Commodity Supply News Shocks: Evidence from Scheduled Announcements

Suraj Kumar

Second draft: October 19, 2026

Abstract

I propose a novel method to estimate agricultural supply news shocks using the U.S. Department of Agriculture's (USDA) crop-specific yield forecasts, which are revised monthly over the agricultural season, kept confidential until release, and driven largely by weather. I orthogonalize these yield revisions with respect to their own lags and lagged real crop prices, to purge the predictable direction of revisions between consecutive months within a season and any expected component already priced into commodity prices. The resulting instrument has a robust first-stage F-statistic above 30. Compared with Adjemian and Jo (2025), the orthogonalized revision shocks have a much more muted impact on real activity and a delayed but significant effect on food and core prices.

1 Introduction

Motivation. Commodity prices have a fundamental effect on the economy. Agricultural commodities such as corn, rice, wheat, and soybeans influence the prices of commonly consumed food products such as bread. In addition, by serving as feed for livestock, they further influence the prices of animal products such as milk, meat, and eggs. Moreover, the food prices consumers observe influence their inflation expectations (D'Acunto et al., 2021). These expectations may then feed into core inflation, which is closely monitored by monetary authorities such as the U.S. Federal Reserve.

While significant attention has been paid to energy commodities and their influence on the economy and monetary policy, I find that similar attention is lacking in the agricultural commodities space. This is especially true of studies that estimate the impact of food commodity price shocks on the economy. Thus, in this study, I examine the impact of adverse agricultural supply news shocks on U.S. inflation and activity after purging the predictable revisions and information already in the price.

I hypothesize that examining only high-frequency changes in agricultural commodity futures is an imperfect way to isolate pure agricultural supply news. The responses of the Baltic Dry Index and oil production in Adjemian and Jo (2025) suggest that the futures-based instrument may still contain demand-side information. Thus, there is a need to better separate demand shocks from supply shocks in the USDA reports.

Method: Orthogonalized Yield Revision. Specifically, I propose measuring agricultural supply news shocks as orthogonalized yield revision shocks. Each month, the USDA publishes the latest expected yield for different crops in the U.S. While weather information is known to the market by the time a WASDE report comes out, the impact of weather on yield is not fully known or priced by the market. Unlike the planting or area harvested information, which may be known beforehand, the yield information is not public before the USDA release. Moreover, the yield revision is unlikely to be driven by any domestic economic factor.[1]

However, USDA yield revisions are not a clean instrument per se, as the USDA has historically revised the yield in the same direction in consecutive months of an agricultural season. Furthermore, the USDA's WASDE reports only come out in the middle of the month, so the market has partly priced the weather information into the agricultural commodity price.[2] This leads to the current month's revision being predictable from the previous month's price (refer to Appendix C).[3] So, in order to remove these two potential sources of endogeneity, I orthogonalize the yield revisions by removing the effect of past revisions and past real commodity prices. The orthogonalized instrument has a robust F-statistic (> 30) and does not suffer from the weak instrument problem.

Results. I then use these yield revisions as an instrument for the real field crop price in a proxy SVAR setup to gauge the macroeconomic impact, in a way quite similar to Adjemian and Jo (2025). I find that adverse yield revision shocks, signaled by a one percent rise in the real field crop price, imply a much more muted response of real activity, with real GDP falling by 0.09 percent at about ten months, and industrial production by 0.1 percent at month three before rebounding into expansion. This one percent adverse supply news shock causes a delayed pass-through to prices: the core price index rises by about 0.02 percent and the food-at-home price index by about 0.09 percent after about two and a half years, a much slower transmission to food and core inflation than implied by the high-frequency changes in

[1]The alternative, which I started out with as the premise for this study, was to use a large language model to analyze the text of the WASDE release to identify more exogenous factors behind the production change estimates for a commodity. This looked promising in the beginning, but later I found that the text does not materially add any new information, as the yield revision shows no sign of being driven by macroeconomic variables (Table 4) or widely used macroeconomic shocks (Table 5). Moreover, for yield revisions, the numerical data are already sufficient. I have extracted the economic drivers behind each USDA revision using the LLM *DeepSeek V4.1 Flash*. However, I was not able to come up with a way to use this structured information to enhance the shocks. Refer to Appendix E for more information.

[2]The market may have perfect information about the weather but not perfect information about how the weather affected crop yields, which the USDA knows from its internal surveys.

[3]I have also tried using the Reuters Eikon analyst expectations about the yield and production estimates for the three main commodities (corn, soybeans and wheat). However, they also suffer from a similar serial correlation, where the analysts are serially surprised in one direction. Refer to Appendix B for more information.

agricultural commodity futures. This is in sharp contrast to the futures shocks of Adjemian and Jo (2025): in my replication, GDP and IP trough at about -0.18 percent after about a year and the food-at-home price peaks at 0.16 percent after a year. Moreover, yield revision shocks imply a fall in the S&P 500 and a rise in the VIX on impact, in contrast to the futures shocks considered in Adjemian and Jo (2025), although these impact responses are imprecisely estimated. The real oil price and the Baltic Dry Index[4] fall in the first year and rise only later, consistent with weaker activity followed by recovering demand; the size of the delayed rise in the oil price, however, remains hard to explain.

Contribution. The study contributes to the existing literature in three ways. First, it suggests a new way to measure agricultural supply shocks, using the USDA's own yield revisions after removing their predictable part. In contrast to Adjemian and Jo (2025), since it is not a price reaction, it is unlikely to pick up any demand-related news that moves prices on the release day. Second, this study builds the shocks on a monthly basis from the forecast revisions of production volumes attributed to the yield instead of decomposing annual actual production into harvest quarters, as done in De Winne and Peersman (2016). Thus, it provides a more real-time update of agricultural supply conditions compared to De Winne and Peersman (2016). Third, compared with Adjemian and Jo (2025), my study estimates a much more muted impact of agricultural supply shocks on real activity, but a significant, delayed pass-through to food and core prices.

Literature Review. The yield revision as a shock is in some ways similar to De Winne and Peersman (2016), who identify the shock from actual harvest volumes, whereas mine is monthly news about upcoming harvests. The shocks of De Winne and Peersman (2016) depend on the production measure, which includes the harvested area. They rely on the lag between planting and harvest to argue that farmers cannot adjust it within the quarter. Since my instrument is a revision to expected yield, it is by design orthogonal to the harvested-area response.[5]

To the best of my knowledge, Adjemian and Jo (2025) is the only study using high-frequency agricultural commodity futures changes as the measure of agricultural commodity shocks on the macroeconomy along the lines of Känzig (2021). Adjemian and Jo (2025) estimate the impact of food commodity shocks on the economy by examining changes in futures prices for agricultural commodities such as corn, rice, wheat, oats, and soybeans on the day of announcements of the USDA's monthly estimates of agricultural supply and demand. They argue that the information provided by the USDA is primarily supply-oriented. They, however, do note that WASDE reports contain estimates of use and exports representing demand factors, but these estimates rely largely on publicly available data and therefore should not be considered "news".[6] Adjemian and Jo (2025) document that adverse

[4]The Baltic Dry Index measures average freight rates for dry-bulk vessels that carry iron ore, coal, grain, and fertilizers across the ocean.

[5]Yield may still respond to economic choices such as input intensity at planting, or abandonment of low-yield fields, so I still need the yield revision to be orthogonal to the other structural shocks.

[6]Adjemian and Jo (2025) consider robustness tests that exclude all crop inventory reports (USDA's Grain Stocks reports) and restrict attention to crop production report months, such as WASDE releases in May, July, August, September, October, and November. They find no significant difference between the robustness

agricultural supply shocks, which are associated with a rise in the real price of field crops, lead to a decline in real GDP and industrial production, and to higher core inflation and food-at-home inflation.[7]

Adjemian and Jo (2025)'s findings suggest that adverse agricultural supply news is associated with an initial increase in both grain exports and the Baltic Dry Index, followed by an eventual decline. They argue that the initial increase in both is driven by higher demand for grain exports, as importers seek to bolster supplies ahead of a poor harvest. This initial increase is followed by a persistent decline as the poor harvest materializes. They also suggest that the decline in oil production is likely driven by higher corn prices caused by a bad harvest. Since corn is a key input into ethanol production, a rise in corn prices can raise ethanol prices and spill over to other energy supply chains, leading to lower energy production. In contrast, De Winne and Peersman (2016) do not find an immediate decline in oil production, but find a decline only after some quarters. They argue that the lagged fall in oil production is driven by weakening global economic activity due to adverse food supply shocks.

Roadmap. I begin with a brief description of the data in Section 2, followed by a brief description of the framework of the proxy SVAR in Section 3. Section 4 details the identification strategy. I present the results of the first stage, the proxy SVAR and the instrument diagnostic checks in Section 5. Finally, Section 6 concludes with a plan for better identification of propagation mechanisms.

2 Data

2.1 USDA Data

The U.S. Department of Agriculture regularly publishes several reports on the state of agriculture, such as the World Agricultural Supply and Demand Estimates (WASDE), Grain Stocks, Prospective Plantings and Acreage.

WASDE. World Agricultural Supply and Demand Estimates reports are the most notable of the USDA reports. They are released monthly, around the second week of the month. They are prepared by the USDA's World Agricultural Outlook Board (WAOB) together with interagency committees. They evaluate the new survey data from the USDA's National Agricultural Statistics Service (NASS), information on foreign markets from U.S. government staff posted abroad, weather reports, satellite imagery and expert judgments. These reports are prepared in secrecy using the latest unpublished survey reports of NASS, which are released at the same time as the WASDE. These reports contain analytical forecasts of the

specifications and their baseline specification.

[7] I replicate their proxy SVAR results in Figure 7 using my own assembled panel. While my data period and variables are the same as theirs, due to data restrictions I have to use a different measure of real GDP than the one they use, and I correct twelve USDA report dates that are misprinted in their Appendix Table A.2; Table A.1 marks them.

balance sheets of major food commodities, such as supply, use, prices and ending stocks, for the U.S. and the world.

For example, the August 2010 WASDE wheat release shows that the production forecast for the 2010–11 agricultural season is raised by 49 million bushels from July, due to higher expected yields in the Northern Plains (see Figure 1). Despite such a large upward revision in supply, demand is also revised upwards, leading to an overall decline in ending stocks of 141 million bushels (see Figure 2). This release shows that a USDA release is usually a mix of both supply and demand factors.

WASDE BALANCE SHEETS. Following Adjemian and Jo (2025), I use the U.S. balance sheets of five field crops: corn, soybeans, wheat, rice and oats. From every WASDE release between December 1987 and September 2026,[8] I record each balance-sheet item for the newest marketing year: area harvested, yield, production, beginning stocks, imports, domestic use, exports and ending stocks. For each item I keep two values, last month's projection and the current month's, so that the revision is computed within a single release.

CROP PRODUCTION. This report contains estimates of harvested area, yield and production by crop and state as of the first of the month. It is released at the same time as the WASDE.

On some release days, WASDE shares the day with other USDA releases such as Grain Stocks, Prospective Plantings or Acreage (refer to Panel B of Figure 3, Appendix A and Table A.1).

2.2 Financial Markets and Macroeconomic Data

COMMODITY FUTURES. I use agricultural futures data from the Bloomberg terminal for five commodities: wheat, corn, rice, oats, and soybeans. For each crop, I use the five generic contracts with the nearest maturities.[9] These commodity futures are traded on the Chicago Board of Trade (CBOT).

MACROECONOMIC DATA. I build a monthly panel of outcome variables to analyze the responses to agricultural supply-side shocks from January 1992 to September 2026. It includes variables such as output, prices, inflation, financial market variables and trade indicators. The real field crop price, the variable instrumented by the supply shocks, is constructed as the weighted average of the NASS monthly prices received for the five crops, with each crop's annual share of U.S. production as its weight, and is then deflated by the U.S. CPI.[10] Refer to Table 1 for more details.

[8] USDA publishes these reports in three formats, so I read them in three ways: machine-readable XML from July 2010, text files from January 1995 to June 2010, and scanned PDFs before 1995.

[9] For corn, as of September 1, 2026, the five contracts by maturity are C 1 Comdty (September 2026), C 2 Comdty (December 2026), C 3 Comdty (March 2027), C 4 Comdty (May 2027), and C 5 Comdty (July 2027). The current Bloomberg export uses the PX_LAST field, which I found to be the same as Bloomberg's settlement field, PX_SETTLE.

[10] The NASS monthly prices for the crops are monthly averages, and the CPI used to construct the real index is the CPI for All Urban Consumers (all items; FRED code CPIAUCSL).

3 Econometric Framework

Here I briefly present the econometric details of the proxy SVAR approach used in this paper. For a more detailed explanation, please refer to Känzig (2021).

Let us assume that we have a reduced-form VAR(L) model:

$$\mathbf{y}_t = \mathbf{b} + \mathbf{B}_1\mathbf{y}_{t-1} + \ldots + \mathbf{B}_L\mathbf{y}_{t-L} + \mathbf{u}_t \tag{1}$$

where L denotes the lag order of the VAR, $\mathbf{y}_t$ is an $n \times 1$ vector of variables, $\mathbf{B}_i$ is an $n \times n$ matrix of coefficients for lag i, $\mathbf{u}_t$ is an $n \times 1$ vector of reduced-form innovations with covariance matrix $\text{var}(\mathbf{u}_t) = \Sigma$, and $\mathbf{b}$ is an $n \times 1$ vector of constants.

I assume that the reduced-form innovations are related to the structural shocks by a linear transformationas follows:

$$\mathbf{u}_t = \mathbf{S}\varepsilon_t \tag{2}$$

where $\mathbf{S}$ is an $n \times n$ nonsingular structural impact matrix, and ε_t is an $n \times 1$ vector of structural shocks. As the structural shocks are assumed to be independent of each other, $\text{var}(\varepsilon_t) = \mathbf{\Omega}$, where $\mathbf{\Omega}$ is a diagonal matrix. From equation (2), we can derive

$$\mathbf{\Sigma} = \mathbf{S}\mathbf{\Omega}\mathbf{S}' \tag{3}$$

Let us assume, without loss of generality, that the agricultural supply news shock is the first shock in the VAR, $\varepsilon_{1,t}$. We aim to estimate the structural impact vector $\mathbf{s}_1$, which is the first column of $\mathbf{S}$.

In an external instruments approach, I assume that the background noise in the surprise or instrument is negligibleto identify the structural impact vector. Let us assume that there is an external instrument z_t; for z_t to be a valid instrument, we require

$$\text{Relevance: } E[z_t\varepsilon_{1,t}] = \alpha \neq 0 \tag{4}$$
$$\text{Exogeneity: } E[z_t\varepsilon_{2:n,t}] = \mathbf{0} \tag{5}$$

where $\varepsilon_{1,t}$ is the agricultural supply shock, and $\varepsilon_{2:n,t}$ is an $(n-1) \times 1$ vector of other structural shocks. Under assumptions $(4) - (5)$, we can identify $\mathbf{s_1}$ as

$$\tilde{\mathbf{s}}_{2:n,1} \equiv \mathbf{s}_{2:n,1}/s_{1,1} = E[z_t\mathbf{u}_{2:n,t}]/E[z_t u_{1,t}] \tag{6}$$

provided that $E[z_t u_{1,t}] \neq 0$. Känzig (2021) notes that $\tilde{\mathbf{s}}_{2:n,1}$ is the population analogue of the IV estimator of $\mathbf{u}_{2:n,t}$ on $u_{1,t}$ using z_t as an instrument. The structural impact vector is $\mathbf{s}_1 = (s_{1,1}, \mathbf{s}'_{2:n,1} s_{1,1})'$. We can set the scale of $s_{1,1}$ by normalization subject to $\mathbf{\Sigma} = \mathbf{S\Omega S}'$. There are two approaches to scaling $s_{1,1}$. First, setting $\mathbf{\Omega} = \mathbf{I}_n$ implies that a unit positive value of $\varepsilon_{1,t}$, equivalent to one standard deviation of the structural shock, has an impact of $s_{1,1}$ on $y_{1,t}$, where we normalize $s_{1,1} = 1$. Alternatively, we can set $\mathbf{\Omega} = \text{diag}(\sigma^2_{\varepsilon_1}, \ldots, \sigma^2_{\varepsilon_n})$, and $s_{1,1} = x$ implies that a unit positive value of $\varepsilon_{1,t}$ has a positive effect of magnitude x on $y_{1,t}$.

4 Identification

In this section, I construct two main instruments to estimate agricultural supply shocks. The instruments must satisfy the relevance (equation 4) and exogeneity (equation 5) conditions stated in Section 3. The first approach relies on high-frequency changes in agricultural commodity futures prices, and the latter on the crop-specific revisions of the yield forecast for an agricultural season, updated monthly by the USDA.

4.1 The Futures Surprise (Benchmark)

I use the changes in agricultural commodity futures around the WASDE announcement to construct a benchmark agricultural supply shock to compare against the WASDE revision. This approach is well established in the literature since Känzig (2021)'s study of oil futures price changes around OPEC announcements. Recently, Adjemian and Jo (2025) have used this approach to estimate the agricultural supply shocks.

For this approach of high-frequency changes in commodity prices to signal supply-side news, it is essential to ensure that the price movements around the USDA's WASDE release are not driven by demand factors. Känzig (2021) argues that, by considering a sufficiently tight window around the OPEC announcement, it is possible to isolate the supply-side factors. For oil, the demand factors driven by global economic conditions are known to market participants and unlikely to change in the tight window around the announcement. Further, OPEC announcements do not affect the economy through an information channel, as monetary policy does, which makes it harder to interpret high-frequency asset price changes as a measure of monetary policy.

Adjemian and Jo (2025) argue that around the USDA announcement, "supply news" dominates the market. They state that, although the USDA reports both demand-side and supply-side components in the release, it is the confidential supply-side news, such as new production or yield estimates, that is not known to the market prior to the release. They indicate that demand-side information, such as commodity use and exports, is based on widely available public data, so it should not be considered news. However, they do note that some USDA announcements also supply confidential demand-side news, such as when the

Grain Stocks reports are published together with the WASDE release; they run robustness checks that exclude such releases from their sample or focus on the reports that only provide new crop information to the market. They find that the results are not significantly different from their baseline results, both qualitatively and quantitatively.

I use agricultural futures data from the Bloomberg terminal for five commodities: wheat, corn, rice, oats, and soybeans. For each crop, I use the five generic contracts with the nearest maturities. These commodity futures are traded on the Chicago Board of Trade (CBOT). Following Känzig (2021); Adjemian and Jo (2025), I construct the USDA announcement surprise by measuring the change in U.S. field-crop futures around the announcement of USDA WASDE reports from 1992 to 2024. Specifically, I estimate individual field-crop surprises as the natural-log difference between the futures settlement price on the announcement day and the settlement price on the last trading day before the USDA report.

$$s_{i,d}^{h} = 100 \times \left(\log F_{i,d}^{h} - \log F_{i,d-1}^{h}\right) \tag{7}$$

where $i \in \{\text{corn}, \text{soybeans}, \text{wheat}, \text{oats}, \text{rice}\}$ indexes the five commodities, $h = 1, \ldots, 5$ indexes the five nearest-expiration futures contracts by maturity, $F_{i,d}^{h}$ denotes the futures price used for commodity i's h-th nearest contract on trading day d, and $d-1$ denotes the last trading day before the report.

The futures contracts of different maturities of a given commodity largely move together and carry overlapping information. Looking across the five maturities allows me to capture a common announcement response rather than relying on a single contract that may be affected by a maturity-specific liquidity or roll issue. I therefore summarize them in a single series using the first principal component (PC1), estimated separately for each commodity:

$$s_{i,d} = \boldsymbol{v}_i' \, \boldsymbol{s}_{i,d}, \qquad \mathbf{s}_{i,d} = \left(s_{i,d}^{1}, \ldots, s_{i,d}^{5}\right)', \tag{8}$$

where $\mathbf{v}_i$ is normalized to unit sum and oriented so that a price increase loads positively.[11]

Thus, the constructed commodity-specific price surprises can be aggregated, using annual U.S. production shares, into a production-weighted WASDE announcement surprise for trading day d as follows:

$$s_d = \sum_i w_{i,y} \, s_{i,d}, \quad w_{i,y} = \frac{Q_{i,y}}{\sum_j Q_{j,y}} \tag{9}$$

[11]Principal component analysis (PCA) summarizes correlated variables with a smaller number of uncorrelated linear combinations. Applied to the five maturity returns $\mathbf{s}_{i,d}$, PC1 is the linear combination with the largest sample variance. Because the five maturities respond to the same announcement, this component captures their common price movement. I normalize $\mathbf{v}_i$ to unit sum and orient it so that a price increase loads positively. With unit-sum loadings, the score is a weighted average of the five returns and is expressed in the same percent units as the individual surprises. In my data, the PC1 factor loads almost equally across the five maturities.

where $Q_{i,y}$ is U.S. production[12] of commodity i in year y (in million metric tons) and the weights are held fixed within the year.

There are various ways to aggregate these daily surprises (trading day d) to a monthly frequency (month t). Adjemian and Jo (2025) follow the convention of summing the daily surprises in a month to obtain the monthly surprise. Gertler and Karadi (2015) suggest an alternative approach when the variables in the monthly VAR are monthly averages instead of end-of-month values. A surprise at the end of a month is expected to have a smaller influence on the monthly average than a surprise at the beginning of the month. This applies here because the NASS price received is a monthly average. Hence, Gertler and Karadi (2015) propose, first, for each day of the month, to cumulate the USDA report-day surprises over the last 31 days (e.g., on May 15, cumulate all surprises since April 15), and second, to average these 31-day sums across the days of the month.[13]In this paper, I refer to these two aggregations as "sum" for summing the daily surprises within the month, and as "GK" for the approach of Gertler and Karadi (2015). I estimate the strength of both AJ's instrument and my proposed instruments under both aggregations.

However, in my view, examining only high-frequency changes in agricultural commodity futures is an imperfect way to isolate agricultural supply news. The futures price moves on the release day capture the market's reaction to everything in the report, including demand-side items such as exports, use and stocks. The immediate rise in the S&P 500 and in the Baltic Dry Index after the shock in Figure 3 of Adjemian and Jo (2025) suggests that futures prices carry some demand-side news as well. In the next subsection, I propose a novel way to estimate these agricultural supply shocks using the USDA's monthly WASDE revisions of a crop's seasonal forecasts of yield, area and production, which measure supply directly from the USDA's own surveys.

4.2 The WASDE Revision

The USDA WASDE release measures total use (TU) of the commodity, such as domestic use (DU) and exports (EX), and the supply components (TS), such as beginning stocks (BS), production (P) and imports (IM), and records the difference between total supply and total use as ending stocks (ES).[14]

[12]For the incomplete 2026 production year, I carry forward the fixed 2025 production shares for all 2026 WASDE events. This avoids using month-by-month 2026 forecasts, which would make the aggregation weights vary with information arriving during the sample; the 2026 weights are therefore provisional and will be revisited once post-harvest estimates are available.

[13]I implement this with the equivalent form given by Gertler and Karadi (2015): cumulate all report-day surprises into a daily series, average it within each month, and take the first difference of these monthly averages. A surprise on day d of a D-day month then enters that month with weight $(D-d+1)/D$ and the next month with weight $(d-1)/D$. The two forms give identical weights when both months have 31 days and differ negligibly otherwise.

[14]By the balance-sheet identity, $ES = BS + P + IM - DU - EX$. The revision in the rest of the balance sheet, $r^R_{i,t} = 100\,[(ES_{i,t} - ES_{i,t-1}) - (P_{i,t} - P_{i,t-1})]/TU_{i,t-1}$, is the ending-stocks revision not explained by production, in percent of last month's total use. It collects the revisions to beginning stocks, imports, domestic use and exports.

I consider using month-to-month revisions in the production, area harvested and yields of different commodities from the WASDE releases as a source of potential instruments. The identification assumption requires that revisions are exogenous to every other structural shock in the month, including domestic and foreign demand shocks; that is, they are uncorrelated with those shocks (equation 5). The production revisions can be disaggregated into parts such as yield revisions and changes in the expected area harvested.[15]

Identification Assumption. I assume that changes in the yields for the agricultural season are mainly driven by weather and other growing conditions, which are likely exogenous to the other structural shocks in the release month. Once a crop is in the ground, farmers have little scope to change its yield: the main inputs, such as seed, fertilizer, and area planted, are chosen at or before planting, which precedes the harvest by three to six months for the spring-planted crops and by about nine months for winter wheat (Figure 3); see also De Winne and Peersman (2016). Farmers can still influence a yield by adjusting some inputs such as in-season fertilizers, fungicide application and timing of harvest, and these may respond to crop and energy prices.[16] However, I regard these margins as small relative to the weather. A second, more serious threat is that the weather can also affect the economy through channels other than crop supply. For instance, heat waves that reduce yields also raise electricity and natural-gas demand, which then feeds into energy prices and the CPI. The yield revision would then be correlated with a weather-driven energy-cost shock, which violates the exogeneity condition in equation (5). In the data, the orthogonalized yield revision is not significantly correlated with any of the published macroeconomic shocks in Table 5, including the oil supply and oil demand shocks.

The harvested-area revision requires a stronger assumption. Harvested area differs from planted area because some fields are abandoned or are put to other uses such as corn cut for silage instead of for grain, or winter wheat used for grazing by cattle. These decisions depend on prices: when the grain price is low relative to the cost of harvesting, or when forage becomes more valuable, farmers harvest less area for grain. It is very likely that early-season revisions of harvested area largely reflect planting-area updates.[17] Thus, the harvested-area revision can pick up demand-driven price movements, and it fails the exogeneity tests in Tables C.7 and C.8. I therefore use the yield revision as my main instrument.

I construct the potential instruments as the production-weighted revisions in an agricul-

[15] Production = Area Harvested × Yield

[16] In-season fertilizers such as nitrogen on corn, wheat and rice, are made from natural gas. So, a higher corn price or a lower fertilizer (or gas) price makes a larger application worthwhile. In deciding whether to apply fungicide against pest activity, which is partly caused by weather, a farmer has to weigh the cost of the application against the potential benefit at the expected crop price. At harvest, farmers can leave the grain to dry in the field, which saves the cost of drying fuel such as propane but raises field losses, or harvest early and dry the grain artificially. Each of these choices has a marginal effect on yield, likely smaller than the effect of the weather already experienced by the crop.

[17] For instance, most harvested-area revisions come in July, when the July WASDE factors in the updated Acreage report released at the end of June. Planted area is driven by spring weather but also by prices expected during planting, so these revisions can signal earlier demand news. The remaining harvested-area revisions come later in the harvest period and may reflect abandonment or the use of the crop for forage or grazing, depending on the crop prices.

tural commodity's harvested area, yield or production. I define the revision $r^X_{i,t}$ in commodity i in WASDE release month t for a WASDE production-related balance-sheet item (X) as:

$$r^X_{i,t} = 100 \ln\left(\frac{X_{i,t}}{X_{i,t-1}}\right) \tag{10}$$

where X represents harvested area (A), yield (Y) or production (P). $X_{i,t}$ denotes the estimate for the new marketing year and $X_{i,t-1}$ denotes the same estimate in the previous WASDE release.[18]

$$r^X_t = \sum_i w_{i,y} r^X_{i,t} \tag{11}$$

where r^X_t represents the production-weighted revisions in an agricultural commodity's production, harvested area or yield, with $w_{i,y}$ defined in the same way as in equation (9). Figure C.1 plots the resulting monthly instruments.

The raw revisions, however, are partly predictable, for two reasons. First, within a season the USDA tends to revise a crop's yield in the same direction in consecutive releases: for instance, in a drought year it may cut the estimate in August, again in September and again in October. Thus, the yield revisions are serially correlated (Figure C.3). However, this persistence does not show up as a seasonal pattern, because the direction of the revisions differs between good years and drought years (Figure C.2). The seasonality in the production revision comes mainly from the area revision, which is revised down in June (Figure C.2).

Second, the real crop price of month $t-1$ predicts the raw yield revision of month t (Table C.5), likely because the USDA releases its report in the middle of the month, by which time the market has partly priced in the observed weather. A revision that its own past or last month's price predicts is not news to the market on the release day, so I remove the predictable part.

4.3 Orthogonalization

Thus, I use an orthogonalization procedure to remove the effect of past revisions and past real commodity prices while constructing the instrument as follows:

$$r^X_t = \alpha^X + \sum_{k=1}^{12} \gamma^X_k r^X_{t-k} + \sum_{k=1}^{12} \delta^X_k p_{t-k} + \tilde{r}^X_t \tag{12}$$

[18] In May, when the new crop year first appears, I set all revisions ($r^X_{i,t}$) to zero. Also, when there is no new estimate, as in some late-winter months such as February and March, $r^X_{i,t}$ is zero because $X_{i,t} = X_{i,t-1}$.

where r_t^X is the five-crop revision of balance sheet item $X \in \{\text{area}, \text{yield}, \text{production}\}$ from equation (11), and p_t is 100 times the log of the seasonally adjusted real field crop price, the same series as in the VAR. I estimate the regression by OLS on 1993M1 to 2024M12 and scale the residual $\tilde{r}_t^X$ to unit standard deviation to serve as the instrument.

Thus, the constructed residual $\tilde{r}_t^X$ contains only the information in the revision that is not predicted by its own past or by the past price. For the yield revisions, it shows no autocorrelation, by design (Figure 6), and no seasonal pattern (Figure 5); however, area revisions and production revisions still show some seasonal pattern.[19] The orthogonalized instruments, scaled to unit standard deviation, are shown in Figure 4.

4.4 Release-day relevance

Table 2 reports the release-day relevance of the orthogonalized revisions at the crop level.[20] The orthogonalized yield revision significantly lowers the commodity futures curve on the day of the release. For instance, a 1% higher yield revision on a WASDE release day leads to the commodity futures price curve moving down by about 0.25% (column 1).[21] The effect is lower when only the yield is considered, with the first principal component moving down by about 0.18% (column 3). Similarly, the orthogonalized production revision also lowers commodity futures prices (column 4). The effect of the orthogonalized component of the yield revision ($\tilde{r}^Y$) on commodity futures remains the same when I add the predictable part ($r^Y - \tilde{r}^Y$), which by itself has no significant effect (column 6).

5 Results

5.1 First Stage

The crucial requirement of the external instruments approach is that the instrument should be correlated with the structural shock of interest but not with other structural shocks. Further, valid inference and reliable results require the instrument to be "strong", with the F-statistic in the first-stage regression of the real field crop price residual from the VAR on the instrument above a certain threshold. Montiel-Olea et al. (2021) recommend reporting the heteroskedasticity-robust first-stage F-statistic and suggest that values above 10 indicate approximately valid coverage of standard 95% confidence intervals.

[19] For example, the orthogonalized area revision $\tilde{r}^A$ still shows a consistent seasonal pattern, such as downward revisions in June and October and upward revisions from February to May. This results in the seasonal pattern in the production revisions.

[20] I estimate a regression similar to equation (12) but for specific crops as $r_{i,t}^X = \alpha_i^X + \sum_{k=1}^{12} \gamma_{i,k}^X r_{i,t-k}^X + \sum_{k=1}^{12} \delta_{i,k}^X p_{i,t-k} + \tilde{r}_{i,t}^X$, where $r_{i,t}^X$ is the revision of commodity i in item $X \in \{\text{yield } (Y), \text{area } (A), \text{production } (P), \text{rest } (R), \text{ending stocks } (ES)\}$ and $p_{i,t}$ is the real NASS price of commodity i.

[21] Table C.2 in Appendix C shows a similar impact on a monthly basis: a 1% higher orthogonalized yield revision lowers the commodity futures price by 0.23%.

Table 3 reports the strength of the selected instruments. First, the AJ futures surprises, the first principal components of the commodity futures, have a much larger robust F-statistic under the GK timing than under the calendar-month sum timing.[22] Second, the orthogonalized yield revision shows a much larger robust F-statistic, above 30 under both the calendar-sum and the "GK" timing, compared with 5.2 and 14.5 for the AJ futures. Further, the yield revision instrument explains 6 percent of the variation in the real field crop price residuals. This evidence suggests that the orthogonalized yield instrument is not a weak instrument.

5.2 Effects on the U.S. Economy and the Global Macroeconomy

I study the impact on the U.S. economy and the resulting effect on the global economy in a way similar to Adjemian and Jo (2025). Figure 7 reports the impulse responses of macroeconomic variables to a one standard deviation commodity futures price shock.[23] Figure 8 reports the corresponding impulse responses to my estimated one standard deviation orthogonalized yield revision shock. Since my estimated structural shocks are based on yield revisions, whereas the AJ shocks are based on commodity futures price changes, the impulse responses are not directly comparable. In order to compare them, I normalize both shocks to have a one percent impact on the real field crop price. Figure 9 shows this comparison.[24] In the following paragraphs, I compare the transmission mechanism implied by the orthogonalized yield revision shocks with that of the AJ shocks. Unless stated otherwise, all numbers refer to the normalized responses in Figure 9.

Prices. Figure 9 shows that, under the orthogonalized yield revision shock, the real field crop price rises slightly for the first three months after impact and then falls to about a third of its impact response within a year. Under the AJ shock, the response keeps building up, to about twice its impact response at six months, before declining; in my sample, however, the AJ path is less precisely estimated, with wide confidence bands. The more immediate price response to the yield revision shock suggests that the current price adjusts quickly to news about future supply. This is what one would expect for a storable commodity: a predictable rise in the price over the next six months could be arbitraged away by buying now and storing. Some delay in the price adjustment is still possible, because the price received by farmers partly reflects sales contracted earlier, so new prices may take some time to show up in the data. This delay, however, cannot explain a near doubling of the response.

On core CPI, yield revision shocks imply a zero response on impact but a gradual build-up over the next two to three years. On the other hand, AJ shocks imply a positive response on impact, followed by a gradual decline. I think core prices should take some time to build

[22] The instruments' strength remains when I use the non-seasonally adjusted data with month dummies.

[23] Refer to Table 1 for the variables considered and their data sources.

[24] This comparison simply scales the impulse responses so that, on impact, the real field crop price moves by 1 percent. I do not reconstruct the bands to account for the uncertainty in the response on impact. This may not be the ideal way to compare the impulse responses to two different shocks, but I use it here as a simple comparison. In Appendix Figures D.8 (calendar-sum timing) and D.9 ("GK" timing), I re-estimate the bootstrap confidence intervals.

up after adverse agricultural supply shocks. This could reflect the delayed pass-through of higher biofuel costs, such as ethanol for gasoline or soybeans for biodiesel, to firms' production expenses. The other likely channel is inflation expectations, as consumers' inflation expectations are driven by the prices of the items they buy most frequently, such as food at home (D'Acunto et al., 2021). Both food-at-home prices and core prices show a significant hump after about two and a half to three and a half years.

The response of livestock producer price inflation is similar under both shocks. As AJ mention, higher expected feed costs incentivize producers to send more cattle to slaughter in the short term, leading to higher short-term supply but a lower herd size in the future, and thus lower long-term supply (Schulz, 2022), which raises livestock inflation in the long run. The response of the ethanol price is also quite similar under both shocks.

The AJ shocks imply a faster rise in the food-at-home price, that is, a quicker transmission of agricultural supply news shocks to the food-at-home price, and consequently a more immediate decline in food-at-home consumption. In contrast, the yield revision shocks imply a much slower transmission to the food-at-home price and consumption. While a more detailed investigation is required to establish the correct response, I conjecture that the delayed response indicates that adverse supply shocks affect multiple agricultural seasons, as lower ending stocks carry over into lower beginning stocks in the next year. This continued shortfall leads to a slower adjustment of retail food prices because of frictions associated with contracts and menu costs.

Real Activity. In Adjemian and Jo (2025), real GDP and industrial production decline on impact. The yield revision shocks provide only weak evidence of an immediate decline on impact. The peak response comes later: industrial production falls by 0.1 percent after a quarter, whereas GDP falls by 0.09 percent after three quarters. As AJ argue, agriculture, food and related industries represent about 5 percent of U.S. GDP, and U.S. food manufacturing comprises roughly 10–13 percent of industrial production, so some impact is expected. Yield revision shocks imply a much quicker rebound in activity than the AJ shocks. A yield revision is news about a single season's harvest, whose effect on activity can be mitigated by drawing down ending stocks or by increasing planting in the next season. This is consistent with the response of the real field crop price, which also returns to about a third of its impact response within a year.

However, the quick rebound in activity does not signal a quick response of consumer prices. Higher field crop prices reach retail food prices, and consequently core inflation through the inflation expectations channel, only with a delay, for two reasons. First, field crops are a very small share of the retail price of food (Whitt, 2025), and processors often buy their inputs forward to hedge against higher input costs. Second, crops serve as cattle feed, so higher crop prices lead to higher meat prices only after a livestock cycle: the livestock PPI first falls, as producers reduce their herds, and rises above zero only after about fifteen months.

Financial Markets. Surprisingly, AJ shocks imply a rise in the S&P 500 on impact and a fall in financial market uncertainty, which is hard to square with the fall in activity indicators such as real GDP or IP when the agricultural outlook is weak. Yield revision

shocks imply the opposite response: the S&P 500 falls and the VIX rises on impact, which aligns with an adverse activity outlook. However, the yield revision responses on impact are less precisely estimated, as zero is not excluded from the 90 percent confidence band for the S&P 500 or from the 68 percent confidence band for the VIX. Yield revision shocks imply a gradual rebound in the S&P 500 and a decline in the VIX as real GDP and IP recover.

Trade Indicators. AJ find an initial rise in grain exports and a consequent rise in the cost of shipping, measured by the Baltic Dry Index. They argue that this is likely driven by importers seeking more supplies given the poor forecast for the agricultural season, leading to more exports in the immediate aftermath of the shock; later, as the poor agricultural supply materializes, grain exports and consequently the shipping cost fall. Yield revisions, however, imply a different response, where both exports and the shipping cost fall on impact. Grain exports quickly recover after about a year, but interestingly the Baltic Dry Index continues to climb after a year.

The continued rise in the Baltic Dry Index may reflect two factors. First, it may reflect a shift in the geography of grain trade: a smaller U.S. crop raises demand for South American supplies, which travel longer routes to Asian importers, raising the demand for dry-bulk shipping. Second, the Baltic Dry Index is also a signal of global industrial activity, and its rise coincides with a rise in U.S. industrial production, which, however, is less precisely estimated.

Oil. AJ shocks imply a much more immediate rise in the real oil price and fall in oil production. In contrast, the yield revision shocks show no immediate fall in oil production on impact, only a gradual decline. Under the yield revision shocks, the real oil price falls before rising above its initial level after a lag of one to two years. These responses are consistent with the responses of industrial production and the Baltic Dry Index, which remain depressed for about a year before recovering to their initial levels, and of real GDP, which recovers more slowly. Thus, recovering demand against sluggish supply from oil-producing countries, such as the OPEC members, could explain the increase in the real oil price.[25] However, a delayed rise of 1 percent in the real oil price in response to a 1 percent rise in the U.S. real field crop price is hard to explain.

5.3 Robustness and Exogeneity Checks

The orthogonalized yield revisions carry no serial correlation (Figure 6), by construction, since each is a residual on its own twelve lags. Table 4 shows that none of the other 13 variables in the VAR Granger-causes them, and the fourteen variables are jointly insignificant ($p = 0.93$); the real field crop price passes by construction, since its lags are removed in the orthogonalization. In my sample, by contrast, GDP Granger-causes the AJ shocks ($p = 0.002$). Table 5 shows that the yield revision shocks are not significantly correlated

[25] On the margin, another channel at play could be the substitution from biofuels to other liquid fuels because of the higher cost of biofuels. However, this channel is likely small for the world oil price, as biofuels are a small share of world liquid fuels and blending limits restrict the substitution between ethanol and gasoline.

with any of the published macroeconomic shocks, while the AJ shocks are correlated with the Global Supply Chain Pressure Index (0.13, $p = 0.02$).

6 Conclusion

In summary, yield revisions point to a much slower transmission mechanism of agricultural supply shocks to inflation, such as food inflation or core inflation, and a minimal impact on activity indicators such as real GDP and industrial production. In contrast, the AJ shocks imply a much quicker transmission to food and core inflation and a more immediate decline in GDP and industrial production. The response of oil is quite puzzling across both shocks, but more so for the AJ shocks, as they lead to an immediate rise in the oil price and a decline in oil production. To better understand the propagation mechanism, I plan to follow Känzig (2021): construct a baseline VAR and then augment it with one variable at a time to study the propagation channels. Moreover, it would be useful to test the impulse responses with an internal instrument method, with the instrument ordered first in the VAR, and, as a robustness check, to plot the impulse responses using either local projections or heteroskedasticity-based identification to correct for any background noise in the shocks.

World Agricultural Outlook Board

waob

ISSN: 1554-9089

World Agricultural Supply and Demand Estimates

United States Department of Agriculture

Office of the Chief Economist	**Agricultural Marketing Service Farm Service Agency**	**Economic Research Service Foreign Agricultural Service**

WASDE - 485	Approved by the World Agricultural Outlook Board	August 12, 2010

WHEAT: U.S. wheat ending stocks for 2010/11 are lowered this month as higher expected exports more than offset an increase in forecast production and lower projected feed and residual use. Production is forecast 49 million bushels higher mostly reflecting higher yields for durum and other spring wheat, especially in the Northern Plains. Winter wheat production is also raised slightly as higher yields in the Northern Plains and Pacific Northwest more than offset reductions in the eastern Corn Belt. Feed and residual use is lowered 10 million bushels as rising values have priced wheat out of feed rations. Exports are projected 200 million bushels higher with declines in foreign production, particularly in the FSU-12, reducing global supplies and making U.S. wheat competitive in key Middle East and North Africa markets. U.S. ending stocks are projected 141 million bushels lower from last month, and down 21 million from 2009/10. The 2010/11 season-average farm price is projected at $4.70 to $5.50 per bushel, up 50 cents on both ends of the range.

Several changes are made to U.S. supply and use estimates for 2009/10 based on the latest U.S. Bureau of Census trade and mill grind data and revisions. Exports are raised 16 million bushels and imports are raised 4 million bushels. Food use is lowered 3 million bushels. Thus, feed and residual is estimated 10 million bushels lower.

Figure 1: The wheat narrative of a WASDE report: WASDE-485, August 12, 2010, page 1

Notes: Each WASDE report opens every crop's balance sheet with a short narrative that gives a cause for each revision. In this release, U.S. wheat production for 2010/11 is raised 49 million bushels, "mostly reflecting higher yields for durum and other spring wheat," and exports are raised 200 million bushels on lower foreign production; together they lower ending stocks by 141 million bushels. Figure 2 shows the same revisions in the balance sheet. Source: USDA World Agricultural Outlook Board, *World Agricultural Supply and Demand Estimates*, WASDE-485 (August 12, 2010), p. 1.

WASDE-485-11 August 2010

U.S. Wheat Supply and Use 1/

Item	2008/09	2009/10 Est.	2010/11 Projections July	2010/11 Projections August
Area		Million acres		
Planted	63.2	59.1	54.3	54.3
Harvested	55.7	49.9	48.3	48.3
Yield per harvested acre	44.9	44.4	45.9	46.9
		Bushels		
Beginning stocks	306	657	973	973
Production	2,499	2,216	2,216	2,265
Imports	127	119	100	100
Supply, total	2,932	2,991	3,289	3,338
Food	927	917	940	940
Seed	78	70	76	76
Feed and residual	255	149	180	170
Domestic, total	1,260	1,137	1,196	1,186
Exports	1,015	881	1,000	1,200
Use, total	2,275	2,018	2,196	2,386
Ending stocks	657	973	1,093	952
CCC inventory	0	0		
Free stocks	657	973		
Outstanding loans	27	51		
Avg. farm price ($/bu) 2/	6.78	4.87	4.20- 5.00	4.70- 5.50
		Million bushels		

Figure 2: The U.S. wheat balance sheet in a WASDE report: WASDE-485, August 12, 2010, page 11

Notes: The 2010/11 projection columns give the July and the August forecast; a revision is the change from one to the next. Highlighted rows: production (red), exports (blue) and ending stocks (yellow). From July to August, harvested area is unchanged at 48.3 million acres and yield is raised from 45.9 to 46.9 bushels per acre, so the whole production revision, from 2,216 to 2,265 million bushels, comes from yield: $r^Y = 100\ln(46.9/45.9) = 2.16$ and $r^P = 100\ln(2{,}265/2{,}216) = 2.19$. Exports rise from 1,000 to 1,200 million bushels and ending stocks fall from 1,093 to 952. Figure 1 shows the narrative that explains these revisions. Source: USDA World Agricultural Outlook Board, *World Agricultural Supply and Demand Estimates*, WASDE-485 (August 12, 2010), p. 11, "U.S. Wheat Supply and Use."

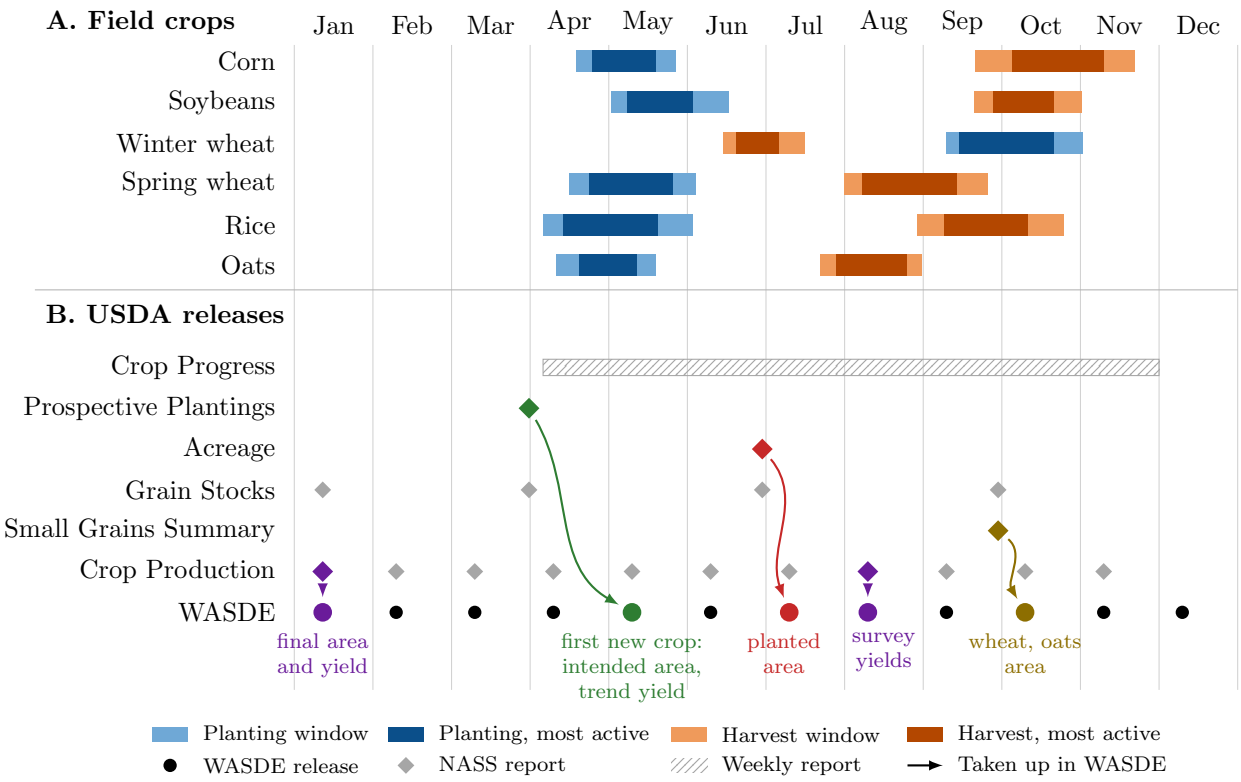


Figure 3: The crop calendar and the USDA release calendar. Panel A gives the usual planting and harvesting dates of the five crops: light bars are the window of each operation, from about 5 to about 95 percent complete, and dark bars its most active period, from about 15 to about 85 percent complete. Winter wheat is planted in the autumn and harvested the following summer. Panel B marks the release day of each USDA report in a typical year: circles are WASDE releases and diamonds are NASS reports. Arrows mark the NASS report whose numbers the next WASDE takes into its new-crop balance sheet. The May WASDE, the first projection for the new marketing year, uses the planting intentions of the March 31 Prospective Plantings report and trend yields. The July WASDE takes planted and harvested area from the June 30 Acreage report, which is published after most corn and soybeans are planted. Survey-based corn and soybean yields first come in the August Crop Production, released at the same time as the WASDE. The October WASDE takes wheat and oat area from the Small Grains Summary of September 30. The January WASDE takes final area and yield from the Crop Production Annual Summary. *Notes:* Crop dates are for the state with the most harvested acres of each crop in 2009 (Iowa for corn and soybeans; Wisconsin for spring oats; Kansas for winter wheat; North Dakota for spring wheat; Arkansas for rice); Table B.1 gives the dates and the months from planting to harvest. WASDE and Crop Production are drawn on the 10th of the month (the 12th in January); the actual day varies between about the 8th and the 12th. Grain Stocks, Prospective Plantings, Acreage and the Small Grains Summary come out on the last business day of the month and Crop Progress weekly on Mondays from early April to the end of November. Sources: USDA NASS, *Usual Planting and Harvesting Dates for U.S. Field Crops*, October 2010; USDA NASS release calendars and the WASDE release schedule, 2026.

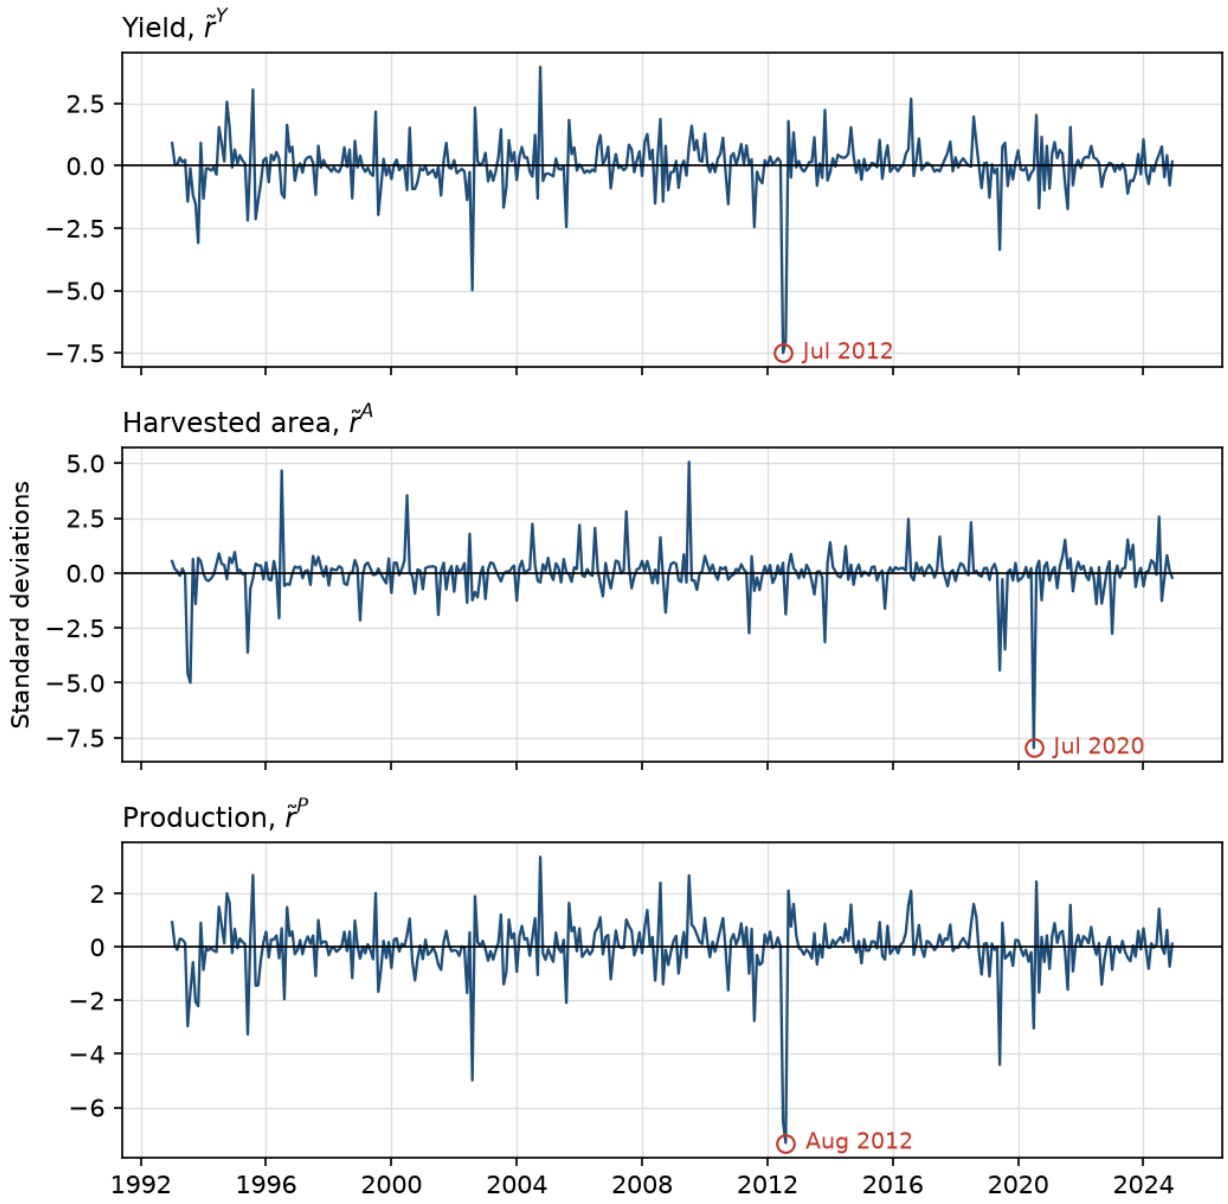


Figure 4: Orthogonalized WASDE revision instruments: yield ($\tilde{r}^Y$, top), harvested area ($\tilde{r}^A$, middle) and production ($\tilde{r}^P$, bottom), calendar-month sum, 1993M01–2024M12, in standard deviations. *Notes:* Each series is the residual w_t of $z_t = c + \sum_{k=1}^{12} g_k z_{t-k} + \sum_{k=1}^{12} h_k p_{t-k} + w_t$, estimated by OLS on 1993M01–2024M12, divided by its standard deviation, where z_t is the five-crop revision of Figure C.1 summed by calendar month (each release counted in full in the month it is published) and p_t the X-13 seasonally adjusted real field crop price of the VAR. Positive means USDA now sees more supply than a month ago, beyond what its own past revisions and the past crop price predict. The circled month is each series' largest absolute value. Source: USDA WASDE.

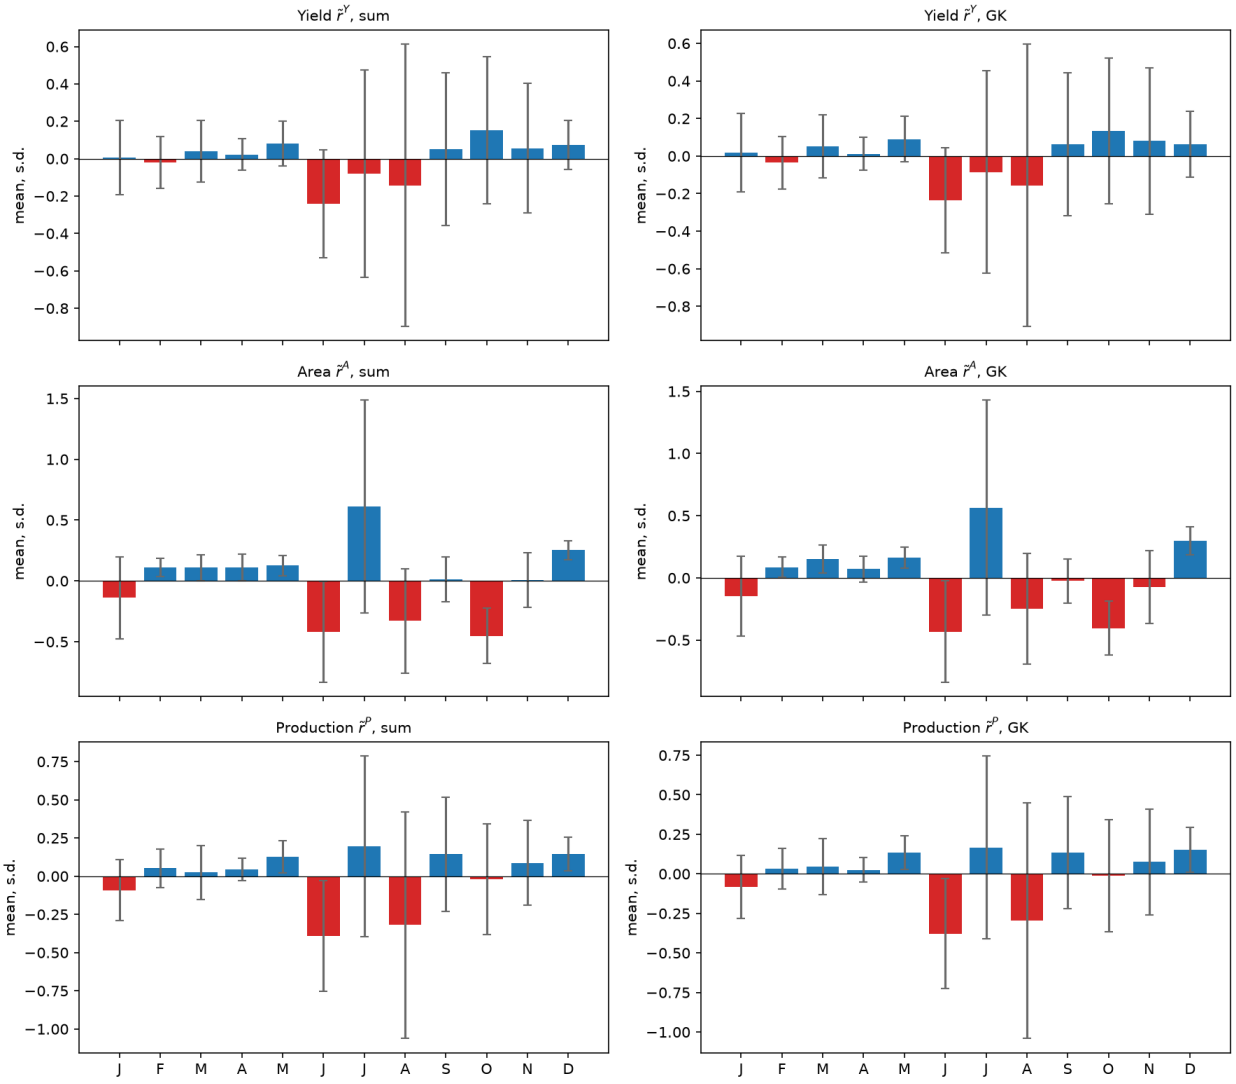


Figure 5: Seasonal pattern of the orthogonalized revision instruments: month-of-year mean of $\tilde{r}^Y$ (top), $\tilde{r}^A$ (middle) and $\tilde{r}^P$ (bottom), calendar-month sum (left) and GK timing (right), 1993M01–2024M12, in standard deviations. *Notes:* Each bar is $\bar{z}_m = \frac{1}{n_m}\sum_y z_{m,y}$ with whiskers $\pm 2\, s_m/\sqrt{n_m}$, as in Figure C.2; $n_m = 32$ years. Blue means more supply on average, red less. The orthogonalization includes twelve own lags, which absorb part of any fixed month pattern, so the absence of a pattern is weaker evidence than its presence. No month of $\tilde{r}^Y$ has $|t| > 2$ under either timing, and a robust F test that the eleven month dummies are jointly zero gives $p = 0.864$ (sum) and 0.836 (GK). The area revision stays seasonal (June and October below zero, December and February–May above; joint $p < 0.001$ under both timings), and the production revision has June below zero and May and December above ($p = 0.204$ and 0.214). Source: USDA WASDE.

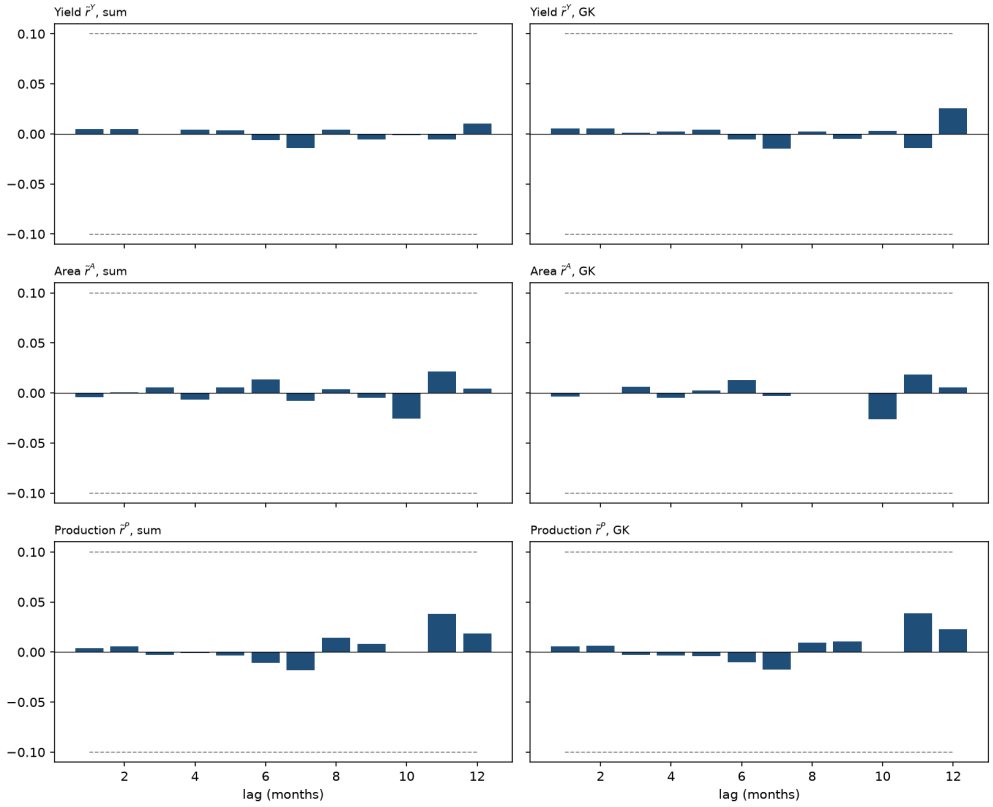


Figure 6: Autocorrelation of the orthogonalized revision instruments $\tilde{r}^Y$ (top), $\tilde{r}^A$ (middle) and $\tilde{r}^P$ (bottom), calendar-month sum (left) and GK timing (right), lags 1–12, 1993M01–2024M12. *Notes:* Sample autocorrelations over the 384 months of the first-stage sample; the dashed lines are the white-noise 95 percent band $\pm 1.96/\sqrt{T} = \pm 0.10$. No autocorrelation of any series lies outside the band (the largest in absolute value is 0.04). This holds by construction: each series is the residual of a regression on its own twelve lags, so it is uncorrelated with them in sample. The figure shows that the GK split of releases across months and USDA's habit of revising in the same direction in consecutive releases (Figure C.3) leave no serial correlation in the instrument; it is not an independent test of exogeneity, which Tables 4 and 5 provide. Source: USDA WASDE.

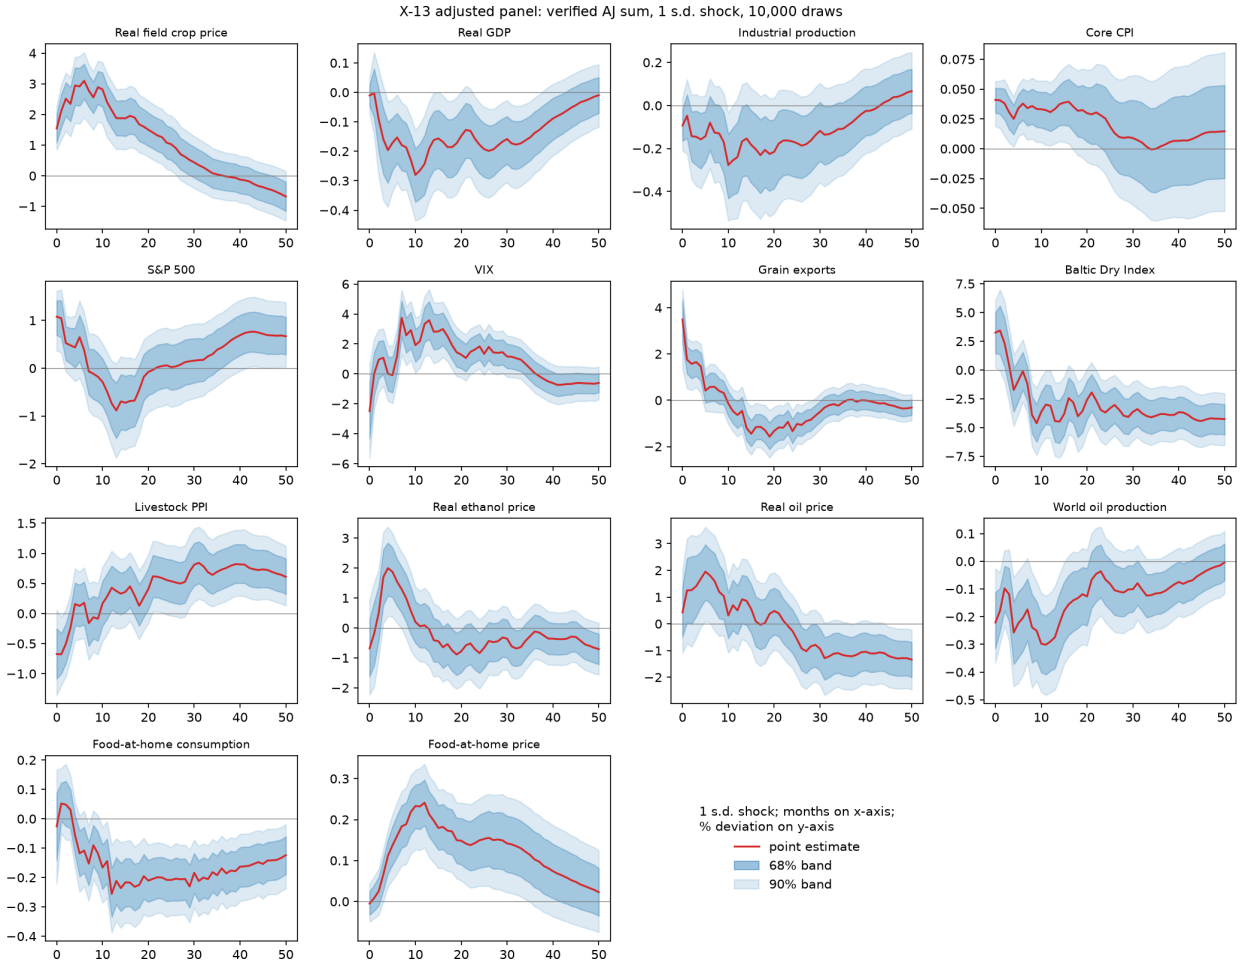


Figure 7: Impulse responses to a one standard deviation agricultural supply news shock: replication of Adjemian and Jo (2025), Figure 3, on the seasonally adjusted panel

Notes: Proxy-SVAR with twelve lags and a constant, fourteen variables in 100 times logs, monthly data 1992M01–2024M12 (identification sample 1993M01–2024M12). The instrument is the monthly sum of production-weighted first principal component surprises in the nearest five maturities of corn, soybean, wheat, rice and oat futures around USDA report dates, using my verified calendar (Table A.2 with twelve report dates corrected to the dates the reports print). All series are seasonally adjusted with X-13. Real GDP is the Brave–Butters–Kelley monthly series, not the S&P Global index used in the original. The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws. Horizons are in months.

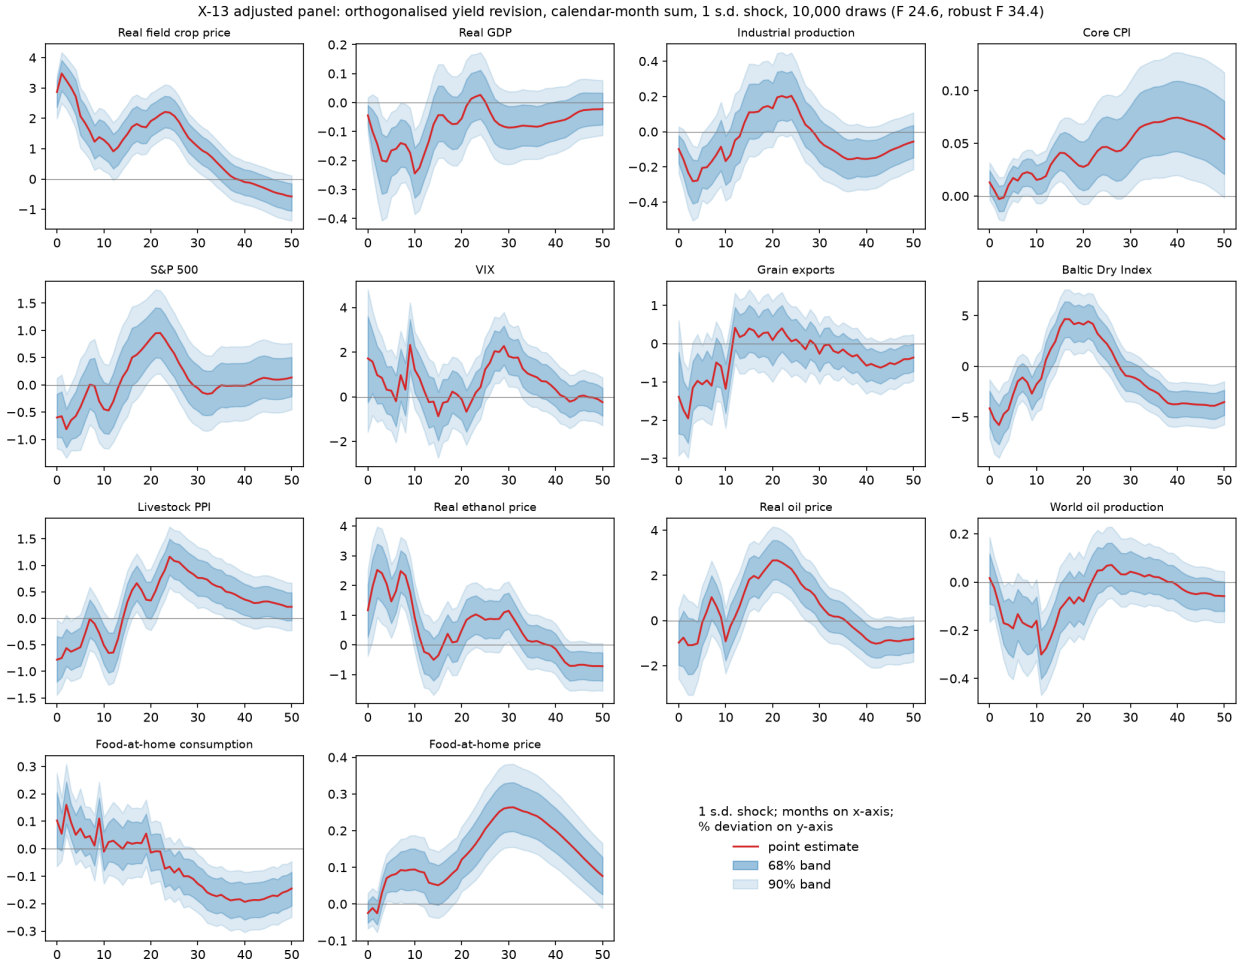


Figure 8: Impulse responses to the orthogonalized yield-revision shock ($\tilde{r}_t^Y$, calendar-month sum), seasonally adjusted panel

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. X-13 seasonally adjusted series and a constant, as Figure 7. Each release enters the calendar month in which it is published, in full, as in Figure 4. The instrument is the orthogonalized five-crop yield revision: the residual $\tilde{r}_t^Y$ of $r_t^Y = \alpha^Y + \sum_{k=1}^{12} \gamma_k^Y r_{t-k}^Y + \sum_{k=1}^{12} \delta_k^Y p_{t-k} + \tilde{r}_t^Y$, with $r_t^Y = \sum_i w_{i,y} r_{i,t}^Y$ and p the seasonally adjusted real field crop price, estimated on 1993M01–2024M12 and scaled to unit standard deviation. The projection removes the part of the revision that the market had already seen: the raw revision is correlated with the crop price residual of the previous month (robust t of -5.4 for the calendar-month sum), the orthogonalized one is not ($t = -0.65$). The shock is signed so that the crop price rises, so it is an adverse yield revision. First-stage F 24.6, robust F 34.4, Newey–West F 25.9. A one standard deviation shock is one standard deviation of the structural shock each instrument identifies, so its size differs from that of Figure 7; Figure 9 compares the two shocks scaled to the same 1 percent rise in the real field crop price on impact. The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage but not the projection. Horizons are in months.

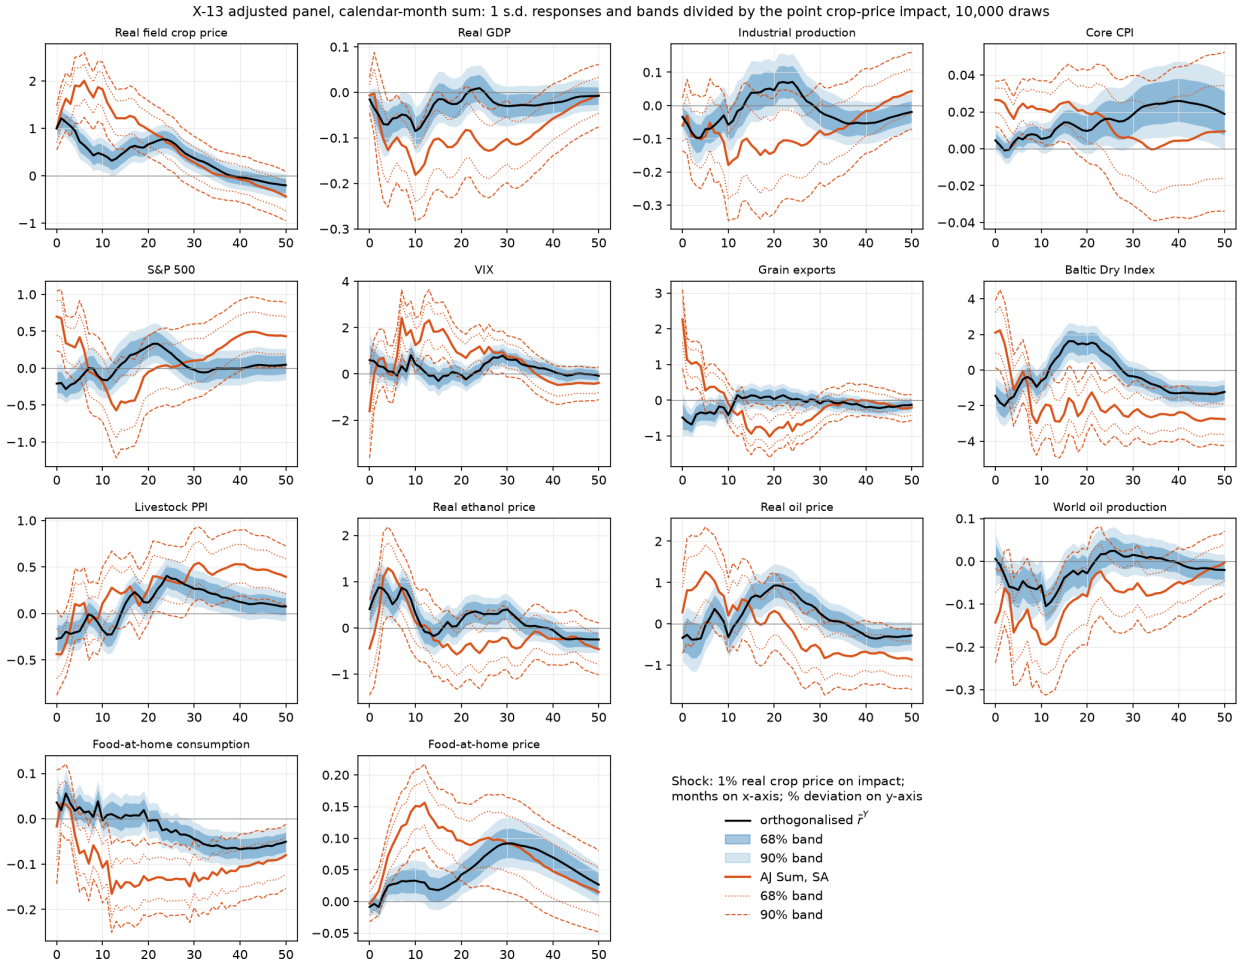


Figure 9: Impulse responses to the orthogonalized yield-revision shock ($\tilde{r}_t^Y$) and to the Adjemian–Jo shock, rescaled to a 1 percent rise in the real field crop price on impact, calendar-month sum, seasonally adjusted panel

Notes: The one standard deviation responses of Figures 8 ($\tilde{r}_t^Y$) and 7 (Adjemian–Jo sum), each divided by one number, its point estimate of the impact response of the real field crop price (2.87 percent for $\tilde{r}_t^Y$, 1.55 percent for Adjemian–Jo), so that both shocks raise the crop price by 1 percent on impact. A one standard deviation shock is one standard deviation of the structural shock each instrument identifies, so its size differs between the two; the rescaling only changes units. The same number divides the point estimate and both bands, so the figure is Figures 8 and 7 on a common scale. The bands treat the impact as known and leave out the uncertainty about it, which is why the crop price band at horizon 0 is not a point; Figure D.8 shows the bands when each bootstrap draw is divided by its own impact, as in Känzig (2021). Twelve lags, the same fourteen variables in 100 times logs, X-13 seasonally adjusted series and a constant, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), each release in the calendar month in which it is published. The black line and the blue regions (dark 68 percent, light 90 percent) are $\tilde{r}_t^Y$ (first-stage F 24.6, robust F 34.4). The red line is the Adjemian–Jo shock, with its 68 percent band dotted and its 90 percent band dashed (first-stage F 6.5, robust F 5.2). Bands from 10,000 moving-block bootstrap draws. Horizons are in months.

Table 1: Variables in the VAR and proxy SVAR

	Variable	Series, source and construction	Frequency and aggregation	Seasonally adjusted
1	Real field crop price	Constructed. NASS monthly prices received for corn, soybeans, wheat, rice and oats, weighted by each crop's annual share of U.S. production in metric tons, and deflated by the CPI for all items (FRED CPIAUCSL). Weights are WASDE-based for 2003–2024 and NASS-based for 1992–2002.	Monthly; weights annual	X-13
2	Real GDP	Brave–Butters–Kelley monthly real GDP, Chicago Fed (FRED BBKMGDP). Substitute for the S&P Global Market Intelligence monthly GDP index used by Adjemian and Jo, which is a paid series.	Monthly	At source
3	Industrial production	Industrial production index, Federal Reserve Board (FRED INDPRO).	Monthly	At source
4	Core CPI	CPI for all items less food and energy, BLS (FRED CPILFESL).	Monthly	At source
5	S&P 500	S&P 500 index, daily close (CEIC series UDIBB).	Monthly mean of daily closes	X-13
6	VIX	CBOE Volatility Index (FRED VIXCLS).	Month-end close	X-13
7	Grain exports	U.S. exports of grains and feeds to the world, in metric tons, USDA Foreign Agricultural Service GATS.	Monthly total	X-13
8	Baltic Dry Index	Baltic Exchange Dry Index, daily (CEIC series 268193601).	Monthly mean of daily values	X-13
9	Livestock PPI	Producer price index for farm products: livestock, BLS (FRED WPU013).	Monthly	X-13
10	Real ethanol price	Fuel ethanol price in dollars per gallon, USDA ERS U.S. Bioenergy Statistics (Table 14), deflated by the CPI for all items.	Monthly	X-13
11	Real oil price	WTI spot crude oil price in dollars per barrel (FRED WTISPLC), deflated by the CPI for all items.	Monthly	X-13
12	World oil production	World crude oil production including lease condensate, thousand barrels per day, U.S. Energy Information Administration.	Monthly	X-13
13	Food-at-home consumption	Personal consumption expenditures on food, chain-type quantity index, BEA (FRED DFXARA3M086SBEA).	Monthly	At source
14	Food-at-home price	Personal consumption expenditures on food, chain-type price index, BEA (FRED DFXARG3M086SBEA).	Monthly	At source

Notes: The sample is 1992M01–2024M12 (396 months), the sample of Adjemian and Jo (2025). The first twelve months supply the VAR lags, so identification uses 1993M01–2024M12. Every variable enters the VAR in 100 times natural logs of the level, in the order listed, so that the impulse responses are percent deviations. The real field crop price is variable 1 because the shock is normalized on it. Variables that are not seasonally adjusted at the source ("X-13") are adjusted with the Census Bureau's X-13ARIMA-SEATS program, and the ones adjusted at the source are used as published. The Baltic Dry Index, S&P 500, grain exports and ethanol series are exported from the vendor and do not update automatically.

Table 2: Commodity futures responses to orthogonalized WASDE revisions

Release-day PC1 futures surprise (%) on each crop's orthogonalized revision	(1) Main	(2) No area	(3) Yield only	(4) Production	(5) Ending stocks	(6) Yield split
Yield revision $\tilde{r}^Y$	-0.254^{***}	-0.251^{***}	-0.177^{***}			-0.177^{***}
	(0.044)	(0.044)	(0.044)			(0.045)
Area revision $\tilde{r}^A$	-0.037					
	(0.028)					
Production revision $\tilde{r}^P$				-0.127^{***}		
				(0.030)		
Rest revision $\tilde{r}^R$	-0.179^{***}	-0.165^{***}		-0.158^{***}		
	(0.036)	(0.035)		(0.037)		
Ending-stocks revision $\tilde{r}^{ES}$					-0.251^{***}	
					(0.035)	
Predicted yield revision $r^Y - \tilde{r}^Y$						-0.116
						(0.107)
Observations	1,726	1,726	1,726	1,726	1,726	1,726
Releases	350	350	350	350	350	350
R^2	0.041	0.039	0.023	0.024	0.046	0.024
Crop fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column estimates $s_{i,t} = \alpha_i + \beta_Y \tilde{r}^Y_{i,t} + \beta_A \tilde{r}^A_{i,t} + \beta_P \tilde{r}^P_{i,t} + \beta_R \tilde{r}^R_{i,t} + \beta_{ES} \tilde{r}^{ES}_{i,t} + \gamma_Y (r^Y_{i,t} - \tilde{r}^Y_{i,t}) + \varepsilon_{i,t}$, where $s_{i,t}$ is the release-day PC1 futures surprise of commodity i in WASDE release month t, in percent, and α_i is a commodity fixed effect. Each column keeps only the regressors with an entry: (1) $\tilde{r}^Y$, $\tilde{r}^A$, $\tilde{r}^R$; (2) $\tilde{r}^Y$, $\tilde{r}^R$; (3) $\tilde{r}^Y$; (4) $\tilde{r}^P$, $\tilde{r}^R$; (5) $\tilde{r}^{ES}$; (6) $\tilde{r}^Y$ and its predicted part $r^Y - \tilde{r}^Y$. The table repeats Table C.1 with each commodity's revision orthogonalized: $\tilde{r}^X_{i,t}$ is the residual of $r^X_{i,t} = \alpha^X_i + \sum_{k=1}^{12} \gamma^X_{i,k} r^X_{i,t-k} + \sum_{k=1}^{12} \delta^X_{i,k} p_{i,t-k} + \tilde{r}^X_{i,t}$, where $r^X_{i,t}$ is the revision of commodity i in item $X =$ yield (Y), area (A), production (P), rest (R) or ending stocks (ES) (zero in months without a revision); the yield, area and production revisions are $100 \ln(X_t/X_{t-1})$, and the rest revision is the change in ending stocks minus the change in production and the ending-stocks revision the change in ending stocks, both as a percent of last month's total use; $p_{i,t}$ is 100 times the log of commodity i's NASS price received, and every coefficient is commodity specific; it is estimated on 1993M01–2024M12 (1,920 commodity-months) and kept in percent, so the slopes compare one for one with Table C.1. Every column pools all five crops with crop fixed effects, on the 1,726 release–crop observations of Table C.1; standard errors clustered by release date in parentheses. For comparison, the raw yield revision alone gives -0.168 (0.043) on the same rows. Table C.2 runs the same split on the monthly series. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Tests on instrument strength: orthogonalized revisions

	AJ futures		Orthogonalized yield revision $\tilde{r}^Y$	
	(1) Sum	(2) GK	(3) Sum	(4) GK
Seasonally adjusted (SA)				
Coefficient	0.333	0.514	−0.635	−0.636
F-statistic	6.46	15.72	24.55	24.58
F-statistic (robust)	5.17	14.51	34.43	35.94
F-statistic (Newey–West)	5.01	13.24	25.90	27.40
R^2	1.66	3.95	6.04	6.04
R^2 (adjusted)	1.40	3.70	5.79	5.80
Observations	384	384	384	384
Unadjusted with month dummies (NSA + D)				
Coefficient	0.376	0.607	−0.579	−0.572
F-statistic	8.70	23.50	21.31	20.74
F-statistic (robust)	7.85	21.02	36.89	35.18
F-statistic (Newey–West)	8.50	17.60	39.85	39.40
R^2	2.23	5.80	5.28	5.15
R^2 (adjusted)	1.97	5.55	5.03	4.90
Observations	384	384	384	384

Notes: First-stage regressions $u_t = a + b\, z_t + e_t$ of the real field crop price residual u_t on each instrument z_t, as in Table C.3. The residual comes from the crop price equation of the fourteen-variable VAR with twelve lags, estimated on 1992M01–2024M12, with the X-13 seasonally adjusted series and a constant (SA) or the unadjusted series with a constant and February–December dummies (NSA + D). Columns (1)–(2) are the Adjemian–Jo futures surprise as a calendar-month sum and in Gertler–Karadi (GK) timing; columns (3)–(4) the orthogonalized five-crop WASDE yield revision $\tilde{r}^Y$ as a calendar-month sum and in GK timing (Figure 4). Each orthogonalized series is the residual of the revision on twelve own lags and twelve lags of the seasonally adjusted real crop price, estimated once on 1993M01–2024M12; the same series enters the SA and the NSA + D VAR. The sample is 1993M01–2024M12 (384 months). Table C.4 reports the orthogonalized production and area revisions, and Table B.4 the same columns on the poll months. Instruments are scaled to unit standard deviation over the regression months, so b is in percent per standard deviation; the futures surprise is a price move (expected $b > 0$), and the revisions are positive when supply is larger (expected $b < 0$). $F = (b/\text{se})^2$; the robust F uses HC1 standard errors and equals the Montiel Olea–Pflueger effective F with one instrument; the Newey–West F uses 5 lags. R^2 in percent. The orthogonalized series is a generated instrument; the standard errors do not account for the first step.

Table 4: Granger causality tests of the orthogonalized instruments

	AJ futures		Orthogonalized yield revision $\tilde{r}^Y$	
	Sum	GK	Sum	GK
Own lags	0.199	0.027	0.986	0.990
Real field crop price	0.092	0.004	0.958	0.939
Real GDP	0.002	0.009	0.684	0.826
Industrial production	0.622	0.717	0.961	0.979
Core CPI	0.171	0.169	0.837	0.890
S&P 500	0.640	0.494	0.409	0.500
VIX	0.899	0.721	0.745	0.786
Grain exports	0.918	0.499	0.876	0.886
Baltic Dry Index	0.356	0.026	0.541	0.574
Livestock PPI	0.569	0.589	0.768	0.811
Real ethanol price	0.221	0.057	0.243	0.271
Real oil price	0.837	0.671	0.391	0.445
World oil production	0.573	0.556	0.526	0.522
Food-at-home consumption	0.061	0.235	0.966	0.978
Food-at-home price	0.833	0.539	0.601	0.626
All 14 variables	0.146	0.016	0.927	0.959

Notes: As Table C.5: for each instrument z_t, $z_t = c + \sum_{k=1}^{12} a_k z_{t-k} + \sum_{j=1}^{14} \sum_{k=1}^{12} b_{jk}\, y_{j,t-k} + e_t$, with y_j the fourteen unadjusted VAR variables; each entry is the p-value of the OLS F test that the row's twelve lags are jointly zero. The sample is 1994M01–2024M12 (372 months) for every column, because the orthogonalized series start in 1993M01 and need twelve own lags. Columns: the Adjemian–Jo futures surprise and the orthogonalized yield revision $\tilde{r}^Y$, each as a calendar-month sum and in GK timing; Table C.7 reports the orthogonalized production and area revisions. p-values below 0.05 are in red. The orthogonalized series are residuals on their own twelve lags and twelve lags of the seasonally adjusted crop price, so the "Own lags" row and the crop price row pass by construction; the other thirteen rows and the joint test are the informative ones.

Table 5: Correlation of the orthogonalized instruments with published macroeconomic shocks

	AJ futures		Orthogonalized yield revision $\tilde{r}^Y$		
	Sum	GK	Sum	GK	Months
Oil supply news (Känzig)	0.001	0.058	−0.006	−0.017	384
	(0.988)	(0.256)	(0.901)	(0.743)	
Oil supply (Baumeister–Hamilton)	−0.045	−0.055	−0.032	−0.023	384
	(0.383)	(0.285)	(0.530)	(0.658)	
Oil consumption demand (BH)	−0.056	−0.035	0.088	0.074	384
	(0.271)	(0.499)	(0.086)	(0.147)	
Oil inventory demand (BH)	0.032	0.043	−0.044	−0.038	384
	(0.534)	(0.396)	(0.390)	(0.454)	
Economic activity (BH)	−0.006	0.013	0.066	0.066	384
	(0.903)	(0.798)	(0.200)	(0.198)	
Monetary policy (BRW)	0.050	0.042	0.019	0.020	367
	(0.344)	(0.422)	(0.719)	(0.706)	
Monetary policy (Bauer–Swanson, orth.)	−0.058	−0.011	0.035	0.044	372
	(0.264)	(0.826)	(0.502)	(0.397)	
Change in fed funds rate	0.008	0.018	−0.009	−0.007	384
	(0.881)	(0.728)	(0.864)	(0.894)	
Carbon policy	−0.089	0.015	0.060	0.042	246
	(0.165)	(0.821)	(0.346)	(0.515)	
Supply-chain pressure (GSCPI)	0.130	0.192	0.050	0.049	324
	(0.019)	(0.001)	(0.366)	(0.376)	
Geopolitical risk (GPR)	−0.050	−0.063	0.048	0.049	384
	(0.329)	(0.218)	(0.346)	(0.336)	
Trade policy uncertainty	0.015	0.044	−0.055	−0.055	384
	(0.766)	(0.390)	(0.282)	(0.283)	
Economic policy uncertainty	0.050	0.070	−0.069	−0.070	384
	(0.330)	(0.171)	(0.174)	(0.171)	

Notes: As Table C.6: Pearson correlation between the instrument and a published shock series, with the p-value of $H_0 \colon \rho = 0$ in parentheses, on the months inside 1993M01–2024M12 that the shock covers (last column). Columns as in Table 4; Table C.8 reports the orthogonalized production and area revisions. A correlation with $p < 0.05$ is in green if positive and in red if negative. BH stands for Baumeister–Hamilton, BRW for Bu–Rogers–Wu, GSCPI for the New York Fed Global Supply Chain Pressure Index and GPR for the Caldara–Iacoviello geopolitical risk index.

References

Adjemian, M. K. and J. Jo (2025). The macroeconomic effects of agricultural supply news. Working paper, University of Georgia. Available from the corresponding author upon request.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). Exposure to grocery prices and inflation expectations. *Journal of Political Economy 129*(5), 1615–1639.

De Winne, J. and G. Peersman (2016). Macroeconomic effects of disruptions in global food commodity markets: Evidence for the United States. *Brookings Papers on Economic Activity Fall 2016*, 183–286.

Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics 7*(1), 44–76.

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

Montiel-Olea, J. L., J. H. Stock, and M. W. Watson (2021). Inference in structural vector autoregressions identified with an external instrument. *Journal of Econometrics 225*(1), 74–87.

Schulz, L. L. (2022, May). Detailing the supply curve for beef. *Ag Decision Maker*, Iowa State University Extension and Outreach. May 22, 2022; as cited in Adjemian and Jo (2025).

Whitt, C. (2025, July). U.S. retail food price data: Frequently asked questions and issues for Congress. CRS Report R48589, Congressional Research Service.

Appendix

A Other USDA Reports

Grain Stocks. The Grain Stocks report gives the quantity of grain in storage in the U.S. on a reference date. It covers the stocks split into on-farm and off-farm storage. On-farm stocks are the producers' reports of the total quantity on their operations, whereas off-farm stocks refer to storage in commercial facilities. The report mainly covers crops such as wheat, corn and soybeans at the sub-national level, with a breakdown of the stocks by the top five states. It is estimated from surveys that run in the first two weeks of December, March, June and September. The results of the December survey are published in the middle of January of the next year, and the results of the other surveys are published at the end of the respective quarter.

Prospective Plantings. This report contains the first official estimate of how many acres farmers intend to plant in the coming season. It comes out at the end of March together with the March Grain Stocks report. It covers the intended planted acreage for the principal crops, by crop and by state.

Acreage. This report contains the acres that farmers have actually planted and expect to harvest. The acres expected to be harvested may differ from the planted acres, as some acres are abandoned, grazed or cut for silage. It is released on the last business day of June along with the June Grain Stocks report.

Table A.1: USDA announcements, 1992–2024

Notes: The event calendar of the Adjemian–Jo shock: the dates on which USDA published WASDE, Grain Stocks (GS), Prospective Plantings (PP) and Acreage (AR), one row per month of 1992M01–2024M12 (396 months, 493 publication dates). Reports published on the same day are one event, since one daily futures price change covers them all. The dates start from Table A.2 of Adjemian and Jo (2025), and each is checked against the release date printed in the archived report itself. A dagger (†) marks the 11 publication dates where the report prints a different date (12 report dates: Grain Stocks and Acreage share 6/29/2001), which replace the Table A.2 date: 6/12/1998 (WASDE; Table A.2: 6/10), 6/29/2001 (GS & AR; Table A.2: 6/28), 8/10/2001 (WASDE; Table A.2: 8/12), 9/14/2001 (WASDE; Table A.2: 9/12), 9/28/2001 (GS; Table A.2: 9/30), 10/12/2001 (WASDE; Table A.2: 10/11), 11/9/2001 (WASDE; Table A.2: 11/12), 12/11/2001 (WASDE; Table A.2: 12/10), 6/12/2002 (WASDE; Table A.2: 6/10), 3/10/2006 (WASDE; Table A.2: 3/13), 9/12/2008 (WASDE; Table A.2: 9/10). The publication was canceled by federal government shutdowns in 2013M10 and 2019M01, and those months have no event. Until April 1994 USDA released the reports after the close of grain futures trading, so the price change is measured on the next trading day; from May 1994 the reports came out before the open (8:30 a.m. ET, noon from 2013), and the change is measured on the publication day. Table A.2's control dates, one week after each publication date and used in its placebo test, are not used here.

Period	Publication date	Reports
1992M01	1/13	WASDE & GS
1992M02	2/11	WASDE
1992M03	3/11; 3/31	3/11 (WASDE); 3/31 (GS & PP)
1992M04	4/10	WASDE
1992M05	5/11	WASDE
1992M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
1992M07	7/9	WASDE
1992M08	8/12	WASDE
1992M09	9/10; 9/30	9/10 (WASDE); 9/30 (GS)
1992M10	10/8	WASDE
1992M11	11/10	WASDE
1992M12	12/10	WASDE
1993M01	1/12	WASDE & GS
1993M02	2/10	WASDE
1993M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
1993M04	4/12	WASDE
1993M05	5/11	WASDE
1993M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
1993M07	7/12	WASDE
1993M08	8/11	WASDE
1993M09	9/9; 9/30	9/9 (WASDE); 9/30 (GS)
1993M10	10/12	WASDE
1993M11	11/9	WASDE
1993M12	12/9	WASDE
1994M01	1/12	WASDE & GS
1994M02	2/10	WASDE
1994M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
1994M04	4/12	WASDE
1994M05	5/10	WASDE
1994M06	6/9; 6/30	6/9 (WASDE); 6/30 (GS & AR)
1994M07	7/12	WASDE
1994M08	8/11	WASDE
1994M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
1994M10	10/12	WASDE
1994M11	11/9	WASDE
1994M12	12/9	WASDE

continued on next page

Table A.1 continued

Period	Publication date	Reports
1995M01	1/12	WASDE & GS
1995M02	2/10	WASDE
1995M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
1995M04	4/11	WASDE
1995M05	5/11	WASDE
1995M06	6/12; 6/30	6/12 (WASDE); 6/30 (GS & AR)
1995M07	7/12	WASDE
1995M08	8/11	WASDE
1995M09	9/12; 9/29	9/12 (WASDE); 9/29 (GS)
1995M10	10/11	WASDE
1995M11	11/9	WASDE
1995M12	12/12	WASDE
1996M01	1/16	WASDE & GS
1996M02	2/9	WASDE
1996M03	3/12; 3/29	3/12 (WASDE); 3/29 (GS & PP)
1996M04	4/11	WASDE
1996M05	5/10	WASDE
1996M06	6/12; 6/28	6/12 (WASDE); 6/28 (GS & AR)
1996M07	7/12	WASDE
1996M08	8/12	WASDE
1996M09	9/11; 9/30	9/11 (WASDE); 9/30 (GS)
1996M10	10/11	WASDE
1996M11	11/12	WASDE
1996M12	12/12	WASDE
1997M01	1/10	WASDE & GS
1997M02	2/12	WASDE
1997M03	3/11; 3/31	3/11 (WASDE); 3/31 (GS & PP)
1997M04	4/11	WASDE
1997M05	5/12	WASDE
1997M06	6/12; 6/30	6/12 (WASDE); 6/30 (GS & AR)
1997M07	7/11	WASDE
1997M08	8/12	WASDE
1997M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
1997M10	10/10	WASDE
1997M11	11/10	WASDE
1997M12	12/11	WASDE
1998M01	1/13	WASDE & GS
1998M02	2/11	WASDE
1998M03	3/12; 3/31	3/12 (WASDE); 3/31 (GS & PP)
1998M04	4/9	WASDE
1998M05	5/12	WASDE
1998M06	6/12[†]; 6/30	6/12 (WASDE); 6/30 (GS & AR)
1998M07	7/10	WASDE
1998M08	8/12	WASDE
1998M09	9/11; 9/30	9/11 (WASDE); 9/30 (GS)
1998M10	10/9	WASDE
1998M11	11/10	WASDE
1998M12	12/11	WASDE
1999M01	1/12	WASDE & GS
1999M02	2/10	WASDE
1999M03	3/11; 3/31	3/11 (WASDE); 3/31 (GS & PP)
1999M04	4/9	WASDE
1999M05	5/12	WASDE
1999M06	6/11; 6/30	6/11 (WASDE); 6/30 (GS & AR)

continued on next page

Table A.1 continued

Period	Publication date	Reports
1999M07	7/12	WASDE
1999M08	8/12	WASDE
1999M09	9/10; 9/30	9/10 (WASDE); 9/30 (GS)
1999M10	10/8	WASDE
1999M11	11/10	WASDE
1999M12	12/10	WASDE
2000M01	1/12	WASDE & GS
2000M02	2/11	WASDE
2000M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2000M04	4/11	WASDE
2000M05	5/12	WASDE
2000M06	6/9; 6/30	6/9 (WASDE); 6/30 (GS & AR)
2000M07	7/12	WASDE
2000M08	8/11	WASDE
2000M09	9/12; 9/29	9/12 (WASDE); 9/29 (GS)
2000M10	10/12	WASDE
2000M11	11/9	WASDE
2000M12	12/12	WASDE
2001M01	1/11	WASDE & GS
2001M02	2/8	WASDE
2001M03	3/8; 3/30	3/8 (WASDE); 3/30 (GS & PP)
2001M04	4/10	WASDE
2001M05	5/10	WASDE
2001M06	6/12; 6/29†	6/12 (WASDE); 6/29 (GS & AR)
2001M07	7/11	WASDE
2001M08	8/10†	WASDE
2001M09	9/14†; 9/28†	9/14 (WASDE); 9/28 (GS)
2001M10	10/12†	WASDE
2001M11	11/9†	WASDE
2001M12	12/11†	WASDE
2002M01	1/11	WASDE & GS
2002M02	2/8	WASDE
2002M03	3/8; 3/28	3/8 (WASDE); 3/28 (GS & PP)
2002M04	4/10	WASDE
2002M05	5/10	WASDE
2002M06	6/12†; 6/28	6/12 (WASDE); 6/28 (GS & AR)
2002M07	7/11	WASDE
2002M08	8/12	WASDE
2002M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2002M10	10/11	WASDE
2002M11	11/12	WASDE
2002M12	12/10	WASDE
2003M01	1/10	WASDE & GS
2003M02	2/11	WASDE
2003M03	3/11; 3/31	3/11 (WASDE); 3/31 (GS & PP)
2003M04	4/10	WASDE
2003M05	5/12	WASDE
2003M06	6/11; 6/30	6/11 (WASDE); 6/30 (GS & AR)
2003M07	7/11	WASDE
2003M08	8/12	WASDE
2003M09	9/11; 9/30	9/11 (WASDE); 9/30 (GS)
2003M10	10/10	WASDE
2003M11	11/12	WASDE
2003M12	12/11	WASDE

continued on next page

Table A.1 continued

Period	Publication date	Reports
2004M01	1/12	WASDE & GS
2004M02	2/10	WASDE
2004M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2004M04	4/8	WASDE
2004M05	5/12	WASDE
2004M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2004M07	7/12	WASDE
2004M08	8/12	WASDE
2004M09	9/10; 9/30	9/10 (WASDE); 9/30 (GS)
2004M10	10/12	WASDE
2004M11	11/12	WASDE
2004M12	12/10	WASDE
2005M01	1/12	WASDE & GS
2005M02	2/9	WASDE
2005M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2005M04	4/8	WASDE
2005M05	5/12	WASDE
2005M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2005M07	7/12	WASDE
2005M08	8/12	WASDE
2005M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2005M10	10/12	WASDE
2005M11	11/10	WASDE
2005M12	12/9	WASDE
2006M01	1/12	WASDE & GS
2006M02	2/9	WASDE
2006M03	3/10†; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2006M04	4/10	WASDE
2006M05	5/12	WASDE
2006M06	6/9; 6/30	6/9 (WASDE); 6/30 (GS & AR)
2006M07	7/12	WASDE
2006M08	8/11	WASDE
2006M09	9/12; 9/29	9/12 (WASDE); 9/29 (GS)
2006M10	10/12	WASDE
2006M11	11/9	WASDE
2006M12	12/11	WASDE
2007M01	1/12	WASDE & GS
2007M02	2/9	WASDE
2007M03	3/9; 3/30	3/9 (WASDE); 3/30 (GS & PP)
2007M04	4/10	WASDE
2007M05	5/11	WASDE
2007M06	6/11; 6/29	6/11 (WASDE); 6/29 (GS & AR)
2007M07	7/12	WASDE
2007M08	8/10	WASDE
2007M09	9/12; 9/28	9/12 (WASDE); 9/28 (GS)
2007M10	10/12	WASDE
2007M11	11/9	WASDE
2007M12	12/11	WASDE
2008M01	1/11	WASDE & GS
2008M02	2/8	WASDE
2008M03	3/11; 3/31	3/11 (WASDE); 3/31 (GS & PP)
2008M04	4/9	WASDE
2008M05	5/9	WASDE
2008M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)

continued on next page

Table A.1 continued

Period	Publication date	Reports
2008M07	7/11	WASDE
2008M08	8/12	WASDE
2008M09	9/12†; 9/30	9/12 (WASDE); 9/30 (GS)
2008M10	10/10	WASDE
2008M11	11/10	WASDE
2008M12	12/11	WASDE
2009M01	1/12	WASDE & GS
2009M02	2/10	WASDE
2009M03	3/11; 3/31	3/11 (WASDE); 3/31 (GS & PP)
2009M04	4/9	WASDE
2009M05	5/12	WASDE
2009M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2009M07	7/10	WASDE
2009M08	8/12	WASDE
2009M09	9/11; 9/30	9/11 (WASDE); 9/30 (GS)
2009M10	10/9	WASDE
2009M11	11/10	WASDE
2009M12	12/10	WASDE
2010M01	1/12	WASDE & GS
2010M02	2/9	WASDE
2010M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2010M04	4/9	WASDE
2010M05	5/11	WASDE
2010M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2010M07	7/9	WASDE
2010M08	8/12	WASDE
2010M09	9/10; 9/30	9/10 (WASDE); 9/30 (GS)
2010M10	10/8	WASDE
2010M11	11/9	WASDE
2010M12	12/10	WASDE
2011M01	1/12	WASDE & GS
2011M02	2/9	WASDE
2011M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2011M04	4/8	WASDE
2011M05	5/11	WASDE
2011M06	6/9; 6/30	6/9 (WASDE); 6/30 (GS & AR)
2011M07	7/12	WASDE
2011M08	8/11	WASDE
2011M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2011M10	10/12	WASDE
2011M11	11/9	WASDE
2011M12	12/9	WASDE
2012M01	1/12	WASDE & GS
2012M02	2/9	WASDE
2012M03	3/9; 3/30	3/9 (WASDE); 3/30 (GS & PP)
2012M04	4/10	WASDE
2012M05	5/10	WASDE
2012M06	6/12; 6/29	6/12 (WASDE); 6/29 (GS & AR)
2012M07	7/11	WASDE
2012M08	8/10	WASDE
2012M09	9/12; 9/28	9/12 (WASDE); 9/28 (GS)
2012M10	10/11	WASDE
2012M11	11/9	WASDE
2012M12	12/11	WASDE

continued on next page

Table A.1 continued

Period	Publication date	Reports
2013M01	1/11	WASDE & GS
2013M02	2/8	WASDE
2013M03	3/8; 3/28	3/8 (WASDE); 3/28 (GS & PP)
2013M04	4/10	WASDE
2013M05	5/10	WASDE
2013M06	6/12; 6/28	6/12 (WASDE); 6/28 (GS & AR)
2013M07	7/11	WASDE
2013M08	8/12	WASDE
2013M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2013M10	—	Canceled (federal government shutdown)
2013M11	11/8	WASDE
2013M12	12/10	WASDE
2014M01	1/10	WASDE & GS
2014M02	2/10	WASDE
2014M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2014M04	4/9	WASDE
2014M05	5/9	WASDE
2014M06	6/11; 6/30	6/11 (WASDE); 6/30 (GS)
2014M07	7/11	WASDE
2014M08	8/12	WASDE
2014M09	9/11; 9/30	9/11 (WASDE); 9/30 (GS)
2014M10	10/10	WASDE
2014M11	11/10	WASDE
2014M12	12/10	WASDE
2015M01	1/12	WASDE & GS
2015M02	2/10	WASDE
2015M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2015M04	4/9	WASDE
2015M05	5/12	WASDE
2015M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2015M07	7/10	WASDE
2015M08	8/12	WASDE
2015M09	9/11; 9/30	9/11 (WASDE); 9/30 (GS)
2015M10	10/9	WASDE
2015M11	11/10	WASDE
2015M12	12/9	WASDE
2016M01	1/12	WASDE & GS
2016M02	2/9	WASDE
2016M03	3/9; 3/31	3/9 (WASDE); 3/31 (GS & PP)
2016M04	4/12	WASDE
2016M05	5/10	WASDE
2016M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2016M07	7/12	WASDE
2016M08	8/12	WASDE
2016M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2016M10	10/12	WASDE
2016M11	11/9	WASDE
2016M12	12/9	WASDE
2017M01	1/12	WASDE & GS
2017M02	2/9	WASDE
2017M03	3/9; 3/31	3/9 (WASDE); 3/31 (GS & PP)
2017M04	4/11	WASDE
2017M05	5/10	WASDE
2017M06	6/9; 6/30	6/9 (WASDE); 6/30 (GS & AR)
2017M07	7/12	WASDE

continued on next page

Table A.1 continued

Period	Publication date	Reports
2017M08	8/10	WASDE
2017M09	9/12; 9/29	9/12 (WASDE); 9/29 (GS)
2017M10	10/12	WASDE
2017M11	11/9	WASDE
2017M12	12/12	WASDE
2018M01	1/12	WASDE & GS
2018M02	2/8	WASDE
2018M03	3/8; 3/29	3/8 (WASDE); 3/29 (GS & PP)
2018M04	4/10	WASDE
2018M05	5/10	WASDE
2018M06	6/12; 6/29	6/12 (WASDE); 6/29 (GS & AR)
2018M07	7/12	WASDE
2018M08	8/10	WASDE
2018M09	9/12; 9/28	9/12 (WASDE); 9/28 (GS)
2018M10	10/11	WASDE
2018M11	11/8	WASDE
2018M12	12/11	WASDE
2019M01	—	Canceled (federal government shutdown)
2019M02	2/8	WASDE
2019M03	3/8; 3/29	3/8 (WASDE); 3/29 (GS & PP)
2019M04	4/9	WASDE
2019M05	5/10	WASDE
2019M06	6/11; 6/28	6/11 (WASDE); 6/28 (PP & AR)
2019M07	7/11	WASDE
2019M08	8/12	WASDE
2019M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2019M10	10/10	WASDE
2019M11	11/8	WASDE
2019M12	12/10	WASDE
2020M01	1/10	WASDE & GS
2020M02	2/11	WASDE
2020M03	3/10; 3/31	3/10 (WASDE); 3/31 (GS & PP)
2020M04	4/9	WASDE
2020M05	5/12	WASDE
2020M06	6/11; 6/30	6/11 (WASDE); 6/30 (GS & AR)
2020M07	7/10	WASDE
2020M08	8/12	WASDE
2020M09	9/11; 9/30	9/11 (WASDE); 9/30 (GS)
2020M10	10/9	WASDE
2020M11	11/10	WASDE
2020M12	12/10	WASDE
2021M01	1/12	WASDE & GS
2021M02	2/9	WASDE
2021M03	3/9; 3/31	3/9 (WASDE); 3/31 (GS & PP)
2021M04	4/9	WASDE
2021M05	5/12	WASDE
2021M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2021M07	7/12	WASDE
2021M08	8/12	WASDE
2021M09	9/10; 9/30	9/10 (WASDE); 9/30 (GS)
2021M10	10/12	WASDE
2021M11	11/9	WASDE
2021M12	12/9	WASDE
2022M01	1/12	WASDE & GS
2022M02	2/9	WASDE

continued on next page

Table A.1 continued

Period	Publication date	Reports
2022M03	3/9; 3/31	3/9 (WASDE); 3/31 (GS & PP)
2022M04	4/8	WASDE
2022M05	5/12	WASDE
2022M06	6/10; 6/30	6/10 (WASDE); 6/30 (GS & AR)
2022M07	7/12	WASDE
2022M08	8/12	WASDE
2022M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2022M10	10/12	WASDE
2022M11	11/9	WASDE
2022M12	12/9	WASDE
2023M01	1/12	WASDE & GS
2023M02	2/8	WASDE
2023M03	3/8; 3/31	3/8 (WASDE); 3/31 (GS & PP)
2023M04	4/11	WASDE
2023M05	5/12	WASDE
2023M06	6/9; 6/30	6/9 (WASDE); 6/30 (GS & AR)
2023M07	7/12	WASDE
2023M08	8/11	WASDE
2023M09	9/12; 9/29	9/12 (WASDE); 9/29 (GS)
2023M10	10/12	WASDE
2023M11	11/9	WASDE
2023M12	12/8	WASDE
2024M01	1/12	WASDE & GS
2024M02	2/8	WASDE
2024M03	3/8; 3/28	3/8 (WASDE); 3/28 (GS & PP)
2024M04	4/11	WASDE
2024M05	5/10	WASDE
2024M06	6/12; 6/28	6/12 (WASDE); 6/28 (GS & AR)
2024M07	7/12	WASDE
2024M08	8/12	WASDE
2024M09	9/12; 9/30	9/12 (WASDE); 9/30 (GS)
2024M10	10/11	WASDE
2024M11	11/8	WASDE
2024M12	12/10	WASDE

B Analyst Poll Surprises

Eikon Commodity Estimates Polls. Reuters polls market analysts for estimates of key demand and supply items for crops such as corn, soybeans and wheat before the release of WASDE reports. The poll covers the market expectation of ending stocks for each of the three commodities. Additionally, it contains the analyst estimates of production and yield for corn and soybeans, and only production estimates for wheat. The production and yield estimates are only polled when the NASS publishes a new production forecast for the crop (Table B.2).[26]

[26]Specifically, analysts' estimates for the production and yield forecasts are available from July to November, and in January of the following year, for corn and soybeans. For wheat, only production forecasts are available, from May to August. The farm calendars of corn and soybeans are slightly different. Both are

Here, I consider a potentially more robust way to remove the demand elements from the USDA announcement. Refinitiv Eikon surveys market participants around the USDA announcements about their expectations of the production, yields and ending stocks of key agricultural commodities such as soybeans, corn and wheat. I suggest measuring agricultural supply-side shocks as surprises relative to the market expectation of crop production and yield.

I propose to further weed out the potential demand component of the price surprises in equation (8) by measuring the difference between the expected yield and production and the actual yield and production estimates at the WASDE release. Reuters polls analysts about expectations for commodities such as corn, soybeans and wheat only. Further, for wheat, only the production surprise (s^P) is available, not the yield surprise (s^Y). Thus, I construct two variants of the instrument as follows:

$$z_t^Y = \sum_{i \in \{\text{corn,soy}\}} w_{i,y} s_{i,t}^Y, \quad \text{shocks: yield} \tag{13}$$

$$z_t^P = \sum_{i \in \{\text{corn,soy,wheat}\}} w_{i,y} s_{i,t}^P \quad \text{shocks: production} \tag{14}$$

where $s_{i,t}^j = 100 \ln(A_{j,i,t}/E_{j,i,t})$, $j \in \{Y, P\}$, measures the surprises in the market estimates ($E_{j,i,t}$) of the yield (Y) or production (P) of commodity i from the published estimates of WASDE ($A_{j,i,t}$) at date t. The z_t^Y instrument only measures the yield-related changes in corn and soybeans, whereas the instrument z_t^P measures the production-related surprises across corn, soybeans and wheat. To measure the impact of these surprises on commodity futures, I estimate the following regression:

$$s_{i,t} = \alpha_i + \beta_Y s_{i,t}^Y + \beta_A s_{i,t}^A + \beta_R s_{i,t}^R + \varepsilon_{i,t}, \tag{15}$$

where $s_{i,t}$ is the PC1 futures-price surprise of commodity i from equation (8) and $s_{i,t}^A = s_{i,t}^P - s_{i,t}^Y$. The surprise in the rest $s_{i,t}^R$ of the WASDE balance sheet is defined as

$$s_{i,t}^R = 100 \, \frac{(A_{ES,i,t} - E_{ES,i,t}) - (A_{P,i,t} - E_{P,i,t})}{TU_{i,t-1}}. \tag{16}$$

The difference between the production surprise ($s_{i,t}^P$) and yield surprise ($s_{i,t}^Y$) estimates the market-implied surprise in the area planted or area expected to be harvested ($s_{i,t}^A$). $s_{i,t}^R$ denotes the surprise in the rest of the components of the WASDE balance sheet normalized by the total use in the previous month's WASDE release ($TU_{i,t-1}$) of the commodity i.[27]

planted in the spring and harvested in the fall. Corn is planted a little earlier (April to May) than soybeans (May to June).

[27]The WASDE measures the total use (TU) of the commodity, such as domestic use (DU) and exports

The Stata estimates of equation (15) are reported in Table B.3. I find that yield surprises and area (harvested or planted) surprises are negatively associated with the change in commodity futures prices on the day of the WASDE release. Positive yield surprises, an indication of more supply than expected by the market, remain negatively associated with the commodity futures price movement when we drop the area or the rest surprise in columns (2) and (3), respectively. The yield surprises explain about one-third of the variation in the commodity futures price movement on the day of the WASDE release. When considering production surprise alone, in column (4), or the surprise in ending stocks in column (5), which contains both supply and demand components, the commodity futures price response remains significant.

Figure B.1 plots the time series of my constructed instruments compared to the AJ instruments. It shows that my constructed instruments display a consistent seasonal pattern. The production and yield surprises are consistently positive in months such as August, September and October, which also happen to be the months in which the WASDE releases new production estimates and in which Reuters polls analysts (Figure B.2). On the other hand, in months with no production or yield update, such as February, March and April, the constructed instruments are weakly negative but statistically insignificant. This seasonal pattern is especially true for the instruments constructed with "GK" timing, but slightly weaker for the "calendar sum" measures. These findings so far indicate that analyst surveys have a consistent bias, which needs to be addressed and may lead to low relevance of the instruments later. I will address this in Section 5.

(EX), and the supply components (TS), such as beginning stocks (BS), production (P) and imports (IM). Reuters only polls analysts on the estimates of yields, production (P) and ending stocks (ES). By the balance-sheet identity, $TU = DU + EX$ and $TS = BS + P + IM$, so $ES = TS - TU = BS + P + IM - DU - EX$. Thus, $\Delta R = \Delta ES - \Delta P$ captures the non-production surprise in the analyst estimates. This may include surprises in the supply components, such as beginning stocks or imports, as well as in the demand components, such as domestic use and exports.

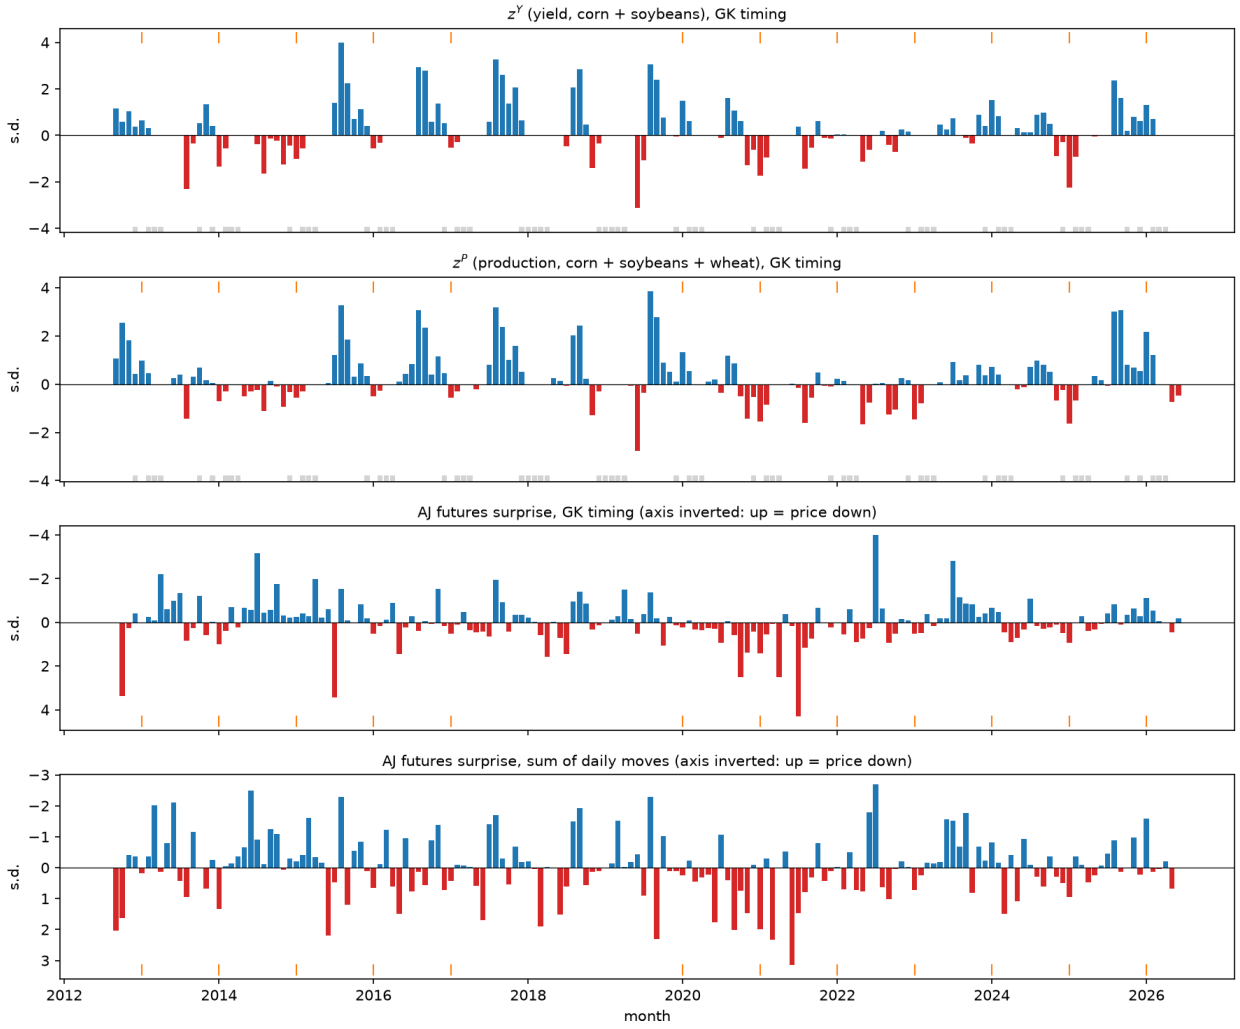


Figure B.1: Monthly poll surprises (z^Y, z^P) and the Adjemian–Jo futures surprise (GK and sum), 2012M09–2026M06, in standard deviations. *Notes:* All four panels use the same sample and are scaled by their own standard deviation over it. In every panel blue means more supply (lower price) and red less supply (higher price). z^Y and z^P are positive when USDA reports more than the Reuters poll mean. The Adjemian–Jo surprise is a futures price change, positive when prices rise, so the two AJ panels are drawn with an **inverted vertical axis**: up is a price fall, that is, more supply. Gray marks at the bottom of the poll panels are months without a yield or production poll, where the poll instruments are zero by construction. Orange ticks mark months whose release shared its day with Grain Stocks, Prospective Plantings or Acreage. Source: Reuters polls, USDA WASDE, daily futures.

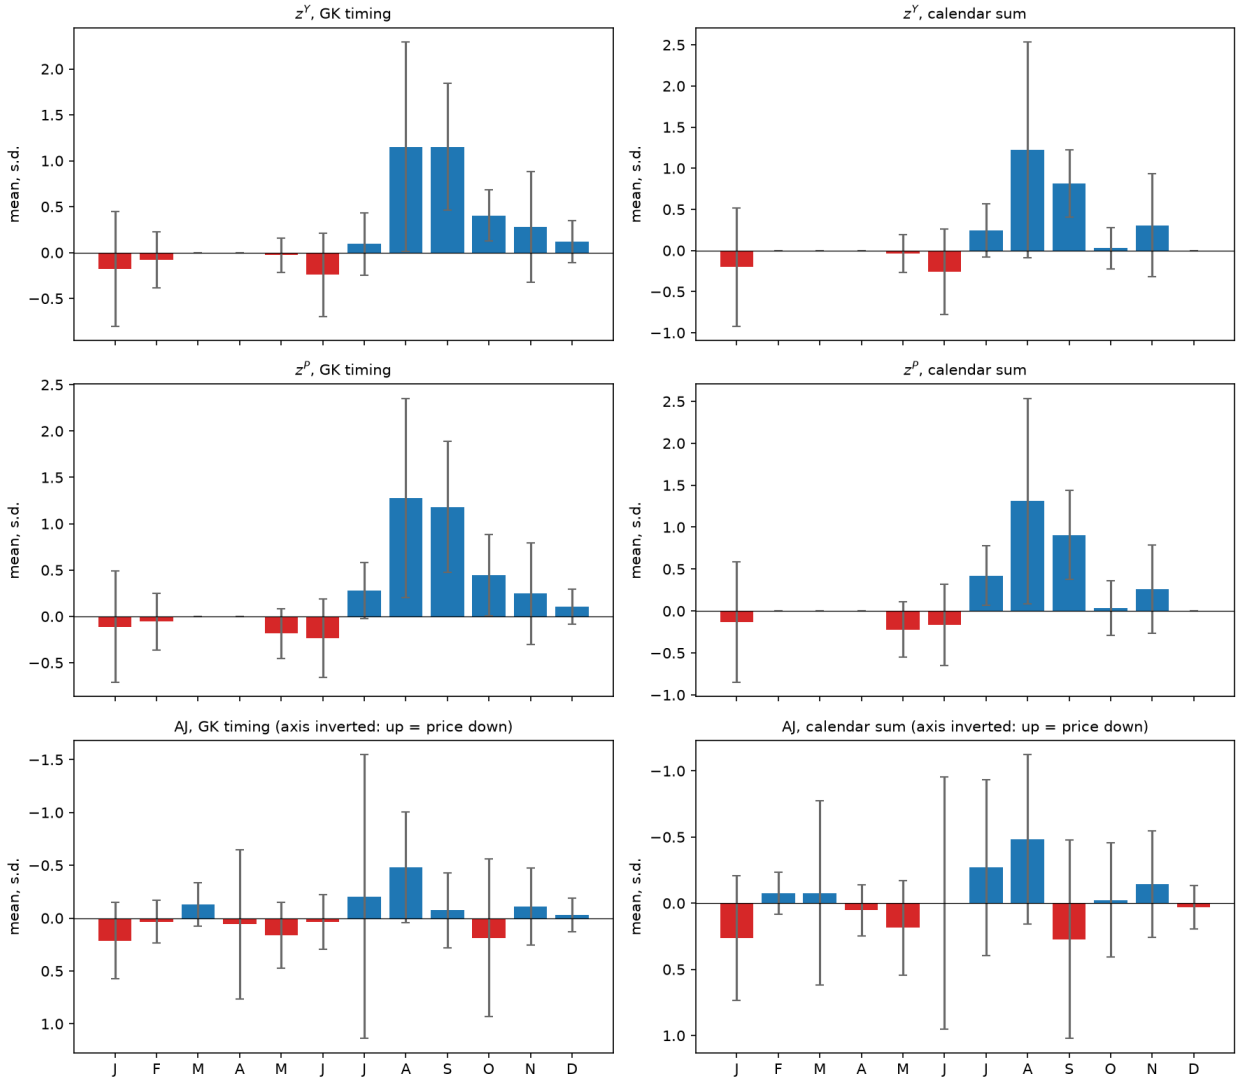


Figure B.2: Seasonal pattern of the instruments: month-of-year mean of z^Y (top), z^P (middle) and the Adjemian–Jo surprise (bottom), under GK timing (left) and the calendar sum (right), 2012M09–2026M06, in standard deviations. *Notes:* Each bar is $\bar{z}_m = \frac{1}{n_m}\sum_y z_{m,y}$, with whiskers $\pm 2\, s_m/\sqrt{n_m}$, where s_m is the standard deviation across years and n_m the number of years (13 or 14). This is the classical standard error of a mean (years independent; no heteroskedasticity or serial-correlation correction, no bootstrap). Blue means more supply (lower price). The AJ panels have an **inverted vertical axis**, so up is a price fall. The February–April and December releases have no yield or production poll, so z^Y and z^P are zero there by construction (under GK, only March and April).

Table B.1: Usual planting and harvesting dates of the five crops

		Planting		Harvest		Months, planting	Marketing
Crop	State	Window	Most active	Window	Most active	to harvest	year
Corn	Iowa	Apr 19–May 26	Apr 25–May 18	Sep 21–Nov 21	Oct 5–Nov 9	5.6	Sep–Aug
Soybeans	Iowa	May 2–Jun 16	May 8–Jun 2	Sep 21–Oct 31	Sep 28–Oct 20	4.7	Sep–Aug
Winter wheat	Kansas	Sep 10–Nov 1	Sep 15–Oct 20	Jun 15–Jul 15	Jun 20–Jul 5	8.8	Jun–May
Spring wheat	North Dakota	Apr 16–Jun 3	Apr 24–May 25	Aug 1–Sep 25	Aug 8–Sep 13	3.6	Jun–May
Rice	Arkansas	Apr 6–Jun 2	Apr 14–May 19	Aug 29–Oct 24	Sep 9–Oct 10	4.8	Aug–Jul
Oats	Wisconsin	Apr 11–May 18	Apr 20–May 11	Jul 22–Aug 30	Jul 28–Aug 24	3.4	Jun–May

Notes: Usual dates in the state with the most harvested acres of each crop in 2009 (spring oats for oats). The window runs from about 5 percent to about 95 percent of the operation complete; the most active period runs from about 15 to about 85 percent. "Months, planting to harvest" is the distance between the midpoints of the two most active periods. The months agree with the U.S. national calendar of the Agricultural Market Information System (AMIS), as used by De Winne and Peersman (2016): planting April–May and harvest September–November for corn, May–June and September–October for soybeans, April–June and August–October for rice, and for wheat planting April–May (spring) and September–October (winter) with harvest June–September. The marketing year is the USDA one; the WASDE publishes the first projection for the new marketing year in May. Source: USDA NASS, *Usual Planting and Harvesting Dates for U.S. Field Crops*, October 2010.

Table B.2: Reuters analyst polls from Eikon used for the USDA surprises

Code	Report	Crop	Quantity polled	Months polled	First poll	Polls
WASDE, U.S. headline marketing year (matched to the WASDE ledger)						
CP	WASDE	Corn	Production, million bushels	Jul–Nov, Jan	2011M01	94
CY	WASDE	Corn	Yield, bushels per acre	Jul–Nov, Jan	2012M09	87
CES	WASDE	Corn	Ending stocks, million bushels	Every WASDE	2010M12	188
SP	WASDE	Soybeans	Production, million bushels	Jul–Nov, Jan	2011M01	93
SY	WASDE	Soybeans	Yield, bushels per acre	Jul–Nov, Jan	2012M09	87
SES	WASDE	Soybeans	Ending stocks, million bushels	Every WASDE	2010M12	186
WP	WASDE	Wheat	Production, million bushels	May–Aug	2011M08	60
WES	WASDE	Wheat	Ending stocks, million bushels	Every WASDE	2010M12	188
WASDE, other quantities (not matched to the ledger)						
CEO, SEO	WASDE old crop	Corn, soybeans	U.S. ending stocks of the previous marketing year	May–Sep	2012M09	71, 70
AWEO	WASDE old crop	Wheat	U.S. ending stocks of the previous marketing year	May–Jun	2013M05	29
WCE, WSE, WAWE	WASDE world	Corn, soybeans, wheat	World ending stocks, million tonnes	Every WASDE	2012M09	167
NASS reports released on separate dates						
CSQ, SSQ, WSQ	Grain Stocks	Corn, soybeans, wheat	U.S. stocks on March 1, June 1, September 1 and December 1	Mar, Jun, Sep, Dec/Jan	2010M12	64
CORN, SOY, AW	Prospective Plantings	Corn, soybeans, wheat	Intended planted area, March survey	March	2010M12	14
APCO, APSB, APWH	Acreage	Corn, soybeans, wheat	Planted area, June survey	June	2011M12	14

Notes: Polls are pulled from Refinitiv Workspace through the Eikon Data API on 3 October 2026. Each series is a Reuters Instrument Code of the form pUS*code*=*statistic*, where the statistic is the median (M), mean (E), high (H), low (L), number of analysts (C) or standard deviation (T). Surprises are computed from the median poll or the mean poll. Eikon stores quantities in bushels, which I convert to million bushels; yields are in bushels per acre and area in million acres. Polls cover corn, soybeans and wheat only, with nothing for rice or oats, and nothing before December 2010. Production and yield are polled only when NASS publishes a new production forecast. The September 2026 observation is the last in the data. Polls are dated at the end of the reference month and are matched to the WASDE release of that month. The number of polls for the Prospective Plantings and Acreage series is 14 because some early years are missing. Eikon's own series for the actual value (USxx=ECI) is not used.

Table B.3: Commodity futures responses to Reuters poll surprises

	Dependent variable: release-day PC1 futures surprise (%)				
	(1) Main	(2) No area	(3) Yield only	(4) Production	(5) Ending stocks
Yield surprise	-1.373^{***}	-1.263^{***}	-0.930^{***}		
	(0.144)	(0.145)	(0.138)		
Area surprise	-0.651^{***}				
	(0.199)				
Production surprise				-0.975^{***}	
				(0.091)	
Rest surprise	-0.925^{***}	-0.667^{***}		-0.668^{***}	
	(0.180)	(0.185)		(0.120)	
Ending-stocks surprise					-0.743^{***}
					(0.111)
Crops	C+S	C+S	C+S	C+S+W	C+S+W
Crop FEs	✓	✓	✓	✓	✓
N	169	169	169	243	552
R^2	0.375	0.341	0.285	0.318	0.179
Releases (clusters)	85	85	85	117	185

Notes: The dependent variable is the release-day PC1 futures surprise from the five nearest maturities. Yield, area, and production surprises are measured in percent; the rest and ending-stocks surprises are measured as a percent of prior total use. Columns (1)–(3) pool corn and soybeans (C+S) and include crop fixed effects: column (1) is the main specification, column (2) drops area, and column (3) includes yield alone. Columns (4)–(5) pool corn, soybeans, and wheat (C+S+W), with crop fixed effects; column (4) includes production and rest surprises, while column (5) includes the ending-stocks surprise alone. Column (4)'s complete-case sample runs January 12, 2011, through August 12, 2026; column (5) uses the full available ending-stocks sample, December 10, 2010, through August 12, 2026. The complete-case sample for columns (1)–(3) runs from September 12, 2012, through August 12, 2026. Column (5) has more observations than columns (1)–(4) because Reuters does not poll analysts for yield, area, and production expectations in every release month; ending-stocks poll observations are available in more months. Standard errors clustered by release date are in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B.4: Tests on instrument strength on the poll months

	AJ futures		Orthogonalized yield revision $\tilde{r}^Y$	
	(1) Sum	(2) GK	(3) Sum	(4) GK
Seasonally adjusted (SA)				
Coefficient	0.057	0.341	−0.621	−0.610
F-statistic	0.09	3.41	11.98	11.53
F-statistic (robust)	0.09	3.53	16.49	15.58
F-statistic (Newey–West)	0.10	3.71	12.87	12.94
R^2	0.06	2.28	7.58	7.32
R^2 (adjusted)	−0.62	1.61	6.95	6.68
Observations	148	148	148	148
Unadjusted with month dummies (NSA + D)				
Coefficient	0.240	0.321	−0.547	−0.530
F-statistic	1.84	3.34	10.12	9.46
F-statistic (robust)	1.67	3.80	15.52	13.64
F-statistic (Newey–West)	1.78	3.83	12.49	11.63
R^2	1.24	2.24	6.48	6.09
R^2 (adjusted)	0.57	1.57	5.84	5.44
Observations	148	148	148	148

Notes: Same regressions and columns as Table 3, on the poll months 2012M09–2024M12 (148 months), the months for which the poll surprises of Table B.2 exist. The VAR residual and the orthogonalized series are those of the full sample (148 months are too few to re-estimate the 25 coefficients of the projection). Instruments are scaled to unit standard deviation over the regression months, so b is in percent per standard deviation; the futures surprise is a price move (expected $b > 0$), and the revisions are positive when supply is larger (expected $b < 0$). $F = (b/\text{se})^2$; the robust F uses HC1 standard errors and equals the Montiel Olea–Pflueger effective F with one instrument; the Newey–West F uses 3 lags. R^2 in percent. The orthogonalized series is a generated instrument; the standard errors do not account for the first step.

C Raw WASDE Revisions

Figure C.1 shows that yield revisions are fairly evenly distributed with a very large negative revision in August 2012 driven by drought (). On the other hand, harvested area revisions show a large negative revision around July 2020, likely driven by Covid-19 restrictions. Figure C.2 plots the monthly average of WASDE revisions for all months; area revisions show a consistent seasonal pattern, with the USDA revising down the area harvested in September, October and November. However, yield revisions show no such consistent seasonal pattern.

Table C.1 shows that a higher yield revision is associated with a lower commodity futures price on the WASDE release day. This effect remains significant when excluding either the area revisions (column 2) or both the area revision and the rest-of-the-balance-sheet revision (column 3). Similarly, production revisions and ending stocks revisions also show a strong and negative correlation with the futures price.

Table C.3 reports the strength of the instruments for all three types of shocks: AJ-based instruments constructed from the daily change in commodity futures prices, instruments constructed from the Reuters analyst poll surprises, and WASDE revision-based instruments. I test instrument relevance on two data panels: seasonally adjusted data, and non-seasonally adjusted data with month dummies.

First, the AJ commodity futures price instrument constructed with "GK" timing (column 2) shows a much larger robust F-statistic than the simple sum of the daily surprises in the month (column 1) across all three panels. This vindicates the Gertler and Karadi (2015) approach of aggregating daily surprises over the simple monthly sum method. Thus, I use the "GK" timing method for the rest of the instruments.

Second, as expected given the strong seasonal pattern observed in the Reuters poll surprise instruments, they are weak in explaining the residual of the real field crop price from the VAR. However, this is more of a sample period issue, as for the same sample period for which the poll-based surprise is available, both the AJ futures and my constructed revisions exhibit a weak instrument problem, with F-statistics consistently below 1 (Panel B).

Third, the revision instruments on the full sample in Panel A are strong, with robust F-statistics above 10. The yield revisions (column 5) exhibit much stronger relevance than the production revisions (column 6).

While "GK" timing makes the instruments much stronger, it also introduces a mechanical correlation between consecutive values of the instrument. This is especially true for the revision instruments. Figure C.3 shows that the AJ futures instrument constructed with GK timing exhibits autocorrelation, while the AJ monthly sum does not. Autocorrelation is a more significant problem in revision instruments constructed with "GK" timing than in those constructed as a monthly sum. Further, it seems that even in the calendar sum instrument, the USDA revises yields and production in the same direction in consecutive releases of a season.

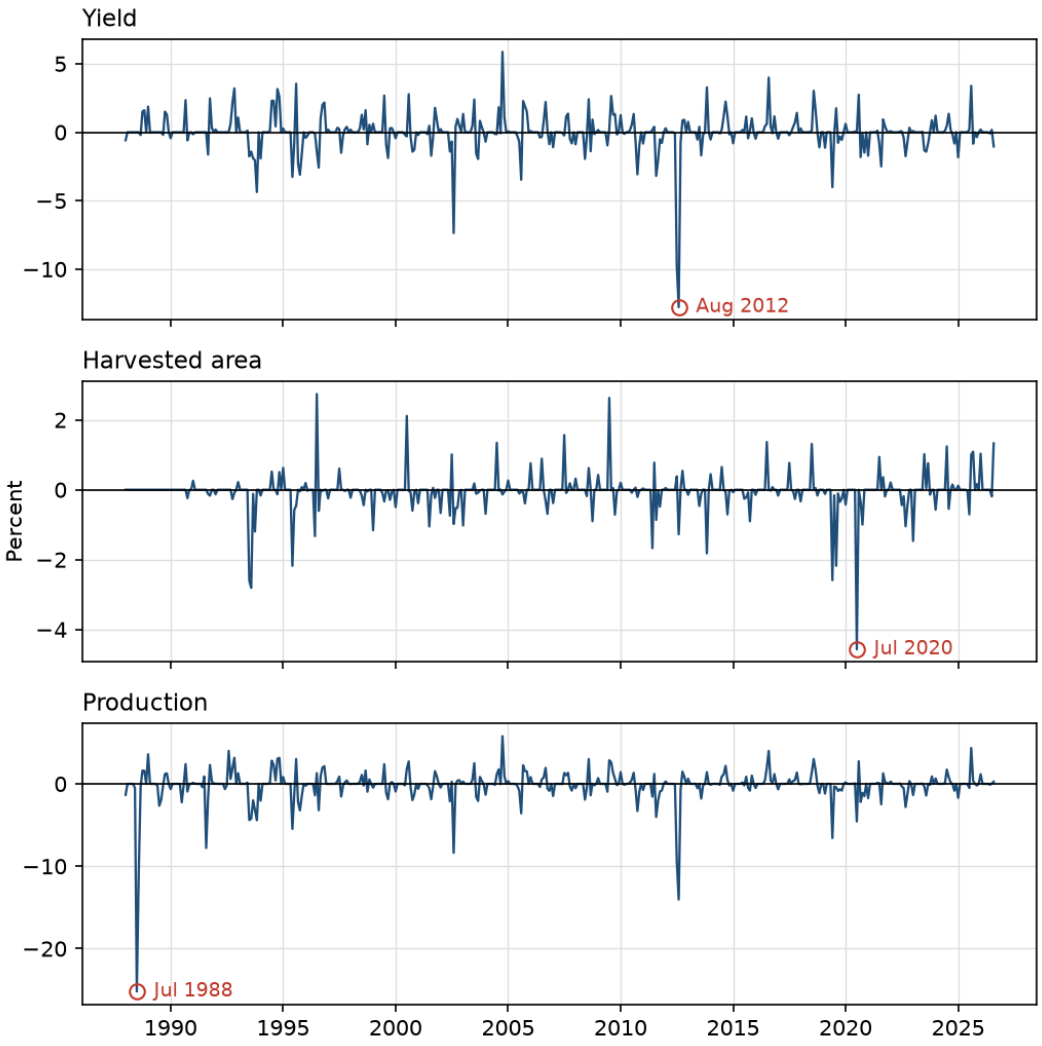


Figure C.1: Monthly WASDE revision instruments: yield (r^Y, top), harvested area (r^A, middle) and production (r^P, bottom), calendar-month sum, 1988M01–2026M08, in percent. *Notes:* For each crop i and release t, the revision is $r^X_{i,t} = 100 \ln(X_{i,t}/X_{i,t-1})$, USDA's month-on-month change in its estimate of the newest marketing year, where $X_{i,t-1}$ is last month's value as reprinted in the same report. The instrument is $z^X_t = \sum_i w_{i,y}\, r^X_{i,t}$ over all five crops (corn, soybeans, wheat, rice, oats), with $w_{i,y}$ crop i's share of U.S. production in metric tons in the release's calendar year y, the same weights as in equation (9), extended to 1987–1991 with the same WASDE rule. The basket is fixed: a release enters only when all five crops have a revision. May releases (the first new-crop projection, not a revision) and, before 1993, the June–August area and yield releases (when the new crop is not yet printed) therefore add nothing, and a month without a release is zero. Each release is counted in full in the calendar month in which it is published. The figure shows the full history of the series; the first stage and the VAR identification use 1993M01–2024M12. Positive means USDA now sees more supply than a month ago. The circled month is each series' largest absolute value. Source: USDA WASDE.

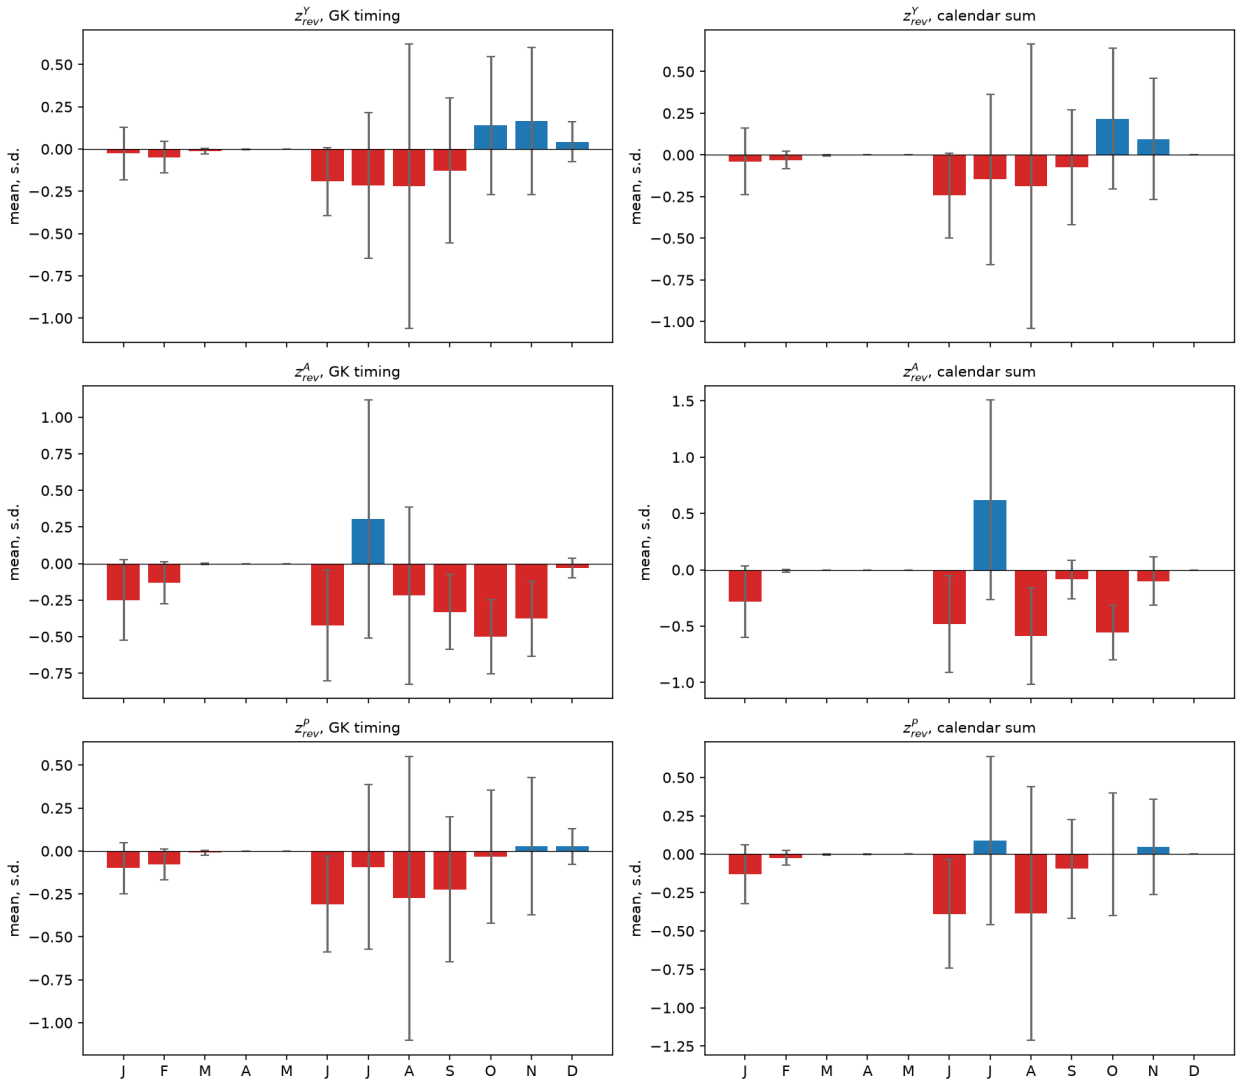


Figure C.2: Seasonal pattern of the revision instruments: month-of-year mean of r^Y (top), r^A (middle) and r^P (bottom), under GK timing (left) and the calendar sum (right), 1993M01–2024M12, in standard deviations. *Notes:* Each bar is $\bar{z}_m = \frac{1}{n_m}\sum_y z_{m,y}$, with whiskers $\pm 2\, s_m/\sqrt{n_m}$, where s_m is the standard deviation across years and $n_m = 32$ the number of years. This is the classical standard error of a mean (years independent; no heteroskedasticity or serial-correlation correction, no bootstrap). Each instrument is standardized over the same sample, the estimation sample of the first stage. Blue means more supply on average (USDA raised its estimate), red less. USDA makes almost no new-crop yield, area or production revisions in the February–May and December releases, so the instruments are (nearly) zero there by construction. Means with $|t| > 2$: r^A in June, September, October and November under GK (June, August and October under the sum), and r^P in June under both timings; none for r^Y. Source: USDA WASDE.

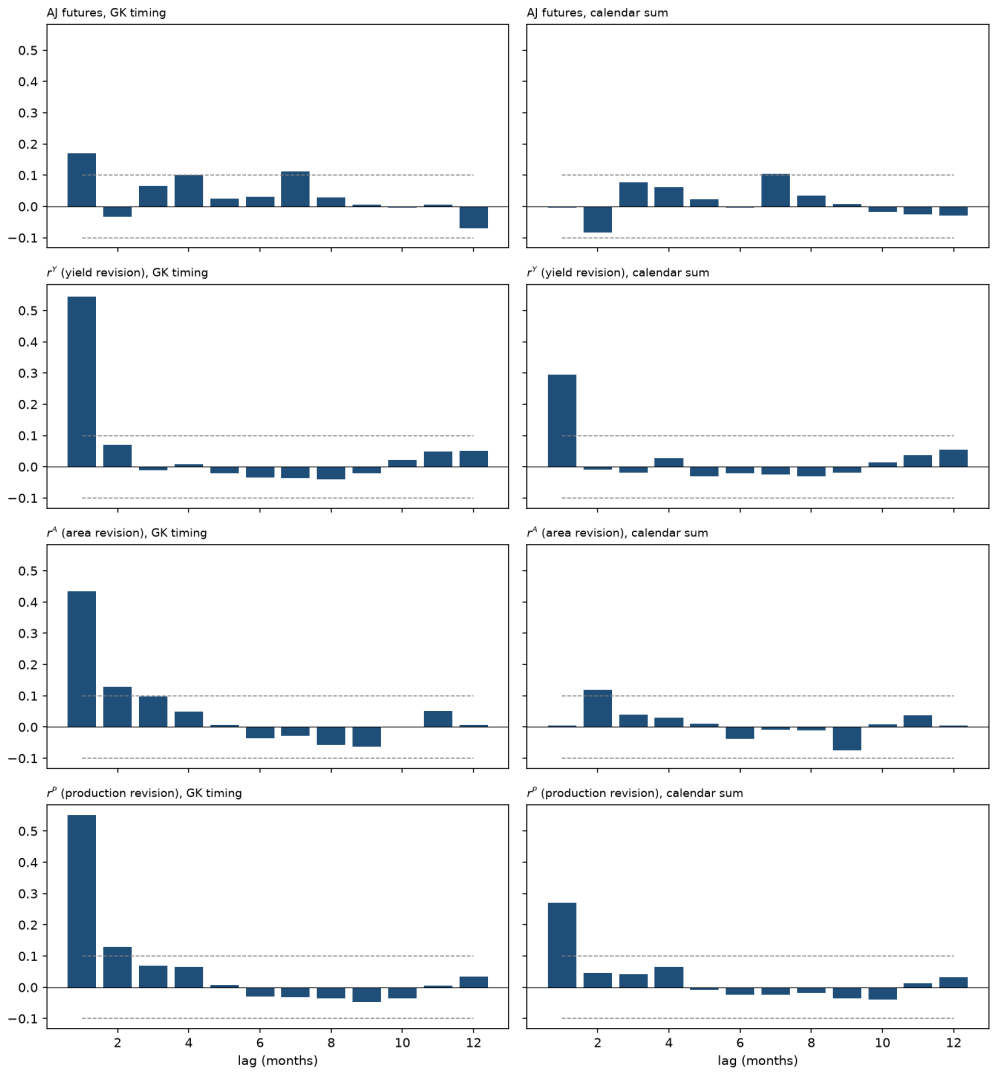


Figure C.3: Autocorrelation of the instruments: the Adjemian–Jo futures surprise (top) and the five-crop WASDE revisions in yield r^Y, harvested area r^A and production r^P, GK timing (left) and calendar-month sum (right), lags 1–12, 1993M01–2024M12. *Notes:* Sample autocorrelations over the 384 months of the first-stage sample; the dashed lines are the white-noise 95 percent band $\pm 1.96/\sqrt{T} = \pm 0.10$. GK timing splits each release between its month and the next, so the GK series has positive first-order autocorrelation by construction; the calendar sum does not split releases, so autocorrelation left in it comes from the news itself. At lag 1 (GK, sum): futures surprise 0.17, 0.00; r^Y 0.54, 0.29; r^A 0.43, 0.01; r^P 0.55, 0.27. USDA revises yield, and hence production, in the same direction in consecutive releases of a season, so r^Y and r^P are serially correlated even without the GK split. Source: USDA WASDE.

Table C.1: Commodity futures responses to WASDE revisions

	Dependent variable: release-day PC1 futures surprise (%)				
	(1) Main	(2) No area	(3) Yield only	(4) Production	(5) Ending stocks
Yield revision r^Y	-0.250^{***}	-0.247^{***}	-0.168^{***}		
	(0.043)	(0.043)	(0.043)		
Area revision r^A	-0.044				
	(0.027)				
Production revision r^P				-0.127^{***}	
				(0.028)	
Rest revision r^R	-0.192^{***}	-0.173^{***}		-0.167^{***}	
	(0.037)	(0.036)		(0.036)	
Ending-stocks revision r^{ES}					-0.239^{***}
					(0.034)
Observations	1,726	1,726	1,726	1,726	1,726
Releases	350	350	350	350	350
R^2	0.046	0.044	0.024	0.028	0.049
Crop fixed effects	Yes	Yes	Yes	Yes	Yes

Notes: The table repeats Table B.3 with USDA's own month-on-month WASDE revision in place of the Reuters poll surprise. The dependent variable is the release-day PC1 futures surprise from the five nearest maturities of each crop, in percent. The yield, area and production revisions are $r^X_{i,t} = 100 \ln(X_t / X_{t-1})$ for the newest marketing year, with last month's value as reprinted in the same report, in percent. The rest revision is the change in ending stocks minus the change in production, and the ending-stocks revision the change in ending stocks, both as a percent of last month's total use. Columns follow Table B.3: column (1) includes yield, area and rest; column (2) drops area; column (3) includes yield alone; column (4) includes production and rest; column (5) includes the ending-stocks revision alone. Every column pools all five crops (corn, soybeans, wheat, rice and oats) and includes crop fixed effects. The sample is the WASDE releases from January 12, 1993, through December 10, 2024 (the first-stage window), with correction vintages excluded; May releases carry no revision for the new marketing year and drop out. The 1,726 observations are 350 releases each for corn, soybeans and wheat, 343 for rice and 333 for oats. Standard errors clustered by release date are in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.2: Monthly futures surprise on the orthogonalized and predicted parts of the WASDE revisions

	Calendar sum			GK timing		
	(1) Main	(2) Yield only	(3) Production	(4) Main	(5) Yield only	(6) Production
Yield, orthogonalized $\tilde{r}^Y$	-0.232^{**}	-0.227^{**}		-0.228^{*}	-0.238^{*}	
	(0.105)	(0.107)		(0.123)	(0.122)	
Yield, predicted $r^Y - \tilde{r}^Y$	0.028	-0.020		-0.075	-0.090	
	(0.238)	(0.234)		(0.122)	(0.119)	
Area, orthogonalized $\tilde{r}^A$	0.385^{**}			-0.196		
	(0.183)			(0.349)		
Area, predicted $r^A - \tilde{r}^A$	-0.993			-0.270		
	(0.810)			(0.429)		
Production, orthogonalized $\tilde{r}^P$			-0.122			-0.225^{**}
			(0.086)			(0.111)
Production, predicted $r^P - \tilde{r}^P$			-0.120			-0.093
			(0.211)			(0.102)
Months	384	384	384	384	384	384
R^2	0.031	0.017	0.007	0.016	0.014	0.016

Notes: Each column estimates $s_t = \alpha + \beta_Y\,\tilde{r}^Y_t + \gamma_Y\,(r^Y_t - \tilde{r}^Y_t) + \beta_A\,\tilde{r}^A_t + \gamma_A\,(r^A_t - \tilde{r}^A_t) + \beta_P\,\tilde{r}^P_t + \gamma_P\,(r^P_t - \tilde{r}^P_t) + \varepsilon_t$ over 1993M01–2024M12 (384 months), keeping only the regressors with an entry: (1) and (4) yield and area; (2) and (5) yield; (3) and (6) production. $s_t = \sum_i w_{i,y}\,s_{i,t}$ is the monthly production-weighted futures surprise of equation (9), in percent, where $s_{i,t}$ is the release-day PC1 futures surprise of commodity i; it is aggregated to the month as a calendar sum (columns 1–3) or with GK timing (columns 4–6). $r^X_t = \sum_i w_{i,y}\,r^X_{i,t}$ is the production-weighted revision in item X = yield (Y), area (A) or production (P), with $r^X_{i,t} = 100\ln(X_t/X_{t-1})$ for commodity i, aggregated with the same timing as s_t. $\tilde{r}^X_t$ is the residual of $r^X_t = \alpha^X + \sum_{k=1}^{12}\gamma^X_k\,r^X_{t-k} + \sum_{k=1}^{12}\delta^X_k\,p_{t-k} + \tilde{r}^X_t$, where p_t is 100 times the log of the seasonally adjusted real field crop price, estimated on the same 384 months; $\tilde{r}^X_t$ and the predicted part $r^X_t - \tilde{r}^X_t$ are kept in percent, so the slopes compare with Table 2. Production is not included beside yield and area, since $\ln P = \ln A + \ln Y$. Heteroskedasticity-robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.3: Tests on instrument strength

	AJ futures		Poll, GK		Revision, GK		Revision, sum	
	(1) Sum	(2) GK	(3) z^Y	(4) z^P	(5) r^Y	(6) r^P	(7) r^Y	(8) r^P
A. 1993M01–2024M12 (384 months)								
Seasonally adjusted (SA)								
Coefficient	0.333	0.514			−0.473	−0.441	−0.584	−0.564
F-statistic	6.46	15.72			13.23	11.43	20.56	19.11
F-statistic (robust)	5.17	14.51			13.22	10.41	28.59	24.37
F-statistic (Newey–West)	5.01	13.24			23.74	18.35	33.88	28.93
R^2	1.66	3.95			3.35	2.91	5.11	4.77
R^2 (adjusted)	1.40	3.70			3.09	2.65	4.86	4.52
Observations	384	384			384	384	384	384
Unadjusted with month dummies (NSA + D)								
Coefficient	0.376	0.607			−0.436	−0.412	−0.555	−0.524
F-statistic	8.70	23.50			11.82	10.48	19.46	17.28
F-statistic (robust)	7.85	21.02			14.21	12.06	26.93	22.57
F-statistic (Newey–West)	8.50	17.60			21.64	17.70	36.31	28.85
R^2	2.23	5.80			3.00	2.67	4.85	4.33
R^2 (adjusted)	1.97	5.55			2.75	2.42	4.60	4.08
Observations	384	384			384	384	384	384
B. Poll months, 2012M09–2024M12 (148 months)								
Seasonally adjusted (SA)								
Coefficient	0.057	0.341	0.058	0.106	−0.099	−0.069	−0.435	−0.357
F-statistic	0.09	3.41	0.10	0.32	0.28	0.14	5.64	3.74
F-statistic (robust)	0.09	3.53	0.10	0.30	0.18	0.09	8.47	4.51
F-statistic (Newey–West)	0.10	3.71	0.12	0.32	0.21	0.09	7.71	3.59
R^2	0.06	2.28	0.07	0.22	0.19	0.09	3.72	2.50
R^2 (adjusted)	−0.62	1.61	−0.62	−0.46	−0.49	−0.59	3.06	1.83
Observations	148	148	148	148	148	148	148	148
Unadjusted with month dummies (NSA + D)								
Coefficient	0.240	0.321	0.002	0.001	−0.225	−0.168	−0.523	−0.398
F-statistic	1.84	3.34	0.00	0.00	1.62	0.89	9.19	5.17
F-statistic (robust)	1.67	3.80	0.00	0.00	1.65	0.91	16.42	7.58
F-statistic (Newey–West)	1.78	3.83	0.00	0.00	2.01	1.05	14.11	5.77
R^2	1.24	2.24	0.00	0.00	1.10	0.61	5.92	3.42
R^2 (adjusted)	0.57	1.57	−0.68	−0.68	0.42	−0.07	5.27	2.76
Observations	148	148	148	148	148	148	148	148

Notes: First-stage regressions $u_t = a + b\,z_t + e_t$ of the real field crop price residual u_t on each instrument z_t. The residual comes from the crop price equation of the fourteen-variable VAR with twelve lags, always estimated on 1992M01–2024M12 (residuals for 1993M01–2024M12). Within each panel, the VAR uses the X-13 seasonally adjusted series with a constant (SA), or the unadjusted series with a constant and February–December dummies (NSA + D). Columns (1)–(2) are the Adjemian–Jo futures surprise as a calendar-month sum and in Gertler–Karadi (GK) timing; columns (3)–(4) the Reuters poll surprises in yield z^Y (corn, soybeans) and production z^P (corn, soybeans, wheat), equation (13)–(14), GK timing; columns (5)–(6) the five-crop WASDE revisions in yield r^Y and production r^P, GK timing; and columns (7)–(8) the same revisions as calendar-month sums, each release counted in full in the month it is published. Panel A uses 1993M01–2024M12 (384 months); the poll surprises start in 2012M09, so columns (3)–(4) are blank there. Panel B uses the full poll sample, 2012M09–2024M12 (148 months), for all eight columns. Instruments are scaled to unit standard deviation over the regression months, so b is in percent per standard deviation; the futures surprise is a price move (expected $b > 0$), and the poll and revision series are positive when supply is larger (expected $b < 0$). $F = (b/\mathrm{se})^2$; the robust F uses heteroskedasticity-robust (HC1) standard errors and equals the Montiel Olea–Pflueger effective F with one instrument; the Newey–West F uses $\lfloor 0.75\,T^{1/3} \rfloor$ lags (5 for 384 months, 3 for 148). F-statistics above 10 indicate strong instruments. R^2 in percent.

Table C.4: Tests on instrument strength: orthogonalized production and area revisions

	Production $\tilde{r}^P$		Harvested area $\tilde{r}^A$	
	(1) Sum	(2) GK	(3) Sum	(4) GK
Seasonally adjusted (SA)				
Coefficient	−0.632	−0.638	−0.106	−0.116
F-statistic	24.24	24.76	0.64	0.77
F-statistic (robust)	32.17	34.97	0.56	0.70
F-statistic (Newey–West)	26.24	29.40	0.70	0.85
R^2	5.97	6.09	0.17	0.20
R^2 (adjusted)	5.72	5.84	−0.09	−0.06
Observations	384	384	384	384
Unadjusted with month dummies (NSA + D)				
Coefficient	−0.557	−0.556	−0.067	−0.078
F-statistic	19.60	19.53	0.27	0.37
F-statistic (robust)	30.67	30.31	0.22	0.31
F-statistic (Newey–West)	30.98	32.27	0.26	0.35
R^2	4.88	4.86	0.07	0.10
R^2 (adjusted)	4.63	4.61	−0.19	−0.17
Observations	384	384	384	384

Notes: Same regressions as Table 3, with the orthogonalized five-crop WASDE revisions in production $\tilde{r}^P$ (columns 1–2) and harvested area $\tilde{r}^A$ (columns 3–4), each as a calendar-month sum and in Gertler–Karadi (GK) timing, on 1993M01–2024M12 (384 months). Each orthogonalized series is the residual of the revision on twelve own lags and twelve lags of the seasonally adjusted real crop price, estimated once on 1993M01–2024M12; the same series enters the SA and the NSA + D VAR. Instruments are scaled to unit standard deviation over the regression months, so b is in percent per standard deviation; the futures surprise is a price move (expected $b > 0$), and the revisions are positive when supply is larger (expected $b < 0$). $F = (b/\text{se})^2$; the robust F uses HC1 standard errors and equals the Montiel Olea–Pflueger effective F with one instrument; the Newey–West F uses 5 lags. R^2 in percent. The orthogonalized series is a generated instrument; the standard errors do not account for the first step.

Table C.5: Granger causality tests of the instruments

	AJ futures	Revision, GK		
	GK	r^Y	r^A	r^P
Own lags	0.014	0.000	0.000	0.000
Real field crop price	0.003	0.005	0.040	0.002
Real GDP	0.021	0.850	0.673	0.930
Industrial production	0.803	0.959	0.788	0.975
Core CPI	0.077	0.824	0.382	0.859
S&P 500	0.497	0.563	0.639	0.628
VIX	0.774	0.713	0.819	0.636
Grain exports	0.554	0.806	0.092	0.580
Baltic Dry Index	0.022	0.542	0.017	0.500
Livestock PPI	0.475	0.833	0.761	0.908
Real ethanol price	0.054	0.135	0.114	0.100
Real oil price	0.545	0.344	0.042	0.342
World oil production	0.492	0.375	0.068	0.140
Food-at-home consumption	0.287	0.966	0.565	0.925
Food-at-home price	0.642	0.730	0.001	0.357
All 14 variables	0.020	0.508	0.000	0.158

Notes: The table follows Table C.1 of Adjemian and Jo (2025). For each instrument z_t, one regression $z_t = c + \sum_{k=1}^{12} a_k z_{t-k} + \sum_{j=1}^{14} \sum_{k=1}^{12} b_{jk}\, y_{j,t-k} + e_t$, where y_j are the fourteen variables of the VAR (unadjusted series) and the twelve lags match the VAR's lag length: 181 coefficients on 384 months, 1993M01–2024M12 (lags reach back into 1992). Each entry is the p-value of the OLS F test that the twelve lags of the row's variable are jointly zero; "Own lags" tests $a_1 = \cdots = a_{12} = 0$, and "All 14 variables" tests every $b_{jk} = 0$. A large p-value means past values do not predict the instrument. The first column is the Adjemian–Jo futures surprise, and the other three the five-crop WASDE revisions in yield r^Y, harvested area r^A and production r^P, all in GK timing.

Table C.6: Correlation of the instruments with published macroeconomic shocks

	AJ futures	Revision, GK			
	GK	r^Y	r^A	r^P	Months
Oil supply news (Känzig)	0.058	−0.086	0.053	−0.061	384
	(0.256)	(0.091)	(0.302)	(0.230)	
Oil supply (Baumeister–Hamilton)	−0.055	0.035	−0.043	0.018	384
	(0.285)	(0.496)	(0.399)	(0.720)	
Oil consumption demand (BH)	−0.035	−0.026	0.134	0.020	384
	(0.499)	(0.617)	(0.009)	(0.698)	
Oil inventory demand (BH)	0.043	0.027	−0.062	0.006	384
	(0.396)	(0.593)	(0.225)	(0.903)	
Economic activity (BH)	0.013	0.082	−0.073	0.050	384
	(0.798)	(0.110)	(0.155)	(0.325)	
Monetary policy (BRW)	0.042	0.013	−0.016	0.006	367
	(0.422)	(0.799)	(0.767)	(0.905)	
Monetary policy (Bauer–Swanson, orth.)	−0.011	0.020	0.000	0.018	372
	(0.826)	(0.704)	(0.999)	(0.724)	
Change in fed funds rate	0.018	0.035	0.039	0.044	384
	(0.728)	(0.495)	(0.447)	(0.387)	
Carbon policy	0.015	−0.049	−0.056	−0.059	246
	(0.821)	(0.448)	(0.386)	(0.357)	
Supply-chain pressure (GSCPI)	0.192	−0.010	−0.029	−0.018	324
	(0.001)	(0.863)	(0.602)	(0.746)	
Geopolitical risk (GPR)	−0.063	0.060	−0.007	0.051	384
	(0.218)	(0.242)	(0.885)	(0.322)	
Trade policy uncertainty	0.044	−0.061	−0.221	−0.126	384
	(0.390)	(0.233)	(0.000)	(0.013)	
Economic policy uncertainty	0.070	−0.114	−0.202	−0.168	384
	(0.171)	(0.025)	(0.000)	(0.001)	

Notes: The table follows Table C.2 of Adjemian and Jo (2025). Each entry is the Pearson correlation between the instrument and a published shock series, with the p-value of $H_0\colon \rho = 0$ in parentheses ($t = r\sqrt{(N-2)/(1-r^2)}$, $N-2$ degrees of freedom). Each correlation uses the months inside 1993M01–2024M12 that the shock series covers, given in the last column. Columns are as in Table C.5: the Adjemian–Jo futures surprise and the five-crop WASDE revisions r^Y, r^A and r^P, all in GK timing. BH stands for Baumeister–Hamilton, BRW for Bu–Rogers–Wu, GSCPI for the New York Fed Global Supply Chain Pressure Index and GPR for the Caldara–Iacoviello geopolitical risk index.

Table C.7: Granger causality tests of the orthogonalized production and area revisions

	Production $\tilde{r}^P$		Harvested area $\tilde{r}^A$	
	Sum	GK	Sum	GK
Own lags	0.838	0.904	0.378	0.359
Real field crop price	0.950	0.921	0.846	0.827
Real GDP	0.817	0.919	0.296	0.328
Industrial production	0.952	0.963	0.583	0.677
Core CPI	0.851	0.878	0.074	0.120
S&P 500	0.573	0.616	0.706	0.688
VIX	0.629	0.787	0.615	0.778
Grain exports	0.763	0.799	0.069	0.118
Baltic Dry Index	0.530	0.609	0.038	0.022
Livestock PPI	0.854	0.934	0.693	0.718
Real ethanol price	0.162	0.229	0.124	0.101
Real oil price	0.327	0.387	0.078	0.073
World oil production	0.325	0.317	0.273	0.242
Food-at-home consumption	0.937	0.914	0.442	0.464
Food-at-home price	0.361	0.271	0.006	0.006
All 14 variables	0.709	0.768	0.000	0.000

Notes: As Table 4: for each instrument z_t, $z_t = c + \sum_{k=1}^{12} a_k z_{t-k} + \sum_{j=1}^{14} \sum_{k=1}^{12} b_{jk}\, y_{j,t-k} + e_t$, with y_j the fourteen unadjusted VAR variables; each entry is the p-value of the OLS F test that the row's twelve lags are jointly zero. The sample is 1994M01–2024M12 (372 months) for every column, because the orthogonalized series start in 1993M01 and need twelve own lags. Columns: the orthogonalized production revision $\tilde{r}^P$ and harvested area revision $\tilde{r}^A$, each as a calendar-month sum and in GK timing. p-values below 0.05 are in red. The orthogonalized series are residuals on their own twelve lags and twelve lags of the seasonally adjusted crop price, so the "Own lags" row and the crop price row pass by construction; the other thirteen rows and the joint test are the informative ones.

Table C.8: Correlation of the orthogonalized production and area revisions with published macroeconomic shocks

	Production $\tilde{r}^P$		Harvested area $\tilde{r}^A$		
	Sum	GK	Sum	GK	Months
Oil supply news (Känzig)	0.005	−0.009	0.038	0.028	384
	(0.924)	(0.865)	(0.457)	(0.589)	
Oil supply (Baumeister–Hamilton)	−0.034	−0.021	−0.022	−0.020	384
	(0.509)	(0.681)	(0.671)	(0.694)	
Oil consumption demand (BH)	0.148	0.131	0.160	0.150	384
	(0.004)	(0.010)	(0.002)	(0.003)	
Oil inventory demand (BH)	−0.094	−0.085	−0.113	−0.105	384
	(0.067)	(0.096)	(0.027)	(0.039)	
Economic activity (BH)	0.037	0.040	−0.070	−0.063	384
	(0.465)	(0.429)	(0.170)	(0.219)	
Monetary policy (BRW)	0.027	0.024	0.034	0.029	367
	(0.600)	(0.652)	(0.512)	(0.580)	
Monetary policy (Bauer–Swanson, orth.)	0.014	0.022	−0.037	−0.033	372
	(0.785)	(0.679)	(0.478)	(0.530)	
Change in fed funds rate	0.022	0.021	0.069	0.064	384
	(0.667)	(0.676)	(0.177)	(0.212)	
Carbon policy	0.023	0.016	−0.095	−0.062	246
	(0.725)	(0.802)	(0.138)	(0.332)	
Supply-chain pressure (GSCPI)	0.019	0.021	−0.061	−0.055	324
	(0.731)	(0.707)	(0.275)	(0.321)	
Geopolitical risk (GPR)	0.035	0.037	−0.010	−0.008	384
	(0.489)	(0.465)	(0.851)	(0.876)	
Trade policy uncertainty	−0.105	−0.105	−0.163	−0.163	384
	(0.039)	(0.040)	(0.001)	(0.001)	
Economic policy uncertainty	−0.118	−0.117	−0.168	−0.166	384
	(0.021)	(0.022)	(0.001)	(0.001)	

Notes: As Table 5: Pearson correlation between the instrument and a published shock series, with the p-value of $H_0 : \rho = 0$ in parentheses, on the months inside 1993M01–2024M12 that the shock covers (last column). Columns as in Table C.7. A correlation with $p < 0.05$ is in green if positive and in red if negative. BH stands for Baumeister–Hamilton, BRW for Bu–Rogers–Wu, GSCPI for the New York Fed Global Supply Chain Pressure Index and GPR for the Caldara–Iacoviello geopolitical risk index.

D Additional Impulse Responses

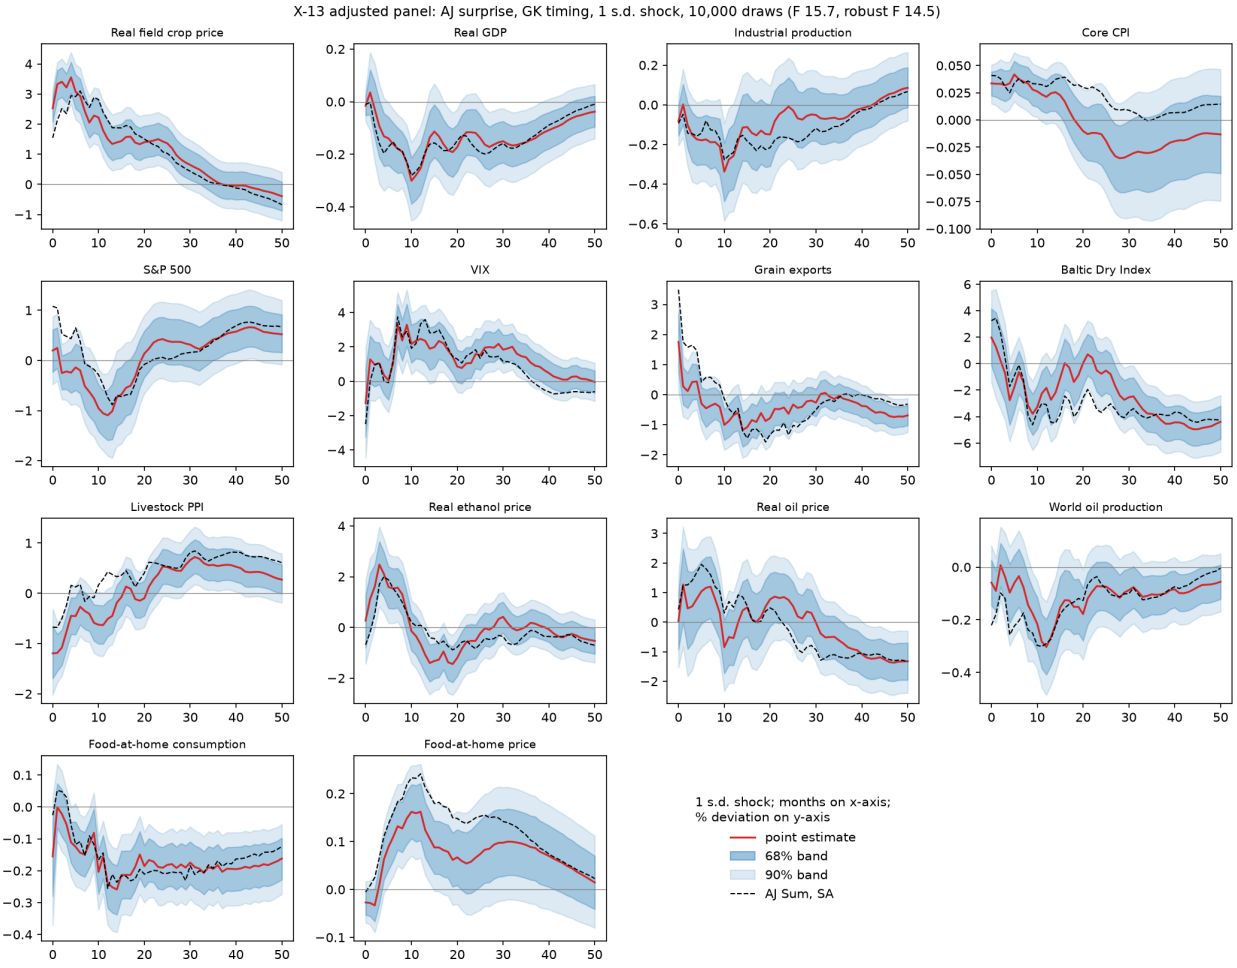


Figure D.1: Impulse responses to a one standard deviation agricultural supply news shock identified with the Adjemian–Jo futures surprise in Gertler–Karadi timing, seasonally adjusted panel

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. X-13 seasonally adjusted series and a constant, as Figure 7. The instrument is the Adjemian–Jo futures surprise of Figure 7 with Gertler–Karadi timing: a release on day d of a D-day month enters that month with weight $(D - d + 1)/D$ and the next with $(d - 1)/D$. First-stage F 15.7, robust F 14.5 (Figure 7: 6.5 and 5.2). The dashed black line is the point estimate of Figure 7 (calendar-month sum). The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage. Horizons are in months.

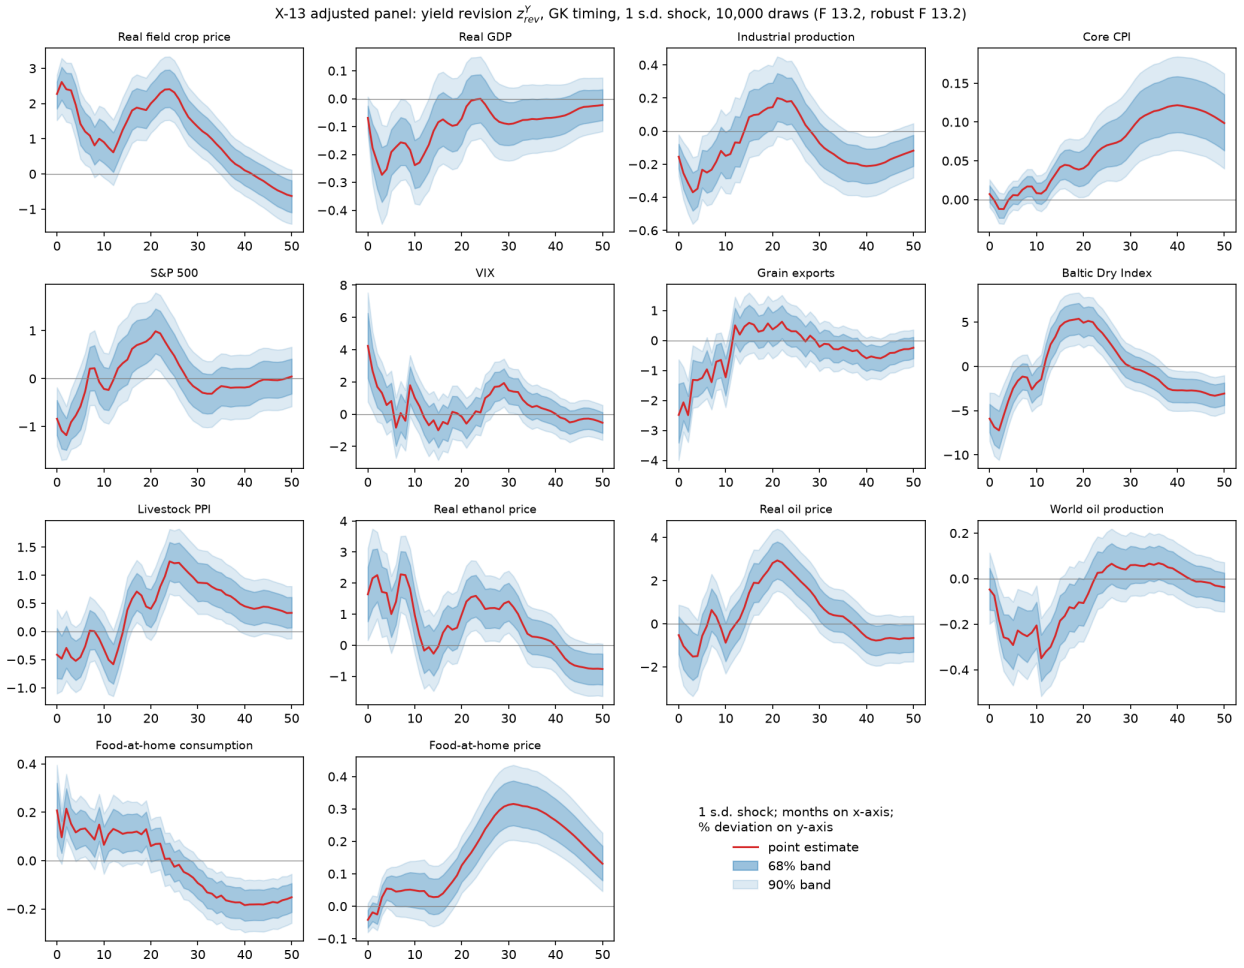


Figure D.2: Impulse responses to a one standard deviation agricultural supply news shock identified with the WASDE yield revision r^Y in Gertler–Karadi timing, seasonally adjusted panel

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. X-13 seasonally adjusted series and a constant, as Figure 7. The instrument is the five-crop yield revision r^Y of Table C.3 (USDA's month-on-month revision of the newest-marketing-year yield, production-weighted over corn, soybeans, wheat, rice and oats); it is positive when USDA sees more supply, and the shock is signed so that the crop price rises, so it is an adverse yield revision. It enters with Gertler–Karadi timing: a release on day d of a D-day month enters that month with weight $(D - d + 1)/D$ and the next with $(d - 1)/D$. First-stage F 13.2, robust F 13.2, Newey–West F 23.7. The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage. Horizons are in months.

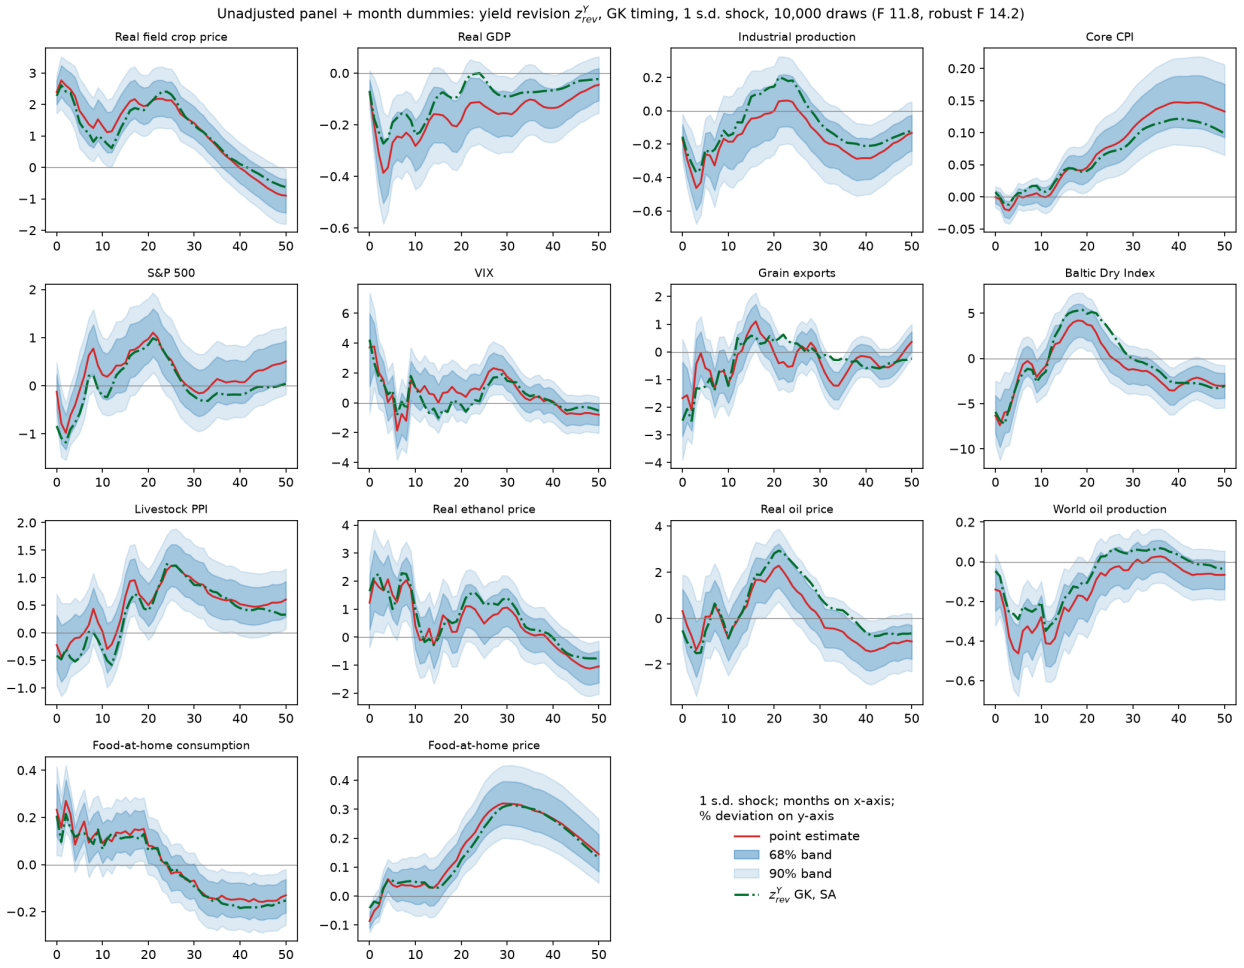


Figure D.3: Robustness: impulse responses to the yield-revision shock (r^Y, Gertler–Karadi timing) on the unadjusted panel with month dummies

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. Unadjusted series with a constant and February–December dummies in every equation (and in every bootstrap draw). The instrument is the five-crop yield revision r^Y of Table C.3 (USDA's month-on-month revision of the newest-marketing-year yield, production-weighted over corn, soybeans, wheat, rice and oats); it is positive when USDA sees more supply, and the shock is signed so that the crop price rises, so it is an adverse yield revision. It enters with Gertler–Karadi timing, as in Figure D.2. First-stage F 11.8, robust F 14.2, Newey–West F 21.6. The dash-dotted dark green line is the point estimate of Figure D.2 (r^Y, Gertler–Karadi timing, seasonally adjusted panel). The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage. Horizons are in months.

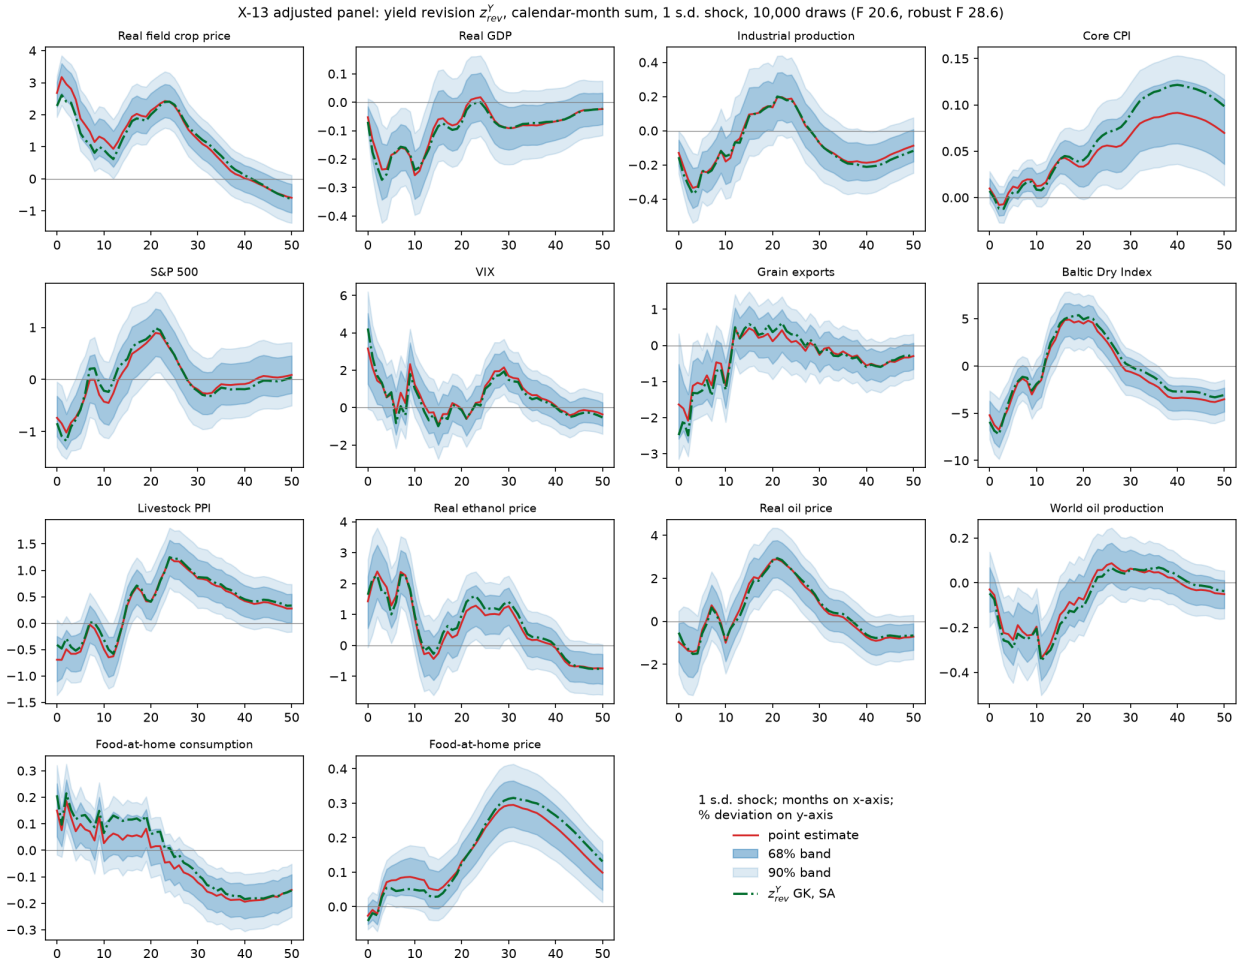


Figure D.4: Robustness: impulse responses to the yield-revision shock with the calendar-month sum (r^Y) instead of Gertler–Karadi timing, seasonally adjusted panel

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. X-13 seasonally adjusted series and a constant, as Figure 7. The instrument is the five-crop yield revision r^Y of Table C.3 (USDA's month-on-month revision of the newest-marketing-year yield, production-weighted over corn, soybeans, wheat, rice and oats); it is positive when USDA sees more supply, and the shock is signed so that the crop price rises, so it is an adverse yield revision. Each release enters the calendar month in which it is published, in full (the timing of Figure 7), so a month without a revision is zero (226 of 384 months are non-zero). First-stage F 20.6, robust F 28.6, Newey–West F 33.9. The dash-dotted dark green line is the point estimate of Figure D.2 (r^Y, Gertler–Karadi timing). The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage. Horizons are in months.

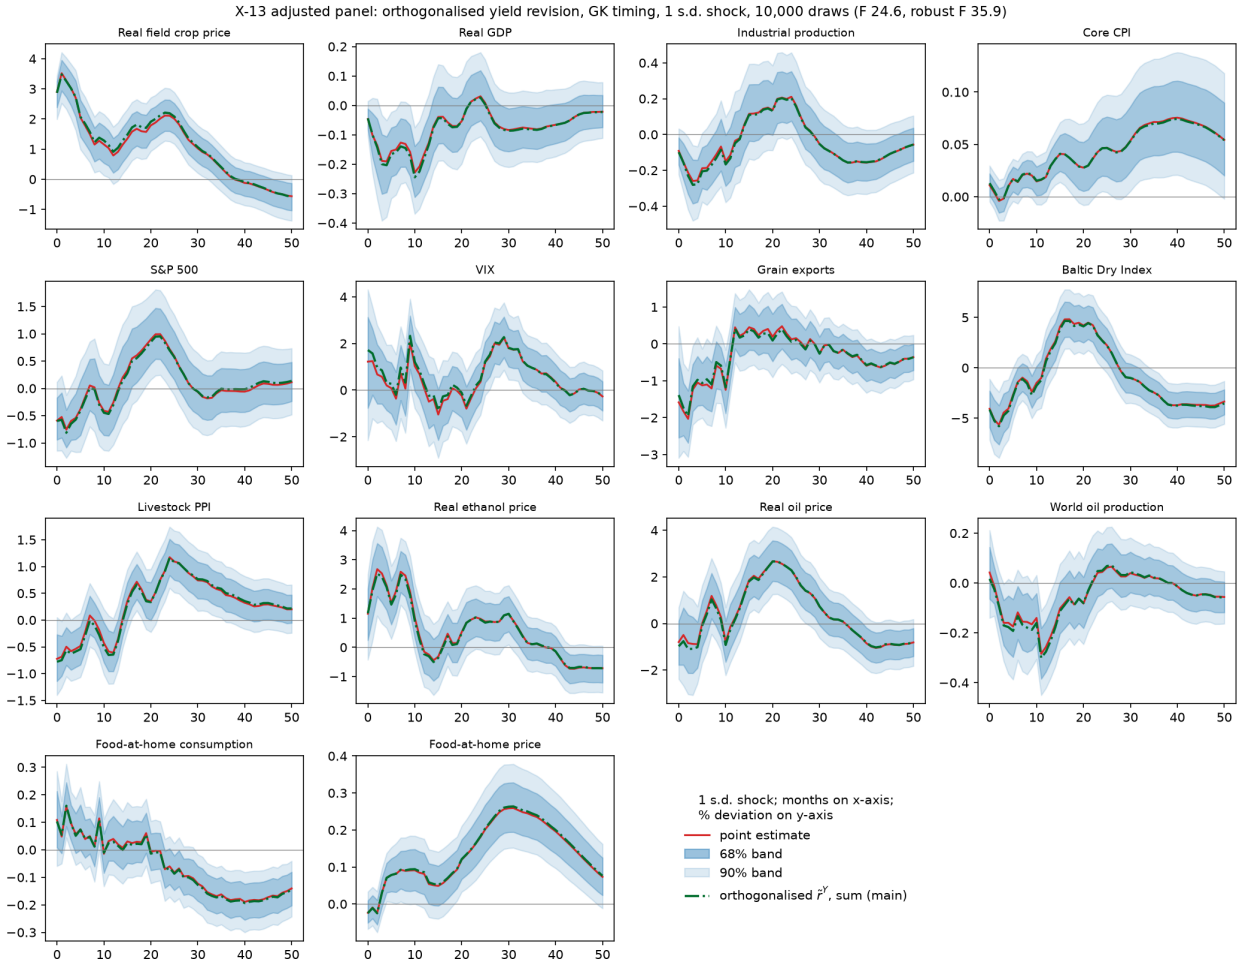


Figure D.5: Impulse responses to the orthogonalized yield-revision shock ($\tilde{r}^Y$, Gertler–Karadi timing), seasonally adjusted panel

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. X-13 seasonally adjusted series and a constant, as Figure 7. Revisions enter in Gertler–Karadi timing, as Figure D.2. The instrument is the orthogonalized five-crop yield revision: the residual w_t of $z^Y_{rev,t} = c + \sum_{k=1}^{12} g_k z^Y_{rev,t-k} + \sum_{k=1}^{12} h_k p_{t-k} + w_t$, with p the seasonally adjusted real field crop price, estimated on 1993M01–2024M12 and scaled to unit standard deviation. The projection removes the part of the revision that the market had already seen: the raw revision is correlated with the crop price residual of the previous month (robust t of -6.2 for Gertler–Karadi timing), the orthogonalized one is not ($t = -0.58$). The shock is signed so that the crop price rises, so it is an adverse yield revision. First-stage F 24.6, robust F 35.9, Newey–West F 27.4. The dash-dotted dark green line is the point estimate of Figure 8 ($\tilde{r}^Y$, calendar-month sum, the main specification). The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage but not the projection. Horizons are in months.

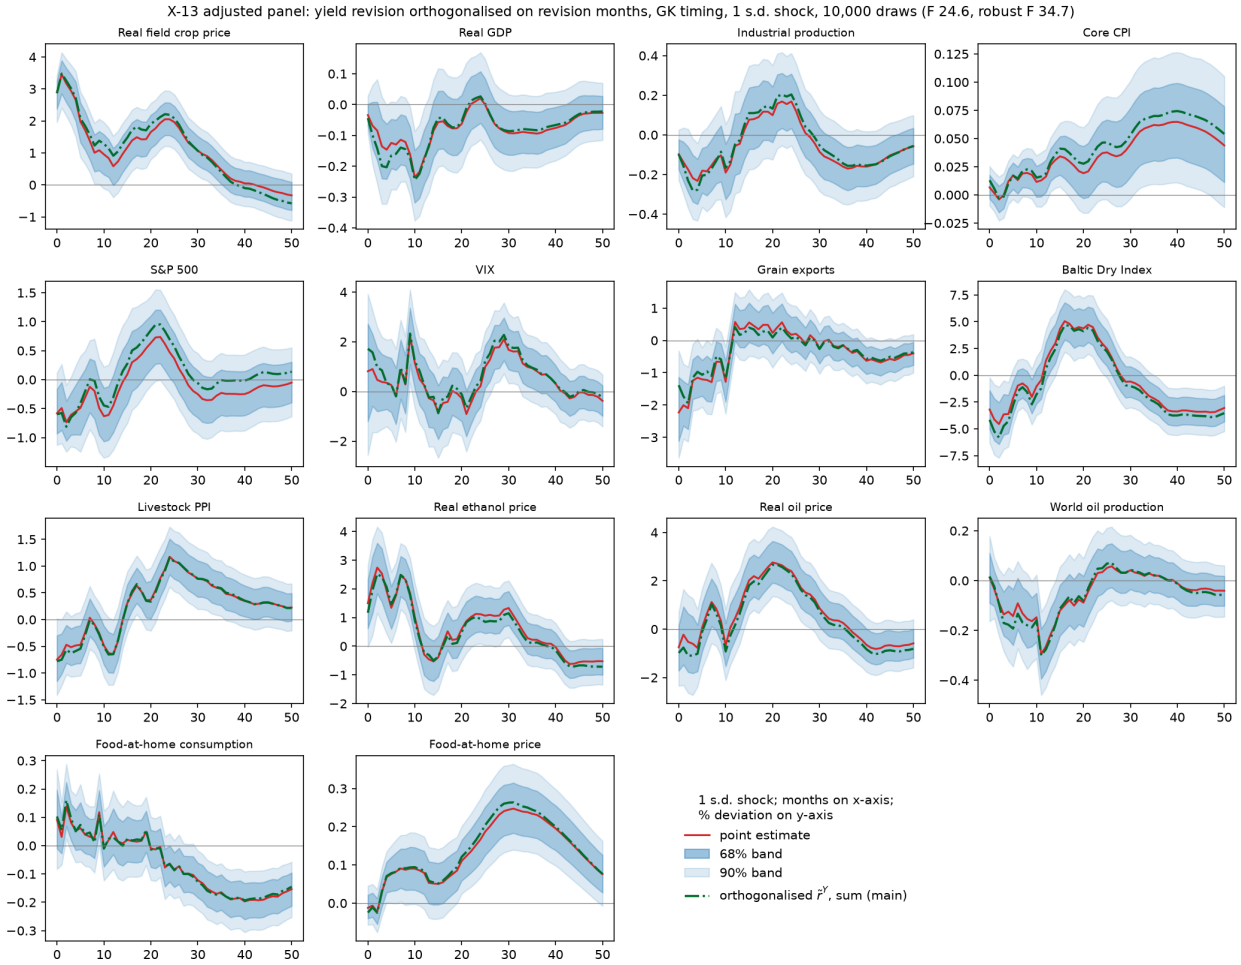


Figure D.6: Robustness: impulse responses to the yield-revision shock orthogonalized on revision months ($\tilde{r}^Y$, Gertler–Karadi timing), seasonally adjusted panel

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. X-13 seasonally adjusted series and a constant, as Figure 7. The instrument is the five-crop yield revision orthogonalized on revision months only: in months with a revision ($z^Y_{rev,t} \neq 0$, 304 of 384) it is the residual w_t of $z^Y_{rev,t} = c + \sum_{k=1}^{12} g_k z^Y_{rev,t-k} + \sum_{k=1}^{12} h_k p_{t-k} + w_t$, estimated on those months, with p the seasonally adjusted real field crop price; in months without a revision it is zero, because the release calendar makes the expected revision zero there. The series is scaled to unit standard deviation. Revisions enter in Gertler–Karadi timing, as Figure D.2. Its correlation with the previous month's crop price residual has a robust t of -0.39. First-stage F 24.6, robust F 34.7, Newey–West F 25.2. The dash-dotted dark green line is the point estimate of Figure 8 (orthogonalized on all months, calendar-month sum, the main specification). The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage but not the projection. Horizons are in months.

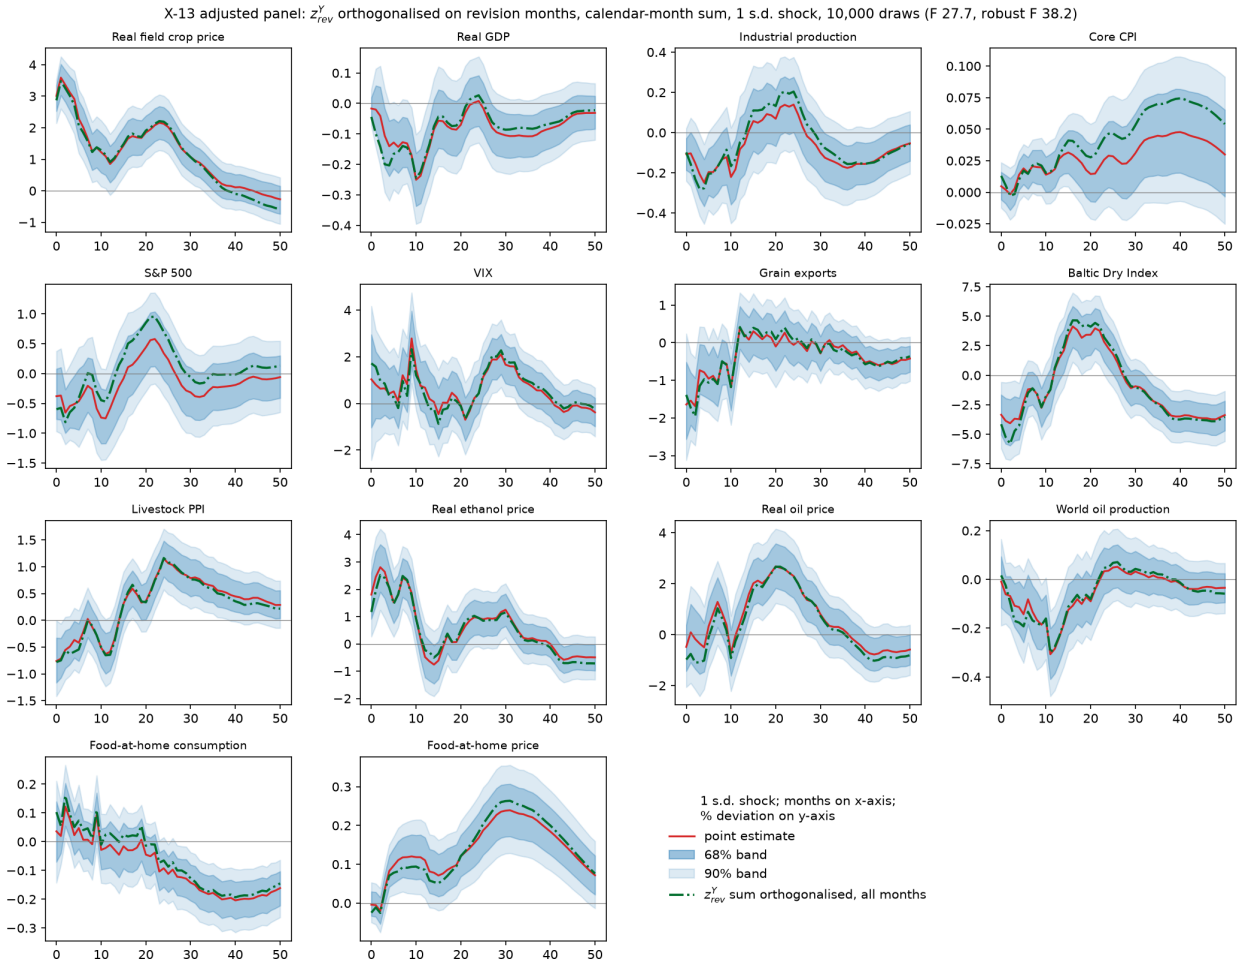


Figure D.7: Robustness: impulse responses to the yield-revision shock orthogonalized on revision months ($\tilde{r}^Y$, calendar-month sum), seasonally adjusted panel

Notes: Proxy-SVAR identical to Figure 7: twelve lags, the same fourteen variables in 100 times logs, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), one standard deviation shock signed to raise the real field crop price on impact. X-13 seasonally adjusted series and a constant, as Figure 7. The instrument is the five-crop yield revision orthogonalized on revision months only: in months with a revision ($z^Y_{rev,t} \neq 0$, 226 of 384) it is the residual w_t of $z^Y_{rev,t} = c + \sum_{k=1}^{12} g_k z^Y_{rev,t-k} + \sum_{k=1}^{12} h_k p_{t-k} + w_t$, estimated on those months, with p the seasonally adjusted real field crop price; in months without a revision it is zero, because the release calendar makes the expected revision zero there. The series is scaled to unit standard deviation. Each release enters the calendar month in which it is published, in full, as Figure D.4. Its correlation with the previous month's crop price residual has a robust t of -0.44. First-stage F 27.7, robust F 38.2, Newey–West F 27.4. The dash-dotted dark green line is the point estimate of Figure 8 (orthogonalized on all months). The solid red line is the point estimate, the dark blue region the 68 percent band and the light blue region the 90 percent band, from 10,000 moving-block bootstrap draws that re-estimate the first stage but not the projection. Horizons are in months.

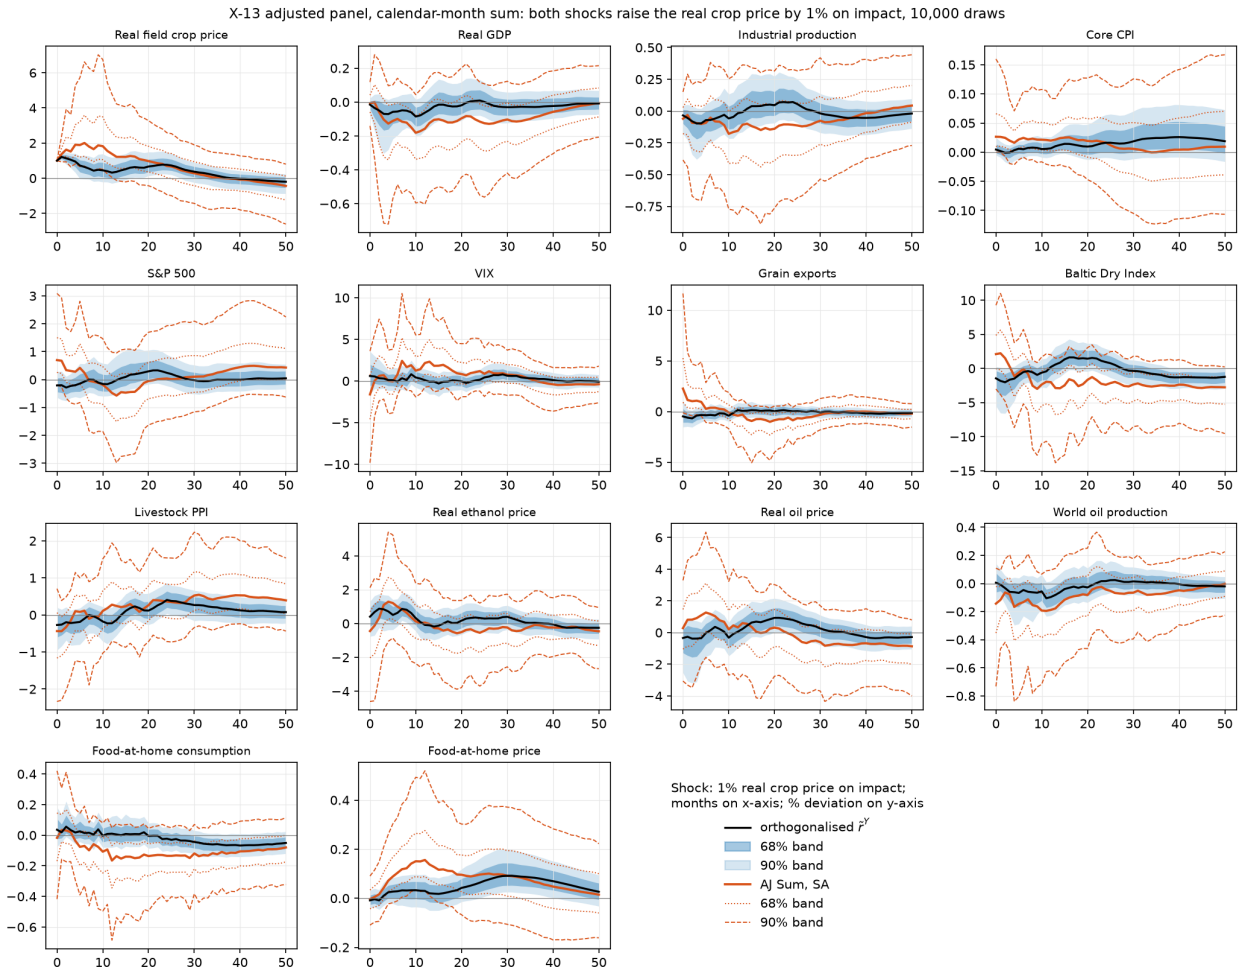


Figure D.8: Impulse responses to the orthogonalized yield-revision shock ($\tilde{r}_t^Y$) and to the Adjemian–Jo shock, both normalized to raise the real field crop price by 1 percent on impact in every bootstrap draw, calendar-month sum, seasonally adjusted panel

Notes: The two proxy-SVARs of Figures 8 ($\tilde{r}_t^Y$) and 7 (Adjemian–Jo sum): twelve lags, the same fourteen variables in 100 times logs, X-13 seasonally adjusted series and a constant, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months), each release in the calendar month in which it is published. A one standard deviation shock is one standard deviation of the structural shock each instrument identifies, so its size differs between the two. Here each shock is instead scaled so that the real field crop price rises by 1 percent on impact, as in Känzig (2021): every response is divided by the impact response of the crop price, $\mathrm{IRF}_j(h)/\mathrm{IRF}_1(0)$, in the point estimate and in each bootstrap draw, so the bands are bands of this ratio. Figure 9 instead divides the 1 standard deviation point estimates and bands by one fixed number, the point estimate of the impact; the lines are the same, the bands differ. The black line and the blue regions (dark 68 percent, light 90 percent) are $\tilde{r}_t^Y$ (first-stage F 24.6, robust F 34.4). The red line is the Adjemian–Jo shock, with its 68 percent band dotted and its 90 percent band dashed (first-stage F 6.5, robust F 5.2); its bands are wide because a weak first stage puts the crop price impact near zero in some draws, which inflates the ratio. These bands overstate the uncertainty: inside the bootstrap, re-estimating the VAR roughly halves every instrument's first-stage F (median across draws 3.0 for Adjemian–Jo and 12.6 for $\tilde{r}_t^Y$). Bands from 10,000 moving-block bootstrap draws; same seed as Figures 8 and 7, so the point estimates are those figures' point estimates divided by their crop price impact (2.87 and 1.55 percent). Horizons are in months.

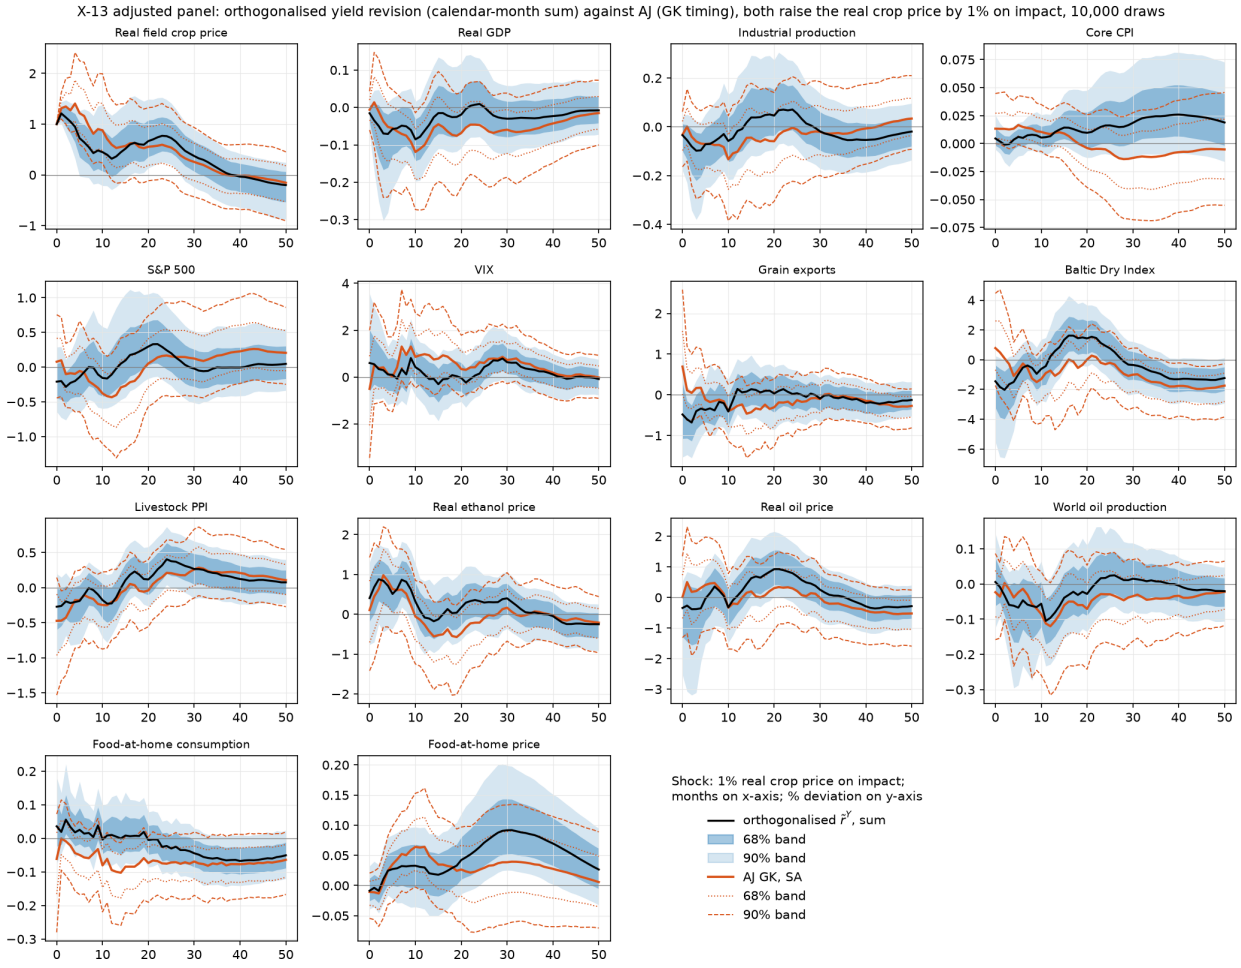


Figure D.9: Impulse responses to the orthogonalized yield-revision shock ($\tilde{r}_t^Y$) (calendar-month sum) and to the Adjemian–Jo shock in Gertler–Karadi timing, both normalized to raise the real field crop price by 1 percent on impact, seasonally adjusted panel

Notes: As Figure D.8, with the Adjemian–Jo surprise in Gertler–Karadi timing (the instrument of Figure D.1) in place of the calendar-month sum, because its first stage is stronger; $\tilde{r}_t^Y$ is unchanged (calendar-month sum, Figure 8). Twelve lags, the same fourteen variables in 100 times logs, X-13 seasonally adjusted series and a constant, monthly data 1992M01–2024M12, identification sample 1993M01–2024M12 (384 months). A one standard deviation shock is one standard deviation of the structural shock each instrument identifies, so its size differs between the two. Here each shock is instead scaled so that the real field crop price rises by 1 percent on impact, as in Känzig (2021): every response is divided by the impact response of the crop price, $\mathrm{IRF}_j(h)/\mathrm{IRF}_1(0)$, in the point estimate and in each bootstrap draw, so the bands are bands of this ratio. The black line and the blue regions (dark 68 percent, light 90 percent) are $\tilde{r}_t^Y$ (first-stage F 24.6, robust F 34.4). The red line is the Adjemian–Jo shock in Gertler–Karadi timing, with its 68 percent band dotted and its 90 percent band dashed (first-stage F 15.7, robust F 14.5); with the stronger first stage its bands are much narrower than in Figure D.8. Bands from 10,000 moving-block bootstrap draws; same seed as Figures 8 and D.1, so the point estimates are those figures' point estimates divided by their crop price impact (2.87 and 2.53 percent). Horizons are in months.

E LLM Text Extraction

The U.S. Department of Agriculture's (USDA) monthly World Agricultural Supply and Demand Estimates (WASDE) releases contain both numerical data, such as demand factors (domestic use for food consumption, seed use, and exports) and supply-side factors (harvested area and yield per harvested acre), and textual data containing a short outlook for major agricultural commodities such as wheat, coarse grains, rice, oilseeds, sugar, cotton, livestock, poultry, and dairy. In addition to using the balance-sheet data on commodity demand and supply-side factors, I use a large language model to analyze the textual release component.

In my analysis, I use the LLM as a structured output generator, where I input a commodity section from one WASDE release, and the LLM generates a structured output broken down by sentence (the prompt is summarized in Figure E.4). For each sentence, the LLM identifies the commodity, channel, direction, causal/noncausal, driver and policy actor. For example, in Figure E.1, the LLM helped to identify that in April 2020, the lower domestic use of wheat was caused by a remeasurement of feed and residual use due to the measurement issue reported in the NASS Grain Stocks report. It also highlights that lower wheat exports are due to competition in the international market.

In my analysis, I plan to use the structured output to create indicator variables that help to identify whether the production revision of the commodity is driven by more exogenous reasons, such as weather, pests and disease, or measurement, or by endogenous decisions, such as policy or external demand.

Figure E.3 shows how the stated drivers map into the balance sheet. Weather labels attach almost entirely to production, measurement labels attach to production, use and stocks, and the endogenous drivers attach mostly to trade and use.

It is very likely that not all the LLM-described reasons related to weather, pests and disease, or measurement are exogenous: as this information is available prior to the WASDE release, the market may have already priced it in. Nevertheless, the LLM provides additional information to help determine whether the cause is exogenous or not.

Figure E.1: April 2020 WASDE, wheat: the commentary labelled sentence by sentence. *Note: the text reproduces the USDA wording of 9 April 2020 (omitted text marked . . .). Each bracket gives the balance-sheet item and direction, then the driver the model reads in the sentence. Teal brackets are exogenous drivers (weather, pest and disease, measurement); violet brackets are endogenous or behavioural drivers (policy, external demand, market competitiveness). Sentences that only restate a number carry no bracket. Labels are from DeepSeek V4.1 Flash, prompt v3.5.*

WASDE release: 9 April 2020 **Commodity topic:** wheat, marketing year 2019/20

U.S. wheat. The outlook is for lower exports, reduced domestic use, and increased ending stocks. The NASS Grain Stocks report, issued March 31, implied **less feed and residual disappearance** than previously estimated [feed and residual ↓ | measurement]. Total feed and residual use is trimmed 15 million bushels to 135 million. Wheat exports are also cut 15 million bushels to 985 million **on a slowing pace and prices that have become uncompetitive in many international import markets** [exports ↓ | market competitiveness]. By class, Hard Red Winter and Soft Red Winter are reduced 10 million and 5 million bushels. The changes result in a 30 million bushel increase in ending stocks to 970 million. The projected season-average farm price is raised $0.05 per bushel to $4.60 **on updated NASS data** [price ↑ | measurement] as well as **surging nearby cash and futures prices** [price ↑ | market competitiveness], **partially resulting from the global COVID-19 pandemic** [price ↑ | external demand].

World wheat. . . . Global exports are lowered 0.9 million tons, led by a 1.5-million-ton reduction for Russia, which was directly offset by an equivalent increase for the EU. The Russia change is based primarily on **newly imposed government export restrictions** [Russia exports ↓ | policy]. The EU is raised on **less competition from Russia** [EU exports ↑ | market competitiveness] as well as expectations of a continued strong pace of exports. . . .

Source: USDA WASDE, 9 April 2020; labels: author's LLM extraction.

Figure E.2: August 2012 WASDE, corn: the 2012 drought in the commentary. *Note: the text reproduces the USDA wording of 10 August 2012 (omitted text marked . . .). Each bracket gives the balance-sheet item and direction, then the driver the model reads in the sentence. Teal is an exogenous driver (here weather) and violet an endogenous one (market competitiveness); a grey bracket marks a change the text presents as following from a change in a different item (induced). Sentences that only restate a number carry no bracket. Labels are from DeepSeek V4.1 Flash, prompt v3.5.*

WASDE release: 10 August 2012 **Commodity topic:** coarse grains, corn, marketing year 2012/13

Supply shock. U.S. feed grain supplies are projected sharply lower again this month with corn production forecast 2.2 billion bushels lower. The forecast U.S. corn yield is reduced 22.6 bushels per acre to 123.4 bushels **as extreme heat and dryness continued, and in many areas worsened, during July across the Plains and Corn Belt** [yield ↓ | weather]. Corn area harvested for grain is also lowered, down 1.5 million acres from the last month's forecast. U.S. corn production for 2012/13 is forecast at 10.8 billion bushels, the lowest since 2006/07. . . .

Market and use response. Imports for 2012/13 are also raised, up 45 million bushels to 75 million, **reflecting strong domestic corn prices and competitively priced foreign supplies** [imports ↑ | market competitiveness]. Total U.S. corn supplies for 2012/13 are projected down 2.0 billion bushels and at a 9-year low. This month's large reduction in U.S. corn supplies [total use ↓ | induced by supply] **and the sharply higher price outlook** [total use ↓ | market competitiveness] are expected to further reduce 2012/13 corn usage. Total use is projected 1.5 billion bushels lower. The biggest reduction again this month is for feed and residual disappearance, projected down 725 million bushels. . . .

Source: USDA WASDE, 10 August 2012; labels: author's LLM extraction.

Figure E.3: Stated drivers of WASDE revisions by affected balance-sheet item. *Note: each cell counts causal LLM labels in 1995–2026 WASDE commentary for wheat, coarse grains, rice and oilseeds (all countries); shading is the share within the driver (row). Exogenous drivers are weather, pest and disease, and measurement; endogenous drivers are policy, external demand and market competitiveness. Drivers labelled other, conflict, trend or assumption, or unstated (about 17 percent of causal labels) are not shown. Production includes yield and area; use includes feed and residual, food, industrial, seed, crush and ethanol; trade is exports and imports.*

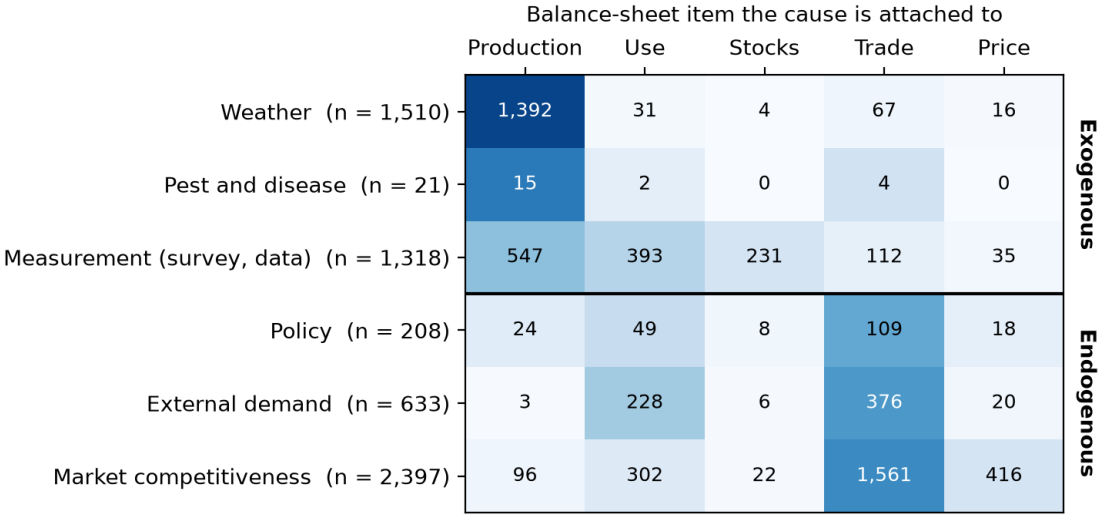


Figure E.4: The sentence-labelling prompt (v3.5). *Note: the figure summarises the system prompt; the verbatim prompt, with its worked examples and output schema, is in the replication package. The model is DeepSeek V4.1 Flash with reasoning enabled. Exogeneity is not asked of the model: it is assigned afterwards in code from the driver and the induced flag.*

1. Input
The whole commentary of one WASDE release, by topic, one numbered sentence each (`E001`, `E002`, ...).
Target topics: wheat, coarse grains, rice, oilseeds. Other topics are read as context only.
The front-page release note, if any, as context.
No numerical revisions, prices or other releases.

→

2. Task
Write one record for every statement in a target sentence: one commodity, one origin, one item, one direction, one driver.
Split a sentence by item, by country and by stated reason.
Write a one-line note quoting the words that decide the role and driver *before* choosing them.
Label only what the text says; otherwise *unstated* or *unknown*.

→

3. Output
One JSON object per call.
Each record carries the fields in panel 4 and a confidence level.
Code checks that every sentence has a record, and retries an empty or unparseable reply.

4. The five labelling axes

Axis	Values and meaning
Item (*channel*)	area planted, area harvested, yield, production, beginning stocks, imports, exports, feed and residual, food, industrial, seed, crush, ethanol, domestic use, total use, ending stocks, trade, price
Sentence type (*role*)	*restatement* (change, no reason); *causal* (a reason is given); *component* (explained by another item, e.g. production lower on yield); *identity* (arithmetic, e.g. stocks from supply and use); *summary*; *source only*; *forecast assumption*
Driver	weather; pest and disease; measurement | policy; external demand; market competitiveness | conflict; trend or assumption; other; unstated
Induced	*yes* only if the text says the change happens because of a change in a different item (exports raised on a larger crop); the parent sentence and item are recorded
Origin	U.S., world, or the country or region as written

5. Three rules that matter most

1. *Naming a report is not a reason.* "The Crop Production report indicates higher yields" has driver *unstated*. Measurement requires that an earlier number was revised, updated or reconciled.
2. *An assumption is not news.* "Trend yield" and "normal weather" are forecast assumptions, never a weather driver.
3. *Arithmetic is not a cause.* "Production lower on reduced yields" is a component, and "with smaller supplies, ending stocks are lower" is an identity. Neither is induced, and weather is not inferred.

6. Two sentences and their records (August 2012, corn)

"The forecast U.S. corn yield is reduced 22.6 bushels per acre to 123.4 bushels as extreme heat and dryness continued, and in many areas worsened, during July across the Plains and Corn Belt."

`corn us yield down role: causal driver: weather induced: no`

"Imports for 2012/13 are also raised, up 45 million bushels to 75 million, reflecting strong domestic corn prices and competitively priced foreign supplies."

`corn us imports up role: causal driver: market_competitiveness induced: no`