33(6), 1177–1185.
- [Sanders and Baker(2012)] Sanders, D. R. and Baker, T. G. (2012). Forecasting corn and soybean basis using regime-switching models. Working paper, Purdue University.

- [Taylor et al.(2006)] Taylor, M. R., Dhuyvetter, K. C., and Kastens, T. L. (2006). Forecasting Kansas wheat basis using historical averages and current market data. *Journal of Agricultural and Resource Economics*, 31(2), 191–211.
- [Tomek and Robinson(1990)] Tomek, W. G. and Robinson, K. L. (1990). *Agricultural Product Prices*, 3rd edition. Cornell University Press, Ithaca, NY.
- [Tyner(2010)] Tyner, W. E. (2010). The integration of energy and agricultural markets. *Agricultural Economics*, 41(s1), 193–201.
- [Working(1949)] Working, H. (1949). The theory of price of storage. *American Economic Review*, 39(6), 1254–1262.