INSTITUT FÜR LEBENSMITTEL- UND RESSOURCENÖKONOMIK

Market expectations, public information and uncertainty resolution in agricultural commodity and equity markets

Dissertation

zur Erlangung des Grades

Doktorin der Agrarwissenschaften (Dr. agr.)

der Landwirtschaftlichen Fakultät

der Rheinischen Friedrich-Wilhelms-Universität Bonn

von

Nguyen Que An Cao

aus

Ben Tre, Viet Nam

Bonn 2023

Erstgutachter:	Prof. Dr. Thomas Heckelei
Zweitgutachter:	Prof. Dr. Michel A. Robe
Drittgutachter:	Prof. Dr. Kathy Baylis
Vorsitzender:	Prof. Dr. Monika Hartmann
Fachnahes Mitglied:	Prof. Dr. Scott H. Irwin

Tag der mündlichen Prüfung: 25. September 2023

Angefertigt mit Genehmigung der Landwirtschaftlichen Fakultät der Universität Bonn

Acknowledgement

There is a proverb saying: *"It takes a whole village to raise a child."* I believe that it holds true for the accomplishment of the Doctorate degree as well. In pursuit of the PhD, I am indebted to the enormous support of many people and communities. With the completion of this thesis, I would like to send each of them my most sincere gratitude.

First and foremost, I would like to thank Prof. Thomas Heckelei for giving me the opportunity to pursue the PhD, and for your enormous understanding and support during all these years. Not only that you taught me the most important foundation of economic and econometric theories, but you also taught me how to become a responsible, diligent researcher and teacher, and to have trust in myself. I also thank you sincerely for holding a space for me to freely develop my research interest, and for doing everything at your capacity to make sure that I can confidently move forward until the end. Thank you, for appreciating the efforts of your PhD students, even in times when they do not lead to remarkable outcomes. You strengthened my belief that it is always possible to be a considerate and kind person, no matter how busy one can be.

I sincerely thank Prof. Michel A. Robe for the opportunity to collaborate for the first time back in 2019, and for all the fruitful collaborations that followed until today. Thank you for the generous funding support for my many research visits in the Department of Agricultural and Consumer Economics (ACES), University of Illinois at Urbana-Champaign, where I had wonderful experiences as a visiting scholar and met great colleagues and friends. I also thank you for various helpful learning experiences, connections and guidance that you gave me as an advisor. Thank you very much for your appreciation for me as your PhD student, and for your valuable advice, especially also in terms of understanding the job market. And, thank you for all the espressos I have taken from your office in ACES!

I am grateful to Prof. Kathy Baylis for her genuine keenness on connecting graduate students to the networks that they belong, as she understands how much we need that in the early stage. Thanks to you, I was connected to ACES, my second "home institution", which plays a crucial part in the success of this dissertation project.

I sincerely thank Prof. Scott Irwin for giving me a root in agricultural commodity markets with his enthusiasm for the field, and his new book titled "Back to the Futures: Crashing Dirt Bikes, Chasing Cows, and Unraveling the Mystery of Commodity Futures Markets". I thank you very much for taking time to give me valuable advice and support, and especially for being a member of my PhD committee, despite your busy schedule.

I would like to send my sincere gratefulness to Prof. Monika Hartmann, who agreed to be the chair of my committee and gave me warm encouragements in the last steps to the finish line.

I am thankful for the administration and support from Bonn International Graduate School (BIGS) – Land and Food of the Faculty of Agriculture, Bonn University. In particular, I thank Dr. Susanne Plattes and Mrs. Regina Kirchner-Bierschenk for their support in various matters during my PhD.

I would like to thank my Master Thesis supervisor and also my first PhD supervisor, Stephanie, for seeding my interest in this research area. But beyond that, thank you wholeheartedly for giving me enormous mental support in the most difficult time of my life, so that I could successfully overcome the grief and prove that I was capable of pursuing a PhD. Thank you for being "persistent" – as you said – in helping me.

A big thank you goes to Bisrat – my wonderful collaborator and big brother – not only for his great contribution to the accomplishment of the last chapter of the thesis, but also for his patience, humble listening and encouragement.

My PhD time would not have been such a wonderful time in my life, had I not been in such a great working group, both in terms of competence and empathy. Thank you very much for all these years, AGPO team – to all those who have departed from our group, those who stay, and those who just joined recently. Thank you for offering me a safe space with full of positive energy and support, warm hugs and laughter. Thank you, Birgit, Kristin and Katrin, my "genii", for taking care of my administrative matters with full commitment. Quite often your care goes beyond what I am allowed to expect. Thank you, Mrs. Goossens, for greatly facilitating my access to all important textbooks that I needed in this process. Thank you, Wolfgang and Andreas, for your readiness to assist me promptly with any IT and computer issues. I thank Arnim, Hanna, Heinz-Peter, Monika, Sebastian, Svetlana and Yaghoob for your advice and support as the experienced post-Docs of the group. Especially thank you, Hugo, for playing the role of a gentle mentor and a big brother to me.

I feel thankful for having my wonderful peer PhD students around to share both the joys and the challenges of the PhD life. Thank you, Anna, Chaoya, Caroline, Christoph, David, Dominik, Elin, El-Mamoun, Eva, Felix, Janine, Julia, Leonie, Nina, Nkoyo, Rienne, Till, Ursula, Uyen, Yannic, and many more. I am particularly grateful for the incredible support of my four PhD angels – Elif, Helena, Linmei and Gowthami – in the run-up to the thesis submission, which made it an unforgettable memory for me.

I am also thankful to the department of ACES as my second "home institution", for the warm welcome, exchanges and supports throughout my visits during the PhD. Thank you, Prof. Peter Christensen, Prof. Philip Garcia, Prof. Margaret Grossman, Prof. Joe Janzen, Prof. Jared Hutchins, Prof. Nick Paulson, Prof. Teresa Serra and Prof. Paul Stoddard. I sincerely thank Jennifer Frank for her great administrative work to facilitate my many trips to visit ACES. My big thank you also goes to Ian Chong, who always made sure that I got the best IT equipment for my work as soon as I arrived – and is still greatly taking care of our project's facilities as of today – beside being a caring friend. Last but not least, thank you my wonderful ACES friends for hosting me with your maximum care and warmth: Prof. Ana Fava, chị Hân and anh Đức Anh, Fred, Gowthami and Arun, Keliang, Joana, Maria and Apostolos, Na, Nico, Noé, Protensia, Qianqian, Richie, Rocío and Ivan, Sarah and Vander, Seo Jeong, Thiago, Yoko, Yuanyuan, and many more.

My thankfulness is extended further to my microeconomics and econometrics students in the last three years. You motivated me to learn and constantly improve my understanding of the basis, through which my research is strengthened. Thank you for giving me the motivation to bring my enthusiasm to class and to joyfully share my knowledge with you.

During these years I spent on a challenging PhD project in Germany, it has been a blessing for me to receive care and love from my German family and friends around. Without that, the challenges of the PhD would have been unbearable. Thank you, Papa Wilfried and Mama Katharina, for giving me a home right on the first day when I arrived in Germany. Thank you, for loving me and being proud of me with the love of true parents, and calming me down when things got tough in all these years. My gratitude is extended further to Mama and Papa Stegen, who always made sure that I never had to spend Christmas time working on the PhD by myself. Thank you, for opening your hearts and arms to me, and for treating me as a member of the family. I am particularly thankful for having around me many good friends who gave me massive supports and cares, which made my PhD time more pleasant. Thank you, my dear Elisabeth, Catha, Isa, Björn, Lennart, Manu, Max, Neus, Nina, Nathalie, Pascal and Sarah, Vanesa, Vicki, and many more. I also thank Jörg and Angelika, for the opportunity to joyfully do handworks in their winery from time to time and refresh my mind after a long time focusing on research, and for your appreciation and care on such occasions.

My whole journey to pursue my academic education in Germany could not have begun, had I not received invaluable support and encouragement from my people at home. They prove that kindness is independent of education, wealth or social status – something I remind myself of every day. I am indebted not only to the material support, but also to unlimited love and trust that they have granted me from the very beginning until this point. I am deeply grateful to my family for encouraging me to keep up and finish the degree even in the time when they needed me the most, and for being proud of me. I especially thank my friends who collected all their savings for my journey to Germany, despite putting themselves at risk. I am thankful for having my soulmate Julie by my side in all these years, who never gives up on the good sides of the human-being in me, and at the same time patiently protect its weak sides. Thank you, for reminding me that the real success is achieved by a beautiful heart. Thank you, my beloved friends who always follow my progress proudly and love me unconditionally: cô Bình An, my 701 family, dì Phượng & dượng Trí, Ngân, Sương, Xuân & anh Quân, Kiều, Trinh, Thúy, Quyên, Quỳnh Anh, Trang, Đức, Đô, Thủy, Phúc, Hai Dương, chị Hảo, and many more.

Last – but never the least, this thesis was completed as part of the honor's promise to my late father – the one who taught me to become a helpful person and take good care of everyone along my way, no matter how far I can go. Had you not taught me that, I wouldn't have become that person who was trustworthy enough for all the help I received. Thank you, not only for giving me a life in flesh and bone, but also for granting me the right attitude to make that life meaningful. Thank you, Cha!

The pursuit of a PhD is a journey. But the completion of it is not a final destination. Therefore, the contribution of all kindnesses to the accomplishment of this dissertation will be extended to any good thing that will come out of it in the future.

Bonn, July 2023

Cao Nguyễn Quế An.

Kurzfassung

Die Hypothese des effizienten Marktes (Efficient Market Hypothesis, EMH) besagt, dass Kursbewegungen aus Änderungen der Markterwartungen resultieren, wenn neue Informationen verfügbar werden. Die vorliegende Arbeit befasst sich mit solchen Erwartungsänderungen im Zusammenhang mit der Veröffentlichung geplanter öffentlicher Bekanntmachungen, die Informationen über die fundamentale Marktdaten enthalten. Konkret untersucht die Arbeit die Rolle von Meldungen des Landwirtschaftsministeriums der Vereinigten Staaten (USDA) bei der Preisfindung an den Agrarrohstoff- und Aktienmärkten in den letzten 15 Jahren. Die drei durchgeführten Studien bieten sowohl neue empirische Erkenntnisse als auch methodische Beiträge zur bestehenden Literatur.

Die erste Studie untersucht umfassend die Entwicklung vorausschauender Unsicherheit und der Stimmung an den Mais- und Sojabohnenmärkten (erfasst durch die optionsimplizierte Volatilität – Ivol) rund um die Tage der Veröffentlichung vier wichtiger Gruppen von USDA-Berichten. Sie zeigt, dass die Berichte eine entscheidende Rolle bei der Reduzierung von Unsicherheit und für die Stimmung auf den Märkten im Vorfeld eines Ereignisses spielen. Der Umfang der Auswirkungen hängt von dem Ausmaß ab, in dem die Berichte den Markt überraschen, sowie von der vorher bestehenden Unsicherheit und Stimmung der Markterwartungen - abgebildet durch verschiedene Merkmale der Verteilung von Analystenprognosen vor dem Ereignis.

Die zweite Studie zeigt, dass – obwohl die Meldungen des USDA keine weitreichenden Auswirkungen auf die US-Aktienmärkte haben – die Berichte signifikante Reaktionen auf die Aktienkurse von Unternehmen des Lebensmittelsektors hervorrufen. Die Richtung und der Umfang der Reaktionen hängen davon ab, wie und in welchem Ausmaß die unerwarteten Nachrichtenkomponenten in den Berichten die erwarteten Cashflows der Unternehmen beeinflusst und ob sich die Auswirkungen auf die erwarteten Cash Inflows (d.h. Umsätze) oder Outflows (d.h. Inputkosten) beziehen.

Im letzten Teil der Arbeit wird eine innovative Methode entwickelt, mit der die *ex-post*-Überraschungskomponente geplanter öffentlicher Bekanntmachungen aus den Markterwartungen vor dem Ereignis herausgelöst werden kann, ohne sich auf Analystenumfragen vor dem Ereignis zu stützen. Die Methode kombiniert die theoretischen Grundlagen der EMH mit der Flexibilität nichtparametrischer Techniken des Maschinellen Lernens (ML), um eine theoretisch konsistente und dennoch hochflexible Methode zur Extraktion von *ex-ante*-Marktüberraschungen bereitzustellen, welche in verschiedenen Marktzusammenhängen angewandt werden kann. Eine Anwendung auf die USDA-Berichte über die Entwicklung und den Zustand der Nutzpflanzenbestände (Crop Progress and Condition Reports, CPCRs) zeigt, dass die Methode in der Lage ist, aus einer großen Anzahl möglicher Vorhersagefehler, die durch den ML-Algorithmus generiert werden, effektiv den besten Proxy für Marktüberraschungen nach der Veröffentlichung zu ermitteln. Die Anwendung zeigt auch, dass die CPCRs eine beträchtliche Menge an neuen Informationen liefern, die über das hinausgehen, was der Markt trotz der jüngsten Fortschritte in der Datenanalyse vorhersieht. Anhand der erheblichen Marktreaktionen auf diese Informationskomponente wird deutlich, dass der Informationsgehalt der Berichte für die Marktteilnehmer nach wie vor wertvoll ist.

Schlüsselwörter: *Marktüberraschungen, Markterwartungen, Marktunsicherheit, Agrarrohstoffe, Aktienmarktreaktionen, Regelmäßige Nachrichten, USDA-Meldungen, Maschinelles Lernen*

Abstract

The Efficient Market Hypothesis (EMH) states that price movements result from changes in market expectations when new information becomes available. This thesis focuses on such revisions of expectations around the releases of scheduled public announcements containing relevant information regarding market fundamentals. Specifically, the thesis investigates the role of USDA announcements in the price discovery processes in agricultural commodity and equity markets in the last fifteen years. The three studies conducted offer both novel empirical findings and methodological contributions to the extant literature.

The first study provides a thorough examination of the evolution of forward-looking uncertainty and sentiment in corn and soybean markets (as captured by Option-implied Volatility – Ivol) around the announcement days of four important groups of USDA reports. It shows that the reports still play a crucial role in resolving pre-event market uncertainty and sentiment. The scope of effects depends on the extent to which the reports surprise the market, as well as the pre-existing uncertainty and sentiment in market expectations – proxied by different characteristics of the pre-event analyst forecast distribution.

The second study reveals that – even though USDA announcements do not have a broad impact on U.S. stock markets – the reports do cause significant reactions in stock prices of food-sector companies. The sign and magnitudes of reactions is determined by how and how much the news component in the reports affects the expected cash-flows of firms, and whether the effect is on expected cash in-flows (*i.e.*, revenues) or out-flows (*i.e.*, input costs).

The last part of the thesis develops an innovative method to tease out the *ex-post* surprise component of scheduled public announcements from the pre-event market expectations without relying on pre-event analyst surveys. The methodology combines the theoretical foundation of EMH and the flexibility of nonparametric Machine Learning (ML) techniques to provide a theoretically consistent yet highly flexible method to extract *ex-ante* market surprises that can be employed in various market settings. An application to the USDA Crop Progress and Condition reports (CPCRs) demonstrates that the framework can effectively identify the best proxy for post-release market surprises among a large set of possible prediction error outcomes generated by ML algorithm. The application also reveals that the CPCRs still provide a substantial amount of new information beyond what the market anticipates, despite recent advancements in data analytics. Through the significant market reactions to this news component, it is evident that the informational content of the reports is still valuable to market participants.

Keywords: *Market Surprises, Market Expectations, Market Uncertainty, Agricultural Commodities, Stock Market Reactions, Scheduled News, USDA Announcements, Machine Learning*

Contents

List of tables

List of figures

Abbreviations

AR	Acreage Report
ARMA	Auto Regressive Moving Average
ATM	At-the-money
BPA	Bushels per Acre
CME	Chicago Mercantile Exchange
CBOT	Chicago Board of Trade
CBOE	Chicago Board Options Exchange
CAPM	Capital Asset Pricing Model
CPCR	Crop Progress and Condition Report
CRSP	Center for Research in Security Prices
CT	Central Time
EA	Earnings Announcement
EL96	Ederington and Lee (1996)
EMH	Efficient Market Hypothesis
EU	European Union
FC	Forecast changes
FE	Fixed-effect
FOMC	Federal Open Market Committee
GARCH	Generalized AutoRegressive Conditional Heteroskedasticity
GS	Grain Stocks
IMV	Implied Variance
IQR	Interquartile Range
IVol	Option-implied return volatility
KW	Kruskal–Wallis
ML	Machine Learning
NASS	National Agricultural Statistics Service
OLS	Ordinary Least Squares
PP	Prospective Plantings

PA	Planted Area
PCA	Principal Component Analysis
PRISM	Parameter-elevation Regressions on Independent Slopes Model
RE	Random-effect
RMWSE	Root Mean Weighted Square Error
SD	Standard deviations
SIC	Standard Industrial Classification
SUR	Seemingly Unrelated Regressions
SPX	Stock Market Index
VIX	CBOE Volatility Index
USDA	U.S. Department of Agriculture
WASDE	World Agricultural Supply and Demand Estimates
USA	United States
XGB	Extreme Gradient Boosting

Chapter 1
Introduction and overview of the thesis

"…These guys had been trading millions of dollars of soybean futures contracts, yet they don't recognize a soybean plant, but that's okay. They didn't need to know what the plant looked like to do their job and do it well. What they did need to know about were crop prices."

(Irwin and Peterson 2023, p. 21).

1.1 Motivation

The original story (McDermott 2016) is about a man nailing a soybean plant to the wall of the trading pit without anyone there recognizing it. This is an illustration of the widely-accepted view that price signals are the most condensed version of relevant information regarding market fundamentals (*e.g.*, Fama *et al.* 1969). But which type of information is relevant to market participants and how the information gets embedded into price signals – which themselves become a secondary type of information – are long-lived subjects of study for many markets. Not only because the knowledge is relevant for various stakeholders – from market participants to academics and regulators – but also because of the dynamic nature of most markets, leading to the need to periodically update the findings for every episode with distinct market structure and environment.

For a long time, U.S. Department of Agriculture (USDA) reports have been considered an important source of information regarding agricultural commodities' supply and demand fundamentals (Adjemian 2012; Adjemian and Irwin 2018; Fortenbery and Sumner 1993; Isengildina, Irwin and Good 2006; Isengildina-Massa *et al.* 2008; Lehecka, Wang and Garcia 2014; McKenzie 2008; McNew and Espinosa 1994; Ying, Chen and Dorfman 2019). These USDA reports, scheduled at different periodicities, aim to provide market participants with timely, reliable and publicly available information, leading to higher information transparency and ultimately more market efficiency. Indeed, numerous studies confirm that the information content of the reports is valuable to market participants through many channels: facilitating price discovery (Adjemian and Irwin 2018; Dorfman and Karali 2015; Hu *et al.* 2020), providing unbiased demand and supply projections (Egelkraut *et al.* 2003; Isengildina-Massa, Karali and Irwin 2020; McKenzie

2008), reducing information asymmetry (Fernandez-Perez *et al.* 2019; Wang, Garcia and Irwin 2014), resolving uncertainty (Isengildina-Massa *et al.* 2008; McNew and Espinosa 1994), and so on. Despite the extensive research already devoted in this area, several important issues remain unaddressed. Unravelling them makes up two main blocks of contributions of this thesis: empirical findings and methodology.

1.1.1 *The role of USDA reports in resolving agricultural commodity markets' uncertainty and impacting sentiment; moving beyond agricultural commodity markets: their impact on stock market valuations*

Both the releasing processes of various USDA reports and the structure of agricultural commodity markets have undergone significant changes in recent time. Most remarkable on the USDA side is the changes in the timing of the releases of many important reports (*e.g.*, World Agricultural Supply and Demand Estimates (WASDE) and Crop Production reports), from "before" to "during" the trading session, and from "with" to "without" pre-release media access (Adjemian and Irwin 2020). More recently, the USDA's National Agricultural Statistics Service (NASS) also makes efforts to supplement the traditional reports with finer-scaled geospatial datasets, of which the weekly Crop Progress and Condition report (CPCR) is one particular example (USDA 2021). Not only can such changes affect the speed and the accuracy of market judgements (Adjemian and Irwin 2020), but the second change in particular may contribute to improving forecasting models that market participants use to form their expectations as well.

On the market side, similar to the markets for other commodities and for financial instruments, the last decade has witnessed rapid developments of agricultural commodity markets in mutually reinforcing aspects. First, coming out of the commodity price booms 2007/08 and 2011/12, there is an increasing connectedness among markets. There is evidence that the interdependencies are strengthened in various dimensions: cross-country food price transmissions (Fernández, González and Rodríguez 2018), among the markets of different commodities (Bonato 2019; Zhang and Broadstock 2020) and between commodity markets and equity markets (Büyüksahin, Haigh and Robe 2010; Büyükşahin and Robe 2014; Cheng, Kirilenko and Xiong 2015; Tang and Xiong 2012). This increased integration potentially results in an expansion of the information set upon which market expectations are formed, and thus alter the ways in which specific news impact the markets on both ends. For example, one could also ask to which extent news about agricultural fundamentals affect stock market prices, and not only vice versa. Second, for major agricultural commodities such as grains and oilseeds, the tendency of geographical diversification continues further on the supply side. Specifically, the United States (USA) now account for only one-third of global corn and soybean exports and oneseventh of global wheat exports, as compared to more than two-thirds (of corn and soybeans) and one-third (of wheat) in the early 1990s (Zulauf *et al.* 2021). At the same time, the export shares of other world players such as Brazil, Argentina, and Ukraine (for corn), Brazil and Argentina (for soybeans), and Russia, the European Union (EU) and Canada (for wheat) have increased. As a consequence, it is natural to question if the relevance of the information about U.S. supply fundamentals released in the USDA reports to the market has changed amid this decentralization tendency.

The way in which markets absorb public information can also be driven by the evolution of other factors, such as climate change or the rapid rise of information technology. On the one hand, recent research shows that as climate change progresses, weather anomalies occur more frequently and more extremely (Coumou and Rahmstorf 2012; Stott 2016). As agricultural production depends heavily on weather conditions, this tendency can introduce more market uncertainty, both in terms of production output (Schlenker and Roberts 2009) and market expectation formation (Schlenker and Taylor 2021). On the other hand, unceasing advancements in computer science and information technology have enabled the "electronification" of major exchange floors– *i.e.*, the replacement of trading pits by computer-based clearing systems. For agricultural commodity markets, this change happened mostly around 2006-07 (Huang *et al.* 2022). At the same time, these advancements have resulted in the use of high-speed, sophisticated market analyzing tools for trading decisions (Karali *et al.* 2019; McKenzie 2008). Consequently, the market can be expected to be more complex in the way information is incorporated into market movements – as more data and techniques are available for market analysis. Indeed, the increasing role of high-frequency trading and algorithmic trading in price discovery and market volatility change is documented by recent literature (Haynes *et al.* 2017; Haynes and Roberts 2015; Raman, Robe and Yadav 2019).

Suppose that those structural changes in market institutions and trading environment lead to some evolution in market participants' information set, as well as in how they process and incorporate such information into trading decisions. Then, it is necessary to revisit the market influences of USDA reports amid these trends. Indeed, several recent studies have looked at the changes in certain aspects (Adjemian and Irwin 2020; Huang *et al.* 2022; Karali, Isengildina-Massa and Irwin 2019; Ying, Chen and Dorfman 2019). However, among the issues discussed above, the literature prior to this thesis lacks empirical answers for two important questions. First, within agricultural commodity markets, how much do USDA reports still help resolve market uncertainty and impact sentiment, as compared to the period preceding these changes (*e.g.*, as documented by Isengildina-Massa *et al.* (2008) and McNew and Espinosa (1994))? Second, given the discussion and general evidence on the strengthened connectedness between markets it seems a relevant question to ask to which extent we can expect that the impact of the news about agricultural fundamentals released in USDA reports goes beyond agricultural commodity markets – in particular, to stock markets? As will be presented in detailed in the next Section, a large part of this thesis is devoted to addressing these crucial questions.

1.1.2 *A novel method of separating anticipated from surprise components in public announcements*

The second block of this thesis focuses on a long-standing puzzle in studying the informational value of public announcements: how to separate the news component of the report from the portion of the report already anticipated by the market. To date, most works in this topic focus on the market impact of the announcements under the basic assumptions of the Efficient Market Hypothesis (*i.e.*, EMH – Fama *et al.* (1969)). Broadly speaking, EMH states that the market should only react to the news in (*i.e.*, the unanticipated component of) an announcement, since its predictable component must have been incorporated into the price before the report's release. However, market expectations are unobservable by nature, and therefore the news is also unobservable. Thus, any study concerned with the market impact of public announcements faces, first and foremost, the challenge of recovering the *ex-ante* market expectations about these announcements.

Announcement surprises are one kind of "shock" to the market. In theory, various models – either parametric or nonparametric – can be considered to capture pre-announcement market expectations. The proxy for the shocks is then derived as the forecast errors of such market expectation models.[1] In practice, however, parametric models are not satisfactory because they require many restrictive assumptions about the functional form of the aggregated expectation formation process, which are not observable to the econometrician. This limitation manifests itself when the underlying relationship between the predicted information and the predictors is complex and highly nonlinear, such as (to take just an example on which this thesis focuses) the relationship between crop condition or yields and the weather. The question is then reduced to how to approximate market expectation nonparametrically. To that effect, it is customary among the extant literature to proxy for *ex-ante* market expectations regarding the announcement using the analyst forecast "consensus" (*e.g.*: Freeman and Tse 1992; Garcia *et al.* 1997; Kinney, Burgstahler and Martin 2002). However, as pointed out by many previous studies (*e.g.*: Chiang *et al.* 2019; Karali, Irwin, and Isengildina-Massa 2019; Rigobon and Sack 2008),

[1] Popular methods to study price discovery such as Information Shares (Hasbrouck 1995) and Component Shares (Harris, McInish, and Wood 2002) are good examples of the parametric approach to separate market expectations and information shocks, though beyond in the context of event study for announcement effects.

this proxy is prone to bias, potentially due to stale information (*i.e.*, expectation revisions between the survey and the announcement) and analysts' incentives to deliberately bias their forecasts (see, *e.g.*, Chiang *et al.* (2019) and the references therein). As a consequence, both market expectations and the market impact of the news (*i.e.*, the difference between the announced information and market expectations) are subject to biases.

Apart from the estimation of the market impact, how much information in a public announcement can be anticipated using the available information prior to the release is a relevant question in its own right. This is particularly helpful for many USDA reports such as the CPCRs, since their usage is not only limited to market participants but may also be relevant to other users outside the market such as crop insurers, policy makers, and agronomists (Beguería and Maneta 2020; Bundy and Gensini 2022; Schnepf 2020; Wardlow, Kastens and Egbert 2006).

These considerations motivate the need to develop an alternative method to disentangle the unpredictable from the predictable content of the reports without resort to analyst forecasts. As much as advancements in information and computing technologies widen market participants' information access and expand their toolkits for analysis, they also provide researchers of market-expectation formation with the same possibilities. One remarkable case is the increasing use of Machine Learning (ML) in the market analysis literature, especially in commodity markets (Gu, Kelly, and Xiu 2020; Ouyang, Wei, and Wu 2019; Roznik, Mishra, and Boyd 2023; Sirignano and Cont 2019; Xu and Zhang 2021). Furthermore, the exploration of such information sources and toolkits for market research, in parallel with their utilization in trading, is essential since the ultimate goal of the researcher is to approximate the true market expectations as closely as possible. This makes up the second contribution of this thesis: an ML-based methodology for decomposing the scheduled public announcement into *ex-ante* market expectations and *ex-post* market surprises. The scope of this contribution goes far beyond agricultural commodity markets, as the resulting framework can be employed for many other markets, ranging from different commodity types to equities, bonds, and foreign exchange.

1.2 Research objective and structure of the thesis

The overarching objective of this thesis is to quantify the news component of public information released in USDA reports, and its impact on agricultural commodity and stock markets for the past fifteen years. To that end, it focuses on three sets of research questions, corresponding to the thesis' three main chapters. This Section presents these research questions and provides an overview of the thesis' structure.

1.2.1 *Research questions*

The thesis contributes to the extant literature by addressing the three following sets of research questions:

(I) How do USDA reports impact uncertainty and sentiment in agricultural commodity markets in present time?

The analysis conducted in this thesis follows prior literature (Bekaert, Hoerova and Lo Duca 2013; Ederington and Lee 1993; Patell and Wolfson 1979) to use option-implied volatilities (Ivol) as proxies for market participant's forward-looking volatility and sentiment, since changes in uncertainty and sentiment[2] map directly into the cost of options-based strategies (Goyal and Adjemian 2020). Options are a financial instrument that is widely traded on commodity derivatives markets. They are designed to protect against, or alternatively to bet on, future fluctuations in commodity prices. Their risk profile is achieved through the right to trade the commodity at a fixed price level at a given future time that is different than the current price of the commodity. Thus, option prices (or "premia") reflect the market valuation of this right, and thereby reflect the degree of forward-looking volatility expected by the market. The more uncertain it is for the market to gauge the future value of the commodity, the more valuable the option is.

Answering this research question involves investigating different aspects of market uncertainty and sentiment around USDA announcement days. Previous works shows that, both theoretically (Ederington and Lee 1996) and empirically for USDA reports (Isengildina-Massa *et al.* 2008; McNew and Espinosa 1994), Ivols drop following scheduled public announcements, indicating the uncertainty resolution effect of the news. Whether this pattern still holds true after 2008 is the starting point of analysis. After finding that there is still a significant drop on the day of announcement, the thesis then extends the question to several novel dimensions to provide a more thorough understanding of the impact beyond previous studies:

(i) How do Ivols evolve on the days before and after the announcement day? In equity market, the literature acknowledges that the impact of news – *e.g.*, earning announcement surprises – can last for several days around the scheduled event-day (Chiang *et al.* 2019). Thus, there is no obvious

[2] In this context, sentiment can be broadly defined as the factors that causes prices movements beyond what can be explained by supply and demand fundamentals. See, *e.g.*, Barberis, Shleifer, and Vishny (1998) and Baker and Wurgler (2006)

reason to assume that Ivols change following USDA announcements should be only a one-day event.

(ii) To what extent do USDA announcements' impact on uncertainty and market sentiment relate to from pre-event market expectations (as reflected in analyst forecasts)?

(iii) On a regular (*e.g.*, daily) basis, forward-looking uncertainty and sentiment in commodity market are positively driven by the uncertainty and sentiment in macroeconomic and broad financial market environment (Adjemian, Bruno, and Robe, 2017; Covindassamy, Robe, and Wallen, 2017; Robe and Wallen 2016). But does this relationship hold on USDA announcement days? That is, how do the uncertainty and sentiment from the broad-market environment and USDA announcements interact and jointly impact agricultural commodity market uncertainty and sentiment on scheduled USDA event days?

(II) To what extent do USDA reports move stock markets in general, and food-sector stocks specifically?

To date, studies about the value of USDA reports examine their impact on a wide range of agricultural commodity markets. But whether their impact ripples beyond those markets and penetrates the equity market – as reflected in the stock prices of publicly traded food-sector firms – has never been answered. As with research question (I), a comprehensive conclusion about this issue requires splitting it into two subproblems and investigating them in sequence:

(i) Do USDA report releases affect U.S. stock market returns and volatility as a whole? Not only is the answer to this question informative in itself, but it also provides the direction for the investigation of the food-sector-specific effect in the remaining part of the analysis. Asset pricing theories maintain that the price of a company's stocks is the total value of expected future cash-flows of that company, discounted to the present time using appropriate risk-adjusted required return rates. Thus, a change in stock price can stem from changes in either cash-flow expectations or in discount rates, or both. Given the small fraction of food-sector stocks in the U.S. stock market, and the fact that many firms do not depend on agricultural prices, it is unlikely that USDA reports could influence the expected cash-flows of the whole stock market. In turn, for the reports' effect to be through discount rates, it must be that market risk (*i.e.*, risk that affects a major part of the whole stock market systematically) would be altered significantly following USDA report releases. This is also unlikely. Indeed, since we find that the overall market does not significantly react to USDA reports, we can exclude this discount rate channel as well. In that case, when examining the causal paths through which USDA news may affect food-sector stock returns and volatility, one can ignore the discount rate pathway and focus solely on the revisions in expected cash-flows at the company level following the news, for firms that should be impacted – *i.e.*, firms from the food sector.

(ii) How do USDA announcements affect firms' excess returns, contingent on the sub-sector in which firms are operated (*e.g.*: farm machinery, restaurant, food retailer, etc.) and the type of information released? Provided that the discount-rate pathway can be excluded, price changes in food-sector firms' stock must be attributed to revisions in market expectations regarding the future cash-flows of these firms. However, the main business activities of food-sector firms are split between farms' input and output markets. Thus, the same news regarding agricultural commodity fundamentals can impact market expectations about firms' future cash-flows in different directions, depending on whether the news have negative or positive implications for firms' input costs or revenues. Furthermore, for a given firm in a given subsector, the magnitude and direction of the impact should intuitively also vary with the type of information (*e.g.*: planted acreage, inventory level, etc.). Hence, without conditioning on these factors, one cannot make robust conclusions about the role of USDA announcements in the stock market for those firms.

(III) How to disentangle the *ex-ante* market expectations and *ex-post* market surprises about a public announcement without resorting to analyst surveys?

While the two previous research questions focus on empirical findings, the last chapter's contribution is mainly methodological. As a practical exercise, it also provides a useful application to the case of CPCRs, which are released at weekly frequency by USDA NASS. Steps to tackle this challenge include:

(i) Building a theoretical framework for market expectation formation, allowing for expectations to be revised not only through incorporating more information but also through updating forecasting models over time.

(ii) Establishing the selection criterion to select the candidates that are most correlated with the true market expectations from a set of predictions.

(iii) Identifying a plausible information set base on which market participants form their expectations about the upcoming crop condition information released in the CPCRs.

(iv) Generating a sufficiently large number of crop condition predictions as candidates for market expectations with a wide range of predicting accuracy using the framework developed in (i) and the information set identified in (iii).

(v) Validating the selected proxies for market expectations about crop condition using extraneous information about the situation of crop production in the sample period.

(vi) Assessing the market impact of the CPCRs for the respective period using the surprise proxies corresponding to the selected market expectation proxies.

1.2.2 *Structure*

The remainder of this introduction chapter is organized as follows. Section 1.3 provides an overview of the data and methodologies employed to answer the research questions and achieve our research objectives, followed by a summary of the key findings of each of the three main chapter in Section 1.4. Section 1.5 concludes and lays on outlook for future research in this topic.

Chapters 2-4 address each of the aforementioned set of research questions/objectives separately, and hence are structured as independent research papers. The research paper presented in chapter 2 is already published as "Cao, A. N. Q. and Robe, M. A. (2022). "Market uncertainty and sentiment around USDA announcements." *Journal of Futures Markets, 42*(2), 250-275. The paper addresses research question (I) by investigating the effect of four main groups of USDA reports (WASDE, GS, PP and AR) on option-implied volatilities in the corn and soybean markets around USDA announcement days. The paper employs both parametric and nonparametric statistical tests and Seemingly Unrelated Regressions (SUR). Chapter 3 constitutes a working paper pre-published as "Cao, A. N. Q., Ionici, O. and Robe, M. A. (2023). "USDA Reports Affect the Stock Market, Too." *SSRN Electronic Journal*. Advance online publication." The paper has received a "Revise and Resubmit" invitation from the Journal of Commodity Markets and is currently under revision. It deals with research questions (II). To that effect, it investigates the impact of news regarding agricultural commodity fundamentals released in the same groups of USDA reports as in chapter 2 but looks at their impact on stock prices – both in general, and for agriculture-related firms in particular. Conclusions are drawn based on the result of a set of statistical tests (both parametric and nonparametric) and unbalanced-panel fixed-effect regressions. Chapter 4 is pre-published as a working paper titled "Cao, A. N. Q., Gebrekidan, B. H., Heckelei, T. and Robe, M. A. (2023). "Market surprises, machine learning and USDA Crop Progress and Condition reports." *SSRN Electronic Journal*." This last chapter is devoted to developing a ML-based framework in order to tease out the *ex-ante* market expectations of, and the corresponding *ex-post* market surprises from, the USDA's CPCRs under the semi-strong form assumptions of EMH. Hence, it answers the research question stated in (III). It blends the flexibility and powerful predicting power of the nonparametric Extreme Gradient Boosting (XGB) algorithm with a theoretical model of market expectation revision – together with the fundamental principle of linear regression with error in variable – to generate a sound method of distinguishing the news and not-news of the crop condition ratings in CPCRs based on post-event market reactions.

1.3 Data and methodologies

To answer these important research questions, the thesis employs a wide range of methods and data. Section 1.3.1 presents the datasets and the motivations behind their usage. Section 1.3.2 provides a brief description of the methodologies employed in every chapter.

1.3.1 *Data*

This thesis focuses on the impact of USDA announcements on futures and option markets of agricultural commodities (*e.g.*: corn, soybeans and wheat), as well as on stock markets. The datasets employed can be divided into four main groups: commodity market data, stock market data, Bloomberg pre-announcement analyst surveys of USDA announcements, and finally the geospatial data for crop condition and the information set used as its predictors. For chapter 2 and chapter 3, the sample period starts from September 2009 and ends in October 2019. For chapter 4, the sample period includes the planting seasons from 2015 to 2021, as the gridded dataset for crop condition is only available for the years since 2015.

(I) Commodity market data are from Bloomberg. Daily constant 90-day Ivols (of corn and soybeans) are used in chapter 2. According to Cui (2012), the Ivol estimates are extracted nonparametrically from at-the-money option prices at the daily market close. In chapter 4, the close-to-open "new-crop" price returns of corn and soybeans are used to evaluate the price movements due to the news released in the CPCRs. Among all corn and soybean contracts traded on CME, new-crop contracts are the contracts with the delivery and expiration dates due right after the harvest (*i.e.*, November for soybeans and December for corn).

(II) Stock market data are mostly from Compustat and CRSP. The study in chapter 3 draws conclusion about the effect of the USDA reports on overall U.S. stock market by comparing the daily returns on the Standard & Poor's 500 stock market index (SPX) – obtained from The Center for Research in Security Prices (CPRS) database. Filtering out the food-sector groups relies on Standard Industrial Classification (SIC) codes, resulting in a list of 154 food-sector firms with detailed information obtained from Compustat. These firms are then categorized further into sub-sectors based on their main activities. The categories include the input side of farms (*e.g.*: machinery, fertilizers, pesticides and agricultural technologies) and the output side of farms (*e.g.*: food processors, livestock producers, biofuel refiners, beverage manufacturers, etc.). The selection of firms into the final sample takes into account various factors such as mergers, acquisition and de-listings. CPRS database is then used to obtain the daily stock returns of the selected firms. For teasing out the proportion of returns stemming from the systematic risk which affects every stock in the market we also need the US Treasury Bill Rates ("T-bill rates") obtained from Bloomberg. The VIX index – a proxy for overall uncertainty and sentiment in equity market used in chapter 2 and chapter 3 – is also obtained from Bloomberg.

(III) Bloomberg pre-announcement analyst surveys of USDA announcements. The answers to research questions (I) and (II) are drawn on an extensive set of important USDA reports – *i.e.*, the World Agricultural Demand and Supply Estimates (WASDE), the Grain Stocks (GS), the Prospective Plantings (PP) and the Acreage (AR) reports. Starting from October 2009, Bloomberg has frequently published the results of its analyst surveys regarding these reports – typically a week before the scheduled release of the reports. In these surveys, analysts are asked to provide predictions of the key information in the upcoming USDA report, similar to the ways surveys are conducted for macroeconomic announcements, for instance the surveys about the federal fund rate announced after Federal Open Market Committee (FOMC) meetings (*e.g.*: Kurov *et al.* 2019). In the absence of an alternative proxy for pre-announcement market expectations, the sample mean or median of the sampled forecasts is widely adopted to represent market consensus expectations – despite its limitation as will be discussed in chapter 4. Resultantly, the analysis in chapter 2 also makes use of the direction of this consensus figure (*i.e.*, decrease or increase from previous period) to capture the general sentiment (*i.e.*, “pessimistic” or “optimistic”) in the market prior to the announcements, and the dispersion of the sampled forecast distribution as a proxy for the dispersion of the forecast distribution in the whole market. Chapter 2 focuses on the reported information for corn and soybeans only, whereas chapter 3 utilizes the information for wheat as well.

(IV) Geospatial data for CPCRs and weather variables as their main predictors. As mentioned previously, beside the traditional format, NASS recently introduced a novel gridded dataset with much finer spatial resolution (*i.e.*, 9x9 kilometers), as compared to the traditional state-level tabular dataset. This dataset opens the opportunity to evaluate the use of fine-scale weather data to predict crop condition, and only then to verify the conjecture that the market can predict accurately the reports in advance (*e.g.*: Bain and Fortenbery 2017). The main goal in building the predicting models for CPCRs is to construct an information set that is (i) highly relevant for crop condition and (ii) as similar as possible to the information set which is available to the markets. Weather variables and previous crop progress and condition are the most prominent candidates for this purpose. Weather variables come from the daily Parameter-elevation Regressions on Independent Slopes Model (PRISM) weather dataset (UCAR 2022). This dataset is at daily frequency and has a 4x4 kilometer spatial resolution. The predicting models fully utilize the dataset with all variables in it, including maximum temperature (Tmax), minimum temperature (Tmin), average temperature (tmean), precipitation (ppt), mean dew point temperature (tdmean), minimum vapor pressure deficit (vpdmin), and maximum vapor pressure deficit (vpdmax). Finally, the aggregation from pixel-level crop condition to an overall crop condition of the main production states in the US resorts to two additional data sources: (i) NASS Quick Stats database for the county-level planted acreage, and (ii) the boundary files from United States Census Bureau (for county and state boundaries) and from Esri ArcGIS Online Platform (for agricultural district boundaries). The information in these datasets is crucial to determine the weight of each geographical unit (*i.e.*, county or state) in the overall crop condition.

1.3.2 *Methodologies*

Across all three studies conducted, theoretical analysis is the starting point and shapes the design of empirical strategies.

For research question (I), the main theoretical consideration is to predict how Ivols change before and after a scheduled announcement. Answering this question requires going back to the idea of using Ivols as a measure of market expectations regarding price volatility of the underlying asset in the remaining life of the option. Thus, the change in Ivols around an announcement ties in with how market expectations about volatility are revised around that time. The theoretical model developed in chapter 2 is a variant of the model of implied variance (IMV) introduced by Ederington and Lee (1996), which is extended to allow for constant-maturity Ivols (*i.e.*, 90-day synthetic) instead of the Ivols of nearby options. From this model, a set of theoretical hypotheses are derived. Next, the study identifies suitable empirical methods to test these hypotheses, including statistical tests and regression analysis. Specifically, statistical tests (*e.g.*: paired-sample tests of mean/median, multiple comparisons across more than two samples) are used to determine the direction of Ivol changes on announcement day, and how it evolves several days before and after the event day (*i.e.,* part (i) of research question (I)). To account for the nonnormality of Ivol changes, both parametric and nonparametric tests are employed. SUR is then used to address part (ii) of the research question simultaneously for corn and soybeans – as it is more efficient compared to Ordinary Least Squares (OLS) due to the existence of cross-equation correlation among regression residuals. Especially, the regression aims to capture the asymmetric effects of report surprises on the magnitude of event-day Ivol changes by separating the surprises into price-bullish and price-bearish surprises. For part (iii), together with the SUR results obtained previously, OLS and different variants of Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) models are used to clarify the contradicting relationship between VIX and commodity Ivols on USDA-announcement *vs.* non-USDA-announcement days.

Likewise, addressing research question (II) starts from the fundamental of asset pricing theory to identify the potential mechanism through which USDA announcements may affect stock market. Next, demand and supply theories offer guidance for hypothetical predictions of the direction and the degree of impact of different types of information on stocks in different agriculture-related subsectors. Econometric analysis is then used to test these hypotheses empirically. The main method to investigate subproblem (i) is paired-sample tests, similar to those employed for part (i) of research question (I). For subproblem (ii), the conclusions are drawn upon both statistical tests and panel regressions, but using excess returns instead of raw returns. Unlike in agricultural commodity markets, stock prices are more strongly driven by factors related to the macroenvironment (*i.e.*, the market risk). Moreover, intuitions suggest that these factors can have long-lasting impact, as opposed to the transitory nature of USDA news generally restricted to crop seasons and reporting cycles. Consequently, there is a risk of endogeneity (when the omitted macroeconomic condition affects agricultural inventory level, for example) and autocorrelation (when the impact of those factors is carried forward to multiple future periods). Hence, to obtain unbiased and precise estimates of the impact of the announcement on food-sector firms, it is crucial to purge the portion of returns as reward for market risk out of the raw stock returns before analyzing the effect of USDA news. To tackle this problem, the study employs Capital Asset Pricing Models (CAPM) (Sharpe 1964; Treynor 1961a, 1961b) with rolling windows to estimate the expected returns induced from market risk for each trading day. The resulting excess returns – *i.e.*, the difference between raw return realization and that expected market-risk reward – capture stock price movements due to factors that are specific to firms, to which USDA news belong. With a rich panel of 154 firms over 151 USDA-report events, unbalanced fixed-effect regressions are then used to efficiently estimate the impact of the news on this excess return component to address part (ii) of the research question. The regression analysis is carried out for the firms on the input-side and output-side of farms separately, as necessitated by previous discussion.

With research question (III), building up a plausible theoretical model of how market participants update their expectations plays a decisive role in designing the predicting models to generate the candidates for market expectation and surprise proxies. In the context of linear regressions with measurement errors in explanatory variables, econometric theories then offer guidance to find the most appropriate selection criterion, which is used to select the best proxies among those candidates. As argued in chapter 4, in the absence of the true market surprises, the selected surprise proxies must be able to explain the more of the variation in post-announcement market returns than alternative surprise measures. Consequently, the R^2 of the univariate linear regression of returns on the surprise candidates is an appropriate criterion to select the best surprise proxy. The combination of these two blocks results in a new the ML routine to tackle the research question. This novel routine is based on the nonparametric Extreme Gradient Boosting (XGB) algorithm, which can handle the large geospatial datasets described in previous Section efficiently.

1.4 Empirical findings

This Section presents three sets of empirical findings of the three studies conducted in the thesis. In chapter 2 and chapter 3, the findings are the direct goals of the analyses answering research questions (I) and (II). In chapter 4, the findings are not the main focus and result from an application of a more general methodological framework, which is the solution to the challenge posed in research question (III). The scope of application of this framework moves beyond the CPCRs, as it can be adopted in various contexts. Nevertheless, altogether, the empirical findings in the three studies shed light on several important angles of the value of USDA announcements that remained uncovered in the extant literature.

1.4.1 *USDA reports still resolve uncertainty and sentiment in agricultural commodity market in the period 2009-2019.*

Consistent with previous studies, the analysis shows that Ivols of both corn and soybeans drop significantly on the days when USDA reports are released. The constant 90-day Ivols in use reflect market expectations about the actual (*i.e.*, realized) volatility in commodity prices within the next 90 days. Thus, the drop indicates that the expectations are revised downward on the event days. As the underlying theoretical model suggests, the main source of this revision is the removal of the event day out of the 90-day window, as it has then become the present day. The drop is thus evidence that the market had expected higher-than-normal realized volatility on the announcement day, which in turns results from the price discovering process. Such situation only occurs when the market as a whole had been confused about the fundamental information of interested and had expected that the released information in the upcoming report would help resolve such confusion. The resolution effect of USDA reports to market uncertainty and sentiment is therefore maintained.

Having confirmed that there is indeed a significant drop in Ivols on the announcement days, the study provides further insights into the development of Ivols around those days, as well as the determinants of the drop. These findings correspond to the three sub-questions of research question (I):

(i) Both corn and soybean Ivols remain lower than normal for a week after the announcement. For corn, there is a slight tendency of increasing Ivol in the run-up to the report releases – though not statistically significant. For soybeans, there is no clear pattern of Ivol change before the releases of the reports. Hence, it is apparent that their uncertainty and sentiment resolution effect lasts longer than just one day.

(ii) The extent of the effect depends on how much news the reports bring to the market – as compared to the pre-event market expectations, proxied by analyst forecast consensus. It is also contingent on the degree of forecast dispersion among the analysts, and the overall sentiment posed by the consensus forecasts. The economic and statistical significance of the effects vary with commodity and the type of report, but the data support a similar direction for both corn and soybean markets. In general, their directions of impact are consistent with what theories suggest. For examples, both price-bullish and price bearish surprises in GS reports (which indicate the current inventory level) seem to drive Ivols upward likely due to stronger implications for the current demand and supply balance, and thus triggering more immediate price adjustments. On the other hand, the surprises caused by WASDE reports do not significantly affect the Ivol changes when they are price-bearish, but the data support upward moving Ivols when they are price-bullish. That said, only when the projected inventory level by the end of the marketing year (which is the key information released in the WASDE) turns out to be lower than what the market expected, the resolution effect is weakened according to the results. This is because low-inventory condition generally implies more volatility to be expected, whereas the effect of abundant stocks is more ambiguous (Baur and Dimpfl 2018; Geman and Smith 2013), and more so in a longer time horizon. Likewise, when the disagreement among the analysts is stronger or when the overall sentiment among their forecasts is more pessimistic (*i.e.*, more price-bullish) compared to the previous period, the market will resort more to the announcement. Consequently, the announcement date can be expected to experience stronger price adjustments, and thus removing it from the Ivol window will cause a larger drop in the Ivol figure on that day.

(iii) While forward-looking volatility in macroenvironment (captured by the VIX index) positively affect forward-looking volatility in agricultural commodity markets, USDA reports mitigate this spillover volatility on the announcement days. Both OLS and GARCH models confirms that the changes of Ivols are positively driven by the changes in VIX on a daily basis. However, on USDA announcement days, the drop due to the scheduled events offsets the spillover effect from macroenvironment.

1.4.2 *The impact of USDA reports moves beyond agricultural commodity markets and reaches the agriculture-related segment of US stock market in the period 2009-2019.*

The findings corresponding to the two subproblems in research question (II) are the following:

(i) The overall US stock market does not react significantly to USDA announcements. Statistical tests on the broad market effects focus on three measures: SPX index returns, the absolute values of them as a measure of realized volatility, and the VIX index as a measure of forward-looking volatility. For all three measures, the tests fail to reject the null hypothesis that there is a significant difference between USDA-event days and non-USDA-event days. Thus, when examining the effect of the reports on the stocks of food-sector firms, it is plausible to exclude the discount risk as a channel through which the effect is materialized.

(ii) Food-sector stocks, on the other hand, react significantly to USDA announcements. As argued in part (ii) of research question (II), the impact of USDA news on the expected cashflows of these stocks is contingent on both the nature of firm activities and the nature of the reported information. For instance, stock returns of firms on the input-side and output-side of farms experience contradicting impacts of the same GS surprises. According to the results, companies that produce farms' input (*e.g.*: fertilizer, machinery, etc.), if the current commodity inventory is lower (higher) than expected, it contributes positively (negatively) to the excess returns, since it implies more (less) input needed to produce a larger (smaller) crop. Conversely, for firms that use agricultural commodities as their input, when the inventory level turns out to be lower (higher) than expected, it implies a potential increase (decrease) in input cost for them, and thus has a negatively (positively) impact on the excess returns. Reports about the planted acreage (*i.e.*, PP and AR) positively affect the excess returns of farm-input stocks when the figures are lower than expected – indicating that the eased cashflow constraints for farm investments (due to higher commodity prices) might be more important for determining farms' increasing input demand than the reduced cropping area in one particular year. As for stocks on the commodity-user side, the results are intuitive: more acreage than expected predicts lower input cost for firms, thereby increase firms' profitability and consequently excess returns, and vice versa. For WASDE reports, the effects are ambiguous for firms on both sides of farms.

1.4.3 *CPCRs still generate substantial surprises to agricultural commodity markets, and the market significantly reacts to its surprises.*

The following empirical findings result from the application of the methodological framework sought in research question (III) to USDA Crop Progress and Condition reports.

(i) The new geospatial dataset of crop condition is an unbiased representation of the state-level, tabular data in the real-time releases of CPCRs. Since we want to use the new high-resolution CPCR dataset to assess the predictability of crop condition, it is important to first verify the compatibility of this dataset to the traditional state-level data. The comparison reveals that the two datasets closely resemble each other, with negligible discrepancies for the most part of the sample period. Thus, it can be concluded with confidence that NASS has properly reproduced the gridded condition layers from the original reports. The gridded data are therefore reliable for the purpose of separating *ex-ante* expectations and *ex-post* surprises of crop condition reports.

(ii) Using an information set that is publicly available to market participants prior to the report releases, XGB algorithm is capable of predicting crop condition with high accuracy. For each corn and soybean crop season, the models generate 12,800 series of crop condition predictions, corresponding to 12,800 unique combinations of six important XGB hyperparameters. The sets of values for these hyperparameters are defined to achieve a broad range of predictive accuracy, as there are no reasons to restrict the possible proxies of market expectations to be among the best predictions. Nevertheless, out of all prediction series for each year, the best prediction series closely approach the actual condition, as pair-sample tests fail to reject the null that the two are drawn from the same population. The second-best predictions are the medians of the prediction sets.

(iii) The best prediction series are, however, not the best proxies for pre-event market expectations, as revealed by post-event market price movements. Despite that the models can generate such highly accurate predictions of crop condition, the log-difference errors of those best prediction series fail to explain the variations in post-announcement returns, as indicated by R^2 of the market return regressions – the selection criterion proposed by the methodological framework. Instead, the prediction series with the errors that yield the highest R^2 are far away from both the best predictions and the median predictions. This result is consistent in all years in the sample for both crops.

(iv) Ex-post price movements do not support the conjecture that the CPCRs no longer surprise the public. The fact that the selected market expectation proxies are distant away from the best condition predictions which are highly accurate makes clear that the CPCRs still surprise the market substantially in the present time. Despite the availability of powerful forecasting tools and the accessibility of high-resolution weather data, the market is unable to fully anticipate the content of the reports before they are released. In extreme cases, the selected proxies indicate that the reports can surprise the market by up to 30 percent deviation from what had been expected. Moreover, additional information about the contradiction between yield forecasts and crop developments in the studied period lends support to the plausibility of these proxies. In general, both corn and soybean markets tend to be overoptimistic in anticipating crop condition.

(v) The new information released in the CPCRs still impact the market significantly. In the period 2016-2021, the analysis suggests that both corn and soybean new-crop future prices still react significantly to the news released in the reports, as captured by the selected surprise proxies. The direction of impact is consistent with supply and demand theories: a better-than-expected crop condition revises the expectations about crop size upward, and is thus price bearish, and vice versa. However, the size of effect is modest. Given that the estimated surprise component of the reports is non-trivial, it can be argued that the small market impact is not due to the predictability of the aggregate crop condition, but is rather due to the transitory nature of crop condition information, which is updated frequently throughout growing season.

1.5 Conclusion and outlook

This thesis starts with the observation that several aspects of the market impact of USDA reports – not only on the markets of the commodities themselves but also on the stock market – remain unexplored. For the latter, the focus is particularly put on companies that either provide input for agricultural production or use agricultural commodities as their input. These questions are addressed by three separate studies, which comprise the next three chapters of the thesis. The studies contribute greatly to the understanding of the market impact of publicly announced information in general, as well as USDA reports in specific, both by empirical findings and methodological innovations.

The first study investigates the role of USDA reports in resolving commodity market uncertainty and sentiment embed in option premia. If the report releases facilitate price discovery process, expected uncertainty and sentiment will be reduced, which can be observed via the drops in Ivol – the option-implied future volatility. The analysis carried out throughout the chapter strongly confirms this effect. It also shows, for the first time, that the effect lasts longer than just one day, as the report releases hold the expected volatility level lower than normal for up to a week afterwards. Furthermore, using pre-report analyst surveys organized by Bloomberg as a proxy for pre-event market expectations, the study finds evidence that the extent of Ivol drops following the report releases is partially determined by the how (*i.e.*, the direction) and how much (*i.e.*, the magnitude) the reported information differs from the analyst consensus, the level of disagreement among the analysts, and the overall sentiment about the current situation (of inventory level) compared to the previous report. Finally, it also suggests that forward-looking volatility in agricultural commodity markets are less prone to spillover effect from the broader financial environment – as reflected in the VIX index – on days when important USDA reports are released.

The second study examines potential impact of USDA reports on US stock market, especially the segment of food-sector companies. Stock returns reflect changes in either discount rates or market expectations about cashflows of firms. The first part of the study verifies that the reports do not significantly affect the whole US stock market. Thus, it also lends support to rule out the possibility that food-sector stocks can be affected by changing discount rate. Therefore, in the second part, theoretical predictions focus on the potential changes in expected cashflows of food-sector firms due to new information released in USDA reports. Results show that, indeed, when conditioning on the type of information released in the reports and whether firms operate on the input or output side of farms, USDA reports significantly affect the excess returns of food-sector companies' stocks. Firms that produce input for farms experience negative (positive) excess returns if the released news predicts lower (higher) input demand for agricultural production. Reversely, USDA news positively (negatively) affects excess returns of firms that use agricultural commodities as their input when the news indicates that commodity supply is more abundant (scarcer), and hence the input costs of those firms are lower.

The first and the second studies share a main limitation regarding the proxies for pre-event market expectations and new information released in the reports. The use of pre-event analyst surveys as a proxy for market expectations is widely adopted in finance literature, but such proxy is highly subject to measurement errors. The third study tackles this challenge by developing a novel ML-based method to separate the news from the anticipated component of public announcements without relying on analyst surveys. This unique framework departs from the traditional approach in previous studies in two major features. One, it reverses the usual procedure of identification: using market prices to identify the *ex-post* surprises and *ex-ante* market expectations. Among a large set of prediction outcomes and their corresponding forecast errors generated by the XGB algorithm, post-released market returns are used to determine the surprise candidate that explains the most variation in them. Thus, it eliminates the sources of bias induced by analyst surveys. Two, the models that produce the best market expectation and surprise proxies are allowed to be updated overtime, reflecting the market's progress in both data utilization and predicting skills. Hence, it reduces further the number of arbitrary restrictions on market expectation formation. An application of the framework on USDA CPCRs shows that, contradicting to the belief that crop condition can be well anticipated by market participants in advance, the reports still bring a substantial amount of new information to the market, and thereby impact post-event futures returns of crops significantly.

Since USDA reports costs tens of "millions of dollars to collect and disseminate" (Karali *et al.* 2019, p. 66), the concern that they have become redundant and thus their continuation should be reconsidered is a valid one. However, the findings in this thesis suggest that such decision should be made with caution. The empirical findings across three studies deliver an important message: USDA announcements are still relevant to market participants, and their scope of impact is not limited to only agricultural commodity markets. In particular, the results of the last study clearly demonstrate that the market as a whole can be considerably flat-footed when anticipating the information released in the reports, even for a type of report whose available information set for expectation formation is publicly and highly accessible.

Going forward, the thesis opens several venues for future research. The methodological framework developed in chapter 4 is highly applicable to many research questions which require an estimate of market expectations. One particular interesting question is to which extend option-implied volatility of agricultural commodities in major exchanges such as CME can help predict food price volatility in low-income countries, where local food prices are potentially dependent on world market fluctuations. In a world with climate change and increasing connectedness among markets, food price volatility can have more severe impacts on global food security. On the one hand, futures and option markets on major exchanges are operated with advanced technologies and highly liquid. On the other hand, local food markets in low-income countries are often faced by low liquidity and high transaction cost due to constraints in market infrastructures. Thus, information asymmetry is a more persistent problem in case of the latter. First, a better understanding of how and which type of information is incorporated into volatility expectations of market participants on option markets (as captured in Ivols) is necessary to determine whether the such volatility expectations are formed efficiently. Then, one can examine the predictive power of Ivols to the realized volatility of given local markets in various predicting models with different information sets and different time horizons. An analysis of these information sets and the ones used by option markets as well as the predictive performances across model variants will shed light on whether and how local markets response to changes in expected volatility. The investigation will potentially improve the predictability of local food price volatility, and thus facilitate efficient coping strategy to mitigate its negative impact. For such study, the insights provided by chapter 2 and the methodology put forward in chapter 4 are particularly relevant.

Finally, this thesis illustrates an important caution regarding the use of ML methodologies for applied works and market analysis specifically, and for causal identification in applied economics in general. The development of the ML-based methodology described in chapter 4 makes clear that the ML toolbox can offer powerful techniques for understanding market mechanisms, especially when highly nonlinear, complex data generating processes are involved. Yet, their usefulness is governed first and foremost by a comprehensive understanding of market theories. ML can efficiently substitute other quantitative methods when the research question involves prediction tasks, but it cannot substitute an appropriate research design giving the prediction task its purpose in the overall identification approach ensuring unbiasedness and consistency of the conclusions. Therefore, what ML methods can and cannot do for a given research problem is an important question to ask, but only after a conceptual framework for the analysis has been defined.

1.6 References

Adjemian, M. K. (2012). "Quantifying the WASDE Announcement Effect." *American Journal of Agricultural Economics*, *94*(1), 238–256. https://doi.org/10.1093/ajae/aar131

Adjemian, M. K. and Irwin, S. H. (2018). "USDA Announcement Effects in Real-time." *American Journal of Agricultural Economics*, *100*(4), 1151–1171.

Adjemian, M. K. and Irwin, S. H. (2020). "The market response to government crop news under different release regimes." *Journal of Commodity Markets*, *19*, 100110. https://doi.org/10.1016/j.jcomm.2019.100110

Bain, R. and Fortenbery, T. R. (2017). "Impact Of Crop Condition Reports on National and Local Wheat Markets." *Journal of Agricultural and Applied Economics*, *49*(1), 97–119.

Baker, M. and Wurgler, J. (2006). "Investor Sentiment and the Cross-Section of Stock Returns." *The Journal of Finance*, *61*(4), 1645–1680. https://doi.org/10.1111/j.1540-6261.2006.00885.x

Barberis, N., Shleifer, A. and Vishny, R. (1998). "A model of Investor Sentiment." *Journal of Financial Economics*, 49, 307–343. https://doi.org/10.1515/9781400829125-015

Baur, D. G. and Dimpfl, T. (2018). "The asymmetric return-volatility relationship of commodity prices." *Energy Economics*, *76*, 378–387.

Beguería, S. and Maneta, M. P. (2020). "Qualitative crop condition survey reveals spatiotemporal production patterns and allows early yield prediction." *Proceedings of the National Academy of Sciences*, *117*(31), 18317–18323.

Bekaert, G., Hoerova, M. and Lo Duca, M. (2013). "Risk, uncertainty and monetary policy." *Journal of Monetary Economics*, *60*(7), 771–788. https://doi.org/10.1016/j.jmoneco.2013.06.003

Bonato, M. (2019). “Realized correlations, betas and volatility spillover in the agricultural commodity market: What has changed?” *Journal of International Financial Markets, Institutions and Money*, *62*, 184–202. https://doi.org/10.1016/j.intfin.2019.07.005

Bundy, L. R. and Gensini, V. A. (2022). “An assessment of USDA corn condition ratings across the U.S. Corn Belt.” *Agronomy Journal*, *114*(1), 601–617. https://doi.org/10.1002/agj2.20973

Büyüksahin, B., Haigh, M. S. and Robe, M. A. (2010). “Commodities and Equities: Ever a “Market of One”?” *The Journal of Alternative Investments, 12*(3), 76–95. https://doi.org/10.3905/JAI.2010.12.3.076

Büyükşahin, B. and Robe, M. A. (2014). “Speculators, commodities and cross-market linkages.” *Journal of International Money and Finance*, *42*, 38–70. https://doi.org/10.1016/j.jimonfin.2013.08.004

Cheng, I.-H., Kirilenko, A. and Xiong, W. (2015). “Convective Risk Flows in Commodity Futures Markets”: *Review of Finance*, *19*(5), 1733–1781. https://doi.org/10.1093/rof/rfu043

Chiang, C.-H., Dai, W., Fan, J., Hong, H. and Jun, T. (2019). “Robust Measures of Earnings Surprises.” *The Journal of Finance*, *74*(2), 943–983. https://doi.org/10.1111/jofi.12746

Coumou, D. and Rahmstorf, S. (2012). “A decade of weather extremes.” *Nature Climate Change*, *2*(7), 491–496. https://doi.org/10.1038/nclimate1452

Cui, C. (2012). “Commodity implied volatility surface.” *Bloomberg OVDV Commodity Volatility Surface White Paper*.

Dorfman, J. H. and Karali, B. (2015). „A Nonparametric Search for Information Effects from USDA Reports.” *Journal of Agricultural and Resource Economics*, *40*(1), 124–143. Retrieved from www.jstor.org/stable/44131280

Ederington, L. H. and Lee, J. H. (1993). “How Markets Process Information: News Releases and Volatility.” *The Journal of Finance*, *48*(4), 1161–1191. https://doi.org/10.1111/j.1540-6261.1993.tb04750.x

Ederington, L. H. and Lee, J. H. (1996). “The Creation and Resolution of Market Uncertainty: The Impact of Information Releases on Implied Volatility.” *The Journal of Financial and Quantitative Analysis*, *31*(4), 513. https://doi.org/10.2307/2331358

Egelkraut, T. M., Garcia, P., Irwin, S. H. and Good, D. L. (2003). “An Evaluation of Crop Forecast Accuracy for Corn and Soybeans: USDA and Private Information

Agencies." *Journal of Agricultural and Applied Economics*, *35*(1), 79–95. https://doi.org/10.1017/S1074070800005952

Fama, E. F., Fisher, L., Jensen, M. C. and Roll, R. (1969). "The adjustment of stock prices to new information." *International Economic Review*, *10*(1), 1–21.

Fernández, A., González, A. and Rodríguez, D. (2018). "Sharing a ride on the commodities roller coaster: Common factors in business cycles of emerging economies." *Journal of International Economics*, *111*, 99–121. https://doi.org/10.1016/j.jinteco.2017.11.008

Fernandez-Perez, A., Frijins, B., Indriawan, I. and Tourani-Rad, A. (2019). "Surprise and dispersion: informational impact of USDA announcements." *Agricultural Economics*, *50*(1), 113–126. https://doi.org/10.1111/agec.12470

Fortenbery, T. R. and Sumner, D. A. (1993). "The effects of USDA reports in futures and options markets." *Journal of Futures Markets*, *13*(2), 157–173.

Freeman, R. N. and Tse, S. Y. (1992). "A Nonlinear Model of Security Price Responses to Unexpected Earnings." *Journal of Accounting Research*, *30*(2), 185. https://doi.org/10.2307/2491123

Garcia, P., Irwin, S. H., Leuthold, R. M. and Yang, L. (1997). "The value of public information in commodity futures markets." *Journal of Economic Behavior & Organization*, *32*(4), 559–570. https://doi.org/10.1016/S0167-2681(97)00013-9

Geman, H. and Smith, W. O. (2013). "Theory of storage, inventory and volatility in the LME base metals." *Resources Policy*, *38*(1), 18–28.

Goyal, R. and Adjemian, M. K. (2020). "The 2019 Government Shutdown Increased Uncertainty in Major Agricultural Commodity Markets."

Grosche, S.-C. and Heckelei, T. (2016). "Directional Volatility Spillovers Between Agricultural, Crude Oil, Real Estate, and Other Financial Markets." In *Food Price Volatility and Its Implications for Food Security and Policy* (pp. 183–205). Springer.

Gu, S., Kelly, B. and Xiu, D. (2020). "Empirical Asset Pricing via Machine Learning." *The Review of Financial Studies*, *33*(5), 2223–2273. https://doi.org/10.1093/rfs/hhaa009

Harris, F. H. d., McInish, T. H. and Wood, R. A. (2002). "Security price adjustment across exchanges: an investigation of common factor components for Dow stocks." *Journal of Financial Markets*, *5*(3), 277–308. https://doi.org/10.1016/S1386-4181(01)00017-9

Hasbrouck, J. (1995). "One Security, Many Markets: Determining the Contributions to Price Discovery." *The Journal of Finance, 50*(4), 1175–1199. https://doi.org/10.1111/j.1540-6261.1995.tb04054.x

Haynes, R., Raman, V., Robe, M. A. and Yadav, P. K. (2017). Algorithmic Trading and Livestock Market Quality. Paper presented at the 2017 NCCC-134 Conference in St Louis, MO. April. 2017.

Haynes, R. and Roberts, J. S. (2015). "Automated trading in futures markets." *CFTC White Paper*.

Hu, Z., Mallory, M., Serra, T. and Garcia, P. (2020). "Measuring price discovery between nearby and deferred contracts in storable and non-storable commodity futures markets." *Agricultural Economics*, *51*(6), 825–840. https://doi.org/10.1111/agec.12594

Huang, J., Serra, T., Garcia, P. and Irwin, S. H. (2022). "To batch or not to batch? The release of USDA crop reports." *Agricultural Economics*, *53*(1), 143–154. https://doi.org/10.1111/agec.12667

Irwin, S. H. and Peterson, D. (2023). "*Back to the futures: Crashing dirt bikes, chasing cows, and unraveling the mystery of commodity future markets*." Mahomet: Ceres Books.

Isengildina, O., Irwin, S. H. and Good, D. L. (2006). "Are Revisions to USDA Crop Production Forecasts Smoothed?" *American Journal of Agricultural Economics*, *88*(4), 1091–1104. https://doi.org/10.1111/j.1467-8276.2006.00918.x

Isengildina-Massa, O., Irwin, S. H., Good, D. L. and Gomez, J. K. (2008). "Impact of WASDE reports on implied volatility in corn and soybean markets." *Agribusiness*, *24*(4), 473–490. https://doi.org/10.1002/agr.20174

Isengildina-Massa, O., Karali, B. and Irwin, S. H. (2020). "Can Private Forecasters Beat the USDA? Analysis of Relative Accuracy of Crop Acreage and Production Forecasts." *Journal of Agricultural and Applied Economics*, *52*(4), 545–561. https://doi.org/10.1017/aae.2020.18

Karali, B., Irwin, S. H. and Isengildina-Massa, O. (2019). "Supply Fundamentals and Grain Futures Price Movements." *American Journal of Agricultural Economics*, *44*(3), 362. https://doi.org/10.1002/ajae.12012

Karali, B., Isengildina-Massa, O. and Irwin, S. H. (2019). "The Changing Role of USDA Inventory Reports in Livestock Markets." *Journal of Agricultural and Resource Economics*, *44*(3), 591–604. Retrieved from https://www.jstor.org/stable/26797577

Karali, B., Isengildina-Massa, O., Irwin, S. H., Adjemian, M. K. and Johansson, R. (2019). “Are USDA reports still news to changing crop markets?” *Food Policy*, *84*, 66–76. https://doi.org/10.1016/j.foodpol.2019.02.005

Kinney, W., Burgstahler, D. and Martin, R. (2002). “Earnings Surprise “Materiality” as Measured by Stock Returns.” *Journal of Accounting Research*, *40*(5), 1297–1329. https://doi.org/10.1111/1475-679X.t01-1-00055

Kurov, A., Sancetta, A., Strasser, G. and Wolfe, M. H. (2019). “Price Drift Before U.S. Macroeconomic News: Private Information about Public Announcements?” *Journal of Financial and Quantitative Analysis*, *54*(1), 449–479. https://doi.org/10.1017/S0022109018000625

Lehecka, G. V., Wang, X. and Garcia, P. (2014). “Gone in Ten Minutes: Intraday Evidence of Announcement Effects in the Electronic Corn Futures Market.” *Applied Economic Perspectives and Policy*, *36*(3), 504–526. https://doi.org/10.1093/aepp/ppu010

McDermott, J. (2016, March 1). “A Eulogy for the Pit Trader.” *MEL Magazine*.

McKenzie, A. M. (2008). „Pre-Harvest Price Expectations for Corn: The Information Content of USDA Reports and New Crop Futures.” *American Journal of Agricultural Economics*, *90*(2), 351–366. https://doi.org/10.1111/j.1467-8276.2007.01117.x

McNew, K. P. and Espinosa, J. A. (1994). “The informational content of USDA crop reports: Impacts on uncertainty and expectations in grain futures markets.” *Journal of Futures Markets*, *14*(4), 475–492. https://doi.org/10.1002/fut.3990140408

Ouyang, H., Wei, X. and Wu, Q. (2019). “Agricultural commodity futures prices prediction via long- and short-term time series network.” *Journal of Applied Economics*, *22*(1), 468–483. https://doi.org/10.1080/15140326.2019.1668664

Patell, J. M. and Wolfson, M. A. (1979). “Anticipated information releases reflected in call option prices.” *Journal of Accounting and Economics*, *1*(2), 117–140. https://doi.org/10.1016/0165-4101(79)90003-X

Raman, V., Robe, M. A. and Yadav, P. K. (2019). “The Third Dimension of Financialization: Electronification, Intraday Institutional Trading, and Commodity Market Quality.” *SSRN Electronic Journal*. Advance online publication. https://doi.org/10.2139/ssrn.3068248

Rigobon, R. and Sack, B. (2008). “8. Noisy Macroeconomic Announcements, Monetary Policy, and Asset Prices: Asset Prices and Monetary Policy.” In J. Y. Campbell (Ed.)

(pp. 335–370). University of Chicago Press. https://doi.org/10.7208/9780226092126-011

Robe, M. A. and Wallen, J. (2016). “Fundamentals, Derivatives Market Information and Oil Price Volatility.” *Journal of Futures Markets*, *36*(4), 317–344. https://doi.org/10.1002/fut.21732

Roznik, M., Mishra, A. K. and Boyd, M. S. (2023). “Using a machine learning approach and big data to augment WASDE forecasts: Empirical evidence from US corn yield.” *Journal of Forecasting.* Advance online publication. https://doi.org/10.1002/for.2956

Schlenker, W. and Roberts, M. J. (2009). “Nonlinear temperature effects indicate severe damages to U.S. Crop yields under climate change.” *Proceedings of the National Academy of Sciences of the United States of America*, *106*(37), 15594–15598. https://doi.org/10.1073/pnas.0906865106

Schlenker, W. and Taylor, C. A. (2021). “Market expectations of a warming climate.” *Journal of Financial Economics*, *142*(2), 627–640. https://doi.org/10.1016/j.jfineco.2020.08.019

Schnepf, R. (2020). “Federal Crop Insurance: Record Prevent Plant (PPL) Acres and Payments in 2019.” *CRS Report. Washington, DC: Congressional Research Service*.

Sharpe, W. F. (1964). Capital Asset Prices: “A Theory of Market Equilibrium under Conditions of Risk.” *The Journal of Finance*, *19*(3), 425–442. https://doi.org/10.1111/j.1540-6261.1964.tb02865.x

Sirignano, J. and Cont, R. (2019). “Universal features of price formation in financial markets: perspectives from deep learning.” *Quantitative Finance*, *19*(9), 1449–1459. https://doi.org/10.1080/14697688.2019.1622295

Tang, K. and Xiong, W. (2012). “Index Investment and the Financialization of Commodities.” *Financial Analysts Journal*, *68*(6), 54–74. https://doi.org/10.2469/faj.v68.n6.5

Stott, P. (2016). “How climate change affects extreme weather events.” *Science (New York, N.Y.)*, *352*(6293), 1517–1518. https://doi.org/10.1126/science.aaf7271

Treynor, J. L. (1961a). “Market value, time, and risk.” *Time, and Risk (August 8, 1961)*.

Treynor, J. L. (1961b). *Toward a theory of market value of risky assets.*

UCAR (2022). “PRISM High-Resolution Spatial Climate Data for the United States: Max/min temp, dewpoint, precipitation | NCAR - Climate Data Guide.” Retrieved from https://climatedataguide.ucar.edu/climate-data/prism-high-resolution-spatial-climate-data-united-states-maxmin-temp-dewpoint

USDA, N. (2021). "Crop Progress and Condition Layers." Retrieved from https://www.nass.usda.gov/Research_and_Science/Crop_Progress_Gridded_Layers/CropProgressDescription.pdf

Wang, X., Garcia, P. and Irwin, S. H. (2014). "The Behavior of Bid-Ask Spreads in the Electronically-Traded Corn Futures Market." *American Journal of Agricultural Economics, 96*(2), 557–577. https://doi.org/10.1093/ajae/aat096

Wardlow, B. D., Kastens, J. H. and Egbert, S. L. (2006). "Using USDA Crop Progress Data for the Evaluation of Greenup Onset Date Calculated from MODIS 250-Meter Data." *Photogrammetric Engineering & Remote Sensing*, *72*(11), 1225–1234. https://doi.org/10.14358/pers.72.11.1225

Xu, X. and Zhang, Y. (2021). "Network analysis of corn cash price co-movements." *Machine Learning with Applications*, *6*, 100140. https://doi.org/10.1016/j.mlwa.2021.100140

Ying, J., Chen, Y. and Dorfman, J. H. (2019). "Flexible Tests for USDA Report Announcement Effects in Futures Markets." *American Journal of Agricultural Economics, 94*(1), 238. https://doi.org/10.1093/ajae/aaz013

Zhang, D. and Broadstock, D. C. (2020). "Global financial crisis and rising connectedness in the international commodity markets." *International Review of Financial Analysis*, *68*, 101239. https://doi.org/10.1016/j.irfa.2018.08.003

Zulauf, C., Brown, B., Orden, D., Schnitkey, G., Swanson, K., Paulson, N. and Coppess, J. (2021). "Concentration and Diversification in World Corn, Soybean, and Wheat Exports." *Farmdoc Daily*, *15*(11).

Chapter 2
Market Uncertainty and Sentiment around USDA Announcements[3]

Abstract: We investigate forward-looking commodity price volatility expectations (proxied by option-implied volatilities or IVols) around scheduled USDA reports. We show that corn and soybean IVols are significantly lower for several trading days after a report. The IVol response to a release depends on agricultural market experts' disagreement and sentiment prior to the USDA report, and on the extent to which the USDA information surprises the market. Whereas commodity IVols are generally positively related to financial-market sentiment and macroeconomic uncertainty (jointly captured by the VIX), this co-movement breaks down on report days—with the VIX and commodity IVols moving in opposite directions.

JEL classification: Q11, G14, G13, G41, Q13

Keywords*:* Commodities, Scheduled News, Forward-Looking Volatility, Surprise, Dispersion, Market Sentiment

---

[3] Chapter 2 is published as Cao, A. N. Q. and Robe, M. A. (2022). Market uncertainty and sentiment around USDA announcements. *Journal of Futures Markets*, *42*(2), 250-275. https://doi.org/10.1002/fut.22283

2.1 Introduction

A vast literature in finance documents that equity and bond markets react to many U.S. macroeconomic announcements—see, *e.g.*, Kurov *et al.* (2019). In the commodity space, similarly, prior work shows that agricultural markets react significantly to scheduled USDA announcements. The latter fact supports the notion that USDA reports contain valuable news and help resolve disagreements among market participants regarding demand and supply fundamentals.

Most of the literature investigates what happens to commodity price levels on USDA event days (*e.g.*, Adjemian, 2012; Karali *et al.*, 2019; Ying, Chen and Dorfman, 2019) or shows how fast the USDA information is impounded into those prices (*e.g.*, Adjemian and Irwin 2018; Lehecka, Wang, and Garcia 2014). As McNew and Espinosa (1994), McKenzie, Thomsen, and Phelan (2007), and Isengildina-Massa *et al.* (2008) note, however, one cannot capture the full impact of the USDA reports without also analyzing how they affect market uncertainty and sentiment.

Documenting that effect, and exploring for the first time its duration and its determinants, is our objective in this paper. We focus on corn and soybeans, because they are the two main U.S. agricultural commodities and also because their growing areas and crop (and, thus, news) cycles broadly overlap. Since changes in uncertainty and sentiment map directly into the cost of options-based strategies (Goyal and Adjemian 2021), our results have important implications not only for policy makers and academics, but also for commodity speculators and for the significant fraction of Corn Belt farmers who use options on futures to alter the exposure of a substantial part of their crop income to commodity price risk (Prager *et al.* 2020).

In equity and bond markets, financial economists and accountants have long used changes in option-implied return volatilities ("IVol") to study the impact of news on forward-looking market uncertainty (Ederington and Lee 1993; Patell and Wolfson 1979). In agricultural markets, two papers by McNew and Espinosa (1994) and Isengildina-Massa *et al.* (2008) use near-dated options-on-futures implied volatilities for the same purpose.[4] We extend that prior work along several dimensions.

First, agricultural markets have evolved massively over the course of the past two decades. Quantitatively, the open interests in corn and soybean options and futures are many times what they were 15 years ago (Robe and Roberts 2019). Qualitatively, changes

[4] In contrast to those studies and our paper, articles that look at market volatility around USDA events focus on realized volatility (captured by the variance equation in GARCH-type models or by the realized sample volatility) rather than forward-looking volatility. An exception is Adjemian *et al*. (2018), who value a missing 2013 WASDE report due to a U.S. government shutdown. Fortenbery and Sumner (1993) is the first study of option prices around scheduled USDA events. See Ying, Chen and Dorfman (2019) for a thorough review of the literature on USDA announcements.

that could materially impact the manner or the extent of market reactions to USDA news include the growth (and, later, the dominance) of electronic and high-frequency trading (Haynes and Roberts 2015; Haynes *et al.* 2017), the demise of the futures trading pits (Gousgounis and Onur 2018), and the influx of ever more sophisticated private forecasting services (Karali *et al.* 2019; McKenzie 2008). Our first contribution is to complement the early literature, which uses data from two decades ago or more (*i.e.*, predating any of those developments), by investigating agricultural IVol responses to scheduled USDA news releases in more modern times.

Intuitively, if "the timing, although not the content, of scheduled announcements is known a priori," then the IVols should already, pre-release, "impound the anticipated impact of important releases on price volatility and (should) decline post-release as this uncertainty is resolved" (Ederington and Lee 1996, p. 513). Using an event-study methodology and data for four different types of USDA announcements in 2009-2019, [5] we find that the commodity IVols fall significantly on the USDA event day—by 2 (soybean) to 2.7 (corn) percent on average in our 2009-2019 sample period. While these decreases are smaller in magnitude than those documented by Ederington and Lee (1996) in interest rate markets, they are similar to the IVol drops found for corn and soybeans two decades ago by Isengildina-Massa *et al.* (2008), using data from 1985 to 2002.[6]

For USDA reports that market observers generally view as the most important (which make up half of our sample of 151 events), we show that the average IVol drop is almost twice as large—3.6 (soybeans) or 5.7 (corn) percent. Furthermore, the IVol remains significantly lower for at least four trading days, and sometimes more than a week, after the event day. These results complement the finding of Adjemian (2012) and Karali *et al.* (2019) regarding the magnitude of commodity futures returns on USDA crop report days: they indicate that, as a group, scheduled USDA reports in recent years remain highly payoff-relevant to agricultural market participants.

Second, our analysis of commodity IVols innovates by recognizing that USDA reports are not released in a vacuum. Precisely, we use regression analyses to establish that the sign and the magnitude of the post-release IVol change depend on agricultural market experts' opinions in the run-up to a release.

Ahead of all major USDA announcements, companies like Bloomberg and Reuters have for over a decade conducted and published surveys of market analysts' expectations regarding the upcoming reports. Those news organizations typically release the details of

[5] We consider the monthly WASDE, quarterly Grain Stocks, annual Prospective Plantings, and annual Acreage reports. McNew and Espinosa (1994) look at crop production reports only; Isengildina-Massa *et al.* (2008) focus on WASDE reports.

[6] Precisely, Isengildina-Massa *et al.* (2008) find that the average magnitude of the IVol decline equals about 3 percent of the annualized IVol level on the day before the report release for corn, and 4 percent for soybeans. As documented in EL96 and in this paper, the IVol on the day before the event is higher than average.

their surveys in the week before the USDA announcement. We argue theoretically and, armed with Bloomberg survey information, provide empirical evidence that the magnitude of the grain and oilseed IVol responses to scheduled USDA announcements is significantly impacted by (i) the gap between the pre-release expert "consensus" forecast and the actual USDA figure (*i.e.*, how big the market surprise is on the event day), (ii) the dispersion of individual expert forecasts around that consensus (which captures disagreements among market experts and, as such, can be seen as a proxy for pre-existing commodity-specific uncertainty), and (iii) the pre-release expert "sentiment."[7]

Looking first at surprises, we document that their effects are significant mostly in the case of inventories-related news contained in the monthly WASDE (prospective) and quarterly (realized) Grain Stocks (GS) reports. With the GS reports, any surprise—whether "bullish" for prices (when USDA figures come in lower than analysts expected) or "bearish" (when the USDA figures are higher than analyst forecasts)—pushes IVols upward significantly (ceteris paribus). In the case of the WASDE reports, the same is true only when the surprise is price-bullish (*i.e.*, when the USDA announces lower future stock levels than the Bloomberg consensus had foreseen)—and, the more bullish the WASDE surprise is for prices, the more the forward-looking volatility increases.[8] Practically speaking, while commodity IVols generally decrease after a USDA report, the decrease is muted (so much so that the forward-looking volatility could actually go up) when the market is caught flat-footed by the USDA—all the more so when the news is bullish for prices.

Next, we look at analyst dispersion, *i.e.*, the extent to which market experts disagree about an upcoming report. Intuitively one would expect that, when experts are "confused" as a group, the USDA news would "settle the market"—resetting participants' expectations and clarifying the path forward. Indeed, for the corn WASDE, the bigger the dispersion of pre-event analyst opinions, the more the IVol drops after the USDA release. The results for other corn reports and for soybeans are statistically insignificant, however, which suggests that dispersion matters less than surprises in agricultural markets.

---

[7] Two recent studies of USDA announcement also consider the possible roles of analyst expectations and uncertainty. Karali *et al.* (2019) use a DCC MGARCH-X model to investigate the role of report surprises on price levels and on the realized variance of agricultural commodity returns. We examine instead the link between prior expert opinions market expectations of future volatility. We look not only at whether expert forecasts were close to the actual release, but also at the extent to which analysts disagreed and at their sentiment (pessimistic or optimistic) prior to the news release. Further afield, Fernandez-Perez *et al.* (2019) examine the link between consensus forecast error and analyst dispersion on futures bid-ask spreads (which acts as a proxy for asymmetric information). Both of those recent studies posit that the price or bid-ask spread changes after USDA announcements can be solely attributed to the reports' informational value, whereas we also control for (i) changes in macroeconomic uncertainty and financial market sentiment around USDA events and (ii) physical market conditions in the runup to the event.

[8] Price-bearish WASDE surprises, in contrast, do not statistically significantly modify the typical commodity IVol response to a scheduled announcement.

Finally, we tease out how the pre-event expert sentiment influences the market's reaction to USDA reports. For each report, we rate the prevailing analyst consensus as "pessimistic" (*resp.* "optimistic") when the median pre-report expert forecast predicts a decrease (*resp.* an increase) in the forecasted USDA variable compared to an objective past reference point. We find, in the case of Grain Stocks reports (but not other USDA reports), a significant negative association between the analyst pessimism prior to the report and the IVol drop on the announcement day. Since we already control for fundamentals-related news (*i.e.*, USDA surprises) and uncertainty (*i.e.*, analyst dispersion) when running the analysis, this result indicates that, when experts had been pessimistic about the actual level of grain inventories, the release of the GS information by the USDA has a stronger market-calming effect. This finding is a novel contribution to a large literature showing the importance of inventories to commodity price dynamics – see, *e.g.*, Bobenrieth *et al.* (2021).

Our third contribution is to show the importance—when assessing the impact of USDA news on agricultural market uncertainty and sentiment—of also controlling for concomitant (*i.e.*, event-day) changes in broad financial market uncertainty and sentiment. Bekaert, Hoerova, and Lo Duca (2013) show that the VIX index (*i.e.*, the Standard and Poor 500 equity-index option-implied volatility) captures jointly the uncertainty about global macroeconomic conditions and the risk aversion among investors. Intuitively, the same should be true in agricultural markets. In essence, insofar as risk aversion affects all asset classes, risk aversion levels in commodity markets should move at least partly in sync with equity-market risk aversion. In the same vein, given that the demand for physical commodities reflects the strength of the economy, uncertainty about the latter should also percolate into agricultural markets.

Consistent with this intuition, Adjemian *et al.* (2017) show that, in the long run, changes in grains and livestock IVols are driven to a significant extent on a day-to-day basis by changes of the VIX index in the same direction. The question we ask here is whether a similar pattern is seen on USDA announcement days—and, thus, if controlling for the VIX helps separate the respective impacts of global *vs.* commodity-specific market uncertainty and sentiment. Surprisingly, we find that the IVol change on USDA announcement days is statistically significantly negatively related to the VIX change on that day. That is, while prior research shows that commodity market sentiment and uncertainty generally move in the same direction as the VIX, we show that this overall pattern is reversed on days when USDA announcements take place. Ceteris paribus, if the VIX increases on the event day, then the IVol drops more that day—and vice-versa.

The chapter proceeds as follows. Section 2.2 extends Ederington and Lee's (1996) theoretical model of market reactions to scheduled announcements, and draws on other literature, to derive testable hypotheses. Sections 2.3 describes the data. Section 2.4 discusses our empirical methodology. Section 2.5 presents our empirical findings. Section 2.6 concludes and discussions possible extensions.

2.2 Theory and hypothesis development

We extend the Ederington and Lee (1996, EL96 for short) model to guide our study of how USDA reports should affect forward-looking commodity market uncertainty and sentiment.

2.2.1 *Predicted IVol change before and after a scheduled USDA announcement*

As noted in the Introduction, EL96 show theoretically that the expectation of future return volatility embedded in the price of a given option (IVol) should, ceteris paribus, first rise in the days leading to a scheduled news release and then fall in the latter's aftermath. In what follows, we adapt their model to predict IVol patterns when using constant-maturity options.

In the EL96 model, the implied variance (IMV) on a given day *t* is the average of the daily expected variances over the remaining life of a given option, starting from day *t+1*. Option traders form expectations using all information available up to day *t*. Thus, the IMV change on the scheduled report day, say *T*, is the sum of two changes:

(i) removing the expected realized variance on the event day *T* from the set of days (until the option's expiration) that was used to calculate the IMV on day *T-1* (because day *T* is now the current day and thus no longer "expected");

(ii) revising the expectation of volatility (or variance) of all the option's remaining days to expiration, starting on day *T +1*.

Regarding the second term (ii) above, EL96 argue that, depending on whether the realized day- *T* volatility is higher or lower than had been expected at *T -1*, market participants will revise upward or downward their expectations of what future realized volatility will be until the option's expiration. However, "rational expectations imply that (...) upward and downward revisions are equally likely and the mean revision across many such scheduled announcements should be approximately zero" (EL96, p.517). We will return to this component in Section 2.2.2; for now, we can focus on the first term.

For the first term (i) above, the core assumption of the EL96 model is that, if there is a scheduled announcement on day T, then on day *T-1* market participants should expect that asset returns will be more volatile than average on day *T*. The intuition is that prices should move a lot to react to the new information, an assumption that is indubitably borne out in the commodity space.[9] Thus, if the maturity date of the option (whose price EL96 use to extract volatility expectations) is fixed, then going from *T-1* to T means removing a higher-than-normal volatility day from the expectation set in (i), which makes the resulting unconditional average IMV fall on day *T*. The same mechanics drive an increase in the IMV in the runup to the event day.

Unlike EL96, we rely for our empirical analysis on constant-maturity (90-day synthetic) options rather than on nearby options (whose time-to-maturity, as in EL96, would instead decrease over time). With constant-maturity options, it is straightforward to show analytically that (keeping the other assumptions of the EL96 model unchanged) the IVol does not increase as the event day approaches but that it still drops in the aftermath the scheduled news release.[10]

In the EL96 model and in our variant thereof, there is no theoretical reason why the post-event IVol decrease should be a one-day affair. First, note that USDA news are incorporated into prices promptly (Adjemian and Irwin, 2018), so the realized volatility increase on which EL96 focus is limited to the event day *T*. Second, insofar as the USDA reports convey large amounts of information to agricultural market participants (Adjemian, 2012), one can show that a given report's impact on part (i) of the IMV in the EL96 model should last for several days—until either new, non-USDA information is released or until a new USDA event day is included in the average (i). Our first testable hypothesis is thus straightforward:

Hypothesis 1: On average, commodity IVols fall on scheduled USDA report days. This decrease remains statistically significant for several business days and is greater for major USDA reports. There is no IVol increase in the run-up to a USDA report.

2.2.2 *Pre-existing commodity-market beliefs and IVol response to USDA news*

As noted in the Introduction, the present paper is the first to ask whether agricultural IVols' responses to scheduled USDA announcements depend on the extents to which market participants are surprised by the information and to which, before the event, experts disagreed about the upcoming release (a proxy for commodity-market uncertainty) and were pessimistic (a proxy for commodity-market sentiment).

[9] For example, Janzen and Bunek (2017) show that the realized volatility of intraday winter wheat futures prices shoots up right after the USDA reports.
[10] One exception is when the 90th day added is also an event day, in which case the IVol should be unaffected by the news release on average. There are very few such cases in our sample. While the low number of such observations makes it difficult to control for this caveat, including those few observations in the sample should bias against our finding an IVol drop on USDA days.

(1) Market consensus and USDA surprise

The "surprise" is the deviation of the information in the scheduled announcement from *pre*-event market expectations. It is the unanticipated shock that leads market participants to revise their expectations *post*-event.

As noted in Section 2.2.2 above, the surprise is zero on average in the EL-96 model, and so part (ii) of the IMV change on the event day *T* is zero. Equity, bond, and forex markets, however, tend to react asymmetrically to "good" *vs.* "bad" news. For example, Braun, Nelson, and Sunier (1995) find significant predictive asymmetry in both the market-wide and the firm-specific components of volatility for various stock portfolios. In a real-time analysis of U.S. dollar spot exchange rates, Andersen *et al.* (2003) report larger surprise-induced conditional-mean jumps when the surprise is bad, compared to the good surprise case. In the same vein, Beber and Brandt (2010) investigate the respective effects of good *vs.* bad macroeconomic news in the U.S. treasury bond market: they find that bond returns react more strongly to bad news than to good news during expansions, and vice-versa during recessions.

There is no reason to believe that commodity markets are any different. One can readily extend the EL96 model to account for the fact that one can sign the surprise insofar as (a) one has data about market expectations regarding the upcoming USDA news and (b) lower-than expected inventories or tighter-than-expected supply/demand balances should boost realized price volatility. There is a long line of research showing theoretically and empirically that commodity prices are more volatile during a scarcity phase than amid conditions of plenty—see, *e.g.*, Geman and Smith (2013) and references cited therein. In the same vein, the theory of storage (Kaldor 1939; Working 1948) states that high levels of commodity inventories help smooth out the impact of demand and/or supply shocks on commodity prices and therefore smooth out price volatility.[11] Defining price-bullish surprises as "tighter commodity supply and/or inventories or higher demand than expected, which should boost prices and volatility," and price-bearish surprises as the opposite, one can extend the EL96 model to derive the following prediction for the event-day IVol change conditional on the surprise:

Hypothesis 2: The IVol response to the USDA news depends on market participants' pre-release expectations. In case of a "bearish" surprise, the IVol should drop more on the event day than it would absent a surprise (*i.e.*, if the market's prior expectations had been met by the content of the announcement). In case

[11] See Baur and Dimpfl (2018) for a recent summary of that literature.

of a "bullish" surprise, the IVol should drop less (and could even increase) post-release than it does on average.[12]

(2) Forecast dispersion

Intuitively, the magnitude of the IVol change after a scheduled USDA report should depend on the level of disagreement between commodity market participants before that announcement. Studies in the equity space document a positive relationship between analysts' forecast dispersion and stock price volatility around firms' earnings announcements. A possible explanation is that forecast dispersion among analysts represents idiosyncratic risk: given that analysts are experts at their forecasted subjects, a high level of dispersion likely reflects uncertainty regarding the subject (Dubinsky *et al.* 2019; Johnson 2004).[13] If this view is shared by those who predicts future price volatility, then we should also expect that the pre-event volatility expectation for the event day T of the IMV in the EL96 model is larger when analyst forecasts are more highly dispersed, causing the IVol to drop more after the USDA event day T is removed from the averaging set—component (i) in Section 2.2.1 above—once the announcement has taken place. Therefore, we have:

Hypothesis 3: Ceteris paribus, the IVol change after a USDA information release is inversely related to the pre-release dispersion of analyst forecasts.

(3) Commodity-specific market sentiment

The EL96 model is predicated upon the Rational Expectation Hypothesis, so that a change in volatility expectations can only be explained by the arrival of new fundamental information. In contrast, the behavioral economics literature suggests that changes in "market sentiment" can also cause a volatility reaction and that sentiment's effect on market volatility may be asymmetrical: in a seminal paper, Barberis, Shleifer and Vishny (1998) develop a theory where "representativeness bias" causes investor to project the most recent news into their future expectation. As such, a negative piece of news is likely to be followed by other negative news, which could imply that more uncertainty ought to be expected for the future (and vice versa).

[12] Hypothesis 2 predicts that the effect of a USDA surprise is asymmetrical. An alternative intuition, which is outside the scope of the EL96 model, is that the IVol response is instead symmetrical: that is, if the market always becomes unsettled whenever it is surprised then, the bigger the surprise is, the smaller the IVol drop should be (or, in extreme cases, the commodity IVol could even increase after any surprise).

[13] Another view brought by the difference-of-opinion school of thought (*e.g.*, Diether, Malloy and Scherbina 2002), posits that forecast dispersion is a result of diverging opinions among market participants, which brings about mispricing once short-sale constraints arise on the market. In the case of commodity futures markets, however, this argument seems moot since there are no short sales constraints.

Proxying commodity-specific sentiment by the degree of analyst optimism/pessimism about the upcoming announced information, we hypothesize that if market analysts are pessimistic before the USDA announcement day *T* about the supply/stock situation, then market participants should expect higher volatility on day *T* (compared to what it would be in the EL96 model), causing the IVol to drop more after day T is removed from the averaging set on the announcement day (*i.e.*, from component (i) in the EL96 model). Therefore, we have:

Hypothesis 4: After controlling for analyst surprise and dispersion, the magnitude of the IVol change depends on pre-release commodity-market analyst sentiment.

2.2.3 *Global macroeconomic environment and commodity IVol response on event day*

Hypotheses 3 and 4 look, respectively, at the possibilities that the pre-announcement commodity-market uncertainty and sentiment could impact the IVol response to the USDA news. In this Section, we turn to the possibility that changes in the macroeconomic environment on the event day T itself may also matter for the commodity IVol behavior that day.

Recent empirical work finds that, for a wide range of commodities, IVols are impacted by the VIX: when the VIX increases, IVols go up—and vice-versa.[14] In order to tease out the impact of scheduled USDA events on corn and soybean IVols, one should therefore control for concomitant changes in the macroeconomic and financial environments:

Hypothesis 5: The IVol response to the USDA news depends on the VIX return on the event day, *i.e.*, on concomitant changes in broad financial market uncertainty and sentiment.

2.3 Data

We examine four groups of scheduled USDA announcements: monthly WASDE, quarterly Grain Stocks (GS), and annual Prospective Plantings (PP) and Acreage (AR) reports. Those are the main reports published by the USDA about the global grain and oilseed markets.

These four sets of reports are released on 15 different USDA announcement days per year (except in 2013 and 2019, when there were only 14 announcement days per year due to U.S. government shutdowns). From September 2009 to October 2019, there are 120 WASDE reports, 41 GS reports (of which 10 overlap with the January WASDE), 10 PP

[14] See Robe and Wallen (2016) in the crude oil space. See Covindassamy, Robe, and Wallen (2017) in the softs space. And see Adjemian *et al.* (2017) and Goyal and Adjemian (2021) in the livestock and grains spaces.

reports, and 10 AR reports. Altogether, we collect a sample of 151 USDA announcement days and the corresponding Bloomberg surveys for 181 reports in total. Table 2.1 synopsizes the characteristics of the reports, including their coverage, frequency and timing, and key information surveyed by Bloomberg.

Table 2.1. USDA Reports Overview

	WASDE	Grain Stocks (GS)	Prospective Plantings (PP)	Acreage (AR)
Frequency	Monthly	Quarterly	Yearly	Yearly
Timing	2nd week of the month	2nd week of January & End of 1st-3rd Quarters	End of March	End of June
Overlap	1st GS (January)	1st WASDE; PP; AR	2nd GS (March)	3rd GS (June)
Information surveyed by Bloomberg	Projected U.S. ending stock of the on-going marketing year	U.S. Ending stock estimates as of 1st Dec, 1st Mar, 1st Jun and 1st Sep	U.S. farmers' planting intention for upcoming crop season	Survey-based estimate of U.S. planted area for current crop season
Baseline for Forecast "Pessimism"	WASDE of previous month	GS of previous year's same quarter	AR of previous year	PP of current year

Note: Table 2.1 describes the 151 USDA reports that we collect for our sample from September 2009 through October 2019. On some dates, the USDA releases more than one report: the third row in the table (labeled "Overlaps") explains which of the WASDE, GS, PP and AR reports overlap. For part of the empirical analysis (see Table 2.5), we include information regarding expert opinions prior to the USDA news release. The information regarding analyst opinions comes from periodic Bloomberg surveys of market experts. The last row of the table indicates the baseline that we use to characterize whether the analyst consensus about upcoming news is optimistic or pessimistic, as explained in Appendix 2.A.1

Since September 2009, Bloomberg has conducted analyst surveys prior to each of these reports. Results of the surveys are released at varying times on Bloomberg News, typically one week before USDA release. The exact timing of the result release is not documented in the survey dataset, so we recover it by tracing back each release on Bloomberg News manually to define the event window for our analysis.

Our Bloomberg survey dataset contains detailed information about the forecasters who participated in each survey. A typical survey summarizes the opinions of about 20 commodity analysts regarding an upcoming USDA announcement. This information allows us to assess the distribution of analyst forecasts and to compute both a “consensus” value (which we set as the median analyst forecasts) and the dispersion of individual analyst forecasts around the consensus.

A few of the USDA reports overlap. Specifically, the PP (March) and AR (June) reports are released together with the second and third GS reports, respectively. As well, the January WASDE and GS reports are released simultaneously. The latter overlap might seem problematic, in that both reports contain information on grain stocks. However, the nature of the information in the two reports is different: the January GS report provides information on what actual stocks were as of December 1st of the previous year, whereas the January WASDE estimates what future stocks should be at the end of the current marketing year.[15]

Since we are interested in forward-looking volatility, we use the constant 90-day IVol for CBOT corn and soybean. To match this maturity choice, we likewise use the CBOE’s constant 90-day Volatility index (VIX) to test Hypothesis 3. All market series, such as the daily VIX, commodity IVols and futures prices, as well as data on USDA announcements and analyst surveys, are retrieved from Bloomberg.[16]

2.4 Methodology

In this Section, we describe the testing strategies for our hypotheses and the construction of the variables needed for that purpose. With Hypothesis 1, we focus on statistical hypothesis testing with the IVol sample around USDA announcements. We examine Hypotheses 2 to 5 using multivariate regressions.

[15] In addition to the WASDE, PP, AR, and GS reports, the U.S. Department of Agriculture also releases NASS Crop Production (CP) reports simultaneously with the WASDE. The NASS CP and the WASDE methodologies to produce crop production estimates are different—see Good and Irwin (2011). Still, the WASDE future stock projections reflect, in part, information about potential crop size (Isengildina-Massa *et al*., 2021). For this reason, our empirical analysis (precisely, the part dealing with the impact of analyst surprise, dispersion, and sentiment on commodity IVol responses to USDA announcements) focuses on the WASDE reports.

[16] A Bloomberg document authored by Cui (2012) details that company’s methodology for extracting forward-looking volatility estimates from at-the-money option prices at the daily market close. Ederington and Guan (2002) and Yu, Lui, and Wang (2010) discuss some of the technical advantages of relying on Bloomberg implied-volatility estimates. One major advantage, in the opinion of the present paper’s authors, is that it makes the analyses easily reproducible.

2.4.1 *Testing Hypothesis 1: Commodity IVols decrease on the announcement day*

(1) Event-day testing. As a first step, we compare the mean and median IVols on the event day *T* against the corresponding values on day *T-1*. Following Isengildina-Massa *et al.* (2008), we use both a parametric paired sample t-test and a nonparametric Wilcoxon signed rank test to account for the nonnormality of the distribution of IVol changes. Denoting the IVol levels on days *T* and *T-1* respectively as $Ivol_T$ and $Ivol_{T-1}$, the common null-hypothesis of these two tests is:[17]

$$H_0\text{: } Ivol_T \geq Ivol_{T-1} \qquad \text{against} \qquad H_1\text{: } Ivol_T < Ivol_{T-1}$$

(2) Event-window extension. Moving beyond the event-day IVol change, we seek a broader picture of how option-implied volatilities behave for five days on either side of the event. Our approach is to perform multiple comparisons of the IVol changes within the event window from a pre-event-window reference. By doing so, we can learn about the timing of any change in the commodity IVol, as well as how persistent these changes are.

Extensions of the t-test and Wilcoxon test that allow comparisons of more than two samples include the parametric one-way ANOVA test to compare group means, and the nonparametric Kruskal-Wallis test to compare group medians. However, they only test the null that all group means/medians are equal, *i.e.*, H_0: $\Delta Ivol_{T-5} = \Delta Ivol_{T-4} = ... = \Delta Ivol_{T+5}$, against the alternative that at least one group has statistically a significantly different mean or median. Without further analysis, it is not possible to know whether each group's mean or median differs from the others. Therefore, we perform multiple comparison procedure using the Turkey-Kramer method based on the result of one-way ANOVA and Kruskal-Wallis test.[18]

(3) Event-window and pre-event-window reference. To capture possible differences between the pre- and post-event IVol change patterns, we consider a window of 5 days before and 5 days after the USDA announcement day. A natural baseline reference to assess within-window IVol changes would be the period just before that 11-day window around the event. One complication is that, because the timing of the Bloomberg analyst surveys varies from one to seven days before a USDA announcement, there can be an overlap between the post-Bloomberg and the pre-USDA periods. To avoid such overlaps,

[17] The difference between the two tests is that the one-sided *t*-test assumes that $\Delta Ivol_T$ (*i.e.*, $Ivol_T - Ivol_{T-1}$) follows a normal distribution with mean 0 and unknown variance under the null-hypothesis, while the Wilcoxon signed rank test only assumes that $\Delta Ivol_T$ is drawn from a continuous distribution that has 0 median and is symmetric around this median under the null. For a detailed description of these two tests, see Isengildina-Massa *et al*. (2008).

[18] An important motivation for using multiple comparisons (rather than simultaneously applying t-tests to every pair of samples) is that the rate of type-I error will be inflated in proportion to the number of pairs of groups being compared simultaneously. Consequently, we can no longer be sure that the probability of incorrectly rejecting the null is no larger than the specified α (Hochberg and Tamhane 1987). The Turkey-Kramer procedure is designed to circumvent this issue by using a studentized range distribution, and adjust the p-values of the pairwise test-statistics accordingly. See, *e.g.*, Stoline (1981) for a review of multiple comparison methods, including the Turkey-Kramer procedure.

we choose as the reference (*i.e.*, baseline) IVol the 5-day average IVol before the Bloomberg survey is released, denoted $\overline{Ivol}$. Figure 2.1 illustrates the timeline and the overlap.

For each day in the window around the event day *T*, we calculate the percentage IVol change as

$$\Delta Ivol_{T+i} = ln\left(\frac{Ivol_{T+i}}{\overline{Ivol}}\right), \text{ where } i = -5, -4, \ldots 5 \tag{2.1}$$

We first apply one-way ANOVA and Kruskal-Wallis tests to see if there is at least one day in the event window when the mean or median $\Delta Ivol_{T+i}$ differs significantly from the others. If the test fails to reject the null, then no further action is needed. Otherwise, we feed the resulting estimated mean (or median) and standard errors into the Turkey-Kramer procedure to compare all possible pairs of $\Delta Ivol_{T+i}$ and $\Delta Ivol_{T+j}$.

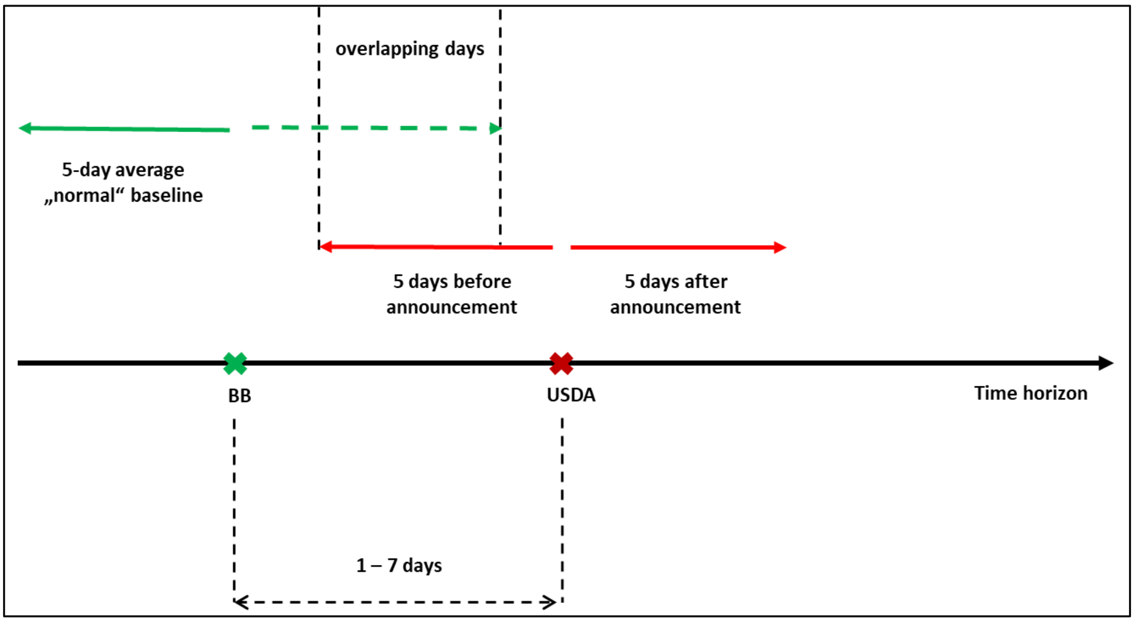

Figure 2.1. Timing of Bloomberg Analyst Surveys and Scheduled USDA Announcements

2.4.2 *Testing Hypotheses 2 to 5: Determinants of the IVol drop*

We regress the event-day commodity IVol change on a set of Bloomberg-survey-related variables (see Hypotheses 2 to 4 in Section 2.2.2), on the VIX return (our proxy for the event-day change in macroeconomic uncertainty and financial market sentiment—see Hypothesis 5 in Section 2.2.3), and on additional control variables (see Item 5 below).

Due to the partial overlaps in the four different reports' respective release schedules, we consider the impact of the four reports on commodity IVols simultaneously.

The information of interest in PP and AR reports is similar in nature (acreage expected to be or actually planted in the current year). Effectively, the AR report is an updated version of the PP report released earlier that year. Thus, we treat the PP and AR reports as a single type of reports, which we call "Planted Area" (denoted PA). This approach has two advantages: it cuts the number of right-hand side variables, and thus increases the number of degrees of freedom; and it helps reduce multicollinearities among the surprise, dispersion, and sentiment variables related to those two annual reports, a problem that stems from their low frequency and the large proportion of observations for these variables that simultaneously have a null value.

Formally, we estimate the following relationship:

$$\Delta Ivol_i = \beta_0 + \sum_j \beta_j S_{i,j} + \sum_j \delta_j D_{i,j} + \sum_j \gamma_j Sentiment_{i,j} + \varphi \Delta VIX_i + \eta Control_i + \varepsilon_i \tag{2.2}$$

where $i = 1,\ 2, \ldots,\ 151$ denotes the i^{th} event day in our eleven-year sample; $j = \{WASDE, GS, PA\}$ denotes the type of report; and, the Δ operator denotes daily close-to-close returns (log difference) on the event day *T* from the previous day *T-1*.

Our variables of interest include:

(1) Surprise, $S_{i,j}$. As in Couleau *et al.* (2020), we use the median Bloomberg analyst forecast as a proxy for the consensus market expectations prior to a USDA announcement. For report j, on the i^{th} event in our sample, we define the "report surprise" as the percentage difference (approximated as a log difference) between the USDA's announced value $A_{i,j}$ and the median forecast value $F_{i,j}$ in the corresponding survey:

$$S_{i,j} = ln\left(\frac{A_{i,j}}{F_{i,j}}\right) \tag{2.3}$$

As discussed in Section 2.2.2, we split the surprises into "bullish" *vs.* "bearish" surprises. A price-bullish surprise $S^-_{I,j}$ occurs if the USDA announces lower stocks (WASDE, GS) or acreage levels (PP, AR) than had been forecasted by the market consensus (hence the negative superscript in our notation); a price-bearish surprise S^+_{ij} captures the opposite situation.

$$S^+_{i,j} = \begin{cases} S_{i,j}, \ if \ S_{i,j} > 0 \\ 0, \ otherwise \end{cases} \quad ; \tag{2.4}$$

and

$$S_{i,j}^{-} = \begin{cases} S_{i,j}, \ if \ S_{i,j} < 0 \\ 0, \ otherwise \end{cases}$$

The regression Equation (2.2) thus becomes:

$$\Delta Ivol_i = \beta_0 + \sum_j \beta_j^- S_{i,j}^- + \sum_j \beta_j^+ S_{i,j}^+ + \sum_j \delta_j D_{i,j} + \sum_j \gamma_j Sentiment_{i,j} + \varphi \Delta VIX_i + \eta Control_i + \varepsilon_i \tag{2.5}$$

By comparing the signs and magnitudes of β_j^- and β_j^+, we can test whether there is an asymmetry in the reaction of grain and oilseed volatility expectations to USDA surprises.

(2) Dispersion $D_{i,j}$. For each forecasted piece of information, we follow prior work—see, *e.g.*, Fernandez-Perez *et al.* (2019) and references cited therein—and calculate dispersion as the ratio of the interquartile range (IQR) to the mean forecast:

$$D_{i,j} = \frac{IQR_{i,j}}{\mu_{i,j}} \tag{2.6}$$

This approach avoids issues related to outliers, unlike the alternative methodology of using the standard deviation of analyst forecasts as a dependent variable.

(3) VIX changes, ΔVIX_i. Grain and oilseed markets are much smaller than equity markets, so we treat the VIX as an exogeneous variable for the purposes of this study.

(4) Expert $Sentiment_{I,j}$. Having controlled for forecasters' expectation (through the surprise), pre-existing commodity-market uncertainty (through dispersion) and global market uncertainty and sentiment (through the VIX), we can test whether the IVol drop on the event day is related to other non-fundamental factors, *i.e.*, to commodity-market "sentiment." We take the "pessimism" of forecasters about an upcoming report as a form of prior market sentiment.[19] We rate a consensus forecast as "pessimistic" when the median predicts a decrease in the forecasted indicator from a reference point. When it predicts an increase, we rate it as "optimistic".[20] To keep things simple, we set

[19] This approach is related to the concept of *"forecast change"* pioneered by Amir and Ganzach (1998). In a corporate finance context, these authors show that the sign of the "forecast change" (defined as the difference between the analysts' earnings forecasts and the previous actual earning of a company) is a significant predictor of the over- or under-reaction in forecasts. Thus, if we find that the pessimistic/optimistic tenor of the market experts' forecasts statistically significantly affects the extent of the USDA-induced IVol drop, then it would be a sign that market sentiment plays a role in how the market reacts to the announcement.

[20] It is important to note that forecast pessimism and forecast surprise need not have the same sign. For instance, the surprise can be "price-bearish" when the USDA releases less "bad" information than what the analysts had predicted.

$Sentiment_{I,j}$ equal to 1 if the median of the analyst forecast for report j released on the i^{th} event day in our sample is pessimistic, and 0 otherwise. The last row of Table 2.1 lists the reference point used to rate the sentiment for each type of report; Appendix 2.A.1 provides additional details.

(5). Regression Equation (2.5) specifies a vector of control variables including day-of-the-week dummies, seasonal dummies, as well as a set of lagged daily returns:

a. *Seasonality*. Every year, IVols in the U.S. corn and soybean markets generally start increasing from April to June, which coincides broadly with the U.S. planting phase (see Appendix 2.A.2 for a visual illustration). To capture this seasonal pattern, we use dummies corresponding to the main development phases of the U.S. crop cycle: planting (April through June), pollination (July and August), and harvest (September through November). The baseline season is the period when the land lays fallow (*i.e.*, December through the following March).
b. *Day-of-the-Week.* We control for the possibility that the IVol reaction to a USDA announcement might differ depending on which specific day of the week the release takes place, by including four weekday dummies (Tuesday to Friday).
c. Lagged returns: for each commodity, we include 1-day lagged returns (using closing prices two days and one day before the i^{th} event day) for the nearby commodity futures, the 90-day commodity IVols, and the VIX.

2.5 **Results**

In this Section, we first provide a summary of the data before presenting the results of our empirical analyses.

2.5.1 *First look at the data*

Table 2.2 reports summary statistics for our main variables of interest, including the levels and returns for the commodity IVol and the VIX, the analyst surprises and dispersions, and the percentage forecast changes (FC)—precisely, the log difference between the median Bloomberg forecast and the corresponding reference point that we use to determine our sentiment variables (see Item 4 in Section 2.4.2). The Table provides values for medians, means, standard deviations (SD), minima and maxima, as well as (in the last column) the counts of negative observations.

There is a clear pattern: the median and the mean of the daily IVol return are negative on the announcement day, for both corn and soybeans. USDA event days with negative IVol returns make up more than three-quarters of the whole sample for each commodity and, across all 151 USDA event days, the IVols fall by 2 (soybeans) to 2.7 (corn) percent on average.

Not all USDA reports are equally important: only a good half of the 151 USDA reports in our sample are considered by market observers to be "big events."[21] On "big" days, the proportion of event days with negative IVol returns jumps from three quarters to almost six sevenths, and the median IVol drop is almost twice as large, averaging 3.6 (soybeans) to 5.7 (corn) percent.

The event-day market surprise is small on average (in absolute terms, less than 0.4 percent of the median forecast for corn and less than 1.1 percent for beans) but its standard deviation is very large. In the case of corn, surprises tend to be bearish for prices and volatility (with all three types of USDA reports), whereas they tend to be bullish for prices and volatility in the case of soybeans.

The pre-event median dispersion of expert forecasts is widest for WASDE reports (11.1 percent of the average forecast for soybeans, and 6.5 percent for corn), followed by forecasts of the quarterly GS reports. The annual PA forecasts exhibit the least dispersion. These patterns hold for both corn and beans.

Turning to expert sentiment ahead of USDA reports, the analyst forecasts are mostly optimistic (the only exception is the corn PA analyst forecasts), *i.e.*, analysts tend to anticipate higher levels of grain stocks or planted areas compared to the (previous) reference point. The forecast change is largest on average in the case of the quarterly GS reports, for both commodities. The magnitude of the change, however, is generally small.

Finally, both the mean and median of the VIX return are small on USDA event days, although the return's standard deviation is large in the sample.

[21] Adjemian & Irwin (2018) list the "big events" for corn and soybeans as the WASDE reports released in January, August, September, October, and November, as well as all the Grain Stocks, Prospective Plantings, and Acreage reports.

Table 2.2. Summary Statistics

	Median	Mean	SD	Min	Max	No. Obs	Obs < 0
VIX	17.65	18.99	5.15	11.85	43.38	2,567	N/A
VIX daily returns (*USDA event days only*)	-0.009	-0.003	0.047	-0.109	0.184	151	88
A. Corn							
IVol, daily	25.68	26.06	7.06	11.41	45.09	2,567	N/A
IVol, all USDA event days †	25.41	25.80	6.73	11.73	43.34	151	N/A
IVol, big-event days ‡	23.84	25.13	7.03	11.73	43.34	82	N/A
Ivol daily returns, all days	-2.8e-4	-2.9e-4	0.03	-0.41	0.36	2,567	1,295
Ivol daily returns, all USDA event days	-0.027	-0.032	0.053	-0.231	0.198	151	119
Ivol daily returns, only big-event days	-0.057	-0.045	0.063	-0.231	0.199	82	69
WASDE surprise	0.004	0.006	0.077	-0.242	0.326	121	52
Grain Stocks surprise	0.002	0.011	0.068	-0.165	0.196	41	20
Planted Area surprise	0.004	0.007	0.018	-0.017	0.055	20	8
WASDE dispersion	0.065	0.083	0.058	0.006	0.253	121	N/A
Grain Stocks dispersion	0.021	0.029	0.024	0.009	0.131	41	N/A
Planted Area dispersion	0.008	0.009	0.004	0.005	0.022	20	N/A
WASDE forecast change§	0.000	-0.005	0.169	-0.621	1.008	121	58
Grain Stocks forecast change	0.010	0.010	0.189	-0.558	0.376	41	15
Planted Area forecast change	-0.001	-0.004	0.026	-0.067	0.043	20	10

Table 2.2 (cont.). Summary Statistics

	Median	Mean	SD	Min	Max	No. Obs	Obs < 0
B. Soybean							
IVol, daily	20.43	20.72	4.63	10.87	37.23	2567	N/A
IVol, all USDA event days†	19.98	20.35	4.40	11.05	31.95	151	N/A
IVol, big-event days‡	19.84	20.03	4.44	11.36	31.95	82	N/A
Ivol daily returns, all days	-0.001	-4.2e-4	0.032	-0.301	0.258	2567	1337
Ivol daily returns, all USDA event days	-0.020	-0.022	0.045	-0.153	0.210	151	114
Ivol daily returns, only big-event days	-0.036	-0.032	0.048	-0.153	0.155	82	67
WASDE surprise	0.000	0.000	0.101	-0.310	0.452	121	55
Grain Stocks surprise	-0.011	0.001	0.091	-0.346	0.265	41	26
Planted Area surprise	-0.004	-0.008	0.021	-0.078	0.034	20	15
WASDE dispersion	0.111	0.125	0.076	0.011	0.401	121	N/A
Grain Stocks dispersion	0.036	0.047	0.030	0.012	0.118	41	N/A
Planted Area dispersion	0.011	0.011	0.006	0.005	0.025	20	N/A
WASDE forecast change§	0.000	0.007	0.146	-0.357	0.747	121	60
Grain Stocks forecast change	0.077	0.093	0.298	-0.623	0.821	41	11
Planted Area forecast change	0.009	0.009	0.022	-0.041	0.053	20	4

Note:
Table 2.2 provides summary statistics for the main variables used in our analysis, including the event-day "surprise" relative to analysts' consensus forecast prior to the event and the "dispersion" of those forecasts around the consensus. The last three rows of each panel are the changes compared to the baselines used to capture analyst sentiments ahead of the forecast. The sample runs from September 2009 through October 2019 and covers 151 USDA reports in that period—see Table 2.1.

† Computed for all 151 Grain Stocks, Prospective Planting, Acreages and WASDE announcement days in the sample.

‡ Corn and soybean "Big-event" days include the WASDE reports in January, August, September, October, and November (not any other), as well as all Grain Stocks, Prospective Plantings, and Acreages report—see Adjemian & Irwin (2018).

§ "Forecast change" is the log difference between (a) the value forecasted by the analysts interviewed by Bloomberg for the upcoming USDA announcement and (b) the reference value. See Table 2.1 for a summary and Appendix 2.A.1 for details.

2.5.2 *Hypothesis 1: IVols decrease on average following a scheduled USDA announcement*

In the first two columns of Table 2.3, we report the test-statistics for one-sided t-test and Wilcoxon signed rank test. For both the corn and the soybean markets, the null can be rejected with a high level of confidence, *i.e.*, there is a statistically significant commodity IVol drop on the USDA announcement day.

In last two columns of Table 2.3, one-way ANOVA and Kruskal-Wallis tests show that, for both commodities, there are at least two days in the 11-day event window whose $\Delta Ivol$ values (where the log difference Δ is computed by reference to the 5-day average IVol value before the latest pre-event Bloomberg survey of agricultural market analysts) are significantly different from each other.[22] We therefore perform the multiple comparison procedure described in Section 2.4.1.

The results of these multiple-comparison tests are visualized in Figure 2.2(a) (corn) and Figure 2.2(b) (soybeans). Table 2.4 reports the *p*-values of the test statistics. For both commodities, the patterns of the log IVol differences between (a) each of the 11 days around the USDA announcement and (b) the 5-day average or "normal" IVol prior to the Bloomberg survey are clearly dissimilar:

- In general, corn and soybean IVols are higher than "normal" on the five days leading up to the announcement but, as predicted by Hypothesis 1, the increase is never statistically significant.[23]
- In sharp contrast to that pre-event behavior, commodity IVols fall significantly on the USDA event day and they remain statistically significantly lower than "normal" for at least four trading days thereafter. Figure 2.2(a) and Figure 2.2(b) show that, for both commodities, the IVol gradually reverts toward its "normal" level. This visual observation is confirmed by the one-sided t-test, as shown in the first column of Table 2.4.[24]

In short, the empirical evidence supports Hypothesis 1. Our above results extend to the past decade early findings in Isengildina-Massa *et al.* (2008) and McNew and Espinosa (1994), that IVols drop significantly on the day of a USDA report release. More importantly, we extend those previous finding by showing that commodity IVols trend upward (though not statistically significantly) for several days before the USDA

[22] For corn, the statistical significance holds for all tests and for all events – small and big. For soybeans, the statistical significance is strongest for the subset of 87 USDA reports that market observers rank as "most important"—see Adjemian and Irwin (2018).
[23] The corn IVol gradually increases for four days before the announcement and reaches its highest *pre*-event level on the day before the USDA event day. The soybean IVol does not exhibit any visible change from the "normal" level prior to the event day. None of those increases, however, is statistically significant.
[24] The Kruskal-Wallis (KW) tests yield similar results. Tables summarizing the KW test result are available by request.

announcement, before dropping significantly on the event day and remaining significantly below “normal” for approximately one week.

Table 2.3. Paired t-Test and Wilcoxon Signed Rank Test Results

	IVol on USDA event day T *vs.* day T-1		**11-day event window around day T**	
	Paired sample t-test	**Wilcoxon signed rank test**	**One-way ANOVA test**	**Kruskal-Wallis test**
A. Corn				
All USDA announcements	**-6.39*****	**-7.24*****	**6.66*****	**83.38*****
Big-event days	**-5.89*****	**-6.05*****	**16.47*****	**187.03*****
Small-event days	**-3.51*****	**-3.58*****	**2.21*****	**16.62***
B. Soybean				
All USDA announcements	**-5.13*****	**-6.17*****	**3.55*****	**58.68*****
Big-event days	**-5.60*****	**-5.27*****	**6.70*****	**102.81*****
Small-event days	**-1.41***	**3.30*****	1.14	8.54

Note: The first two columns of Table 2.3 present the two-sample parametric (Student t) and nonparametric (Wilcoxon signed rang) test statistics for $H_0: \ Ivol_T \geq Ivol_{T-1}$. The two rightmost columns show the results of one-way analysis of variance (ANOVA) and Kruskal-Wallis tests for $H_0: \Delta\, Ivol_{T-5} = \Delta Ivol_{T-4} = \cdots = \Delta Ivol_{T+5}$, with. For the t-tests, Table 2.3 reports left-sided t-values; for the Wilcoxon tests, the left-sided *z*-values. For the one-way ANOVA and Kruskal-Wallis tests, the *F*- and chi-square statistics are reported. For both commodities, we run each test for all USDA announcements together, and also separately for “big-event” days and “small-event” days (see Table 2.2 for the definition of “big” and “small” USDA events).
Statistical significance is denoted using * (10 percent), ** (5 percent), and *** (1 percent).

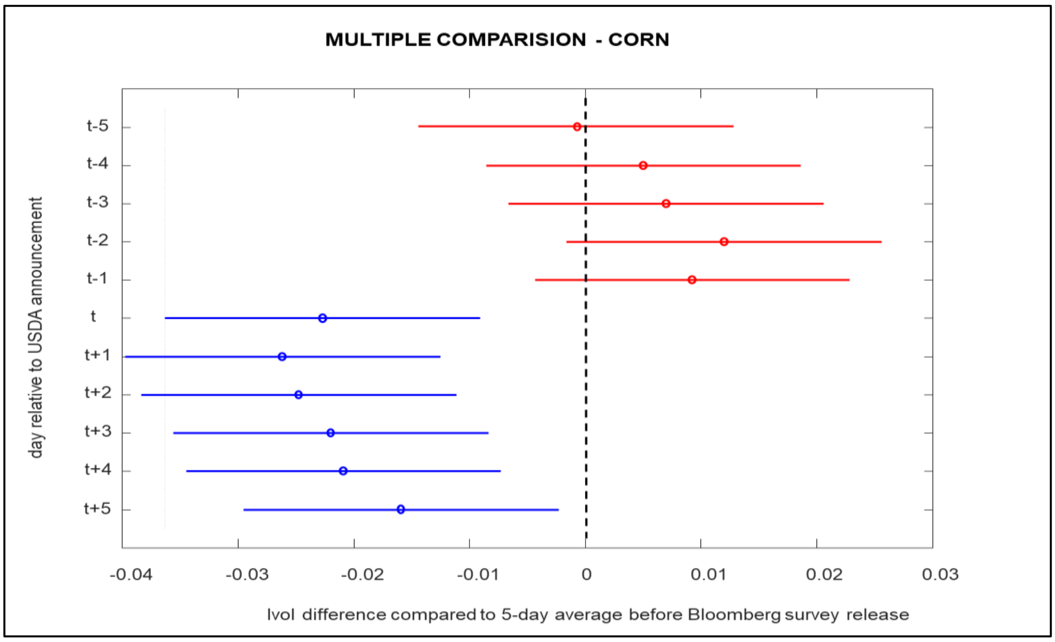

(a) Corn

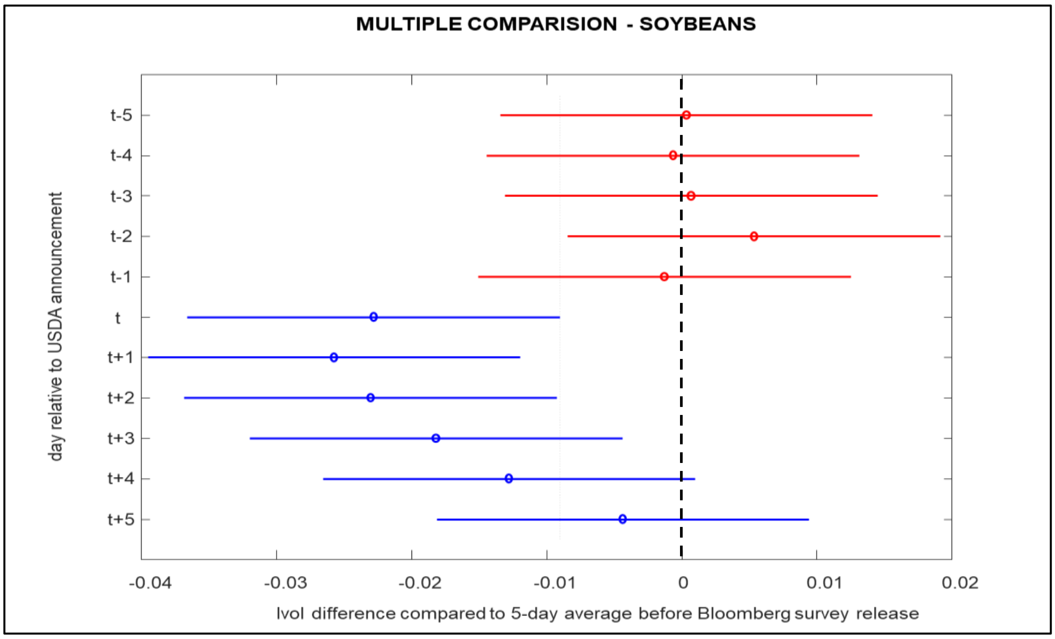

(b) Soybeans

Figure 2.2. Daily IVol Changes (*vs.* 5-day average IVol Prior to the Latest Bloomberg Survey)

Note: The circles in Figure 2.2(a) (corn) and 2.2(b) (soybeans) show the mean estimates of $\Delta\ IVol_{T+i}$ (i = -5, -4, …, +5) for five days before and five days after 151 USDA scheduled announcement days (day t) from September 2009 to October 2019. For each day, we compute log differences between (a) the 90-day commodity option-implied volatility (IVol) at the market close on day $T+i$ and (b) the average IVol for the five days before the most recent *pre*-event Bloomberg survey (which typically takes place five to seven days before the USDA event day). For each day, the colored bars represent the estimated 95-percent confidence intervals. If the confidence intervals of two days overlap each other, then the difference between IVols on those two days are not statistically significant. Likewise, if a colored bar crosses the zero dashed line, then the IVol on that day is not statistically significantly different from the average IVol in the five days before the Bloomberg survey. Test results (bars) for each of the five days before a USDA announcement are plotted in red; the bars for the announcement day T and for the next five trading days appear in blue.

Sources: USDA, Bloomberg and authors' computations.

Table 2.4. IVols Around USDA Announcements – Multi-Day Comparison Tests

Panel A: Corn

	$\overline{Ivol}$	$T-5$	$T-4$	$T-3$	$T-2$	$T-1$	T	$T+1$	$T+2$	$T+3$	$T+4$	$T+5$
$T-5$	0.62											
$T-4$	**0.05****	1.00										
$T-3$	**0.02****	1.00	1.00									
$T-2$	**<0.00*****	0.92	1.00	1.00								
$T-1$	**0.03****	0.98	1.00	1.00	1.00							
T	**<0.00*****	0.26	**0.04****	**0.02****	**<0.00*****	**0.01*****						
$T+1$	**<0.00*****	**0.10***	**0.01***	**<0.00*****	**<0.00*****	**<0.00*****	1.00					
$T+2$	**<0.00*****	0.15	**0.02****	**0.01*****	**<0.00*****	**<0.00*****	1.00	1.00				
$T+3$	**0.01*****	0.30	**0.05***	**0.03****	**<0.00*****	**0.01*****	1.00	1.00	1.00			
$T+4$	**0.01*****	0.38	**0.08***	**0.04****	**0.01*****	**0.02****	1.00	1.00	1.00	1.00		
$T+5$	**0.06***	0.79	0.32	0.20	**0.04****	**0.10***	1.00	0.98	0.99	1.00	1.00	

Panel B: Soybeans

	$\overline{Ivol}$	$T-5$	$T-4$	$T-3$	$T-2$	$T-1$	T	$T+1$	$T+2$	$T+3$	$T+4$	$T+5$
$T-5$	0.54											
$T-4$	0.61	1.00										
$T-3$	0.51	1.00	1.00									
$T-2$	0.14	1.00	1.00	1.00								
$T-1$	**0.07***	1.00	1.00	1.00	1.00							
T	**<0.00*****	0.20	0.25	0.18	**0.04****	0.30						
$T+1$	**<0.00*****	**0.08***	0.11	**0.07***	**0.01****	0.14	1.00					
$T+2$	**<0.00*****	0.19	0.24	0.17	**0.04****	0.28	1.00	1.00				
$T+3$	**0.02****	0.53	0.61	0.50	0.18	0.67	1.00	1.00	1.00			
$T+4$	**0.07***	0.91	0.94	0.89	0.57	0.96	0.99	0.92	0.98	1.00		
$T+5$	0.33	1.00	1.00	1.00	0.99	1.00	0.54	0.31	0.52	0.88	1.00	

Note: Table 2.4 shows the *p*-value matrix of ANOVA-based multiple comparison tests comparing (a) the IVol values on the days in the event window (columns 2 to 12 covering days *T-5* to *T+5*) and paired *t*-tests (column 1) to (b) the average IVol value in the five days before the Bloomberg survey of grain and oilseed market analysts. Each cell in columns 2 to 12, denoted p_{ij}, reports the *p*-value for$H_0: \Delta\, Ivol_{T+i} = \Delta Ivol_{T+j}$, with $i, j = -5,\ -4, \ldots,\ 5$ and $i \neq j$. In the first column, each cell reports the *p*-value for a one-sided *t*-test of each $Ivol_{T+I}$ against the average IVol on the five trading days before the Bloomberg survey, denoted $\overline{Ivol}$. For the days before the USDA announcement day *T* (*i.e.*, from *T-5* to *T+5*), the null hypothesis is that the IVol on that day is larger than $\overline{Ivol}$, indicating an increase in commodity implied volatility. In contrast, the null for the days after USDA announcement (*i.e.*, from $T+1$ to $T+5$) is that the mean IVol on that day is smaller than $\overline{Ivol}$, indicating a drop in implied volatility following the USDA report release. Panel A shows the result for corn; Panel B, for soybeans. Statistical significance is denoted using * (10 percent), ** (5 percent), and *** (1 percent).

2.5.3 *Hypotheses 2 to 4: Pre-event expert forecasters' influence on the IVol response*

Equation (2.5) could be estimated by Ordinary Least Squares (OLS) for each commodity separately, since the residuals are not serially correlated across announcement days.[25] However, because we find significant cross-equation residual correlations, Seemingly Unrelated Regressions (SUR)—as proposed by Zellner (1962)—are more efficient.[26] We therefore focus on the results obtained with the SUR method.[27]

Table 2.5 summarizes our estimations of Equation (2.5) jointly for corn and soybean markets using the SUR estimator. Heteroskedasticity-consistent standard errors are reported in brackets. To compare the predictive importance of the explanatory variables, we also report standardized regression coefficients.

Our result shows that the effects of market surprise and analyst dispersion vary across markets and across reports, in terms of signs as well as magnitudes.

(1) The role of report surprises.

Given how we compute surprises, price-bearish surprises are positive while bullish surprises are negative. Hence, a negative coefficient for bearish surprises ($\beta_j^+ < 0$) indicates an IVol decrease, whereas a positive value of β_j^+ indicates an IVol increase. The opposite holds for a bullish surprise β_j^-. Hypothesis 2 therefore implies that β_j^+ and β_j^- should both be negative.

[25] Breusch-Godfrey statistics for serial correlation tests, using a maximum of 30 lags, are 22.239 and 37.398 for corn and soybeans OLS residuals, respectively. With these test statistics, we cannot reject the null hypothesis of no serial correlation of any order up to 30 across event days for both commodities.
[26] The cross-equation (*i.e.*, between corn and soybeans) residual correlation is 0.353. The Breutsch-Pagan test statistic is 18.867, which strongly rejects the diagonality (*i.e.*, null-covariance) of the corn and soybean IVol residual variance-covariance matrix.
[27] OLS results are available upon request.

Table 2.5. Forecast Surprise, Analyst Dispersion, Sentiment, and Commodity IVol Changes†

	Corn		Soybean	
	Unstandardized	Standardized ‡	Unstandardized	Standardized‡
Constant	**-0.055***** **(0.017)**	-0.416 (0.301)	**-0.037**** **(0.018)**	-0.169 (0.345)
Bearish§ WASDE Surprise	-0.081 (0.103)	-0.073 (0.084)	0.073 (0.068)	0.091 (0.091)
Bullish§ WASDE Surprise	**-0.363*** **(0.208)**	**-0.261***** **(0.077)**	**-0.263**** **(0.112)**	**-0.336***** **(0.090)**
Bearish§ Grain Stocks Surprise	**0.711***** **(0.296)**	**0.383***** **(0.119)**	**0.266***** **(0.100)**	**0.192*** **(0.109)**
Bullish§ Grain Stocks Surprise	**-0.614*** **(0.318)**	**-0.221**** **(0.106)**	**-0.278***** **(0.076)**	**-0.201**** **(0.093)**
Bearish§ Planted Area Surprise	-1.598 (1.662)	-0.189 (0.131)	-0.257 (0.942)	-0.018 (0.083)
Bullish§ Planted Area Surprise	-1.267 (2.460)	-0.059 (0.105)	**1.451**** **(0.617)**	**0.236*** **(0.108)**
WASDE Dispersion	**-0.158**** **(0.068)**	**-0.182*** **(0.101)**	-0.070 (0.061)	-0.131 (0.117)
Grain Stocks Dispersion	-0.311 (0.450)	-0.104 (0.136)	0.280 (0.212)	0.164 (0.117)
Planted Area Dispersion	1.290 (5.310)	0.082 (0.177)	1.351 (1.467)	0.134 (0.134)
WASDE Sentiment ¶	0.005 (0.006)	0.102 (0.167)	0.010 (0.007)	0.229 (0.169)
Grain Stocks Sentiment¶	**-0.060***** **(0.021)**	**-1.124***** **(0.334)**	**-0.042**** **(0.016)**	**-0.938**** **(0.410)**
Planted Area Sentiment¶	-0.018 (0.029)	-0.340 (0.440)	-0.018 (0.025)	-0.402 (0.549)
VIX returns	**-0.135**** **(0.061)**	**-0.120*** **(0.074)**	-0.088 (0.069)	-0.093 (0.078)
Observations	151		151	
R^2	0.359		0.283	
F Statistic (df = 23; 127)	**3.078*****		**2.196*****	
Cross-equation residual correlation	0.353			
Breusch-Pagan test of diagonality	**18.867*****			

Note: Table 5 reports the Seemingly Unrelated Regression (SUR) estimation results for corn and soybeans. Heteroskedasticity-consistent standard errors are reported in brackets. Variables are described in Table 2.2.

Sample period: September 2009 to October 2019.

† For both commodities, the dependent variable is IVol daily returns on the day of USDA announcement.

‡ The standardized coefficients report the number of standard deviations change in the dependent variable associated with one standard deviation change in the independent variables, except for dummy variables.

§ Given the way in which we compute the surprises, bearish surprises are positive and bullish surprises are negative. Hence, a negative bearish surprise coefficient ($\beta_j^+ < 0$) indicates an IVol decrease, whereas a positive β_j^+ indicates an IVol increase, while holding all other factors constant.

Table 5 shows that, for both corn and soybeans, the bullish-surprise coefficients (β_j^-) are indeed statistically significantly negative for both WASDE and GS announcements.[28] In other words, a lower-than-predicted level of commodity inventories (whether actual in a GS report, or projected in a WASDE) brings about a smaller-than-average IVol drop (or even an outright IVol increase) post-announcement. Precisely, if the USDA corn (*resp.* soybean) stock projection in a WASDE comes in one percent under the median Bloomberg forecast, then the result is *ceteris paribus* a 0.36 (*resp.* 0.26) percent increase in the 90-day corn (*resp.* soybean) IVol on the USDA announcement day. For GS reports, the corresponding numbers are 0.61 and 0.28 percent.

In contrast, the estimated bearish-surprise coefficient β_j^+ do not appear to support the part of Hypothesis 2 predicting that IVols should drop more strongly following a bearish report surprise. First, we find that bearish WASDE and PA surprises have statistically insignificant impacts on the post-event IVol change. Put differently, apart from the IVol drop due to the removal of expected high volatility on the announcement day *T* (as discussed in Section 2.2.1), the market's corn and soybean volatility expectations (IVols) are not significantly revised further downward due to a higher-than-expected projected stock level or planted area. Second, bearish GS surprises are actually followed by a significant IVol increase in corn and soybean IVols. This finding implies that any GS surprise, whether bearish or bullish for prices, drives IVols upward (by similar amounts in both cases).

(2) The role of forecast dispersion.

For most reports, we do not find statistically significant coefficients for analyst dispersion. The only exception is for the corn WASDE. Consistent with Hypothesis 3, we find that the pre-event WASDE forecast dispersion significantly predicts the post-event IVol change in the corn market. Insofar as more disagreement among analysts (who could be traders too) implies greater corn market uncertainty prior to the WASDE release, the negative corn WASDE dispersion coefficient implies that the information released in the report becomes the new market consensus and resolves that uncertainty. Quantitatively, we find that a one-percent increase in WASDE forecast dispersion around the mean analyst forecast contributes to a statistically significant 0.16 percent decrease in the corn IVol (other things held equal). This finding complements the conclusion of Karali *et al.* (2019), that USDA reports remain valuable even in the presence of private forecasts.

[28] For corn, bullish PA report surprises are not statistically significant. The effect of bullish soybean PA surprises is negative and significant—an unexpected result. Bearish PA surprises have statistically insignificant impacts on the event-day IVol return, for both commodities.

(3) The role of forecaster sentiment.

As discussed in Section 2.2.2, we expect coefficients of our forecast sentiment dummies to be negative. This is indeed what we find for GS reports. *Ceteris paribus*, the corn (*resp.* soybean) IVol drops an extra 6.0 (*resp.* 4.2) percent when a majority of forecasters are pessimistic (sentiment dummy = 1) about the corn (*resp.* soybean) inventory level in an upcoming GS report, compared to when a majority are optimistic or neutral (sentiment dummy = 0). Both percentage drops are equivalent to about one standard deviation of the IVol change, either for corn and soybeans.

The Efficient Market Hypothesis (EMH) states that fundamentals-related information should be incorporated into market expectations as soon as it becomes available to market participants. Hence, having controlled for new fundamental-related information (proxied by surprises) and for uncertainty (proxied by dispersion), IVol change on day *T* should not be significantly predicted by information available prior to that day. Given how we construct the sentiment variable (see Appendix 2.A.1) and given the timing of the Bloomberg survey's release, the pessimistic/optimistic nature of the median analyst forecast is already known to market participants at the latest by day *T-1*. Therefore, the implication of our finding a significant GS sentiment coefficient is that market sentiment (not just market fundamentals) plays a role in how commodity IVols react to USDA announcements.

2.5.4 *Hypothesis 5: Macroeconomic uncertainty and financial market sentiment.*

For both corn and soybeans, the coefficient of the VIX return is negative. It is statistically significant for corn. All other things equal, a one-percent VIX increase on the USDA event day is associated with a 0.14 percent decrease in corn IVol. As noted in the development of Hypothesis 5 (see Section 2.2), prior work documents empirically that daily VIX and commodity IVol returns are positively correlated.

One possible interpretation of our surprising VIX finding is that it lends additional support to the argument that the USDA information is the "new market consensus." Given that commodity IVol changes are generally positively driven by VIX changes, it must be that, on those few days when the USDA announcements take place, the USDA news helps mitigate the VIX spillover.

In order to verify empirically our conjecture that the influence (on grain and oilseed IVols) of financial market uncertainty and sentiment is reduced on USDA event days, we run an additional analysis of the relation between VIX and commodity IVol returns for all days in sample period.[29] For each commodity, we run the following regression across all 2,567 days in our sample:

$$\Delta IVol_t = \beta_0 + \beta_1 \Delta VIX_t + \beta_2 D_{USDA,t} + \beta_3 \Delta VIX_t * D_{USDA,t} + \beta_4 \Delta IVol_{t-1} + \varepsilon_t \quad (2.7)$$

in which $\Delta IVol_t$ is the daily log-difference of IVol, ΔVIX_t is the daily log-difference of VIX, and $D_{USDA,t}$ is a dummy variable set equal to 1 when day t is a USDA announcement day.

Analogously to Goyal and Adjemian (2021), we first use simple OLS to estimate Equation (2.7). To account for conditional heteroskedasticity, we also estimate Equation (2.7) with standard GARCH (*i.e.*, sGARCH) and exponential GARCH (*i.e.*, eGARCH). For the sGARCH and eGARCH models, our diagnostic tests indicate that a GARCH(1,1) with ARMA(1,1) is sufficient for corn, while soybeans require a GARCH(1,1) with ARMA(3,1).

Table 2.6 summarizes our regression results. As in earlier studies, we find that the coefficient of ΔVIX_t (β_1) is positive for both commodities.[30] Likewise, β_2 is consistently negative and highly significant for both commodities in all three specifications, which reinforces the conclusion that USDA reports reduce commodity IVols. However, the effect of the VIX return on the IVol return is reversed on USDA announcement days: $\beta_3$is negative across all specifications for both corn and soybeans—with statistically significant values for corn (all models) and soybeans (eGARCH model). Moreover, in all cases, the absolute size of β_3 is much larger than that of β_1, leading to a negative net effect of the VIX change on the IVol change on USDA announcement days.

In sum, our results show that, while in general grain and oilseed IVol returns are positively related to changes in macroeconomic uncertainty and sentiment (jointly captured by the VIX), this relationship does not hold on USDA announcement days. This empirical finding points to the need for theoretical work to understand why the value placed by agricultural market participants on the commodity-specific (but consensus-making) information of the USDA reports seems to increases in the level of financial market uncertainty.

[29] We run the regressions in first differences to ensure that all series are stationary.

[30] As in Table 2.6, the VIX regression coefficient is statistically significant for corn but not for soybeans. This lack of significance may reflect the conjunction of two facts. One, historical decompositions by Adjemian *et al.* (2017) show that, in contrast to other determinants of agricultural IVols, macroeconomic uncertainty and financial market sentiment matter the most during periods of elevated financial stress. Two, with the exception of August 2011, there is no major VIX spike in our sample period (2009-2019).

Table 2.6. VIX Impact on Commodity IVols—Regular Days *vs.* USDA Event Days

	Corn			Soybeans		
	OLS	sGARCH	eGARCH	OLS	sGARCH	eGARCH
Constant	**0.002***** **(0.001)**	**0.002***** **(0.001)**	**0.002***** **(0.000)**	0.001 (0.001)	0.000 (0.001)	**0.001***** **(0.000)**
VIX	**0.032**** **(0.014)**	**0.043***** **(0.013)**	**0.047***** **(0.006)**	0.003 (0.014)	0.011 (0.015)	0.013 (0.009)
USDA announcement	**-0.034***** **(0.004)**	**-0.033***** **(0.004)**	**-0.032***** **(0.003)**	**-0.023***** **(0.004)**	**-0.023***** **(0.003)**	**-0.022***** **(0.003)**
VIX*USDA announcement	**-0.152*** **(0.088)**	**-0.159**** **(0.079)**	**-0.166***** **(0.021)**	-0.094 (0.072)	-0.053 (0.079)	**-0.057***** **(0.010)**
Lagged daily IVol log-difference	-0.045 (0.057)	**-0.814***** **(0.204)**	**-0.767***** **(0.017)**	0.032 (0.038)	**0.844***** **(0.072)**	**0.860***** **(0.043)**
AIC	-10153	-4.1388	-4.2140	-10339	-4.1420	-4.1585
BIC	-10118	-4.1114	-4.1798	-10304	-4.1101	-4.1243
Wald/LM tests	**17.375*****	0.4641	0.5078	**11.316*****	3.203	3.860

<u>*Notes*</u>: Table 2.6 provides estimates of the daily impact of VIX changes on commodity IVol changes on USDA *vs.* non-USDA days. Daily models cover the period from August 17, 2009 to October 31, 2019. For both eGARCH and sGARCH models, we estimate a GARCH(1,1) with ARMA(1,1) for corn, and a GARCH(1,1) with ARMA(3,1) for soybean; the choice of model is based on diagnostic tests as in Goyal and Adjemian (2021). For the regression coefficients, heteroskedasticity-consistent standard errors are reported in brackets. Wald test statistics are reported for the OLS models; weighted ARCH LM test statistics are reported for the GARCH models at lag 7. In all cases, statistical significance is denoted using * (10 percent), ** (5 percent), and *** (1 percent).

2.5.5 *Discussion: Information value of Bloomberg surveys*

As a final check of our results, we revisit the Bloomberg analyst surveys that precede a scheduled announcement. We ask two questions. One, do the surveys themselves contain new information—in which case, they might influence commodity returns or IVols prior to a USDA information release (which in turn would impact our measure of the IVol change on the announcement day). Second, is the median analyst survey a biased predictor of the USDA numbers—in which case, our measure of surprise could be affected?[31]

To answer these questions, we adopt and adapt the approach proposed by Balduzzi, Elton, and Green (1998) to study analyst forecasts of corporate earnings (see Appendix 2.A.3 for details). We regress the value of each USDA figure on (i) the median forecasted value in the Bloomberg survey and (ii) the IVol and price returns between the survey publication day and the USDA day. Table 2.7 in Appendix 2.A.3 summarizes the results. They suggest that the Bloomberg analyst forecasts are informationally valuable to market participants: the regression coefficient for the median forecast is significantly different from 0 in all specifications. Table 2.7 also suggests that the median forecast is an unbiased predictor of the USDA information: none of the intercepts is significantly different from 0, and the median forecast coefficients are not significantly different from 1.[32] Most importantly for our analysis, however, the BEG3 specifications in Table 2.7 show that the IVol return coefficients are not statistically significant after we control for the price returns. We conclude that market participants do not significantly revise their volatility expectations between the Bloomberg survey and the USDA announcement day, which assuages the concern of measurement errors.

2.6 **Conclusion**

We provide novel evidence on the impact of scheduled USDA information releases on forward-looking volatilities (IVols) in agricultural markets. We document that, for up to five trading days after the release of a scheduled USDA report (WASDE, Grain Stocks, Prospective Plantings, and Acreage), corn and soybean IVols are significantly lower than they had been a week before the release. The USDA reports' uncertainty-resolution power is substantial for both commodities.

[31] Karali, Irwin, and Isengildina-Massa (2019) raise a measurement-error concern in the case of surveyed analysts' crop production forecasts.
[32] In the case of soybean Grain Stock reports, the F-statistic for the joint hypothesis that the intercept is equal to 0 and the median forecast coefficient is equal to 1 is rejected at 10% level of significance. However, the size of these coefficients is very close to 1: we cannot reject, using t-tests, the hypothesis that the median forecast coefficients equal 1. We therefore conclude that, even if there is a bias in the case of analysts' soybean GS forecasts, it is very small.

The benefit of the USDA information in resolving market uncertainty is impacted by analyst disagreement and sentiment ahead of the report, and by the extent to which the market is surprised on the announcement day. Those three variables, however, do not have uniform impacts on the USDA reports' uncertainty resolution. Notably, for WASDE and Grain Stock reports, event-day surprises that are bullish for prices tend to boost commodity IVols. The effect of price-bearish surprises is more muted, except for Grain Stock reports. The impact of disagreement among market experts in the run-up to a report is usually insignificant, except for the corn WASDE (with greater pre-USDA dispersion boosting the IVol drop on the event day).

Sentiment also matters—both commodity experts' sentiment ahead of the release and changes in broad financial market sentiment on the event day. One, in the case of the Grain Stocks reports, we document that the calming effect of USDA news is larger when market analysts had been pessimistic about stock levels. Two, while commodity IVols are in general positively related to broad financial-market sentiment and macroeconomic uncertainty (jointly captured by the VIX index), we show that this co-movement surprisingly breaks down on USDA report days—with the VIX and commodity IVols moving in opposite directions on that day.

Our findings offer both practical and policy implications for market participants and policy makers. First, they show that the USDA information has value and impacts market volatility expectations. Second, short-run hedging and other derivatives-market positioning around USDA announcements could be improved by considering the IVol forecast-to-announcement patterns that we document, leading to more efficient pricing and risk management in the long run. Finally, public programs involving price volatility, such as crop insurance (Sherrick 2015) or USDA season-average price forecasts that incorporate forward-looking volatility—as advocated by Adjemian, Bruno and Robe (2020)—should also benefit from our conclusions.

Our findings suggest several venues for further research. First, most our empirical predictions are theoretically grounded in an extension of the Ederington and Lee (1996) model of implied volatility around scheduled public announcements. While our empirical analysis provides strong support for most of those predictions, it also points to the need for more theoretical work to better understand (i) why analyst surprises regarding grain inventories boost (*ceteris paribus*) the market's post-USDA-report volatility expectations when the surprise is bearish for prices and (ii) why the generally positive relationship between VIX returns and commodity IVols reverses on USDA event days.

Second, our paper focuses on commodity IVols that can be used as forecasts of future realized volatility (Egelkraut, Garcia, and Sherrick 2007). The IVols on which we rely are derived from the most liquid, at-the-money, options. Options on agricultural commodities, however, are unique in that out-of-the-money call options are usually more expensive than puts (Norland 2019). In other words, agricultural options exhibit positive skew. Given that the underlying returns in these markets generally do not exhibit positive skewness, the likely explanation is market structure: food buyers appear more willing to pay a premium for upside protection than farmers seem ready to pay for downside protection. A natural question is what happens to the volatility skew around USDA events. We leave this question for further research.

Acknowledgement

We are grateful to Mike Adjemian (our CEMA discussant), Siyu Bian, Annabelle Couleau, Adrian Fernandez, Todd Hubbs, Scott Irwin, Olga Isengildina-Massa, Kris Jacobs, Mark Jekanowski, Andy McKenzie, Hong Miao (our JPMCC discussant), Esen Onur, Dan Prager, and especially to Scott Mixon, Thomas Heckelei, and Hyun-Soo Choi (our APAD discussant), for helpful discussions and suggestions. We thank online conference participants at the 2020 Meeting of the Agricultural and Applied Economics Association (AAEA), the 2021 Meeting of the Commodity and Energy Markets Association (CEMA), the 2021 Meeting of the Asian Pacific Derivatives Association (APAD), the 2021 Research Symposium of the JP Morgan Center for Commodities (JPMCC), and online seminar participants at the US Commodity Futures Trading Commission (CFTC), the USDA's Economic Research Service (ERS), the CME Group, Centrec Consulting Group (Urbana, IL), and the Universities of Wisconsin, Arkansas, and Illinois, for helpful comments. This paper was partly supported by a NIFA grant under Multistate Project Number S1072. Michel A. Robe also gratefully acknowledges the financial support received in his capacity as The Clearing Corporation Foundation Professor in Derivatives Trading at the University of Illinois. While this paper was written, Michel A. Robe also served as a part-time consulting senior economist at the CFTC; no CFTC resources or confidential information were used for this project. The views expressed in this article are those of the authors only and may not be attributed to the CFTC, the Commissioners, or other CFTC staff. The authors are solely responsible for any error or omission.

2.7 References

Adjemian, M. K. (2012). "Quantifying the WASDE Announcement Effect." *American Journal of Agricultural Economics*, *94*(1), 238–256. https://doi.org/10.1093/ajae/aar131

Adjemian, M. K., Bruno, V. G. and Robe, M. A. (2020). "Incorporating Uncertainty into USDA Commodity Price Forecasts." *American Journal of Agricultural Economic*, 102(2): 696–712. https://doi.org/10.2139/ssrn.3414585

Adjemian, M. K., Bruno, V. G., Robe, M. A., and Wallen, J. (2017). "What Drives Volatility Expectations in Grain and Oilseed Markets?" Paper presented at the 2017 Annual Meeting of the Agricultural and Applied Economics Association, July. Available online at: https://ageconsearch.umn.edu/record/258452/files/Abstracts_17_05_24_09_18_42_30__199_136_105_95_0.pdf

Adjemian, M. K. and Irwin, S. H. (2018). "USDA Announcement Effects in Real-time." *American Journal of Agricultural Economics*, *100*(4), 1151–1171. https://doi.org/10.1093/ajae/aay018

Adjemian, M. K., Johansson, R., McKenzie, A. and Thomsen, M. (2018). "Was the Missing 2013 WASDE Missed?" *Applied Economic Perspectives and Policy*, *40*(4), 653–671. https://doi.org/10.1093/aepp/ppx049

Amir, E. and Ganzach, Y. (1998). "Overreaction and Underreaction in Analysts' Forecasts." *Journal of Economic Behavior & Organization*, *37*(3): 333–347. https://doi.org/10.1016/S0167-2681(98)00092-4

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Vega, C. (2003). "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange." *American Economic Review*, *93*(1), 38–62. https://doi.org/10.1257/000282803321455151

Baker, M. and Wurgler, J. (2006). "Investor Sentiment and the Cross-Section of Stock Returns." *Journal of Finance*, *61*(4), 1645–1680. https://doi.org/10.1111/j.1540-6261.2006.00885.x

Balduzzi, P., Elton, E. J. and Green, T. C. C. (2001). "Economic News and Bond Prices: Evidence from the U.S. Treasury Market." *Journal of Financial and Quantitative Analysis*, 36(4), 523-543. https://www.jstor.org/stable/2676223

Barberis, N., Shleifer, A. and Vishny, R. (1998). "A model of Investor Sentiment." *Journal of Financial Economics*, *49*, 307–343. https://doi.org/10.1515/9781400829125-015

Beber, A. and Brandt, M. W. (2010). "When It Cannot Get Better or Worse: The Asymmetric Impact of Good and Bad News on Bond Returns in Expansions and Recessions." *Review of Finance, 14*(1), 119–155. https://doi.org/10.1093/rof/rfp006

Bekaert, G., Hoerova, M. and Lo Duca, M. (2013). "Risk, Uncertainty and Monetary Policy." *Journal of Monetary Economics, 60*(7), 771–788. https://doi.org/10.1016/j.jmoneco.2013.06.003

Bobenrieth, E. S. A., Bobenrieth, J. R. A., Guerra, E. A., Wright, B. D. and Di Zeng (2021). "Putting the Empirical Commodity Storage Model Back on Track: Crucial Implications of a 'Negligible' Trend." *American Journal of Agricultural Economics,* 103(3): 1034-1057. https://doi.org/10.1111/ajae.12133

Braun, P. A., Nelson, D. B. and Sunier, A. M. (1995). "Good News, Bad News, Volatility, and Betas." *Journal of Finance, 50*(5), 1575–1603. https://doi.org/10.1111/j.1540-6261.1995.tb05189.x

Couleau, A., Serra, T. and Garcia, P. (2020). "Are Corn Futures Prices Getting 'Jumpy'?" *American Journal of Agricultural Economics* 102(2): 569-588. https://doi.org/10.1002/ajae.12030

Cui, C. 2012. "Commodity Implied Volatility Surface." Bloomberg OVDV White Paper.

Diether, K. B., Malloy, C. J. and Scherbina, A. (2002). "Differences of Opinion and the Cross Section of Stock Returns." *Journal of Finance, 57*(5), 2113–2141. https://doi.org/10.1111/0022-1082.00490

Dubinsky, A., Johannes, M., Kaeck, A. and Seeger, N. J. (2019). "Option Pricing of Earnings Announcement Risks." *Review of Financial Studies, 32*(2), 646–687. https://doi.org/10.1093/rfs/hhy060

Ederington, L. H. and Lee, J. H. (1993). "How Markets Process Information: News Releases and Volatility." *Journal of Finance, 48*(4), 1161–1191. https://doi.org/10.1111/j.1540-6261.1993.tb04750.x.

Ederington, L. H. and Lee, J. H. (1996). "The Creation and Resolution of Market Uncertainty: The Impact of Information Releases on Implied Volatility." *Journal of Financial and Quantitative Analysis, 31*(4), 513. https://doi.org/10.2307/2331358

Egelkraut, T.M., P. Garcia and B.J. Sherrick. 2007. "The Term Structure of Implied Forward Volatility: Recovery and Informational Content in the Corn Options Market." *American Journal of Agricultural Economics* 89(1): 1–11. https://www.jstor.org/stable/4123558

Fernandez-Perez, A., Frijins, B., Indriawan, I. and Tourani-Rad, A. (2019). "Surprise and Dispersion: Informational Impact of USDA Announcements." *Agricultural Economics, 50*(1), 113–126. https://doi.org/10.1111/agec.12470

Fortenbery, T. R. and Sumner, D. A. (1993). "The Effects of USDA Reports in Futures and Options Markets." *Journal of Futures Markets, 13*(2), 157–173. https://doi.org/10.1002/fut.3990130204

Geman, H. and Smith, W. O. (2013). "Theory of Storage, Inventory and Volatility in the LME Base Metals." *Resources Policy, 38*(1), 18–28. https://doi.org/10.1016/j.resourpol.2012.06.014

Glosten, L. R., Jagannathan, R. and Runkle, D. E. (1993). "On the relation between the expected value and the volatility of the nominal excess return on stocks." *Journal of Finance*, *48*(5), 1779–1801. https://doi.org/10.1111/j.1540-6261.1993.tb05128.x

Good, D. and S. Irwin (2011). "USDA Corn and Soybean Acreage Estimates and Yield Forecasts: Dispelling Myths and Misunderstandings." *Marketing and Outlook Briefs* 11(02): 1-15. https://farmdoc.illinois.edu/assets/marketing/mobr/mobr_11-02.pdf

Gousgounis, E. and Onur, E. (2018). "The Effect of Pit Closure on Futures Trading." *Journal of Commodity Markets, 10,* 69–90. https://doi.org/10.1016/j.jcomm.2017.11.002

Goyal, R. and Adjemian, M. K. (2021). "The 2019 Government Shutdown Increased Uncertainty in Major Agricultural Commodity Markets." *Food Policy* (102), forthcoming. https://doi.org/10.1016/j.foodpol.2021

Haynes, R., Raman, V., Robe, M. A. and Yadav, P. K. (2017). "Algorithmic Trading and Livestock Market Quality." CFTC Paper presented at the 2017 NCCC-134 Conference, April. https://legacy.farmdoc.illinois.edu/nccc134/pdf/NCCC-134_2017_Program.pdf

Haynes, R. and Roberts, J. S. (2015). "Automated Trading in Futures Markets. *CFTC Paper*." https://www.cftc.gov/sites/default/files/idc/groups/public/@economicanalysis/documents/file/oce_automatedtrading.pdf

Hochberg, Y. and Tamhane, A. (1987). *Multiple Comparison Procedures*. NY: John Wiley & Sons.

Isengildina-Massa, O., Irwin, S. H., Good, D. L. and Gomez, J. K. (2008). "Impact of WASDE Reports on Implied Volatility in Corn and Soybean Markets." *Agribusiness, 24*(4), 473–490. https://doi.org/10.1002/agr.20174

Isengildina-Massa, O., X. Cao, B. Karali, S.H. Irwin, M.K. Adjemian, and R.C. Johansson (2021). "When Does USDA Information Have the Most impact on Crop and Livestock Markets?" *Journal of Commodity Markets,* 22(1): 100-137. https://doi.org/10.1016/j.jcomm.2020.100137

Janzen, J. P. and Adjemian, M. K. (2017). "Estimating the Location of World Wheat Price Discovery." *American Journal of Agricultural Economics*, *99*(5), 1188–1207. https://doi.org/10.1093/ajae/aax046

Janzen, J. P. and Bunek, G.D. (2017). "Does Public Information Facilitate Price Consensus? Characterizing the Effect of USDA Crop Reports Using Realized Volatility." University of Montana Working Paper presented at the 2015 NCCC-134 Conference. Updated, June 2017. https://legacy.farmdoc.illinois.edu/nccc134/conf_2015/pdf/Bunek_Janzen_NCCC_134_2015.pdf

Johnson, T. C. (2004). "Forecast Dispersion and the Cross Section of Expected Returns." *Journal of Finance*, *59*(5), 1957–1978. https://doi.org/10.1111/j.1540-6261.2004.00688.x

Kaldor, N. (1939). "Speculation and Economic Stability." *Review of Economic Studies, 7(1),* 1–27. https://doi.org/10.2307/2967593

Karali, B., Irwin, S. H. and Isengildina-Massa, O. (2019). "Supply Fundamentals and Grain Futures Price Movements." *American Journal of Agricultural Economics*, 102(2): 548-568. https://doi.org/10.1002/ajae.12012

Karali, B., Isengildina-Massa, O., Irwin, S. H., Adjemian, M. K. and Johansson, R. (2019). "Are USDA Reports still News to Changing Crop Markets?" *Food Policy 84 (April), 66-76*. https://doi.org/10.1016/j.foodpol.2019.02.005

Kurov, A., A. Sancetta, G. Strasser, and M. Halova Wolfe (2019). "Price drift before US macroeconomic news: Private information about public announcements?" *Journal of Financial and Quantitative Analysis*, 54(1), 449-479. https://doi.org/10.1017/S0022109018000625

Lehecka, G., Wang, X. and Garcia, P. (2014). "Gone in Ten Minutes: Intraday Evidence of Announcement Effects in the Electronic Corn Futures Market." *Applied Economic Perspectives and Policy*, *36*(3), 504–526. https://doi.org/10.1093/aepp/ppu010

McKenzie, A. M. (2008). "Pre-Harvest Price Expectations for Corn: The Information Content of USDA Reports and New Crop Futures." *American Journal of Agricultural Economics*, *90*(2), 351–366. https://doi.org/10.1111/j.1467-8276.2007.01117.x

McKenzie, A. M., Thomsen, M. and Phelan, J. (2007). “How do you Straddle Hogs and Pigs? Ask the Greeks!” *Applied Financial Economics*, 17(7), 511–520. https://doi.org/10.1080/09603100500428230

McNew, K. P. and Espinosa, J. A. (1994). “The Informational Content of USDA Crop Reports: Impacts on Uncertainty and Expectations in Grain Futures Markets.” *Journal of Futures Markets*, *14*(4), 475–492. https://doi.org/10.1002/fut.3990140408

Ng, V. K., & Pirrong, S. C. (1994). “Fundamentals and Volatility: Storage, Spreads, and the Dynamics of Metals Prices.” *Journal of Business*, 67, 203–230. https://www.jstor.org/stable/2353103

Nofsinger, J. R. and Prucyk, B. (2003). “Option Volume and Volatility Response to Scheduled Economic News Releases.” *Journal of Futures Markets*, *23*(4), 315–345. https://doi.org/10.1002/fut.10064

Norland, E. (2019). "The Rationale for Ag Options’ Upward Skew." Working Paper, CME Institute, July. Available online at: https://www.cmegroup.com/education/featured-reports/the-rationale-for-ag-options-upward-skew.html

Patell, J. M. and Wolfson, M. A. (1979). “Anticipated Information Releases Reflected in Call Option Prices.” *Journal of Accounting and Economics*, *1*(2), 117–140. https://doi.org/10.1016/0165-4101(79)90003-X

Pirrong, S. C. (2011). *Commodity Price Dynamics: A Structural Approach*. Cambridge, UK: Cambridge University Press.

Prager, D., Burns, C., Tulman, S. and MacDonald, J. (2020). “Farm Use of Futures, Options, and Marketing Contracts.” *ERS Economic Information Bulletin* No. 219, October 2020. Available online at: https://www.ers.usda.gov/webdocs/publications/99518/eib-219.pdf?v=1633

Sherrick, B. (2015). “Understanding the Implied Volatility Factor in Crop Insurance.” *Farmdoc Daily*, *5*, Feb. https://farmdocdaily.illinois.edu/2015/02/understanding-implied-volatility-crop-insurance.html

Stoline, M. R. (1981). “The Status of Multiple Comparisons: Simultaneous Estimation of all Pairwise Comparisons in one-way ANOVA Designs.” *The American Statistician*, *35*(3), 134–141. https://doi.org/10.1080/00031305.1981.10479331

Tetlock, P. C. (2007). “Giving Content to Investor Sentiment: The Role of Media in the Stock Market.” *Journal of Finance*, *62*(3), 1139–1168. https://doi.org/10.1111/j.1540-6261.2007.01232.x

Working, H. (1948). “Theory of the Inverse Carrying Charge in Futures Markets.” *Journal of Farm Economics*, *30*(1), 1–28. https://www.jstor.org/stable/pdf/1232678.pdf

Ying, J., Chen, Y. and Dorfman, J. H. (2019). “Flexible Tests for USDA Report Announcement Effects in Futures Markets.” *American Journal of Agricultural Economics, 94*(1), 238. https://doi.org/10.1093/ajae/aaz013

Zellner, A. (1962). “An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias.” Journal of the American statistical Association, 57(298), 348-368. https://doi.org/10.1080/01621459.1962.10480664

2.8 Appendices

2.A.1 Baselines for "Pessimism" in Forecasts

Based on the nature of the forecasted information in each report, we define their respective reference points as follows:

1. WASDE. The forecast we use is also the most frequently surveyed information—the projected U.S. ending stock of the current marketing year. Every month, the USDA updates the projections in the WASDE in light of demand and supply developments. As reference point or baseline, we therefore choose the actual value in previous month's report.

2. GS reports. The USDA estimates U.S. ending stocks as of the end of the previous quarter. Due to the seasonality of crop production and demand, grain and oilseed inventories also fluctuate seasonally. We therefore use the same quarter of the previous year as the reference point. When forecasters predict a lower (*resp.* higher) stock level than at the same time in the prior year, we call them "pessimistic" (*resp.* "optimistic") about the inventory situation.

3. PP and AR reports. The Prospective Plantings and the Acreage reports both provide information on the planting area of the current crop year. By construction, the AR report is the updated version of the PP report for the same crop year. For the AR report, we therefore proceed as for the WASDE reports, and use the earlier information (in the PP report) to determine if the estimates in the later report (AR) is lower ("pessimistic") or higher ("Optimistic"). For the PP report, we define "optimistic" or "pessimistic" analyst sentiment by comparing the current year's planting intentions (in the PP report) to the planted area in the previous year (in the prior year's AR report).

2.A.2 Seasonality of Agricultural Option-Implied Volatilities (IVol), 2009-2019

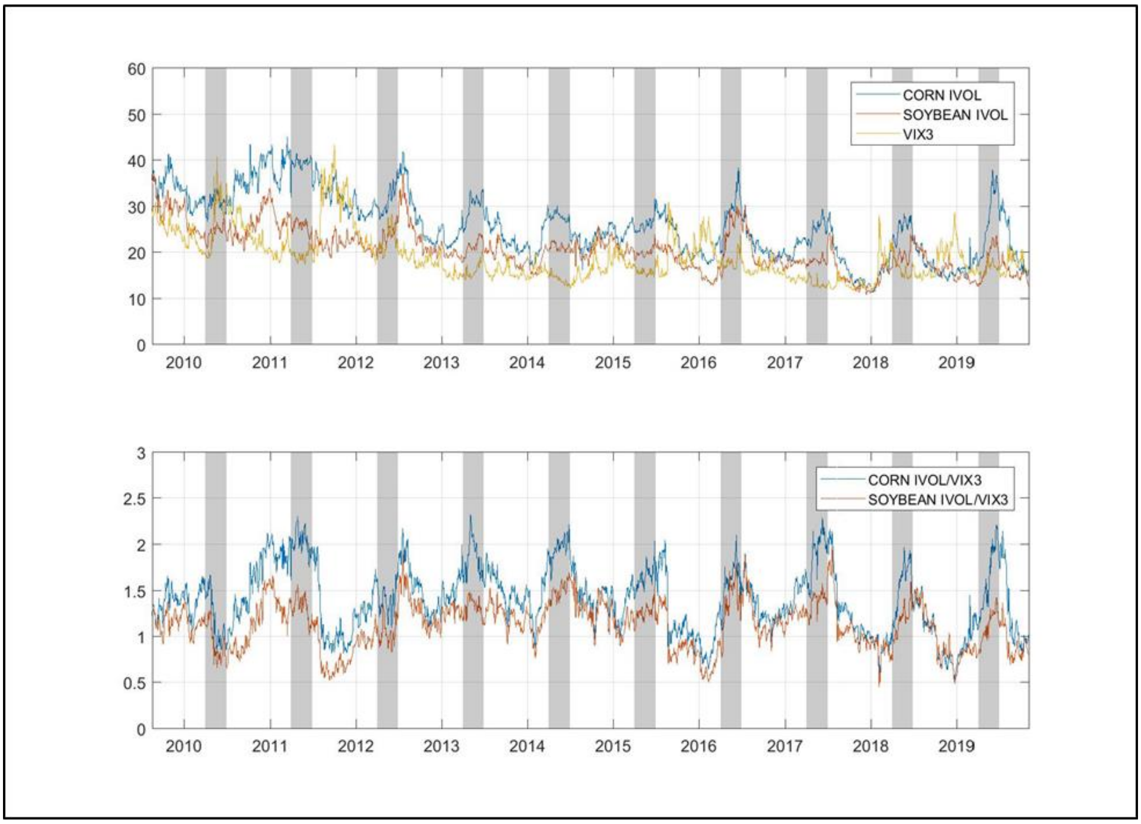

Note: The above Figures plot the daily values of the 3-month VIX and of the forward-looking return volatilities that are embedded in the prices of (synthetic) constant-maturity 90-day at-the-money options on agricultural futures (corn, IVC3; soybean, ICS3) from August 2009 through September 2019. The top panel plots the actual values in percent. The bottom panel expresses the commodity IVols in terms of the contemporaneous VIX. The spring season is shaded to better highlight the seasonality of the commodity IVols (Source: Bloomberg).

2.A.3 Informational Value of the Bloomberg-Survey Analyst Forecasts

Balduzzi, Elton and Green (1998) propose the regression equation (hereafter, the BEG equation):

$$A_{i,j} = \alpha_0 + \alpha_1 F_{i,j} + \alpha_2 R_i + \varepsilon_{i,j} \tag{2.A.3.1}$$

where $A_{I,j}$ and $F_{I,j}$ are, respectively, the actual and the forecasted values of indicator j for the announcement day i, and R_I is the cumulative market return from the day when Bloomberg releases the survey result to the announcement day. Several hypotheses can be tested with this regression:

- If α_1 is significantly different from zero, then the forecast contains information;
- If α_0 is not significantly different from zero and α_1 is not significantly different from 1, then the forecast is unbiased;

- If α_2 is significantly different from zero, then market expectations have been revised between the forecast day and the announcement day. In this case, new information arrives in the market after the forecast.

Table 2.7. BEG Regressions

	BEG1		BEG2		BEG3	
	WASDE	GS	WASDE	GS	WASDE	GS
A. Corn						
Constant (α_0)	27.29	38.02	-8.67	43.69	10.66	50.58
	(31.47)	(43.91)	(35.42)	(56.46)	(33.08)	(44.31)
$F_{i,j}$	**0.99*** **	**1.00*** **	**1.00*** **	**1.00*** **	**0.99*** **	**1.00*** **
	(0.02)	**(0.01)**	**(0.02)**	**(0.01)**	**(0.02)**	**(0.01)**
R_i	**-1856.50*** **	**-2473.96*** **			**-1694.21*** **	**-2620.53*** **
	(361.82)	**(522.37)**			**(374,65)**	**(526.74)**
$\Delta Ivol_i$			**-516.56*** **	109.63	-284.21	386.72
			(190.35)	(347.59)	(183.60)	(278.30)
Observations	120	41	120	41	120	41
R^2	0.96	1.00	0.96	1.00	0.96	1.00
F Statistic (H_0: $\alpha_0 = 0$; $\alpha_1 = 1$)	0.51	0.55	0.04	0.53	0.94	1.23
B. Soybeans						
Constant (α_0)	5.25	11.13	-5.14	9.33	5.09	10.25
	(6.18)	(9.96)	(7.42)	(9.98)	(6.84)	(10.12)
$F_{i,j}$	**1.00*** **	**0.99*** **	**1.01*** **	**0.99*** **	**1.00*** **	**0.99*** **
	(0.01)	**(0.01)**	**(0.02)**	**(0.01)**	**(0.01)**	**(0.01)**

Table 2.7 (cont.) BEG Regressions

	BEG1		BEG2		BEG3	
	WASDE	GS	WASDE	GS	WASDE	GS
B. Soybeans						
R_i	**-787.92*****	-168.14			**-785.49*****	-152.64
	(130.67)	(208.62)			**(138.62)**	(211.51)
$\Delta Ivol_i$			**-103.42***	-71.58	-3.00	-63.80
			(58.61)	(95.64)	(55.03)	(96.85)
Observations	120	41	120	41	120	41
R^2	0.98	1.00	0.97	1.00	0.98	1.00
F Statistic (H_0: $\alpha_0 = 0;\ \alpha_1 = 1$)	0.69	**3.19***	0.27	**3.29****	0.49	**3.23***

Note: Table 2.7 reports the result of three different versions of the BEG equation. For each version, we run a regression for the monthly WASDE ending stock forecasts and one for the quarterly Grain Stocks estimates (there are too few observations for the regression to make sense in the case of the annual PP and AR reports). The original BEG regression is presented in the first two columns of Table 2.7 ("BEG1"); in the next two columns ("BEG2"), R_I is replaced by $\Delta Ivol_I$ (*i.e.*, returns of commodity IVol from the forecast day to the announcement day). The last two columns ("BEG3") include both R_I and $\Delta Ivol_I$ in one regression. Finally, the last row of each panel shows the F-statistic of the joint hypothesis test for H_0: $\alpha_0 = 0$; $\alpha_1 = 1$ – implying the analyst forecasts are unbiased. In all cases, statistical significance is denoted using * (10 percent), ** (5 percent), and *** (1 percent).

Chapter 3
USDA Reports Affect the Stock Market, Too[33]

Abstract: We document that the stock prices of food-sector firms react to USDA news. The economic and statistical significance of the effect depends on the commodity, type of scheduled USDA report, and direction and extent to which the USDA information surprises the market. Individual stock price responses to USDA news differ between firms on the input-side *vs.* firms on the output-side of agricultural (farm) production, based on which component of the firm's cash-flow expectations (costs or revenues) and which variable (commodity price or expected firm output) is impacted by the news. Planted Area surprises have the largest effect for both subsets of firms (ag-as-inputs and ag-as-output), followed by Grain Stocks news—with the effects having the expected sign. In contrast, WASDE surprises have very modest and mixed impacts on food-sector stock returns. Our findings establish that USDA announcements have an impact well beyond their recognized relevance to commodity markets.

JEL classification: G12, G14, Q02, Q11

Keywords*:* Commodity news, Stock market reactions, USDA Announcements

---

[33] Chapter 3 is pre-published as a working paper titled Cao, A. N. Q., Ionici, O. and Robe, M. A. (2023). USDA Reports Affect the Stock Market, Too. *SSRN Electronic Journal*. Advance online publication. https://doi.org/10.2139/ssrn.4391882. The paper has received a "Revise and Resubmit" invitation from the Journal of Commodity Markets and is currently under revision.

3.1 Introduction

The U.S. Department of Agriculture (USDA) has long gathered information on physical conditions in agricultural markets. It disseminates those data, free of charge, through numerous reports. This activity, of course, is not costless. Yet, the "actual amount of agricultural production and marketing information that is mandated by U.S. law is small relative to the actual information produced by USDA" (Pruitt *et al.* 2014, p.25). Given that taxpayers foot the bill for the USDA's reports, the latter's usefulness is therefore open to reexamination every time "a new farm bill is created or when economic conditions prompt lawmakers to contend with growing budget deficits" (Ellison and Lusk 2011, p.1).[34]

In such an environment, a key question is whether the information published by the USDA is valuable—and who values it. The extant literature provides evidence of value by documenting that agricultural futures and options markets react significantly to the news contained in various USDA reports.[35] That prior research investigates their impact on commodity prices (*e.g.*, Adjemian 2012; Goyal and Adjemian 2021; Karali *et al.* 2019; Ying, Chen, and Dorfman 2019) or volatility expectations and market sentiment (*e.g.*, Cao and Robe 2022; Isengildina-Massa *et al.* 2008; McNew and Espinosa 1994). In contrast, the present paper asks whether USDA news ripple beyond commodity markets. Specifically, we investigate for the first time whether USDA news also affect the stock market as a whole and, if not, whether they move the stock prices of publicly listed companies in the "food" sector.

Equity prices are net present values of expected future company cash-flows, discounted at the appropriate risk-adjusted required rates of return. USDA news can therefore, in theory, affect share prices if they alter the expectations of equity market participants regarding the stream of future corporate cash-flows or if they impact the rate at which investors discount that stream.

---

[34] The concern that any given USDA report might be discontinued is more than theoretical. To wit, just three years ago (in February 2020), the USDA's National Agricultural Statistics Service (NASS) announced that it would "no longer publish county-level estimates for dry edible beans, flaxseed, hay (alfalfa and other), potatoes, sugarbeets, sugarcane, sunflower (non-oil and oil varieties) and tobacco" because funding had not been renewed to cover the "collection cost for the surveys used to gather the data used for county level estimates" (USDA 2020).

[35] The same is true of EIA (U.S. Department of Energy) announcements in energy markets. In the same vein, a large body of work in financial economics establishes empirically that U.S. macroeconomic announcements impact equity and bond prices significantly. See, *e.g.*, Kurov *et al*. (2019) for a review of that work.

Given that food-sector companies only make up a small fraction of the U.S. stock market,[36] there is little reason to expect more than a trivial impact on the broad stock market from news about agricultural commodities.[37] Therefore, after providing empirical evidence that USDA reports indeed do not significantly impact the U.S. stock market as a whole, we ignore discount rates and focus on the cash-flow channel.

On the cash-flow side, intuition suggests several reasons why one could expect agricultural prices to impact future earnings of companies in the food sector. On the one hand, in the case of companies that make production factors for the agricultural sector (*e.g.*, fertilizer producers, farm machinery manufacturers, agricultural technology developers, etc.), demand for their products may be affected if a USDA report has implications for planted acreage—either directly (in the annual Prospective Plantings or Acreage reports) or indirectly (for example, if the news' price implications of a World Agricultural Supply and Demand Estimates or WASDE report are sufficient to bring about a change in market expectations of future acreage allocated to food production). Furthermore, regardless of acreage considerations, suppose that financial constraints limit farmers' ability or willingness to purchase capital goods. Then, insofar as the USDA news' implications are bullish for commodity prices and imply a relaxation of those constraints, one should expect farm equipment purchases (machinery, investments in technology) to increase—thus boosting manufacturers' revenues and (assuming their costs are not impacted much) earnings.

On the other hand, for companies in the food transformation sector such as mills, beverage makers, biofuel producers, etc. (*resp.* for restaurant and grocery chains), agricultural commodities (*resp.* products derived from them) are an input. USDA news might thus be expected to affect commodity users' costs in the opposite direction of the first set of companies. Thus, the ultimate impact on those companies' future earnings depends (a) on the degree to which they hedge against commodity price fluctuations, and (b) if they do not, on the extent to which their competitive position allows them to pass through an increase in input costs to their own customers. If the latter's demand is not perfectly inelastic, then the impact of the cost increase should dominate.

Our question, then, is an empirical one: using data from 2009-2019 (before the COVID pandemic), we ask whether the stock prices of publicly traded companies from different sub-sectors (food processors such as Coca Cola and Kraft-Heinz; farm machinery

[36] In 2019, for example, less than one sixteenth (6.5 percent) of all the firms that make up Standard and Poor's S&P 500 stock market index belonged to the "food" sector (as defined by the SIC codes used to delineate the sample in the present paper).

[37] In contrast, for the energy sector, there is evidence of feedback between crude oil and equity index prices—see, *e.g.*, Huang (1996) and Kilian and Park (2009).

producers like John Deere and Caterpillar; fertilizer manufacturers; biofuel producers; restaurants chains; grocery stores; etc.) react significantly to USDA announcements.

Answering this question requires teasing out the news contained in each USDA report, as well as accounting for the extent to which the changes in companies' share prices on USDA event days might simply be echoing broad U.S. stock market movements that have nothing to do with the USDA news. For the first task, we exploit the fact that, ahead of all major scheduled USDA announcements, companies like Bloomberg and Reuters have, for more than a decade, published surveys of commodity analysts' expectations regarding the upcoming reports. Those news organizations typically release the details of their surveys in the week before a USDA news event, which allows us to compute the market surprise on the event day. For the second task, we compute, for each stock, excess returns on the USDA announcement-days. We do so using a rolling multi-month window leading up to ten days before a USDA announcement (precisely, before a scheduled WASDE, Grain Stocks, Prospective Plantings, or Acreage report).

Our sample contains 154 publicly traded companies between September 2009 and October 2019. The starting year reflects the availability of Bloomberg data on commodity analyst surveys, which we use to compute the USDA surprises. We choose the ending year to predate the start of the COVID-19 pandemic. In constructing our dataset, we account carefully for mergers, spin-offs, acquisitions, and de-listings. The resulting sample is unbalanced, and our empirical methodology accounts for that fact.

We consider the stock price reactions for each firm individually and for six sub-sectors (fertilizer and pesticide producers; farm machinery and technology firms; biofuel producers; food processors and beverage firms; restaurant chains and catering firms; food retailers and supermarket chains). Our results yield insights for agribusinesses and financial analysts in the food sector, and they provide a novel measure of USDA reports' value to market participants.

Section 3.2 proposes testable hypotheses. Sections 3.3 describes the data. Section 3.4 discusses our empirical methodology. Section 3.5 presents our results. Section 3.6 discusses policy implications of our findings. Section 3.7 concludes the paper and outlines possible venues for further research.

3.2 Hypothesis development

In this Section, we propose several hypotheses regarding how scheduled releases of USDA news about three key U.S. agricultural commodities (corn, soybean, wheat) could affect the U.S. stock market as a whole (returns, Section 3.2.1), the stock prices of U.S.

companies that operate in the "food industry" broadly defined (excess returns on USDA event days, Section 3.2.2), and USDA-days excess returns for different sub-sectors of that broad industry (Section 3.2.3).

3.2.1 *U.S. Stock Market Returns and Volatility on USDA Event Days*

As noted in the Introduction, food-sector firms make up a small fraction of the overall stock market in the United States, both in terms of the number of companies and in terms of total market value. Furthermore, even if the food sector did make up a large share of the US stock market, that very sector comprises both companies whose stock price should react positively to a USDA surprise (*e.g.*, a WASDE report that is price-bullish for grains and oilseeds) and also companies whose share price should move in the opposite direction in response to the same event; therefore, on average across all stocks, the broad stock market returns should be the same on USDA scheduled announcement days and non-event days. Thus:

Hypothesis 1a: $E(R_M \mid \text{USDA day}) = E(R_M \mid \text{non-event day})$

where, R_M denotes the daily rate of return on the market portfolio (proxied by Standard and Poor's value-weighted S&P 500 stock market index).

In the same vein, insofar as the food sector does not drive the U.S. economy, we conjecture that the overall stock market volatility (both realized volatility and forward-looking uncertainty) to be indistinguishable on USDA event days and on non-event days. Thus, we have:

Hypothesis 1b: $E(|R_M| \mid \text{USDA day}) = E(|R_M| \mid \text{non-event day})$

Hypothesis 1c: $E(VIX \mid \text{USDA day}) = E(VIX \mid \text{non-event day})$

where, we use the absolute daily return on the S&P 500 index $|R_M|$ as a measure of realized stock market volatility and the S&P 500 option-implied volatility index, *VIX*, as the measure of forward-looking stock market volatility.

3.2.2 *Average Food-Sector Stock Returns on USDA vs. non-Event Days*

USDA reports often contain surprises regarding agricultural quantities that can move commodity prices substantially (Adjemian 2012). Some surprises are bullish for agricultural prices; some are price-bearish. Therefore, on average across all USDA events, the average excess returns for food-sector firms should be zero under the joint hypothesis that (i) the CAPM is well specified and (ii) the expected market return (R_M) is the same on event and non-event days (*i.e.*, Hypothesis 1 holds):

<u>Hypothesis 2:</u> there are no significant difference in food-sector companies' average excess stock returns on USDA scheduled announcement days *vs.* non-event days.

where, excess stock returns are defined by reference to the Capital Asset Pricing Model (CAPM).

3.2.3 *Sub-Sector Stock Returns and USDA news*

Notwithstanding Hypothesis 2, as long as Hypothesis 1a holds, we should be able to sign the excess returns for firms that belong to different food-related industry sub-sectors—based on whether the USDA surprise has implications for the total acreage planted with wheat, corn, or beans, and/or whether the surprise is price-bullish or price-bearish for agricultural commodities.

The starting point of our analysis is the fact that the quantity or demand/supply balance surprise in a USDA report has mirror implications for commodity prices, as shown in Table 3.1. For instance, lower (*resp.* higher) than expected grain stocks constitute a bullish (*resp.* bearish) commodity price signal:[38]

Table 3.1. Price-bullish vs. -bearish USDA Surprises

	WASDE	**Grain Stocks**	**Prospective Plantings**	**Acreage**
Quantity Surprise	+ / −	+ / −	+ / −	+ / −
Commodity Price Impact	− / + (Bearish / bullish)	− / + (Bearish / bullish)	− / + (Bearish / bullish)	− / + (Bearish / bullish)

Then, the connection between the excess return on a company's share and a quantity and/or commodity price surprise is obtained by identifying the impact of each surprise in the equation that defines a company's stock price as the present value of all future profits that accrue to its shareholders, properly discounted:

$$S_i = \sum_{t=1}^{\infty} \frac{\Pi_{i,t}}{(1+k)^t} = \sum_{t=1}^{\infty} \left(\frac{Rev_{i,t}(A_t, P_t) - Cost_{i,t}(A_t, P_t)}{(1+k)^t} \right) \tag{3.1}$$

where, S_I is the stock price of company i, k is the risk-adjusted discount rate for its periodic stream of profits $\Pi_{I,t}$, and $Rev_{I,t}(A_t, P_t)$ and $Cost_{I,t}(A_t, P_t)$ capture the reality

[38] See Cao and Robe (2022) and references cited therein

that company i's period *t* revenues and costs may be impacted by the acreage A_t devoted to growing a given agricultural commodity and the price P_t of the latter.

As discussed in the Introduction, if we assume that the discount rate k from Equation (3.1) is not impacted by USDA surprises, then the impact of USDA news on food-sector company excess returns (for companies that are not fully hedged against that possible impact) should depend solely on the numerator. In Sections 3.2.3, we develop hypotheses regarding how that impact varies by sub-sector.

Intuitively, the impact of USDA news on a company's stock price should depend on whether it uses grains and oilseeds for its own output, or whether it produces (or contributes to the production) of the commodities themselves. For this reason, we investigate separately the impact of USDA news on stock prices for "farm input-side" and "farm-output side" (Section 3.2.3) firms. For the same reason, we eliminate from our final sample four food-sector companies that are vertically integrated and, hence, cannot easily be classified in either category.

(1) "Farm Input-side" Companies: Machinery, Fertilizers, and Agricultural Technology

As a first-order approximation, the product development and production costs of companies that manufacture inputs into the production of agricultural commodities (*e.g.*, farm machinery, new technologies for agriculture, fertilizers) should not be affected by changes in agricultural prices (at least in the short run). The impact of USDA news, if any, thus must come from the revenue side, $Rev_{I,t}(A_t, P_t)$—potentially through both acreage A_t and commodity price P_t.

On the acreage side A_t, the more surface used for growing food, the greater the need for farm inputs—seeds, fertilizers, and even tractors or automation technology. Companies that make such inputs should therefore be positively impacted by unexpected increases in planted acres. Hence, higher than expected figures in the WASDE, Prospective Plantings (PP), or Acreage reports (AR) should be good for their stock prices, as should lower than expected grain stocks (insofar as the latter may provide the impetus for additional planting—in the Northern or Southern hemisphere, depending on the time of year).

On the commodity price side, P_t, the effect of USDA news' on those three types of companies' revenues should come from the possible relaxation of financial constraints that might curtail farmers' purchases of capital goods and production inputs. Surprises that are bullish for commodity prices and imply a relaxation of those constraints, then, should boost those companies' revenues. Together, those observations lead to our third hypothesis, which is detailed in Table 3.2:

Hypothesis 3: The effect of Grain Stock (GS) surprises on the share prices of companies that manufacture inputs for crops is unequivocally positive (*resp.* negative) when the surprise is price-bullish (*resp.* bearish). In contrast, the net effect of USDA surprises depends on the relative magnitudes of the A_t and P_t impacts in the case of the WASDE, PP and AR reports.

Table 3.2. USDA Surprises and "Farm Input-side" Company Stock Returns

	WASDE	Grain Stocks	Prospective Plantings	Acreage
USDA Quantity Surprise	+ / −	+ / −	+ / −	+ / −
Commodity Price Impact	− / + (Bearish / bullish)	− / + (Bearish / bullish)	− / + (Bearish / bullish)	− / + (Bearish / bullish)
Revenue Impact (A_t)	+ / −	− / +	+ / −	+ / −
Revenue Impact (P_t)	− / +	− / +	− / +	− / +
Stock Price Impact	?	− / +	?	?

(2) "Farm Output-side" Firms: Transformers, Distributors, Food Retailers, Restaurants

In this Section, we turn to companies that use agricultural commodities as inputs—whether directly (in the food processing sector including grain mills, brewers, or vegetable oil, meat, alcohol, biodiesel producers, etc.) or more indirectly (as would be the case for processed food wholesalers, distributors, retailers, or restaurants).[39]

In Section 3.2.3 we argued that, in the case of "input-side" companies (*i.e.*, those associated with the production of agricultural commodities), both acreage and price surprises have distinct implications for share prices—and the transmission channel is the impacts on corporate revenues. In contrast, the exposure of firms that use agricultural commodities, process them, or distribute the resulting food and products comes from the

[39] Four of the 158 firms in our sample (ADM, The Andersons Inc., Bunge Ltd., and The Seaboard Corp.) have activities tied to both the production and the transformation of grains and oilseeds. Furthermore, in their business as commodity merchandisers, those four companies are mostly exposed to calendar spread prices (*i.e.*, the net costs of storage) but not to outright commodity prices (Robe and Roberts 2019). It is therefore not clear whether (let alone how) their stock prices should react to USDA news. For this reason, we therefore exclude them from our analysis.

cost side of the corporate profit equation. The ultimate impact on profits (and, hence, on share prices) therefore depends on the price elasticity of the demand for those firms' respective products. In what follows, we assume that none of those demands is perfectly inelastic, so that any agricultural commodity price increases resulting from USDA surprises cannot be fully passed through to customers.

Hypothesis 4: the share prices of companies that use agricultural commodities experience negative (*resp.* positive) excess returns following a price-bullish (*resp.* bearish) USDA surprise. The absolute value of that excess return increases with the elasticity of the demand for those firms' products.

It is worth noting here that, historically, "increases in agricultural commodity prices have contributed little to U.S. retail food price increases, because of the small cost share of agricultural products in food prices" (Baumeister and Kilian 2014). Therefore, the impact of USDA commodity price surprises on the share prices of restaurant chains, catering firms, and food distributors and retailers (groceries) should be directionally similar to, but lower in magnitude than, that of firms in the transformation sector. Our empirical methodology, however, does not allow us to gauge the intensity of the price reactions: Hypothesis 4 therefore covers all those firms. Table 3.3 summarizes Hypothesis 4.

Table 3.3. USDA Surprises and "Farm Output-side" Food Company Stock Returns

	WASDE	**Grain Stocks**	**Prospective Plantings**	**Acreage**
USDA Quantity Surprise	+ / −	+ / −	+ / −	+ / −
Commodity Price Impact	− / + (Bearish / bullish)	− / + (Bearish / bullish)	− / + (Bearish / bullish)	− / + (Bearish / bullish)
Cost Impact $(C_{(I,t)})$	− / +	− / +	− / +	− / +
Stock Price Impact	+ / −	+ / −	+ / −	+ / −

Practically, one could run a *t*-test for Hypotheses 3 and 4 as long as one could construct an aggregate index of the USDA surprises. Since there is no obvious way to construct a "surprise index" across commodities (or across different reports whenever two or more reports are published simultaneously), the way to test Hypotheses 3 and 4 is to run regressions using commodity- and report-specific variables. To carry out those regressions, we split the sample into sub-groups (after dropping the four commodity merchandisers) between “ag-as-output” and “ag-as-input” firms.

3.3 Data

Our analysis requires financial market data and information on USDA announcements. Section 3.3.1 describes the construction of our sample of food-sector firms and the information that we gather for each firm, as well as other financial market data. Section 3.3.2 describes the data that we gather regarding USDA announcements.

3.3.1 *Stock Market Data*

We use Compustat to identify all the food-sector companies that were listed on the NYSE, AMEX or NASDAQ markets at some point between September 2009 and October 2019. We create a list of the pertinent SIC codes and use it to construct a sample comprising a wide range of companies—firms that use agricultural products as inputs (food processors, livestock producers, biofuel refiners, beverage manufacturers, restaurant or catering chains, grocery chains and food distributors) and firms that produce inputs for farmers (machinery, fertilizers, ag technology, pesticides, …). We carefully account for mergers, acquisitions, spin-offs, and de-listings (for example, if two firms merge, then they are considered separately before the merger and jointly thereafter). As discussed in Section 3.2.3, we drop four grain merchandising firms from the sample.[40] We also drop companies for which not all quarterly reports or reporting dates are available in the period when they are publicly listed.[41] Our final sample comprises 154 distinct entities.[42]

We collect the daily stock returns for the chosen stocks from CRSP between January 2009 and October 2019. For benchmarking using a 3-month or 6-month CAPM model, we also extract between January 2009 and October 2019 (i) CRSP data regarding the daily returns on the S&P 500 stock market index (SPX) and (ii) Bloomberg daily data on the 90-day and 180-day T-bill rates.

[40] See especially footnote 38.

[41] We employ Compustat for the earnings announcement dates for the period 2006-2019. We match Compustat and Bloomberg information to obtain a list of firm-quarter observations with earnings announcement time stamps.

[42] We use SEC filings, as well as data gathered from Wikipedia, Investopedia, and Google searches, to reconstruct a continuous time series for the current tickers regardless of mergers, acquisitions, splits, and other corporate events.

3.3.2 *Data on USDA Announcements*

We examine four sets of periodic USDA reports for the three main US agricultural commodities (corn, soybeans, and wheat): the monthly WASDE, the quarterly GS reports, and the annual PP and AR reports.

These four categories of reports are released on a total of 15 pre-announced dates each year (except in 2013 and 2019, when there were only 14 announcement days due to U.S. government shutdowns). Using the same tallying as in Cao and Robe (2022, p.256), "from September 2009 to October 2019, there are 120 WASDE reports, 41 GS reports (of which 10 overlap with the January WASDE), 10 PP reports, and 10 AR reports."

For each report, we also collect the results of a Bloomberg pre-event survey of corn, soybean, and wheat market analysts. Bloomberg has conducted such expert surveys since September 2009. Bloomberg News typically publishes the results of its surveys one week before the corresponding USDA event.[43] The Bloomberg survey information contains detailed information about the forecasters who participated in each survey. A typical survey summarizes the opinions of about 20 commodity analysts regarding an upcoming USDA announcement.

The Bloomberg information allows us to assess the distribution of analyst forecasts and to compute both a "consensus" value (which we set as the median analyst forecasts). Table 3.8 in the Appendix, reproduced with permission from Cao and Robe (2022), summarizes the characteristics of the 151 reports in our sample—including their coverage, frequency and timing, and key information surveyed by Bloomberg.[44] Appendix Table 3.9 provides summary statistics regarding the USDA "news" for each event (*i.e.*, the surprise defined as the difference between the actual USDA figure and the market analysts' most recent pre-release Bloomberg consensus forecast).

Altogether, our final sample encompasses 154 companies on 2,560 trading days (Table 3.10 in the Appendix lists their stock tickers, by farm's input/output side and sub-sector) and 151 USDA news events for four types of reports: WASDE, GS, AR, and PP.

3.4 **Methodology**

In this Section, we describe our approach to achieve three main tasks: (i) tease out the excess returns of food-sector company stocks using the CAPM model, so as to isolate sample-firm returns due to the USDA report release from the part of the returns due to

[43] As noted by Cao and Robe (2022), "the exact timing of the result release is not documented in the survey dataset, so (one must) recover it by tracing back each release on Bloomberg News manually."

[44] For more details about the figures of interest in each report, see Cao and Robe (2022).

overall stock market price movements on USDA event days; (ii) constructing proxies for the news component in the USDA reports; and (iii) testing our hypotheses using these measures.

3.4.1 *Estimating excess returns*

Financial asset returns are potentially affected by two types of risk: systematic risk and idiosyncratic risk. While the former affects every stock in the market and thus cannot be diversified, the latter is specifically related to a certain asset (or sector) and thus can be diversified by combining different assets (or sectors) into a portfolio. Since food-sector companies only account for a small fraction of the U.S. stock market broadly defined, the risk tied to USDA-announcements should be considered as sector-specific and thus irrelevant to the market risk premium.

According to the Capital Asset Pricing Model (CAPM) (Sharpe 1964; Treynor 1961a, 1961b), the expected asset returns as a reward for non-diversifiable market risk can be estimated over a certain period of time using the relation

$$R_{i,t} = R_f + \beta_i\left(R_{M,t} - R_f\right) + \varepsilon_{i,t} \tag{3.2}$$

where:

$R_{i,t}$ is the periodic return on company *i*

R_f is the risk-free interest rate

$R_{M,t}$ is the periodic return on the market portfolio

$\varepsilon_{i,t}$ is a residual or excess return component that is orthogonal to the market risk.

In our empirical analysis, we use the S&P 500 stock index returns as proxies for the market returns and U.S. T-bill rates as proxies for the risk-free interest rates.[45]

For each USDA event day, we must separate price movements stemming from the undiversifiable market risk from the potential impact of USDA announcement on stock returns. To do so, we first obtain an estimate of β_I in Equation (3.2) for each stock *i* and each trading day in our sample. We use two rolling windows (of either 90 or 180 days) to do so. We end each regression 10 days prior to the day for which we need to compute an excess return, in order to abstract from food-sector stock price movements that might result from the publication of Bloomberg surveys of commodity analysts (which takes

[45] 3-Month T-bill for the 90-day window, and 6-Month T-bill for the 180-day window

place up to 7 business days prior to a scheduled USDA report). Second, we use the relevant β_I estimate (as well as the contemporaneous risk-free interest and market returns) to predict each stock's return on each given trading day. Finally, we subtract those expected returns from the actual stock returns to generate a daily time series of excess returns $\varepsilon_{I,t}$ for each stock *i*.

A large literature in finance and accounting shows that another component of individual stock returns is the price change around the company's earnings announcement (EA) (*e.g.*, Kross and Schroeder 1984; Nichols and Wahlen 2004). On the one hand, for each company in our sample, we could exclude the subset of EA dates that overlap with the USDA dates when performing statistical tests. On the other hand, overlapping EA-USDA event days account for a very small fraction of the observations in our sample (typically ranging from 0-2%, with a maximum of 7% of the active sample of a stock). For this reason, we retain all observations (including the overlapping ones) and add a dummy control variable that takes the value 1 when a USDA event day overlaps with an EA day.

3.4.2 *Measuring the news component of USDA reports*

Like macroeconomic and corporate announcements, scheduled USDA reports contain anticipated and unanticipated (*i.e.*, news or "surprise") components. Including the part of its content that had been expected by market participants could bias the estimate of a report's market impact. Therefore, to unbiasedly assess the effect of a report on stock returns, we must first extract that report's news component.

As noted in Section 3.3, we use the Bloomberg analyst surveys before each USDA report to assess market experts' expectations regarding the upcoming report. Specifically, for each report-date, we follow Cao and Robe (2022) and compute the surprise as the log-difference between the actual USDA announced figure and the median ("consensus") Bloomberg forecast. The same procedure is applied for three commodities: corn, soybeans, and wheat—which are the most important crops grown in the United States. Appendix Table 3.9 provides summary statistics regarding the consensus forecasts and USDA news for each USDA event.

3.4.3 *Testing methodology*

In this Section, we present the methodologies that we employ to test the five hypotheses presented in Section 3.2.

(1) Testing hypotheses 1a to 1c and 2: Compare excess returns on USDA announcement days vs. non-USDA announcement days

Assuming that the market returns and food-sector companies' estimated excess returns are all normally distributed, these hypotheses can be tested using two-sample *t*-test on the USDA days *vs.* non-USDA days. Alternatively, the nonparametric Kruskal-Wallis (KW) test can also be applied to account for the case of nonnormality, including the market forward-looking volatility (proxied by the VIX index). For each test, the null hypothesis is that there is no difference in the mean (*t*-test) or median (KW test) values on USDA days *vs.* "normal" days.

In the case of all food companies' excess stock returns (*i.e.*, for testing hypothesis 2), one key issue for the test is the choice of the "normal" benchmark for non-USDA daily returns. To avoid any potential price movement in the run-up to, and also just after, the announcement day t, we compute the average excess returns of three trading days before day *t-2* (that is, from day *t-5* to day *t-3*) and three trading days after day t+2 (*i.e.*, from day *t+3* to day *t+5*). When any of these normal days happens to be an EA day as well, we replace it by one day further backward or forward. For consistency, we apply the same procedure to test Hypotheses 1a, 1b and 1c.

(2) Testing Hypotheses 3 and 4: Regression analysis

We carry out a standard panel data analysis to assess the role of USDA news on the excess returns of food-sector company stocks on USDA days. As discussed in Section 3.2.3, we run the regression for two subsets of excess returns: "farm input-side" (Hypothesis 3) and "farm output-side" (Hypothesis 4) companies. To capture the differentiated effects of commodity price-bearish *vs.* price-bullish surprises in the USDA reports, we split the surprise variable into commodity price-bearish (denoted S_t^+) and commodity price-bullish (denoted S_t^-) surprises. For each of our two subsets of food sector-linked firms, the general regression equation is thus given by:

$$ER_{i,t} = \mu + \gamma^+ S_t^+ + \gamma^- S_t^- + \delta Control_{i,t} + \alpha_i + v_{i,t} \tag{3.3}$$

where $i = 1, 2, \dots, 154;\ t = 1, 2, \dots, 2560; and$:

$ER_{i,t}$ is the excess returns of stock *i* on trading day *t*;[46]

μ is the overall intercept;

[46] $ER_{i,t}$ in Equation (3.3) is the estimate of $\varepsilon_{i,t}$ in Equation (3.2).

S_t^+ is a column vector of the commodity price-bearish news components ("surprises") of all commodities of interest for all USDA announcements released on day t.[47] Since we are interested in three commodities (corn, soybeans, wheat) reported in three types of USDA reports (WASDE, GS, and Planted Area[48]), S_t^+ has 9 elements in total, each of which is

cross-sectionally invariant (*i.e.*, for each day t, all stocks have the same vector S_t^+);

S_t^- is a column vector of the nine commodity price-bullish news components that we construct analogously to S_t^+;

γ^+ is the row vector of coefficients that capture the marginal effect of the USDA price-bearish surprises on stock excess returns;

γ^- is the row vector of coefficients that measure the marginal effect of the USDA price-bullish surprises on stock excess returns;

$Control_{i,t}$ is a vector of control variables that capture other time-varying characteristics of firms (the firm's market capitalization the month before, plus a dummy variable indicating whether the USDA event day is also an earnings announcement day for that firm);

$\alpha_i + v_{i,t} \equiv \zeta_{i,t}$ is a composite error that contains a firm-specific, time-invariant unobserved effect α_I and a zero-mean, homoscedastic, non-serially-correlated disturbance v_{it}.

We are interested in the estimates of the influence of USDA report surprises, which are captured in the vectors $\widehat{\gamma^-}$ and $\widehat{\gamma^+}$.

Since not all the companies in our samples have full-length return observations throughout the sample period (September 2009 to October 2019),[49] our sample constitutes an unbalanced panel of stock returns with sample lengths varying from stock to stock.[50]

When it comes to the question of whether firm fixed-effect (FE) or random-effect (RE) estimators should be used to estimate Equation (3.3), two arguments favor FE. Firstly, as

[47] For each type of report, for those trading days when there is no scheduled report, the values of both surprise variables are set equal to zero.

[48] By construction, the Acreage report (released end of June) is a mid-season update for the Prospective Planting report (released end of March). Therefore, to keep the model parsimonious, we follow Cao and Robe (2022) and group these two reports into one single group called "Planted Area."

[49] Some companies start to trade publicly, or are delisted, after September 2009. Furthermore, some firms are taken over or merge at some point in our sample.

[50] There are 86 out of 154 companies with full-length sample, *i.e.*, with 2,560 trading day and no missing values.

discussed, we hypothesize that the extent to which a company's stock price reacts to USDA information depends on the sub-industry to which the company belongs. For example, the implication of an unexpected increase in wheat acreage should be different for manufacturers of fertilizer *vs.* mills or restaurant chains. Since the sub-industry to which it belongs is a time-invariant characteristic of a company, its effect can be captured in α_I, together with other time-invariant characteristics that may affect how a specific stock generally reacts to news. The possible correlation between α_I and other time-varying USDA-related explanatory variables will cause a bias in $\widehat{\gamma^-}$ and $\widehat{\gamma^+}$ and, therefore, should be eliminated—which supports using FE estimators. Secondly, for unbalanced panels, the latter usually produce more robust estimates (Wooldridge 2020, p. 447).[51]

The panel methodology applied in our regressions draws heavily on Wooldridge (2020). We carry out our empirical analysis using the MATLAB Panel Data Toolbox developed by Álvarez, Barbero, and Zofío (2017).

3.5 **Results**

In this Section, we present the results for each of the tasks described in Section 3.4.

3.5.1 *CAPM model estimates of excess returns*

For each firm and trading day, we estimate a CAPM model with two different rolling windows (90 and 180 trading days) and two estimation methods (OLS and maximum-likelihood estimation with expectation maximization algorithm to account for missing data (Dempster, Laird and Rubin 1977). Figure 3.1 plots the resulting excess returns from these four models, for each of our 154 firms on the x-axis, on 151 USDA event days (blue) *vs.* the normal baseline on non-USDA days (orange) on the y-axis.

Looking at Figure 3.1, we see qualitatively similar results for the four models. There is a clear pattern that excess returns tend to be larger in absolute value on USDA days (blue dots), whereas for non-USDA days they are concentrated more closely around zero (orange dots). In all four models, however, the average (not absolute) excess returns seem to be distributed around 0.

[51] In addition to those theoretical arguments, we find empirical evidence to support the FE model in our setting. Specifically, in Tables 3.6 and 3.7 below, we estimate both the FE and RE models and we show, using Hausman (1978) and Mundlak (1978) tests, that the FE estimator is indeed preferred to the RE estimator. Regardless, the two estimators yield similar numerical estimates.

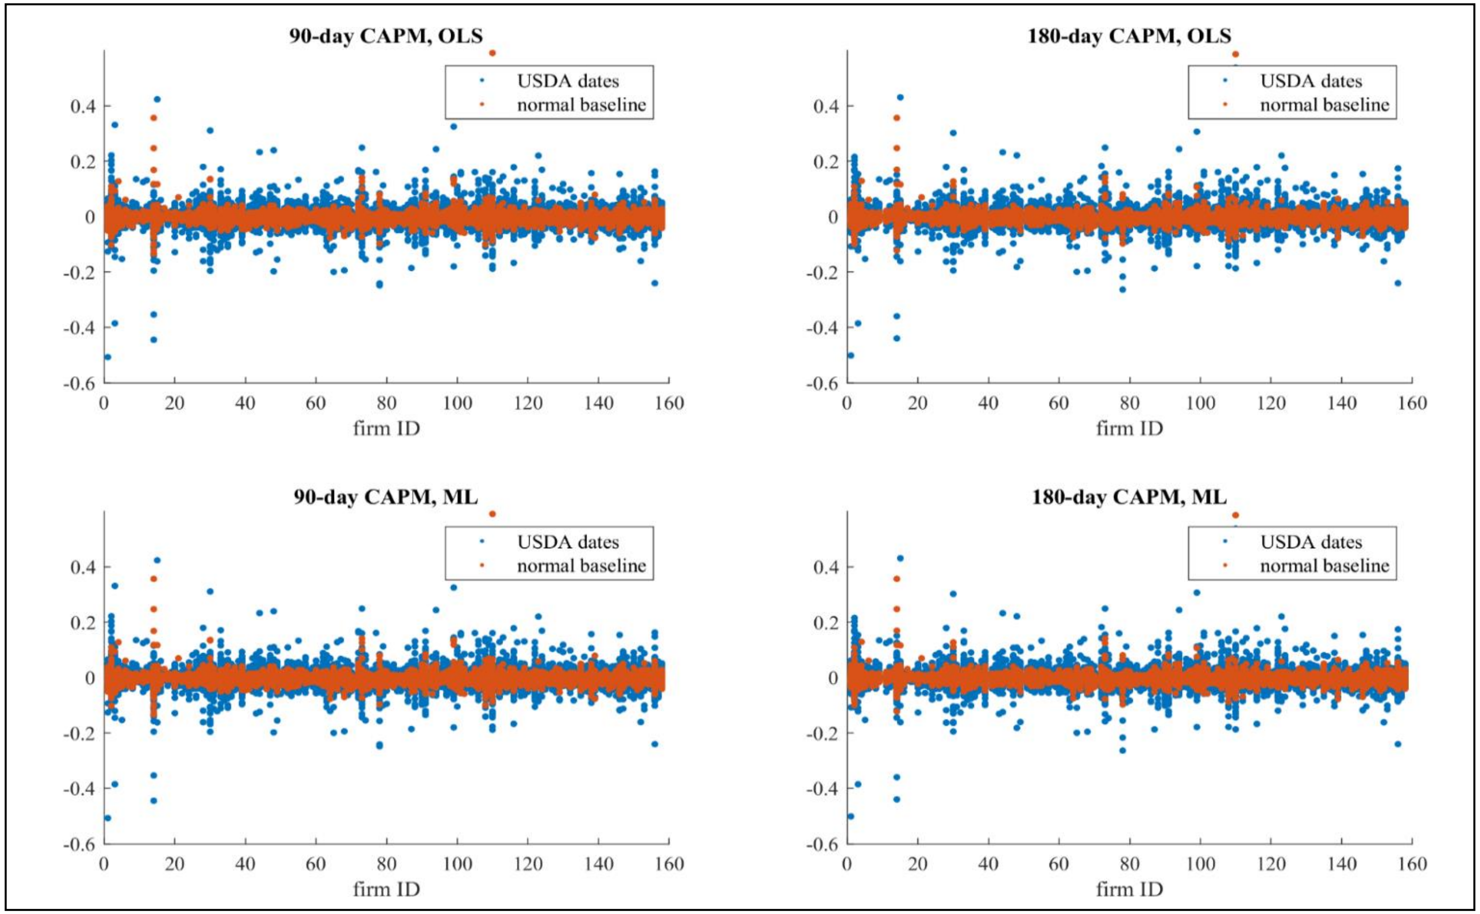

Figure 3.1. Alternative Estimations for Excess Returns Plotted by Firm ID Number

Table 3.4 provides summary statistics for the raw and the estimated excess returns on both USDA and non-USDA event days. The Table focuses on the 180-day OLS estimation for brevity, as the empirical results are similar across these four specifications.[52] Table 3.4 indicates that, both for the raw and for the estimated excess returns, the difference between the returns on USDA event days *vs.* on normal days is minimal on average, as both are close to zero and have large variance.

Table 3.4. Summary Statistics – Food Sector Raw and Excess Returns on USDA-announcement Days

	Median	Mean	SD	Min	Max	Jarque-Bera test statistics	No. non-missing Obs
A. Raw returns							
All trading days	0.0003	0.0006	0.0264	-0.4915	3.1123	**$1.5\text{x}10^{10***}$**	308,393
"Normal" baseline	0.0000	0.0005	0.0395	-0.4894	2.6182	**$1.0\ \text{x}10^{9***}$**	18,184
USDA days	0.0014	0.0019	0.0268	-0.4894	1.1830	**$4.4\ \text{x}10^{7***}$**	18,184
B. 180-day CAPM excess returns, all firms							
All trading days	-0.0015	-0.0016	0.0257	-0.5014	3.1304	**$1.9\text{x}10^{10***}$**	308,393
"Normal" baseline	-0.0010	-0.0016	0.0120	-0.1225	0.5860	**$1.1\text{x}10^{8***}$**	18,184
USDA days	-0.0012	-0.0011	0.0262	-0.5014	1.1587	**$4.8\text{x}10^{7***}$**	18,184

Statistical significance code: *** 0.01 **0.05 *0.10

[52] Since the 180-day rolling window generally requires more non-missing return data to provide plausible estimates, 6 are dropped out of the initial sample of 160 firms, resulting in a total number of 154 firms in the final sample.

Table 3.4 (cont.). Summary Statistics – Food Sector Raw and Excess Returns on USDA-announcement Days

	Median	Mean	SD	Min	Max	Jarque-Bera test statistics	No. non-missing Obs
C. 180-day CAPM excess returns, Input firms							
All trading days	-0.0000	0.0003	0.0284	-0.5014	3.1304	**$1.8x10^{10***}$**	54,496
"Normal" baseline	0.0003	0.0003	0.0145	-0.0754	0.5860	**$9.2\ x10^{7***}$**	3,212
USDA days	-0.0006	-0.0002	0.0278	-0.5014	0.5373	**$9.1\ x10^{5***}$**	3,212
D. 180-day CAPM excess returns, Output firms							
All trading days	-0.0019	-0.0020	0.0250	-0.4875	2.5848	**$4.1x10^{9***}$**	253,897
"Normal" baseline	-0.0013	-0.0020	0.0114	-0.1225	0.3561	**$6.0\ x10^{6***}$**	14,972
USDA days	-0.0014	-0.0013	0.0259	-0.4398	1.1587	**$5.5\ x10^{7***}$**	14,972

Statistical significance code: *** 0.01 **0.05 *0.10

3.5.2 *Hypotheses 1a to 1c and 2*

Table 3.5 summarizes the testing results for Hypotheses 1a to 1c and 2, using paired sample t-tests and Wilcoxon signed-rank tests. Since the Jarque-Bera test statistics (in Table 3.4) consistently reject the normality of the return distributions for both raw and excess returns, we focus in what follows on the results of the nonparametric Wilcoxon signed rank tests.

As Panel A in Table 3.5 shows, we cannot reject the null hypothesis that the S&P 500 index returns are not statistically significantly different on USDA announcement days *vs.* on non-announcement days. This result is consistent with our intuition in Section 3.2.1, that any potential significant movements of food-sector stocks on USDA announcement days should either be (1) cancelled out across all food-sector firms or (2) too small to move the stock market as a whole.

Panels B and C of Table 3.5 show that, in general, there is no significant difference in stock market volatility between USDA event and non-event days, consistent with Hypotheses 1b and 1c. Using the absolute returns of the S&P 500 index, $|R_{SPX}|$, as the proxy for realized market volatility, Panel B shows that we cannot reject the null hypothesis that the realized volatility of the stock market is no different on USDA days compared to normal days. Likewise, Panel C shows that we cannot reject the null hypothesis that the VIX (*i.e.*, the stock market's forward-looking volatility) is not affected by USDA events.

Figure 3.2 plots the distribution (by firm) of *p*-values of Wilcoxon signed rank test statistics of USDA-days excess returns against the "normal" average excess returns, as discussed in Section 3.2.2. Consistent with what is suggested by the summary statistics, Figure 3.2 shows that the p-values of the test are "spread evenly" between 0 and 1. Altogether, there are only 15 stocks (out of 154) for which the test's *p*-value is lower than 0.1. Without conditioning on the news content of the reports, at this stage, the statistical evidence suggests that USDA events have very little impact on food-sector companies (except for a small number of firms)—as predicted by Hypothesis 2.

Table 3.5. Statistical Test Results for Hypotheses 1a to 1c and 2

	Null hypothesis	Alternative hypothesis	Paired sample t-test	Wilcoxon signed rank test
Panel A. Hypothesis 1a	$E(R_{SPX}$ | USDA) = $E(R_{SPX}$ | non-USDA)	$E(R_{SPX}$ | USDA) ≠ $E(R_{SPX}$ | non-USDA)		
Test-statistics			0.664	0.942
p-value			0.508	0.346
Conclusion			Cannot reject the null	Cannot reject the null
Panel B. Hypothesis 1b	$E(\|R_{SPX}\|$| USDA) = $E(\|R_{SPX}\|$| non-USDA)	$E(\|R_{SPX}\|$| USDA) ≠ $E(\|R_{SPX}\|$| non-USDA)		
Test-statistics			0.464	-0.528
p-value			0.643	0.598
Conclusion			Cannot reject the null	Cannot reject the null
Panel C. Hypothesis 1c	E(*VIX* | USDA) = E(*VIX*| non-USDA)	E(*VIX* | USDA) ≠ E(*VIX*| non-USDA)		
Test-statistics				-0.813
p-value			Not applicable	0.416
Conclusion				Cannot reject the null
Panel D. Hypothesis 2	$E(ER_i$ | USDA) = $E(ER_i$ | non-USDA)	$E(ER_i$ | USDA) ≠ $E(ER_i$ | non-USDA)		
Test-statistics			See Figure 3.2	See Figure 3.2
p-value				
Conclusion			Cannot reject the null	Cannot reject the null

Statistical significance code: *** 0.01 **0.05 *0.10

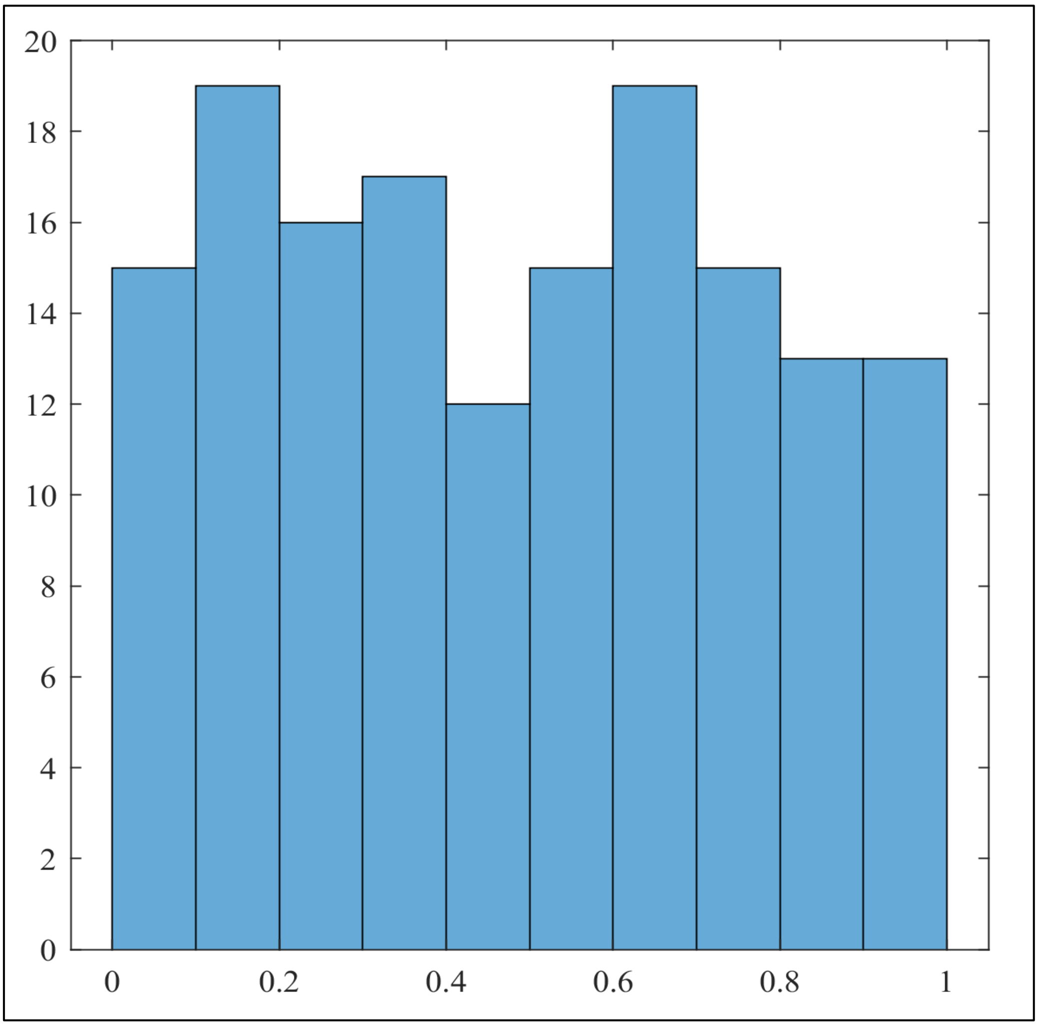

Figure 3.2. Histogram of *p*-Values of Wilcoxon Signed Rank Test Statistics of USDA Excess Returns against the "Normal" Average Excess Returns

3.5.3 Testing Hypotheses 3 and 4

This Section focuses on the results of the unbalanced fixed-effect panel regressions described in Section 3.4.3. This part of the analysis is designed to tease out the different dimensions of USDA reports' effect on firms' excess returns.

(1) Hypothesis 3: "Farm input-side" companies

Table 3.6 reports the estimates of $\widehat{\gamma^-}$ (*resp.* $\widehat{\gamma^+}$), *i.e.*, of the effect of different types of USDA reports' price-bullish (*resp.* bearish) surprises on the USDA days' excess stock returns for the 26 farm-input suppliers in our sample. Heteroskedasticity-robust standard errors are reported. According to Hypothesis 3 (see Table 3.2), the signs of the coefficients for WASDE and Planted Area surprises are ambiguous and to be determined empirically. In contrast, for GS reports, $\widehat{\gamma^+}$ should be negative for price-bearish surprises to have a negative impact on those firms' excess returns. Likewise, $\widehat{\gamma^-}$ should also be negative for price-bullish GS surprises to have a positive impact on the excess returns of farm input producers.[53]

Table 3.6 shows that the WASDE surprises, as expected from the above discussion, have a mixed effect on the excess returns across three commodities. Moreover, that impact is generally small. Starting with wheat, we see that a one percent WASDE surprise (from the pre-event Bloomberg consensus forecast) for ending-stock projections only brings about a (statistically significant) 0.046 percent negative excess return (compared to non-USDA days) when the surprise is price-bearish. But a statistically significant average negative excess return of similar magnitude (0.063 percent) is also observed when the wheat surprise is price-bullish. For corn, in contrast, a price-bullish WASDE surprise causes a statistically significant positive excess return (an average 0.067 percent per one percent ending-stocks surprise), whereas the effect of a price-bearish surprise is insignificant. Soybean ending-stock projection surprises do not significantly move excess returns of firms in our sample in any direction. From those varied results, we conclude that it is ambiguous whether the cash-flow expectation adjustment (which causes the excess returns, as shown in Equation (3.1), and is induced by WASDE surprise) is driven more by the adjustment in acreage expectation or in the commodity output price expectation.

In contrast to the mixed results for WASDE, Table 3.6 shows that the GS surprise coefficients that are statistically significant are all negative, consistent with Hypothesis 3. First, a price-bearish corn GS surprise (which means the measured realized inventory level is higher than expected) yields 0.062 percent negative excess returns per one percent of higher-than-expected stock level. Second, the impact of a price-bullish wheat stocks announcement is even stronger: 0.22 percent negative excess returns per one percent of lower-than-expected stock level. The rest of the GS coefficients are not statistically significant.

The Planted Area surprise coefficients in Table 3.6 are the largest in magnitude among the three types of announcements. It is apparent that price-bearish planting surprises significantly precede negative excess returns (almost 0.4 percent for corn), while price-bullish acreage surprises cause positive excess returns (nearly 0.5 percent for corn and 0.2 percent for soybeans). Unlike the case of WASDE, here the cash-flow expectations

[53] By construction, bullish surprises only take negative values, so they must be multiplied by a negative coefficient to yield positive excess returns.

through the commodity price channel P_t clearly dominates the acreage (*i.e.*, quantity) channel A_t (see Table 3.2).

Table 3.6. Unbalanced Panel FE Estimates of the Effect of USDA Surprises on Food-sector Stock Excess Returns on USDA Announcement Days, Farms' Input Producer Firms†

	Corn		Soybeans		Wheat	
	Coefficients	Robust standard errors	Coefficients	Robust standard errors	Coefficients	Robust standard errors
WASDE Bearish Surprise	0.027	0.033	-0.019	0.014	**-0.046****	**0.019**
Grain Stocks Bearish Surprise	**-0.062*****	**0.018**	-0.020	0.018	-0.001	0.005
Planted area Bearish Surprise	**-0.357*****	**0.010**	0.139	0.171	0.150	0.127
WASDE Bullish Surprise	**-0.067*****	**0.016**	-0.010	0.007	**0.063****	**0.029**
Grain Stocks Bullish Surprise	0.037	0.027	0.033	0.021	**-0.222****	**0.081**
Planted area Bullish Surprise	**-0.486****	**0.218**	**-0.227****	**0.089**	0.060	0.147
EA day dummy‡	-0.0001	0.0014				

Statistical significance code: *** 0.01 **0.05 *0.10

No. stocks (N):	26
No. USDA days (T):	151
No. non-missing daily return observations (n):	54,496
F Statistic of individual effects:[§]	**12.96*** **
Wooldridge's test statistics:[§§]	0.005
Hausman test statistics:[⁋]	4.09
Mundlak's test statistics:[⁋⁋]	**179.07*** **

<u>Note</u>: Table 3.6 reports the unbalanced panel FE estimation results for the subset of stocks of companies which produces farms' input, including machinery, fertilizer, pesticides, and biotechnology. Heteroskedasticity-consistent standard errors are reported.

[†] For all report-commodity combinations, the independent variables of interest are the surprises (percentage deviation of the actual USDA information from the Bloomberg median forecast). Given the way these independent variables are measured, their regression coefficients characterize the "elasticity" of excess returns (also measured in the percentage terms) to the correspondent change in the interest variables on the day of USDA announcement.

[‡] To avoid misunderstandings, the EA day dummy is a stand-alone explanatory variable in the regression equation, without being interacted with any other variable.

[§] *F*-test for the null hypothesis that all unobservable individual effects, *i.e.*, α_I, are not significantly different from zero. The test statistics follows an *F*-distribution with (*N-1, n-N-K*) degrees of freedom under the null, with *K* being the number of time-varying regressors in the model.

[§§] Wooldridge's test is used to test the null hypothesis that there is no serial correlation among the residuals. The test statistics follows an *F*-distribution with (*1, N*) degrees of freedom under the null.

[⁋] Hausman test is used to test the null hypothesis that there is no significant difference between RE and FE estimated coefficients. Thus, a rejection of the null means RE estimator is consistent. FE is estimator is consistent under both the null and the alternative hypothesis (Wooldridge 2020). The test statistics follows a Chi-squared distribution with K degrees of freedom under the null, with *K* being the number of time-varying regressors in the model.

[⁋⁋] Mundlak test tests the null hypothesis that individual means are zero, which imply that RE model should be preferred. The test statistics follows a Chi-squared distribution with *K* degrees of freedom under the null.

(2) Hypothesis 4: “Farm output-side” companies

Analogously to Table 3.6, Table 3.7 reports the FE coefficients $\widehat{\gamma^-}$ and $\widehat{\gamma^+}$ estimated using the subset of 128 “farm-output side” firms (*i.e.*, companies that use farms’ agricultural production as their input). Those include food/beverage processors, biofuel producers, catering/restaurant chains, and food retailers/distributors. According to Hypothesis 4 (see Table 3.3), we expect the coefficients for all three USDA report types to be positive, such that a price-bearish USDA surprise will cause positive excess returns, and a price-bullish one will cause negative excess returns.

For the majority of the statistically significant coefficients in Table 3.7, this turns out to be the case. For example, a one percent bearish Planting Area surprise is estimated to yield about 0.3 percent positive excess returns (for corn and soybeans), whereas a bullish Planting Area surprise of the same size causes about 0.2. percent negative excess returns (for soybeans and wheat). As with the farms’ input producers, here Planted Area news also has the largest impact on commodity buyer stocks.

For GS reports, we find that only price-bearish corn stock surprises significantly move excess returns upward. However, the size of effect is very modest: 0.025 percent excess returns per one percent of higher-than-expected actual corn stocks.

Finally, as for “input-side firms”, WASDE surprises again have a mixed and very subtle impact on the excess returns of the present subgroup of food company shares. Price-bearish news about soybean and wheat projected ending-inventories bring about statistically significant excess returns (as expected from Table 3.3), but only by 0.012 and 0.032 percent, respectively. In contrast, corn price-bearish WASDE surprises bring about 0.023 percent negative excess returns for “farm output-side” firms—which is not predicted by Hypothesis 4. However, that same-size corn surprise, when price-bullish, also moves the excess returns downward to almost the same extent: 0.014 percent, which is consistent with Hypothesis 4. Another unexpected effect is the wheat price-bullish surprise, which cause a positive excess return of 0.06 percent on average.

Table 3.7. Unbalanced Panel FE Estimates of the Effect of USDA Surprises on Food-sector Stock Excess Returns on USDA Announcement Days, Farms' Output Buyer Firms†

	Corn		Soybeans		Wheat	
	Coefficients	Robust standard errors	Coefficients	Robust standard errors	Coefficients	Robust standard errors
WASDE Bearish Surprise	**-0.023*****	**0.005**	**0.012*****	**0.004**	**0.032*****	**0.007**
Grain Stocks Bearish Surprise	**0.025*****	**0.007**	-0.008	0.010	0.021	0.023
Planted area Bearish Surprise	**0.272*****	**0.073**	**0.275*****	**0.054**	-0.142	0.104
WASDE Bullish Surprise	**0.014*****	**0.005**	0.002	0.004	**-0.060*****	**0.011**
Grain Stocks Bullish Surprise	0.003	0.014	-0.000	0.005	-0.021	0.032
Planted area Bullish Surprise	-0.094	0.071	**0.161*****	**0.046**	**0.246*****	**0.053**
EA day dummy‡	-0.000	0.001				

Statistical significance code: *** 0.01 **0.05 *0.10

No. stocks (N):	128
No. USDA days (T):	151
No. non-missing daily return observations (n):	253,897
F Statistic of individual effects:[§]	**21.29*****
Wooldridge's test statistics:[§§]	0.79
Hausman test statistics:[⁋]	**92.72*****
Mundlak's test statistics:[⁋⁋]	**959.80*****

Note: Table 3.7 reports the unbalanced panel FE estimation results for the subset of stocks of companies who are the buyers of farms' output, including food/beverage processors, biofuel producers, catering/restaurant chains and food retailers/distributors. Heteroskedasticity-consistent standard errors are reported.

[†] For all report-commodity combinations, the independent variables of interest are the surprises (percentage deviation of the actual USDA information from the Bloomberg median forecast). Given the way these independent variables are measured, their regression coefficients measure the "elasticity" of excess returns (also measured in the percentage terms) to the correspondent change in the interest variables on the day of USDA announcement.

[‡] To avoid mis understanding, the EA day dummy is a stand-alone explanatory variable in the regression equation, without being interacted with any other variable.

[§] F-test for the null hypothesis that all unobservable individual effects, *i.e.*, α_I, are not significantly different from zero.

[§§] Wooldridge's test is used to test the null hypothesis that there is no serial correlation among the residuals.

[⁋] Hausman test is used to test the null hypothesis that there is no significant difference between RE and FE estimated coefficients. Thus, a rejection of the null means RE estimator is consistent. FE is estimator is consistent under both the null and the alternative hypothesis (Wooldridge 2020).

[⁋⁋] Mundlak test tests the null hypothesis that individual means are zero, which imply that RE model should be preferred.

3.5.4 *Robustness*

The results so far focus on results with excess returns based on a CAPM estimated using OLS over a 180-day rolling window. As mentioned earlier, we also estimated the CAPM excess returns using different rolling window lengths (*i.e.*, 90 days instead of 180 days) and performed the analysis on both. Results are similar for both, except that the standard errors tend to be larger for the 90-day variant—both in Equation (3.2) and Equation (3.3). We find qualitatively similar results using maximum-likelihood estimators.

Despite the trivial fraction of USDA event days that overlap with earning announcement days, we control for this confounding factor by including in Regression Equation (3.3) a dummy variable that takes the value 1 if a firm announces its earning on that trading day. Results (reported in Table 3.6 and Table 3.7) show that, for the firms in our sample in the 2009-2019 period, an earnings announcement drift (as suggested by earlier literature) is not visible.

3.6 **Policy implications**

The magnitude of the USDA news' impact that we find for food-sector equities is smaller than the price reaction identified in commodity futures markets (see, *e.g.*, Adjemian 2012)—but that is as it should be: equity prices discount cash-flows for numerous years (over the course of which the relevance of a given USDA report will fade), while agricultural futures prices reflect the current supply-demand balance of a storable but ultimately perishable commodity. Regardless, the point of our analysis is to show that there is an equity-market impact and that it is statistically significant.

An obvious policy implication is that USDA announcements matter to the broad economy. That is, while "plenty of research investigates whether USDA are influential on commodity futures markets from a variety of angles" (Ying, Chen, and Dorfman 2019, p. 991), our paper identifies a new dimension along which to measure and quantify the value of this public information that costs tens of "millions of dollars to collect and disseminate" (Karali *et al.*, 2019, p. 66). Precisely, we provide empirical evidence that the USDA's public reports do more than to "facilitate the efficient functioning" of, and to "reduce information asymmetries and facilitate the policy and program formation, operation and evaluation processes" in, commodity markets (*ibid.*, p.67): in fact, USDA reports' informational impact ripples beyond their immediate reflection in the commodity space, with USDA news reflected in the prices of long-term assets such as food-sector equities.

A vast literature in finance and accounting investigates the impact of earnings surprises. In a Journal of Finance study using a period (1983-2015) that partly overlaps with ours, Chiang *et al.* (2019) find that a one standard deviation increase in their measure of surprise generates a 1.5 percent cumulative abnormal stock return three days after a quarterly earnings announcement. This kind of magnitude comes to about 0.5 percent per day, and thus it explains both the attention paid to analyst forecasts by investors and the care taken by corporate insiders to provide markets with corporate guidance.[54] While the kind of surprise that we investigate is different, the one-day excess returns that we find following USDA grain stocks (*resp.* acreage) surprises are 0.22 percent (*resp.* 0.5 percent) for some commodities. A second implication of our findings is that corporate insiders and investors should thus pay attention to USDA news. In line with this recommendation, equity research teams who track firms in the "food/ags," "clean energy," "ferts/chems," "machinery" or related sub-sectors pay close attention to the same USDA reports that we analyze in the present paper, and they provide information to their clients regarding scheduled USDA releases.[55]

A third policy implication directly follows from the above. It relates to the importance of security measures to prevent information leakage prior to the release of scheduled agency reports.[56] Market authorities typically focus on commodity derivatives trading venues in the case of EIA and USDA reports. Yet, options trade on many of the stocks in our sample, and a simple back-of-the-envelope computation shows that a trader could earn a juicy one-day return of over ten percent on an at-the-money (ATM) one-month stock option following a 0.5 percent stock price reaction to an Acreage surprise.[57] This reality suggests

[54] Many firms manage market expectations of their earnings—see Bertomeu *et al*. (2021) and references therein—and a majority of companies start issuing earnings guidance right after they first go public (Allee *et al*. 2020).

[55] To take just one example, the October 12, 2022 OPCO Equity Research report on Green Plains Inc., an ethanol producer in our sample, refers to the September 2022 WASDE report in to forecast demand for the company's output.

[56] See Adjemian and Irwin (2018) for a discussion of the USDA lockup and other safety procedures around report releases. See also Huang, Serra, and Garcia (2021) for a discussion of how analyses that (unlike the present paper) rely on very high-frequency data must account for possible intraday information leakage.

[57] The 10 percent figure assumes 20 percent stock price volatility, 30 days to option expiration, no dividends, and 5 percent interest rate; Appendix Table 3.11 give returns for various combinations of parameter values (volatility, time to maturity). We compute the option returns using a plain Black-Scholes model, as the percentage difference in option price resulting from a 0.5 percent change in the underlying stock price. The delta of an ATM option is always one half, so the dollar price change is half the underlying stock price change; however, because the option price is much less than the stock price, the percentage return is much higher for the option than for the stock. Put differently, option positions are "levered"—which is why unscrupulous traders use options to trade on ill-gotten information (see, *e.g.*, Augustin, Brenner, and Subrahmanyam (2019), Fishe and Robe (2004), and references cited in those papers).

that market regulators may have to look further afield (namely, in equity markets) for possible malfeasance involving USDA information.

3.7 Conclusion

Our paper provides the first investigation of stock market reactions to scheduled USDA reports. As our starting point, we separate food-sector firms' daily raw returns into event-day market-risk reward and excess return, and focus our analysis on the latter. To that end, we compute excess returns for 154 publicly-listed U.S. food-sector companies in our sample, using 180-day rolling windows, between 2009 and 2019. Both parametric and nonparametric tests show that, when averaged across all types of reports and across all grain and oilseeds, the prices of food-sector stocks in general do not react significantly to USDA surprises.

In contrast, we document that food-sector stocks do respond, once one controls for the commodity, type of USDA report, and direction and extent to which the USDA news surprises the market. In particular, stock price responses to USDA news differ between firms on the input-side *vs.* firms on the output-side of agricultural (farm) production, contingent on which component of firm's cash-flow expectations (costs or revenues, price or quantity) is impacted by the news.[58]

Our empirical results support several intuitive hypotheses. On the one hand, for firms that supply inputs to farms (fertilizer, pesticides, machinery, or technology), a positive deviation from expected future commodity supply delivered in the WASDE and in the Prospective Planting or Acreage reports should be a "good news" (as more input will be needed). At the same time, the same news also implies expected falling agricultural commodity prices, thus tightening farmers' cash constraint for further investments. In that sense, this kind of surprise should be "bad news" for those firms too. In contrast, a higher-than-expected actual inventory level (in a Grain Stocks report) ought to consistently reduce those firms expected future cash-flows, given that plantings elsewhere in the world should adjust downward for the next harvest. On the other hand, for firms that use agricultural commodities as inputs, we should expect a more consistent pattern of stock price reactions for all types of USDA reports, given that news in all these reports affects the expectation of such firms' margins through the same channel (namely, input costs). That said, a larger-than-expected commodity supply implied by any of these reports should in turn imply a cost reduction for these firms so long as the demand for their own

[58] For this set of analyses, we exclude from the sample vertically-integrated, publicly-listed agribusinesses that cannot be easily classified as either input- or output-side firms.

products is not perfectly inelastic, and thus leads to an upward revision of firms' expected cash-flows.[59]

Overall, our empirical results support these hypotheses statistically in the case of Planted Area and Grain Stock news. Planted Area surprises have the largest absolute effect on returns (with up to 0.5 percent excess returns) for both subsets of firms (ags-as-inputs and ags-as-output), followed by the GS reports. In both cases, the effects have the expected sign. In contrast, for both subsets of firms, we find that WASDE surprises have a very modest (and mixed) impact on stock returns.

Our study is the first to shed light on the question of whether, and how, USDA report releases have an impact beyond commodity markets in general, and the stock prices of food-sector firms specifically. Our findings have important policy implications, which we spell out in Section 3.6. They also suggest additional paths for further research. First, our findings are at the daily frequency. Research in agricultural economics, however, shows that the commodity futures-market impact of USDA news takes places withing ten minutes of the announcement (*e.g.*, Lehecka, Wang, and Garcia 2014). It would be interesting to revisit our analysis using intraday data, as the high frequency would sharpen the computation of excess returns. Second, our analysis abstracts from the possibility that some companies may hedge their exposure to commodity price fluctuations. Future research could investigate the relationship between firms' hedging decisions and the extent to which their stock prices react to USDA news.

[59] The reverse holds true for a negative news.

Acknowledgement

We are grateful to Mike Adjemian, Alexander Belozertsev (our CEMA discussant), Phil English, Thomas Heckelei, Scott Irwin, Kristen Owen, Fei Qin (our AAEA discussant), and Hugo Storm for helpful discussions. We thank session participants at the 2022 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management (NCCC-134 in St. Louis, MO), the 2022 annual meeting of the Commodity and Energy Markets Associations (CEMA in Chicago, IL), and the 2022 annual meeting of the Agricultural and Applied Economics Association (AAEA in Anaheim, CA) for useful comments on a first draft titled "USDA Announcements and the Stock Prices of Food-Sector Companies." Cao gratefully acknowledges financial support from the Bonn International Graduate School at the University of Bonn's Agricultural Faculty. Robe's contribution was partly supported by an ACES Academy for Global Engagement seed grant and by a NIFA grant under Hatch Multistate Project Number S1072. Robe also gratefully acknowledges the financial support received in his capacity as The Clearing Corporation Foundation Professor in Derivatives Trading at the University of Illinois. While this paper was written, Robe served as a part-time consulting senior economist at the CFTC; no CFTC resources or confidential information were used for this project. The views expressed in this article are those of the authors only and may not be attributed to the CFTC, the Commissioners, other CFTC staff, or the United States government. The authors are solely responsible for any error or omission.

3.8 References

Adjemian, M.K. (2012). “Quantifying the WASDE Announcement Effect.” *American Journal of Agricultural Economics*, 94(1), 238–256. https://doi.org/10.1093/ajae/aar131

Adjemian, M.K. and S.H. Irwin (2018). “USDA Announcements Effects in Real Time.” *American Journal of Agricultural Economics*, 100(4), 1151–1171.

Allee, Kristian D., Theodore E. Christensen, Bryan S. Graden, and Kenneth J. Merkley. (2020). “The Genesis of Voluntary Disclosure: An Analysis of Firms’ First Earnings Guidance.” *Management Science*, 67(3), 1914–1938.

Augustin, P., M. Brenner, and M.G. Subrahmanyam (2019). “Informed Options Trading.” *Management Science*, 65(12), 5697–5720.

Álvarez, I. C., Barbero, J. and Zofío, J. L. (2017). A Panel Data Toolbox for MATLAB. *Journal of Statistical Software*, 76(6), 1–27. https://doi.org/10.18637/jss.v076.i06

Baumeister, C., and L. Kilian (2014). Do Oil Price Increases Cause Higher Food Prices? *Economic Policy*, 29(80), 691–747.

Bertomeu, Jeremy, Edwige Cheynel, Edward Xuejun Li , and Ying Liang. (2021). “How Pervasive Is Earnings Management? Evidence from a Structural Model.” *Management Science*, 67(8), 4643–5300.

Cao, A. N. Q. and Robe, M. A. (2022). Market uncertainty and sentiment around USDA announcements. *Journal of Futures Markets*, 42(2), 250–275.

Chiang, C.-H., W. Dai, J. Fan, H. Hong, and J. Tu. (2019). “Robust Measures of Earnings Surprises.: *Journal of Finance*, 74(2), 943–983. https://doi.org/10.1111/jofi.12746

Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1), 1–22.

Ellison, B.D., and J.L. Lusk (2011). “Taxpayer Preferences for USDA Expenditures.” *Choices*. Quarter 2. Accessed online, January 21, 2023: http://choicesmagazine.org/choices-magazine/submitted-articles/taxpayer-preferences-for-usda-expenditures

Fishe, R.P.H. and M.A. Robe (2004). “The Impact of Illegal Insider Trading in Dealer and Specialist Markets: Evidence from a Natural Experiment.” *Journal of Financial Economics*, 71(3), 461–488.

Goyal, R. And M.K. Adjemian (2021). "The 2019 Government Shutdown Increased Uncertainty in Major Agricultural Commodity Markets." *Food Policy*, 102, 102064

Hausman, J. A. (1978). "Specification Tests in Econometrics." *Econometrica*, 46(6), 1251–1271.

Huang, J., T. Serra, and P. Garcia (2021). "The Value of USDA Announcements in the Electronically Traded Corn Futures Market: A Modified Sufficient Test with Risk Adjustments." *Journal of Agricultural Economics*, 72(3), 712-734.

Huang, R. D., Masulis, R. W. and Stoll, H. R. (1996). Energy shocks and financial markets. Journal of Futures Markets, 16(1), 1–27.

Isengildina-Massa, O., Irwin, S. H., Good, D. L. and Gomez, J. K. (2008). "Impact of WASDE Reports on Implied Volatility in Corn and Soybean Markets." *Agribusiness*, 24(4), 473–490. https://doi.org/10.1002/agr.20174

Karali, B., Irwin, S. H. and Isengildina-Massa, O. (2019). "Supply Fundamentals and Grain Futures Price Movements." *American Journal of Agricultural Economics*, 102(2), 548-568.

Karali, B., Isengildina-Massa, O., Irwin, S. H., Adjemian, M. K. and Johansson, R. (2019). "Are USDA Reports still News to Changing Crop Markets?" *Food Policy*, 84 (April), 66-76.

Kilian, L. and C. Park. (2009). "The Impact of Oil Price Shocks on the U.S. Stock Market." *International Economic Review*, 50(4), 1267-1287. https://www.jstor.org/stable/25621509

Kross, W. and Schroeder, D. A. (1984). "An Empirical Investigation of the Effect of Quarterly Earnings Announcement Timing on Stock Returns." *Journal of Accounting Research*, 22(1), 153–176.

Kurov, A., A. Sancetta, G. Strasser, and M. Halova Wolfe (2019). "Price Drift before US Macroeconomic News: Private Information about Public Announcements?" *Journal of Financial and Quantitative Analysis*, 54(1), 449-479.

Lehecka, G.V., Wang, X. and Garcia, P. (2014). "Gone in Ten Minutes: Intraday Evidence of Announcement Effects in the Electronic Corn Futures Market." *Applied Economic Perspectives and Policy*, 36: 504-526. https://doi.org/10.1093/aepp/ppu010

McNew, K. P. and Espinosa, J. A. (1994). "The Informational Content of USDA Crop Reports: Impacts on Uncertainty and Expectations in Grain Futures Markets." *Journal of Futures Markets*, 14(4), 475–492. https://doi.org/10.1002/fut.3990140408

Mundlak, Y. (1978). "On the Pooling of Time Series and Cross Section Data." *Econometrica*, 46(1), 69–85.

Nichols, D. C. and Wahlen, J. M. (2004). "How do Earnings Numbers Relate to Stock Returns? A Review of Classic Accounting Research with Updated Evidence." *Accounting Horizons*, 18(4), 263–286.

Pruitt, J. Ross, Glynn T. Tonsor, Kathleen R. Brooks, and Rachel J. Johnson (2014). "End User Preferences for USDA Market Information. *Food Policy*, 47, 24–33.

Robe, M. A. and Roberts, J. S. (2019). Who holds positions in agricultural futures markets. Available at SSRN 3438627.

Sharpe, W. F. (1964). "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *Journal of Finance*, 19(3), 425–442.

Treynor, J. L. (1961a). "Market Value, Time, and Risk." Unpublished manuscript, August 8.

Treynor, J. L. (1961b). "Toward a Theory of Market Value of Risky Assets." Unpublished manuscript.

USDA, NASS (2020). "USDA NASS Discontinues County Level Estimates for Selected Crops." Accessed March 2, 2023: https://www.nass.usda.gov/Newsroom/Notices/2020/02-13-2020.php

Wooldridge, J. M. (2020). *Introductory Econometrics* (7th Edition): Cengage Learning.

Ying, J., Chen, Y. and Dorfman, J. H. (2019). "Flexible Tests for USDA Report Announcement Effects in Futures Markets." *American Journal of Agricultural Economics*, 94(1), 238. https://doi.org/10.1093/ajae/aaz013

3.9 Appendix

Table 3.8. USDA Reports Overview

	WASDE	Grain Stocks (GS)	Prospective Plantings (PP)	Acreage (AR)
Frequency	Monthly	Quarterly	Yearly	Yearly
Timing	2nd week of the month	2nd week of January & End of 1st-3rd Quarters	End of March	End of June
Overlap	1st GS (January)	1st WASDE; PP; AR	2nd GS (March)	3rd GS (June)
Information surveyed by Bloomberg	Projected U.S. ending stock of the on-going marketing year	U.S. Ending stock estimates as of 1st Dec, 1st Mar, 1st Jun and 1st Sep	U.S. farmers' planting intention for upcoming crop season	Survey-based estimate of U.S. planted area for current crop season

Note: Table 3.8, drawn from Cao and Robe (2022), describes the 151 USDA reports that we collect for our sample from September 2009 through October 2019. On some dates, the USDA releases more than one report: the third row in the Table (labeled "Overlaps") explains which of the WASDE, GS, PP and AR reports overlap. For part of the empirical analysis (see Section 3.2.3), we include information regarding expert opinions prior to the USDA news release. The information regarding analyst opinions comes from periodic Bloomberg surveys of market experts.

Table 3.9. Summary Statistics – Bloomberg Surveys and USDA News

	Median	Mean	SD	Min	Max	No. Obs	Obs < 0
A. Corn							
WASDE surprise	0.004	0.006	0.077	-0.242	0.326	121	52
Grain Stocks surprise	0.002	0.011	0.068	-0.165	0.196	41	20
Planted Area surprise	0.004	0.007	0.018	-0.017	0.055	20	8

	Median	Mean	SD	Min	Max	No. Obs	Obs < 0
B. Soybean							
WASDE surprise	0.000	0.000	0.101	-0.310	0.452	121	55
Grain Stocks surprise	-0.011	0.001	0.091	-0.346	0.265	41	26
Planted Area surprise	-0.004	-0.008	0.021	-0.078	0.034	20	15

Table 3.9 (cont.). Summary Statistics – Bloomberg Surveys and USDA News

	Median	Mean	SD	Min	Max	No. Obs	Obs < 0
C. Wheat							
WASDE surprise	0.006	0.007	0.039	-0.139	0.138	121	45
Grain Stocks surprise	0.012	0.007	0.029	-0.074	0.055	41	16
Planted Area surprise	0.004	-0.001	0.016	-0.038	0.018	20	9

Note: Table 3.9 provides summary statistics for the event-day "USDA surprise" relative to analysts' consensus forecast prior to the USDA scheduled event. The sample runs from September 2009 through October 2019 and covers 151 USDA reports in that period—see Table 3.8. The information in Table 3.9 for corn and soybeans is similar to Table 2 in Cao and Robe (2022).

Table 3.10. Sample Firms

Sector	Sub-Sector	Stock Tickers
Farm's input-side Firms	Farm Machinery & Technology (15)	AGFS, BLT, TITN, ALG, ARTW, CAT, DE, LNN, SANW, TSCO, TTC, TWI, UAVS, AG/AGCO, CNH/CNHI
	Fertilizers & Pesticides (11)	DD, DOW, DWDP, SYT, AVD, CF, FMC, MBII, MON, MOS, RKDA
Farm's output-side Firms	Beverages & Food Processors (60)	AQB, BETR, BNNY, BUD, GMCR, HNZ, KFT, KRFT, MDLZ, PF, POST, RAH, SFD, BGS, BRID, CAG, CALM, CPB, DAR, DEO, FARM, FLO, FRPT, GIS, HAIN, HRL, HSY, JBSS, JJSF, K, KHC, KO, LANC, LW, LWAY, MED, MGPI, MKC, NBEV, PEP, PPC, REED, RMCF, SAFM, SAM, SJM, SMPL, STZ, SXT, TAP, THS, TR, TSN, TWNK, BREW/HOOK, CPO/INGR, EAST/ESDI, FIZ/FIZZ, HANS/MNST, DPS/KDP
	Catering & Restaurant Chains (43)	BOBE, CPKI, ARKR, ARMK, BDL, BJRI, BLMN, CAKE, CBRL, CHUY, CMG, DENN, DNKN, DPZ, DRI, EAT, FAT, FRGI, GTIM, JAX, LOCO, LUB, MCD, NATH, NDLS, PBPB, PLAY, PZZA, RRGB, RUTH, SBUX, SHAK, STKS, TACO, TAST, TXRH, WEN, WING, YUM, BBQ/DAVE, DIN/IHP, JACK/JBS, PZZI/RAVE
	Food Retailers & Distributors (17)	HTSI, KR, CASY, CHEF, COKE, CORE, HFFG, IFMK, NGVC, PFGC, SFM, SVU, SWY, UNFI, USFD, WMK, WFM/WFMI
	Biofuels (8)	AMRS, GEVO, GPP, AMTX, FF, GPRE, REGI, REX/RSC
Integrated	Merchandisers (4)	ADM, ANDE, BG, SEB

Table 3.11. At-the-money Option Returns

stock price	Volatility		1-month call	1-month put	3-month call	3-month put
			Panel A: At-the-Money option prices			
100	15		$ 1.94	$ 1.52	$ 3.64	$ 2.39
	20		$ 2.51	$ 2.10	$ 4.61	$ 3.37
	25		$ 3.09	$ 2.67	$ 5.60	$ 4.36
	30		$ 3.66	$ 3.24	$ 6.58	$ 5.34
			Panel B: Return from a 0.5% underlying stock price rise			
100.5	15		14.7%	-14.1%	8.2%	9.0%
	20		11.1%	-10.6%	6.3%	6.5%
	25		8.9%	-8.4%	5.1%	5.1%
	30		7.5%	-7.0%	4.3%	4.2%
			Panel C: Return from a 0.5% underlying stock price drop			
99.5	15		-13.5%	15.6%	-7.8%	9.0%
	20		-10.4%	11.4%	-6.1%	6.5%
	25		-8.5%	8.9%	-5.0%	5.1%
	30		-7.2%	7.3%	-4.2%	4.2%

Note: Table 3.11 shows the returns that would accrue to the holder of an at-the-money stock option, from a 0.5 percent change in the underlying stock price. Panel A shows the prices of at-the-money calls and puts on a $100 stock, for various values of the volatility on the underlying asset (percent, annualized) and of the options' times to maturity. Panel B shows the option returns from a stock price increase, Panel C, from a stock price decrease.

Chapter 4
Market surprises, machine learning and USDA Crop Progress and Condition reports[60]

Abstract: Traditionally, a predefined surprise proxy (such as the consensus errors of analyst forecasts) is used to estimate the market impact of public announcements. We instead use the post-event price movements to tease out what the market consensus must have been, and to estimate the event-day surprises. Our empirical analysis focuses on the USDA's weekly Crop Progress and Condition reports (CPCRs), which we show can be forecasted using weather "big" data. Departing from conventional Machine Learning (ML) approaches, we create a new ML routine to incorporate important features of a market expectation model under the Efficient Market Hypothesis' semi-strong form. We find that the market often overestimates the condition of both crops by about 5-6%, with occasional spikes up to 22%. Moreover, these surprise estimates suggest that the reports still cause significant post-release market reactions, though of small magnitudes.

***JEL classification*:** C53, G14, Q02, Q11

***Keywords*:** Market Surprises, Market Expectations, Machine Learning, Crop Condition, Commodities, Scheduled News, USDA announcements

---

[60] Chapter 4 is pre-published as a working paper titled Cao, A. N. Q., Gebrekidan, B. H., Heckelei, T. and Robe, M. A. (2023). Market surprises, machine learning and USDA Crop Progress and Condition reports. *SSRN Electronic Journal*. Advance online publication. https://ssrn.com/abstract=4515193.

4.1 Introduction

Public announcements regarding market fundamentals greatly facilitate the price discovery process, and thereby improve market efficiency (Adjemian 2012; Baur and Orazem 1994; Brunnermeier 2005; Ederington and Lee 1993; Garcia *et al.* 1997). In the spirit of the Efficient Market Hypothesis (EMH, Fama *et al.* 1969), most studies to date focus on estimating a post-release market impact (*e.g.*, on price returns, volatility, etc.) to assess the extent to which such announcements bring new information to the market. This new information component (*i.e.*, the market surprise) is the difference between the announced figure and the *ex-ante* market expectation about that figure. Unfortunately, the market expectations are unobservable *ex-ante*, and thus so are the surprises *ex-post*. As a remedy, most studies to date use a pre-event consensus forecast (Chiang *et al.* 2019) as a proxy for *ex-ante* market expectations. However, this proxy is prone to serious measurement errors, which may be either correlated or uncorrelated with the true market expectations.[61] Consequently, the resulting estimates of market impact using such proxy are subject to biases, and one might not be able to draw correct conclusions about the market impact of the reports. Additional efforts have been devoted to developing new estimators (*e.g.*, Karali, Irwin, and Isengildina-Massa 2019; Rigobon and Sack 2008) as well as to constructing alternative analyst-forecast-based expectations (*e.g.*, Chiang *et al.* 2019; Hirshleifer, Lim and Teoh 2009). Still, all these studies stick to analyst forecasts as the starting point of the surprise computation, and thus the risk of getting biased estimates of market impact persists – especially the one induced by incentive-driven biases.

These considerations have two implications. One, to see how much new information a public announcement brings to the market requires to examine the surprise component itself in magnitude and pattern, and not only the estimated market impact. Two, to that effect, an analyst-forecast-free method of extracting the surprises must be considered. These are the two main goals of this paper.

Towards both ends, we develop a novel nonparametric framework to extract the surprises of the announcement content without relying on analyst forecasts. The intuition underlying our approach is simple: regardless of whether the true surprise causes a large market reaction or not, it should be the most correlated with price movements compared to other proxies with measurement errors, provided that the measurement errors satisfy classical error-in-variable assumptions. Hence, in the absence of the true news, we can select the proxy that explains the most variation in the price returns among a set of candidates, when the set is large and able to generate unbiased proxies for the news. Critically, rather than computing some pre-defined "consensus forecast error" using

[61] Bartolini (2008) provides a helpful discussion of the problems associated with using such a proxy.

analyst surveys and taking its impact on prices as the impact of the news, as is done in much of the extant accounting, financial, and agricultural economics literatures, we let the post-event price signals speak out and choose the proxy that best explains its movements. In our theoretical framework, based on which the surprise candidates are generated, we allow for the expectation formation process to change over time to incorporate potential structural changes affecting the way in which market participants access and process information.

For our analysis, we choose a report that is widely known to market participants, and that can be forecasted using objective, highly accessible and publicly available information. For these reasons, we focus on a weekly report that (we show) can be forecasted unbiasedly using high-resolution weather data: the Crop Progress and Condition Report (CPCR) published weekly by the U.S. Department of Agriculture's (USDA) National Agricultural Statistics Service (NASS).

The CPCR's main role is to keep market participants up to date on the progress and the condition of major U.S. crops (like corn and soybeans) during their respective growing seasons. Despite its subjective, survey-based nature, the CPCR is a timely, comprehensive, and accurate assessment of crop development across the main U.S. production areas (Beguería and Maneta 2020). As long as a crop's condition during the growing period determines the final yield, news regarding the development of crops are a crucial determinant of agricultural prices (Boudoukh *et al.* 2007; Schnepf 2006; Stevens 1991). Thus, one would expect that the reports are valuable to market participants and, therefore, that CPCR surprises have significant post-release impact on the market. It is true that some monthly or quarterly USDA reports, such as the World Agricultural Supply and Demand Estimates (WASDE) and grains stock reports, command more attention in the academic literature. Furthermore, some of the literature to date (Bain and Fortenbery 2017; Dorfman and Karali 2015; Lehecka 2014; Ying, Chen and Dorfman 2019) suggests that CPCRs impact commodity prices less than other USDA reports do, and that the impact has been falling over time as the market has gotten better at anticipating CPCR's content. This said, a number of other studies demonstrate how the CPCRs can be used effectively for early in-season yield forecasting (Beguería and Maneta 2020; Irwin and Hubbs 2018; Kruse and Smith 1994) and as a ground-truth validation source for remote-sensing-based crop mapping (Gao *et al.* 2017; Wardlow, Kastens and Egbert 2006; Worrall, Rangarajan and Judge 2021). As we submitted at the beginning of this introduction, these two seemingly contradictory sets of findings suggest that the surprise component itself must be investigated in order to properly assess the contribution of the CPCRs to market participants and other user groups.

In 2021, NASS introduces the "county-level" gridded CPCRs dataset for the years starting from 2015, in additional to the traditional state-level data. The present paper is the first to deploy this fine-scale dataset, together with weather data at high-resolution (both spatially and temporally), to produce a large set of predicted conditions of corn and soybeans during 2016-2021. We choose ML to accomplish this task due to its flexibility in dealing with complex relationships (such as the relationship between weather and crop condition/yield) as well as its capacity to handle such large amount of data efficiently. However, we depart from conventional, built-in ML cross validation routines to select the "best predicting model," and create a new routine that can best reflect our theoretical framework and produce numerous candidates for market expectations in the way that is similar to the expectation formation process. The development of this routine is a second, technical contribution of our paper.

Precisely, our routine is developed on the basis of Extreme Gradient Boosting (XGB) – a powerful nonparametric ML algorithm. We create a hyperparameter space that contains 12,800 combinations of six important XGB hyperparameters, and let the models produce 12,800 surprise series candidates for each year-crop in our sample period. Among these candidates, we choose the ones that explain the most variations in the post-release price movements.

Results show that our models can generate highly accurate predictions for both U.S. crop condition (*i.e.*, less than 0.5 percent away from the actual release on average). Consistent with our theoretical prediction, however, we find that these best predictions are consistently rejected by market returns as being appropriate proxies for the market's expectations. Instead, using the post-event price movements to identify the latter, our algorithm picks the forecasts that are far away from the best predictions and the median predictions as the best proxies for the market expectations. Our empirical analysis indicates that both corn and soybean markets tend to be five to six percent overoptimistic (in log-difference terms) about crop condition during our sample period. In extreme cases, actual corn and soybean condition can be 18 percent (*resp.* 22 percent) below what the market had been expecting. This is a third contribution of our paper (to the agricultural economics literature), beside the broader methodological contributions (to the financial economics literature) of using a tailored ML routine and correlations between forecast errors and event-day price reactions to tease out the market surprise.

We provide evidence that the market surprises derived from our models can plausibly characterize the developments of crop cycles in the sample period. For example, the overoptimistic expectations in 2021 are consistent with the fact that yield expectations in both markets were always at records that year – according to historical statistics of yield surveys provided by NASS Quick Stats database, despite the continuously worsening crop condition reports as the year progressed. Putting our proxies together for the whole period 2016-2021, we find that the sign and magnitude of the surprises that we identify are consistent with the theory: a report of better (*resp.* worse) than expected crop condition is price-bearish (*resp.* bullish). Remarkably, market reactions to the crop condition news are statistically significant during this period, though with small size.

Our work contributes to the extant literature in at least three different ways. First, we demonstrate that machine-learning routines, when developed rigorously based on market theories, have a great potential to be an effective method to derive a parameter-free, analyst-forecast-free proxy for the surprise component of public consensus information, in particular the USDA reports. We are not aware of any other work in which supervised machine-learning routines are employed to separate the unanticipated component of public consensus information. Second, to the best of our knowledge, we are also the first to effectively exploit market returns to identify the underlying surprise without having to resort to analyst forecasts.[62] Third, we provide empirical evidence that contradicts the notion that the CPCRs nowadays provide only a limited amount of new information to the public. This may change the perception of various user groups of the reports, as well as how they incorporate the information content of the reports into many relevant estimates.

The remainder of the paper proceeds as follows. Section 4.2 describes the CPCRs, as well as the weather and market data used in our analyses, which provides the reader with the necessary context for corn and soybean markets. Section 4.3 discusses several important theoretical aspects of the market expectation formation model, in the context of crop condition forecasting under the semi-strong form EMH. Section 4.4 builds on that discussion and introduces the design of the new ML routine. Section 4.5 discusses our empirical findings. Section 4.6 concludes.

4.2 The new geospatial CPC dataset: an opportunity for ML analysis

The CPCR is a survey-based, weekly report that has been produced by NASS since 1986. It is the most comprehensive and frequently published information source for U.S. main crops' progress and condition. The report is released every first business day of the week, with information available for each crop throughout its growing season.

[62] Our approach has some connection to the principal component analysis (PCA) exercise of Rigobon and Sack (2008). But, differently from them, we construct an explicit information set for market expectations and do not merely rely on market price data. Thus, our model outcomes can be interpreted more intuitively than the PCA factors. Yet, our model selection approach does not prevent us to draw valid inferences about the market impact of this surprise component.

Here we briefly introduce the newly available "county-level" gridded crop condition dataset[63] developed by NASS and its potential advantages compared to the traditional state-level data. This new dataset is an important motivation for our use of ML methods in estimating the unanticipated report content. The weather dataset as well as commodity futures market data are also described in this Section.

4.2.1 *Two innovations for a gridded crop condition dataset*

Figure 4.1 illustrates the difference of the new format (left) *vs.* the traditional format (right) of the crop condition dataset. Until 2020, CPCRs were only published at the state level, using the standard format shown in panel (a). Though the surveys have always been carried out at county level, one goal of the state-level aggregation is to protect farmers' confidentiality (USDA 2021). For each state-crop, the CPCR reports the proportions of acreage in five rating categories summing up to 100%: very poor (VP), poor (P), fair (F), good (G), and excellent (EX) conditions. The crop condition is reported only for the crop's main production states.

Deriving a reliable proxy for *ex-ante* market expectations of this dataset is challenging for several reasons. First, unlike quantitative information such as grain stocks or demand and supply estimates, a crop's condition is a qualitative metric and there is no unique approach of quantifying it. To our knowledge, there exists no public guidance on how to classify a certain crop into the five aforementioned categories. This ordinal structure and the lack of reference are critical obstacles for parametric inferences of pre-event expectation of the CPCR, as such inferences would typically require restrictive distributional and functional assumptions on the data generating processes.

Second, the CPCRs are highly spatially aggregated; in contrast, crop condition predictors (in particular, weather variables) are available at much higher spatiotemporal resolutions. Weather variables interact spatially, temporally, and cross-sectionally (*e.g.*, between temperature and precipitation) in a complex fashion (Schlenker and Roberts 2009; Westcott and Jewison 2013). Furthermore, the production of corn and soybeans spans a wide geographic area of the USA characterized by heterogeneous growing conditions. In such a setting, arbitrary model specifications and data aggregation – as would be required in traditional parametric analyses – are likely to cause a serious loss of information.

Finally, while nonparametric proxies of market consensus derived from pre-report analyst surveys seem promising in reducing bias (Chiang *et al.* 2019), they were not available in

[63] Both the traditional and the new formats of the CPCR contains two main indicators: (i) crop progress and (ii) crop condition. In this paper, we only focus on the information about crop condition. For more discussion about the crop progress, see Hubb and Irwin (2020a) and Lehecka (2014).

the case of CPCR until recently, likely due to the lack of common understanding of the crop rating criteria as discussed above.

In response to the rising demand for finer-scale analysis (USDA 2021), NASS introduced in 2021 weekly gridded layers for both progress and condition data of corn, soybeans, and wheat. NASS made the information available retroactively, going back to 2015.[64] A visualized example of this dataset is provided in panel (b) of Figure 4.1. Though the dataset is produced using the same surveys that are used for the state-level reports, NASS has made efforts to bring it much closer to county level. The new spatial resolution is 9x9 kilometers, which potentially reduces the loss of information due to spatial aggregation. Furthermore, instead of discrete categories, the new dataset provides continuous indices of the crop condition based on equally-spaced integer values attached to each category/development phases. [65] Though it is not clear how the numerical values assigned to the categories/phases were chosen, it has the potential to be an official reference for how market agents could combine the categorical information. [66] With these two important innovations, the gridded presentation of the CPCRs opens an opportunity for more flexible nonparametric methods to exact the surprise component of the report content. In this paper, we propose a highly flexible ML approach to accomplish this task.

[64] According to NASS, it is unlikely that the gridded data will be produced for the years before 2015. Consequently, our current sample only contains seven years from 2015 to 2021, which is relatively short compared to previous studies. Also, this dataset does not replace but complements the traditional CPCR format, which is still regarded as the official format of the reports.

[65] See the dataset documentation provided by NASS at:
https://www.nass.usda.gov/Research_and_Science/Crop_Progress_Gridded_Layers/CropProgressDescription.pdf

[66] Another caveat according to NASS is that the fine-scale dataset might be subject to some discrepancies from the original figures when aggregated back to state- and country level crop condition. This is due to smoothing and interpolation in the production of the fine-scale gridded dataset, as the data collecting and processing process was initially designed for state-level reports only USDA (2021). If the difference is large enough, it might introduce serious measurement errors into the new dataset and consequently bias the estimated market expectations. Thus, it is necessary to assess the numerical differences between the two data formats, as will be discussed in Section 4.5.1.

4.2.2 *Weather and market data*

With recent rapid developments in information technologies, weather data have become publicly accessible with much finer spatial and temporal resolutions and in ever more timely fashion. A natural question is to which extent these developments have contributed to the predictability CPCRs content.

Corn Condition by Percent

	VP	P	F	G	EX
CO	26	22	28	20	4
IL	2	5	28	51	14
IN	4	9	26	47	14
IA	9	15	29	41	6
KS	7	12	27	39	15
KY	0	2	8	49	41
MI	3	8	32	46	11
MN	1	3	16	52	28
MO	2	4	16	61	17
NE	6	13	22	42	17
NC	6	10	32	42	10
ND	2	8	31	47	12
OH	4	11	36	43	6
PA	9	16	40	25	10
SD	2	4	15	67	12
TN	1	3	22	59	15
TX	5	14	39	31	11
WI	2	4	17	48	29
18 Sts	5	9	25	46	15
Prev Wk	5	9	25	46	15
Prev Yr	4	10	30	45	11

(a) Traditional crop condition report format

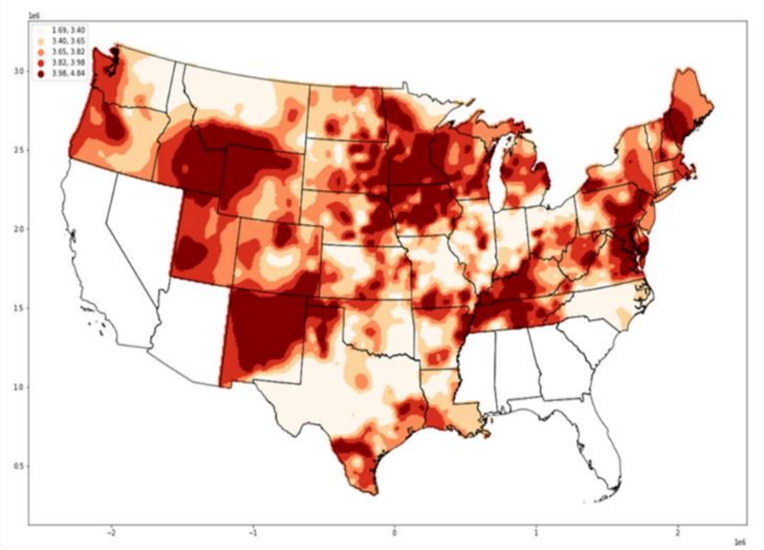

(b) New crop condition report format

Figure 4.1. Gridded, "County-level" Crop Condition (left) *vs.* Traditional, Ordinal State-level Crop Condition Formats (right)

In this paper, we approach predictability using the daily Parameter-elevation Regressions on Independent Slopes Model (PRISM) weather dataset, which has a 4x4 kilometer spatial resolution. This dataset was developed, and is maintained by Oregon State University's PRISM Climate Group. Alongside Daymet, it is one of the two most widely-used sources of weather and climate data for the US continent in studies about the impact of climate and weather on agriculture (Mourtzinis *et al.* 2017; UCAR 2022). The daily data comprise seven weather variables: maximum temperature (*Tmax*), minimum temperature (*Tmin*), average temperature (*tmean*), precipitation (*ppt*), mean dew point temperature (*tdmean*), minimum vapor pressure deficit (*vpdmin*), and maximum vapor pressure deficit (*vpdmax*).[67]

[67] For more detailed description about the dataset, see: https://climatedataguide.ucar.edu/climate-data/prism-high-resolution-spatial-climate-data-united-states-maxmin-temp-dewpoint

In our sample period (2015-2021), all trading of CBOT corn and soybean futures is electronic after July 2015, and pit trading's share of the total CBOT futures trading volume was small in the runup to pits' closure (Gousgounis and Onur 2018). On the CME's Globex (*i.e.*, the electronic platform used for CBOT futures trading), each trading day has two sessions. The "U.S. night" session starts at 7PM Central Time (CT) on U.S. calendar day $t-1$ (which is formally called the "day t open", because it is already calendar day t in Asia) and ends at 7:45AM CT on U.S. calendar day t. The "U.S. day" session that starts at 8:30AM CT and closes at 1:20PM CT on U.S. calendar day t. For the purposes of our price analyses, we follow the approach adopted by Lehecka (2014). We use the "day t close"-to-"day $t+1$ open" price returns to capture the market impact of the CPCRs, given that the latter are released at 4PM on calendar day t (*i.e.*, at a time when the U.S. grain and oilseed futures markets are closed). We use close-to-open future returns of "new-crop" (i.e., December) corn and (November) soybean contracts traded on the Chicago Mercantile Exchange (CME). As their name suggests, these contracts expire shortly after the harvest of the current crops. Thus, news about current crop condition is directly relevant to the market valuation of these contracts. We use open and close (settlement) prices as reported on Bloomberg.

As discussed in Sections 4.4.3 and 4.5.1 below, we also use the Acreage Planted data from NASS Quick Stats (USDA 2017) to calculate county- and state-level crop production shares when it comes to the national crop condition surprises.

4.3 Theoretical framework

The futures prices of agricultural commodities reflect the aggregate expectations about what their spot prices will be at a certain point in the future. Futures price movements thus follow changes in expectations about market fundamentals. For a CPCR to cause a market reaction at the time of its release, it must be that market participants rely on the reports to revise their crop yield expectations. This is because once a crop has been planted, its condition during the planting season is a crucial determinant of the end-of-season crop size.[68] As a crop's condition changes continuously during the growing season, traders will revise their yield expectations accordingly, which will be reflected in price movements. If this updating process is efficient, then the price movement caused by each CPCR will only reflect traders' reaction to the unanticipated portion of the information released in that report – since all the previously publicly available information has already been incorporated into the market prices beforehand. This is the

[68] For a thorough discussion about the relationship among weather, crop condition and expected yield, see Bundy and Gensini (2022) and the references therein.

basic assumption of semi-strong form EMH. Based on this intuition, we present a simple model of crop condition market expectation revision.

Suppose that the market's expectation of crop condition C_t in period t can be approximated by some model $f(.)$, using publicly available information I_t at time t. The market surprise S_t is then the forecast error of this forecasting process,

$$S_t = C_t - f(I_t) \tag{4.1}$$

Over time, market expectations are revised as new information arrives. But a change in market expectations can also be due to changes in the way the market processes the same information, such as advancements in data analyzing techniques and computing power or in the structure of the data-generating processes.[69] Thus, we should allow this expectation formation process to be updated after some periods. To see why this is important, suppose that in period $t+1$, $f(.)$ is no longer the best approximation to market expectations, but is instead replaced by some other function $h(.)$. The surprise of this period is then the forecast error of the updated process

$$S_{t+1} = C_{t+1} - h(I_{t+1}) \tag{4.2}$$

At the same time, the new crop condition can be thought of as what was there previously plus some development ΔC_{t+1},

$$C_{t+1} = C_t + \Delta C_{t+1} \tag{4.3}$$

The previous crop condition C_t is, obviously, known to everyone at time $t+1$. Thus, it should be fully incorporated into the information set I_{t+1} to form an optimal forecast for the next period, $h(I_{t+1})$, under the semi-strong form EMH. Hence, the prediction task in period $t+1$ concentrates on ΔC_{t+1}, and the surprise component S_{t+1} must be part of it:

$$\Delta C_{t+1} = \frac{dh(.)}{dI} \Delta I_{t+1} + S_{t+1} \tag{4.4}$$

with

[69] For example, the relationship between crop condition and weather can also changes due to advancements in biotechnologies, or climate change.

$$\Delta I_{t+1} = I_{t+1} - I_t \tag{4.5}$$

being the new information that arrives between t and $t + 1$. This explains why the entire week-to-week change ΔC_{t+1} is not a good proxy for the market surprise, as noted by Lehecka (2014).

Recursively substituting (4.3), (4.4), and (4.1) into (4.2) yields

$$h(I_{t+1}) \ = \ f(I_t) \ + \ S_t \ + \ \frac{dh(.)}{dI} \ \Delta I_{t+1} \qquad (4.6).$$

Equation (4.6) shows that the updated process of expectations embraces three components: the formerly ideal process $f(I_t)$, the unexpected error S_t of that former process, and finally the revised expectation due to both new information ΔI_{t+1} and changing the model. The last term on the right-hand side – as intuitions suggest – implies that in the process of updating the market expectation model from $f(.)$ to $h(.)$ in period $t + 1$, $h(.)$ must be used to process all the data available at time t in retrospect.[70] Thus, errors from the fit of $h(I_t)$ should not explain any market reaction at time t.

On the other hand, if we were to continue using $f(.)$ for period $t + 1$ when that functional form is no longer the best approximation to market expectations, then

$$S_{t+1}^{f} \ = \ C_{t+1} - f(I_{t+1}) \qquad (4.7).$$

Substituting the relationships in (4.2) and (4.6) into (4.7) and then rearranging yields

$$S_{t+1}^{f} = \left[\frac{dh(.)}{dI} - \frac{df(.)}{dI}\right] \Delta I_{t+1} + S_t + S_{t+1} \qquad (4.8).$$

Equation (4.8) shows that the consequence of continuing to use (4.8) when the expectation formation process has changed is that the pseudo-surprise S_{t+1}^{f} differs from the true surprise S_{t+1} by two components: the disagreement between the two models on the revised expectation due to new information ΔI_{t+1}, and the surprise from the previous period S_t. In other words, even if S_t is truly unanticipated (*i.e.*, no bias in expectations), if we keep on using $f(.)$ while it is no longer the best representation of market

[70] In case of nonoptimal model updating, the refit for period t can be thought of as some weighted average of the updated model $h(I_t)$ and the old one, $f(I_t)$.

expectations, it will be carried forward to the future estimates. The same intuition holds true for individual forecasters: if one does not learn from one's previous forecast errors when revising the forecasting model for future periods, then the mistake will persist.

In what follows, we show that nonparametric ML approaches, when properly designed, can be an effective way to incorporate such updating process in the estimation of market expectations and surprises.

4.4 Methodology: an innovative ML routine

This Section focuses on ML's capability to tease out the surprises contained in the USDA's crop condition reports. First, we discuss our motivation for choosing a nonparametric ML method over traditional parametric analyses. Despite ML's advantages, there are important reasons why we cannot use a standard ML approach to find the crop condition surprises. To see why, we provide a brief overview of the standard ML approach. From there, we point out the need for modifications to make it suitable for extracting the crop condition surprises, and we show how these modifications are incorporated into our modeling framework.

4.4.1 *Why do we need ML?*

Apart from greater data availability as introduced in Section 4.2, our rationales for the use of a nonparametric ML approach mainly come from its strength in dealing with complex, nonlinear relationships in the data generating process (Storm, Baylis and Heckelei 2020). Theory is quite silent when searching for the appropriate functional forms. This is especially true for the complex market expectation formation, as well as the relationship between weather and crop yield.

First, as demonstrated by Chiang *et al.* (2019), an ideal measure of the surprise is a highly nonlinear aggregating function of individual forecasts, which depends on many unknown parameters. Consequently, it is impossible to come up with a theoretically motivated parametric measure of the surprise, which necessitates the use of a nonparametric approach.

Second, in the agricultural context, weather variables make up an essential part of the information set based on which end-of-season supply expectations are formed after the crop has been planted. The processes governing impacts of the different weather characteristics on crop conditions are highly complex and highly nonlinear across space and time, leaving market participants with infinitely many different ways to approximate them. That complexity weakens further any theoretical justification for using a specific functional form of aggregated expectations.

Finally, as our theoretical framework makes clear, under the EMH any previous public information should be incorporated into the information set used for the current period, including past forecast errors. Since predictive accuracy can be improved either by adding more data or adapting to a more appropriate modeling structure, the best model used for some previous periods is not necessarily the best one for the next periods. Allowing for such configuration updating requires a great deal of flexibility, which is limited in parametric models.

4.4.2 *A standard ML approach with cross-validation*

In the first place, ML is designed specifically for prediction. Predictive accuracy is its ultimate objective. This objective is accomplished through finding the balance between the in-sample (*i.e.*, training set) good fit and the out-of-sample (*i.e.*, test set) generalizability. In ML terminology, we want a model that produces the smallest errors possible, not only in training but also in test sets.[71]

Different procedures have been developed for this purpose, mostly as standardized, built-in routines. Figure 4.2 illustrates one of the most popular ones: the train-validation-test split approach. For a given prediction task, one needs to predefine the relative sizes of the in- and out-of-sample subsets, the ML algorithm, the objective function, and the hyperparameter space. Each point in this hyperparameter space is a specific configuration of the ML algorithm. Tuning these configurations regulates various aspects of the algorithm's operation, such as the rate of learning or the degree of complexity allowed to guard against overfitting, and find the model with the best predictive performance. Depending on the algorithm used and how wide the set of possible values for each hyperparameter is, this space can contain a very large number of hyperparameter combinations.

The general idea of this process is to split the training set (*i.e.*, the subset of data used for learning the underlying relationships between the outcome variable and the predictors)[72] into smaller subsets, or folds. For each possible hyperparameter combination, each fold will take turns serving as a semi-final test set (often called the validation set), while the model is trained (*i.e.*, fitted) on all remaining folds. The average value of the objective function across all these validation sets is then computed. This average, often referred to as the cross-validation error, is used to determine the best configuration across the hyperparameter space.

[71] For more intuitive explanations of this tradeoff, see Storm, Baylis and Heckelei (2020).

[72] In this paper, we only focus on the supervised ML category. The unsupervised ML methods – in which an outcome or target variable is absent – is not referred to since they are not suitable for our task.

Nowadays, this whole process of trying different combinations of hyperparameters usually happens within built-in routine,[73] also called hyperparameter tuning. To optimize computational and storage resources, these routines will only return the model configuration that produces the smallest average loss (largest average gain) across all validation folds, depending on the choice of the objective function. Finally, the winning model configuration is used to predict on the final test set, where the ultimate interest of the task normally concentrates. In the context of our paper, one example is to use the forecast errors in the test set as a proxy for the crop condition surprises during that period – provided that the test set spans across a consecutive period and is representative of a full geospatial unit of interest such as county, state, or country – and make statistical inference for the market impact within this test set.

Though this a powerful and widely-used approach to achieve high predictive accuracy, as we explain in the next Section, it cannot be applied mechanically to our problem, *i.e.*, to estimating the crop condition surprise.

4.4.3 *Tailored ML routine and selection criterion*

The fact that the standard ML routines only focus on the best predictive model turns out to be a critical shortcoming for our purpose. Our objective is not to come up with the best crop condition forecast, as there is no reason to assume that market expectations are always the best forecasts.

It is true that coming up with the most accurate forecast should be the goal of individual market participants. But since some will do better than others, the aggregate (market) expectation need not be the winner of them all. In other words, the surprises and the forecast errors produced by the best predicting model may not be the same. Because the standard cross-validation described above only returns the model with the best predicting performance, it is not possible to examine other candidates to see if they produce forecasts that better approximate the true market expectations. This intuition prompt us to "open the black box" of build-in validation routines, *i.e., to explicitly evaluate the forecasts and forecast errors produced by every candidate in the hyperparameter space – not just the best one.*

[73] In Python language, examples are the packages GridSearchCV and RandomizedSearchCV.

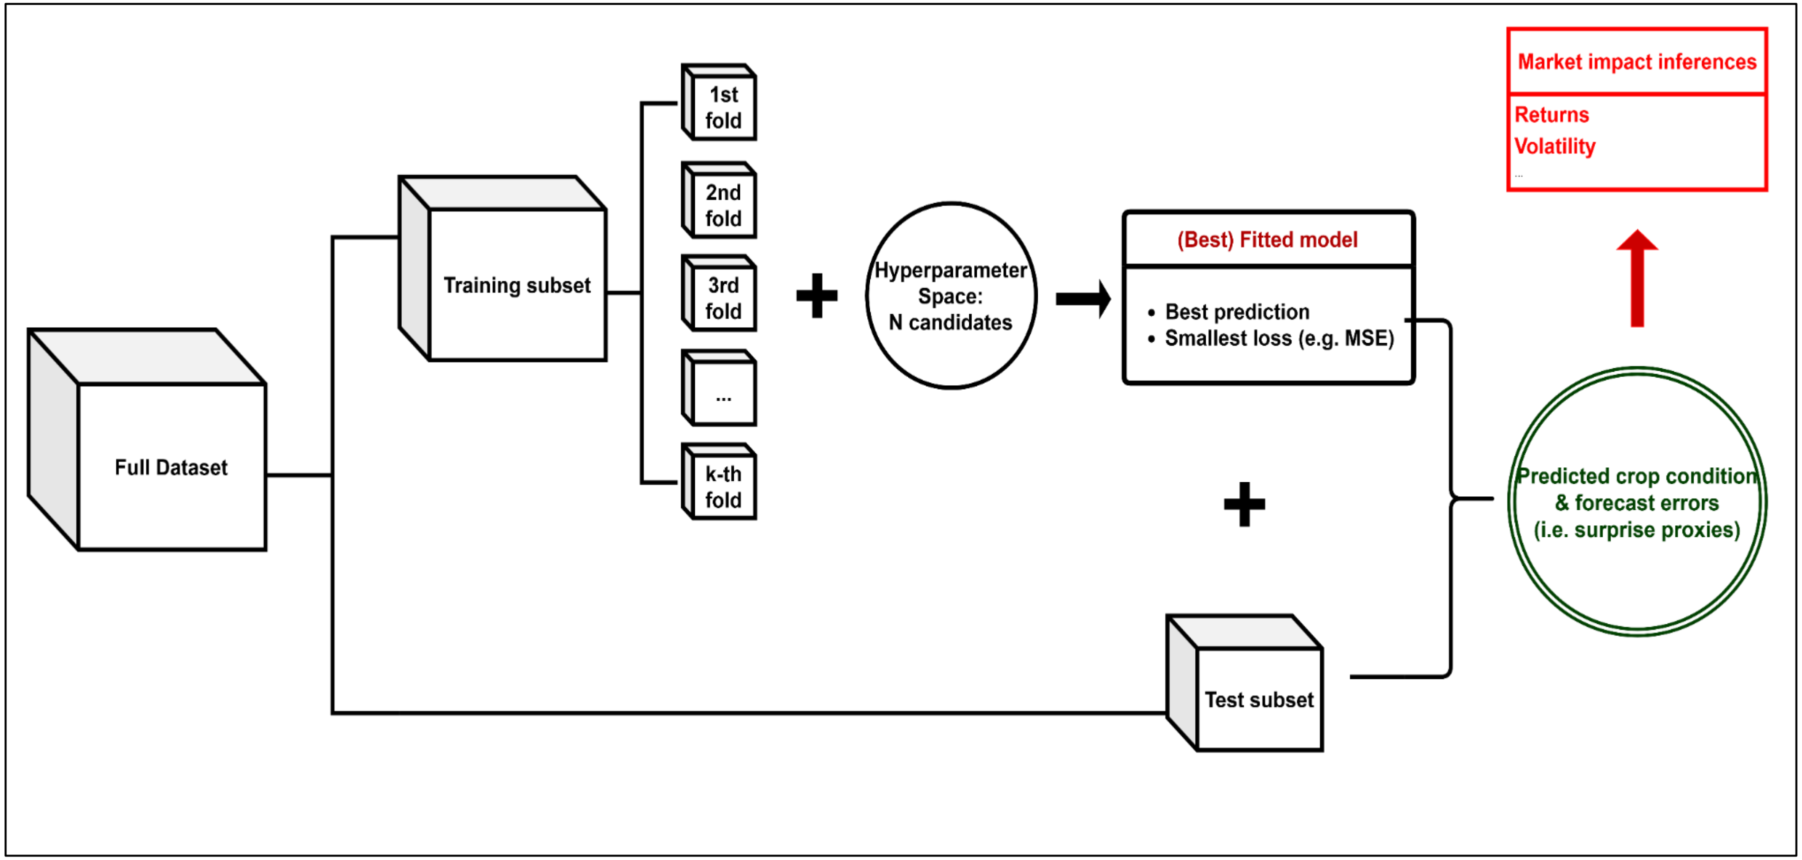

Figure 4.2. An Example of a Standard ML Workflow with K-fold Validation

The core question that follows directly is, how do we know if a given series of weekly crop condition forecasts closely resembles the market expectations of that period? In other words, how do we know if a given series of forecast errors is more correlated with the true market surprises than the other series? Obviously, an average value is not informative since the variation in it is discarded. Because the true market expectations are unobservable, we must rely on some other observable process that can be plausibly assumed to be correlated with them. But, as Bartolini (2008, p. 364) precisely points out, *"there is no instrument correlated with the true news that can be used to address this problem: if there were one such instrument, z_t^* would not be news anymore."*[74] This leaves us with the only possibility: the post-release market returns. *To access the quality of a proxy for the announcement surprises, we propose to consider its relationship with the observable post-announcement market reactions*.

This conclusion leads us to a selection criterion beyond the possibilities provided by conventional ML approaches: the ability of the resulting forecast errors to explain the observed market returns. Suppose that a ML model predicts crop condition as $g(I_t)$ and has a forecast error X_t, *i.e.*,

$$X_t = C_t - g(I_t) \qquad (4.9).$$

The idea is to find $g(I_t)$ that is most similar to $f(I_t)$ such that X_t is closest to S_t in Equation (4.1). But market expectation $f(I_t)$ is not observable, and neither is the market surprise S_t. However, under the maintained assumption that the market consensus is an unbiased forecast, it is natural to assume that $\mathrm{S_t}$ is distributed with mean zero and a finite variance σ_S^2.

For a ML algorithm designed to produce an unbiased forecast of the crop condition (and due to the regularization in most ML algorithms), its forecast errors X_t should be distributed with mean zero and a finite variance σ_X^2. Following Karali, Irwin, and Isengildina-Massa (2019), we acknowledge that this forecast error differs from the true market surprise S_t by a random measurement error component η_t,

$$X_t = S_t + \eta_t \qquad (4.10).$$

As there is no reason to believe that this measurement error is correlated with the true surprise S_t, we can assume that it fits the classical error-in-variables assumptions and is

[74] z_t^* refers to the true surprise, equivalent to S_t in our notations.

distributed with mean zero and variance σ_η^2. Post-announcement market returns are defined as

$$r_t = lnP_{t,1} - lnP_{t,0} \qquad (4.11),$$

where $P_{t,1}$ and $P_{t.0}$ are market prices after and before the announcement of period t, respectively. Assuming that the impact of the news released in the announcement on market returns is orthogonal to other factors, its marginal effect on market returns can be captured by the coefficient β_1 in the simple linear regression equation

$$r_t = \beta_0 + \beta_1 S_t + \varepsilon_t \qquad (4.12),$$

where ε_t satisfies the usual assumptions.[75] The proportion of the variation in returns explained by this news component is then measured by the coefficient of determination of Equation (4.12), denoted R_S^2.

$$R_S^2 = 1 - \frac{\hat{\sigma}_\varepsilon^2}{\hat{\sigma}_r^2} \qquad (4.13),$$

with $\hat{\sigma}_\varepsilon^2$ and $\hat{\sigma}_r^2$ being the estimated sample variance of the regression residuals ε_t and r_t, respectively. Analogously, for any proxy X_t we can run the regression

$$r_t = b_0 + b_1 X_t + v_t \qquad (4.14)$$

and obtain the corresponding coefficient of determination R_X^2,

$$R_X^2 = 1 - \frac{\hat{\sigma}_v^2}{\hat{\sigma}_r^2} \qquad (4.15).$$

Using asymptotic theory, it can be shown that[76]

$$plimR_X^2 = R_S^2 \frac{\sigma_X^2 - \sigma_\eta^2}{\sigma_X^2} \qquad (4.16).$$

[75] In equity space, this equation is known as “POSTCAR regression” Chiang *et al.* (2019).

[76] This result is consistent with Karali, Irwin, and Isengildina-Massa (2019). Readers should read that article for a more general discussion about the consequence of measurement errors on the estimated impacts of USDA announcements in agricultural markets.

That is, under classical error-in-variables, the coefficient of determination of the market impact regression with a given news proxy is smaller than the one obtained using the true surprises by a fraction of measurement error's variance in total variance of the observed proxy. Thus, as a proxy gets closer to the true surprise, σ_η^2 becomes smaller. Consequently, this difference should become narrower. Though one can never completely eliminate the measurement errors in any proxy for the true surprises due to its latent nature, a measure most correlated with the true surprise must have the highest explanatory power to the corresponding market movement outcomes among a set of candidates. Therefore, *choosing the candidate that yields the highest R_X^2 in regression Equation (4.12) will pull us close to the true surprise*.[77]

It is important to point out that our approach does not jeopardize the assessment of the market impact of the news, so long as the slope coefficient estimates in Equation (4.12) are not involved in the selection criteria. The true market surprise variable S_t, with the highest R^2 as shown, can have an insignificant impact β_1 on market returns over a given time interval, and so can the proxy X_t.

Until here, it is clear that both the workflow and the objective function for model selection must be different from those in standard ML approaches. But further modifications must be incorporated to make sure that the proxy X_t is generated in a similar manner to the expectation formation process. In the context of crop condition expectations, neither the time nor the spatial dimensions of the data can be neglected. As it is natural with crop development and market expectation formation, ensuring that we do not use the data of the future to predict the past is critical.[78] Furthermore, consistent with our theoretical framework, allowing for the possibility of updating the model over time is important. For example, suppose that we are interested in drawing conclusion about the market surprise during the period 2020-2021. Using a standard approach, we would have all the data of

[77] To get a sense of how this relationship holds for small samples, we perform a simple Monte Carlo simulation with 10,000 repetitions for a sample size of 20 observations, resembling an average crop cycle when the CPCRs are available. Each time, a noise component is added to the pre-defined surprise variables before a regression of the market returns on this noisy proxy is run. The experiment reveals that the probability of getting an $\hat{R}_X^2$ that is higher than the "true" $\hat{R}_S^2$ is around 3%.

[78] In principle, methods to account for the time dimension of the data are available within conventional ML frameworks (see, *e.g.*, Oliveira, Torgo, and Santos Costa (2021), for a thorough review of cross-validation schemes in spatiotemporal settings). However, these treatments still leave a gap between the periods used to fit the data and the periods in the final test set, namely the periods used in the validation set. From expectation revision perspective, what works well for the validation set may not work that well for the test set. More importantly, without the possibility to evaluate every model candidate separately, they can only improve temporal forecast accuracy, not market expectation replication.

this period as the test set, and the training set would contain all available data before 2020. But by making those choices, the data that we would use for the year 2021 would be restricted to be smaller than the information set available to market participants, which also includes the year 2020. At the same time, that approach would exclude the possibility that a different model configuration might perform better after having the 2020 data included in the training set, as compared to the one selected from a training set that ends in 2019. Hence, ideally, incorporating the most recent data and evaluating the models must be done simultaneously.

Likewise, as market participants ultimately care about the implication of crop condition on the overall crop size of the US, it is natural to weigh the crop condition of a geographical unit by the production share of that unit when assessing the national figures. Since the loss (gain) in built-in ML routines is the arithmetic mean across all observations, it washes out the heterogeneous importance of across time and space, and is thus likely to be a biased estimator of the nationally aggregated surprises.[79]

To sum up, to have a ML workflow function in a similar way to how market expectations are formed as suggested by theories, we cannot adapt a standard workflow that is primarily designed to achieve the highest possible predictive power. Figure 4.3 summarizes the modified workflow, discussed in this Section, that addresses these concerns.

[79] Introducing customized loss (gain) functions is also possible with available ML routines. However, to our best knowledge, the latter do not allow for dynamic updates of the weights across different parts of the dataset. This is necessary because the production share of each geographical unit changes from year to year. Even within a year (when the production share can be argued to be invariant), crop planting does not happen at the same time across all production areas. Using fixed weights will likely result in biased estimates of the national condition. We discuss these issues in detail in Section 4.5.1.

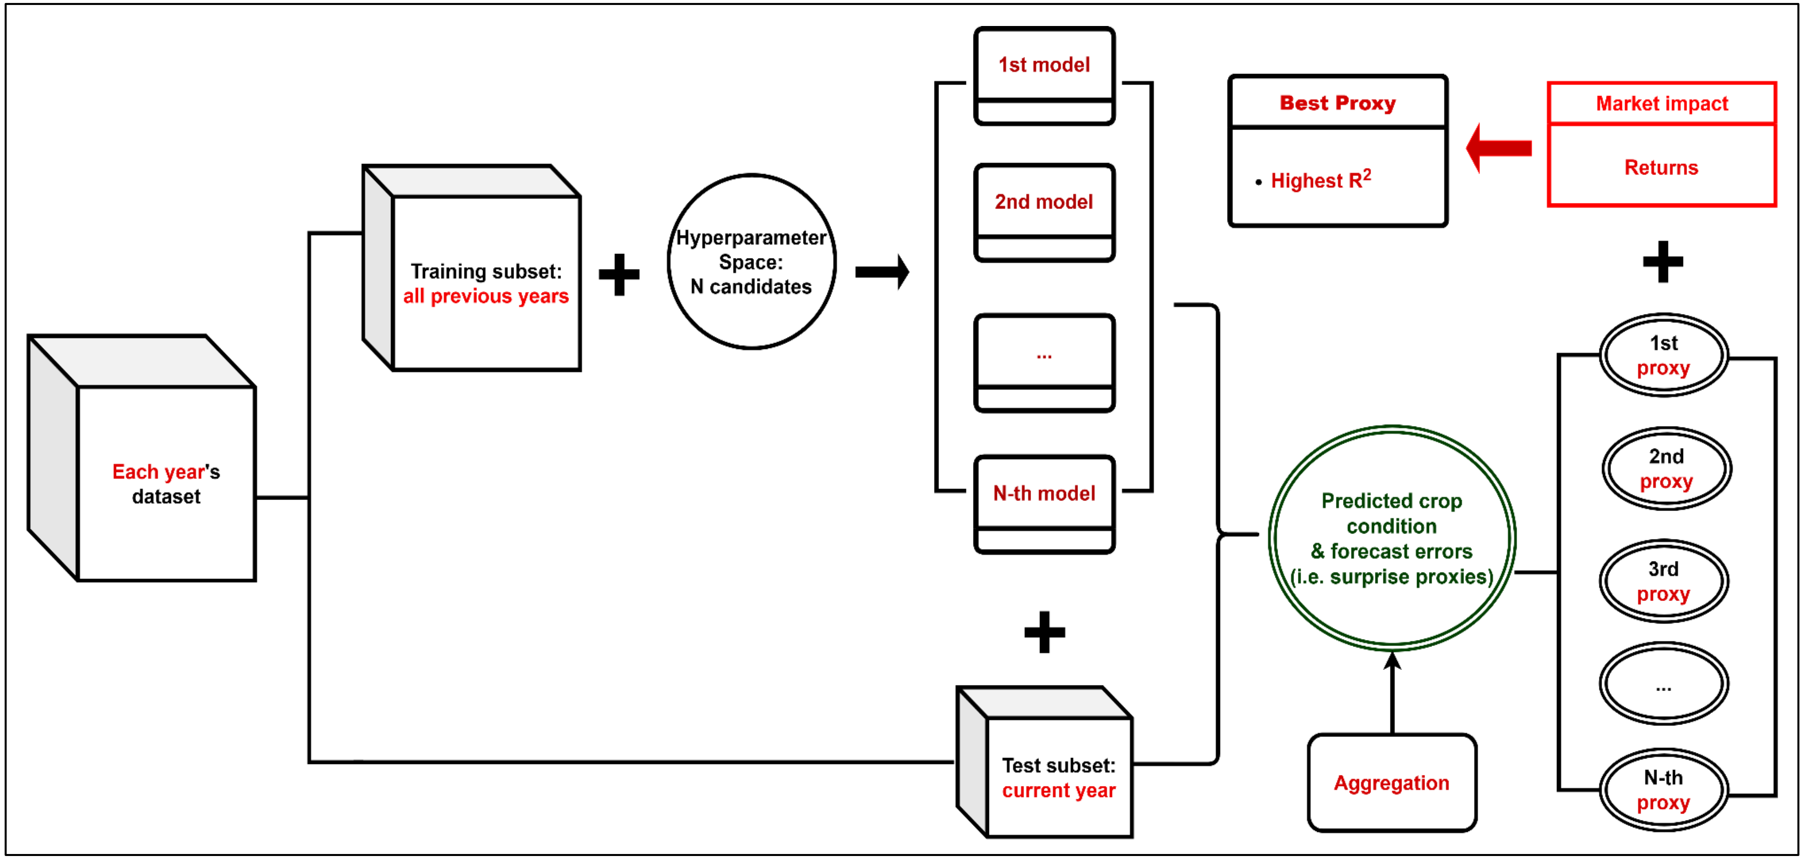

Figure 4.3. Modified ML Workflow to Estimate the Surprise Component of Crop Condition Reports

4.5 Empirical implementation and results

Our ultimate goal is to assess the surprise component of the crop condition content in the CPCRs. To that end, we proceed in four steps as presented in Figure 4.4, including: (i) – data preprocessing, (ii) – consistency check of the new condition dataset, (iii) – extracting the surprise component, and (iv) – assessing the impact of the surprises on corn and soybean new-crop futures prices. This Section describes the procedures we follow and our findings for these steps. Since both CPCR gridded layers and PRISM data are in geospatial format, we follow a standard workflow of preprocessing geospatial data for ML applications. Hence, for abbreviation, this step is not described in detail here.[80] A list of predictors used to predict pixel-level crop condition is provided in Appendix 4.A.1.

4.5.1 *Consistency check: how reliable is the gridded dataset?*

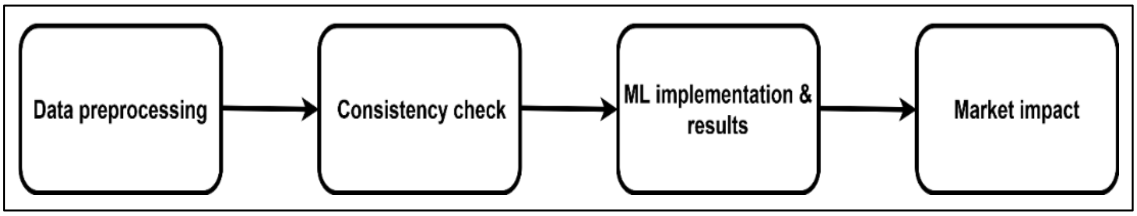

Figure 4.4. Implementation Steps

Before using the gridded CPCR dataset for our analysis, it is important to verify that this newly developed dataset is not a biased representation of the original dataset, which is the one released to the markets in real time. In this Section, we therefore check whether interpolation and smoothing introduce a serious inconsistency between the crop conditions reported by the two datasets.

As described in Section 4.2.1, the two datasets are different in two important attributes: the numerical format, and the level of spatial aggregation. Hence, for comparison we must bring them to the same expression. The former can be easily done using the formula provided in NASS documentation of the new gridded dataset (USDA 2021) to convert the ordinal format of the original datasets to the continuous form as used in the gridded dataset. For a given week-geographical unit (*e.g.*, county, state, or all 18 production states)[81], this formula is given by

[80] Nevertheless, the Python scripts used for this step are available on our documentation homepage.

[81] For corn, 18 production states include CO, IL, IN. IA, KS, KY, MI, MN, MO, NE, NC, ND, OH, PA, SD, TN, TX, WI: For soybeans, they are AR, IL, IN, IA, KS, KY, LA, MI, MN, MS, MO, NE, NC, ND, OH, SD, TN, WI.

$$Condition = 5 * EX + 4 * G + 3 * F + 2 * P + 1 * VP \quad (4.17),$$

where Condition is a continuous index that can take values ranging from one to five. The five abbreviated terms represent the crop shares in five categories, ranging from Excellent (EX) down to Very Poor (VP).

The second reconciliation, with respect to the spatial dimension, is more complicated. Though NASS documentation provides a helpful overview of how the 9x9 km weekly layers are created, for confidentiality reasons it is not meant to provide a precise procedure with all data used. Hence, reverse-engineering the state or national condition from these layers requires a rigorous treatment.

Figure 4.5 summarizes the procedure that we develop for this purpose. Its significant development beyond NASS guidance is the use of agricultural district boundaries to identify the "Other (combined) counties" (*i.e.*, counties that do not appear separately under their names in the database, but instead are merged in one group within each district). These counties make up a nontrivial share of the total corn and soybean acreage planted within the main production states.[82] Omitting them from the aggregation would likely cause a bias in the aggregated crop condition for 18 states. Another important calibration is our dynamic weighting scheme to determine the contribution of each geographical unit (*e.g.*, county or state) to the aggregate crop condition. The timing of crop cycles varies a lot across a large production area in the United States. Thus, assuming a constant influence of each county to the 18-state condition throughout the growing season would be unrealistic. For each state-week, we first exclude the counties with no reported condition in that week. Then, we take the annual acreage shares of the remaining counties and scale them up, such that they always sum up to one for a given state-week. We apply the same method when aggregating crop condition from state level to 18-state level.

In Figure 4.6, we plot the resulting 18-state (national) condition indices from the two datasets for corn (panel a) and soybeans (panel b). g_{cc} (in blue, solid) denotes "gridded crop condition", whereas o_{cc} (in orange, solid) denotes "original crop condition". The two series are plotted using the left-hand-side vertical axis. In addition, $ln\ (g_{cc}/o_{cc})$ (in red, dotted) is the log-difference between the two series, plotted using the right-hand-side vertical axis.

[82] In 2019, for example, they account for almost 21% and 22 % corn and soybean acreage planted in 18 production states of each crop, according to NASS Quick Stats.

During the period 2015-2021, there are 153 (*resp.* 135) weeks when gridded and traditional corn (*resp.* soybean) condition data are both available.[83] Despite the many challenges with identification described above, the overall discrepancies between the two resulting indices are minimal, which is an assurance that the gridded dataset is well-developed and is reliable for our analysis. It also proves that our aggregation scheme is generally appropriate.

[83] For some unknown reason, soybean gridded condition is not available for week 32 in 2018, and thus is dismissed from the analysis.

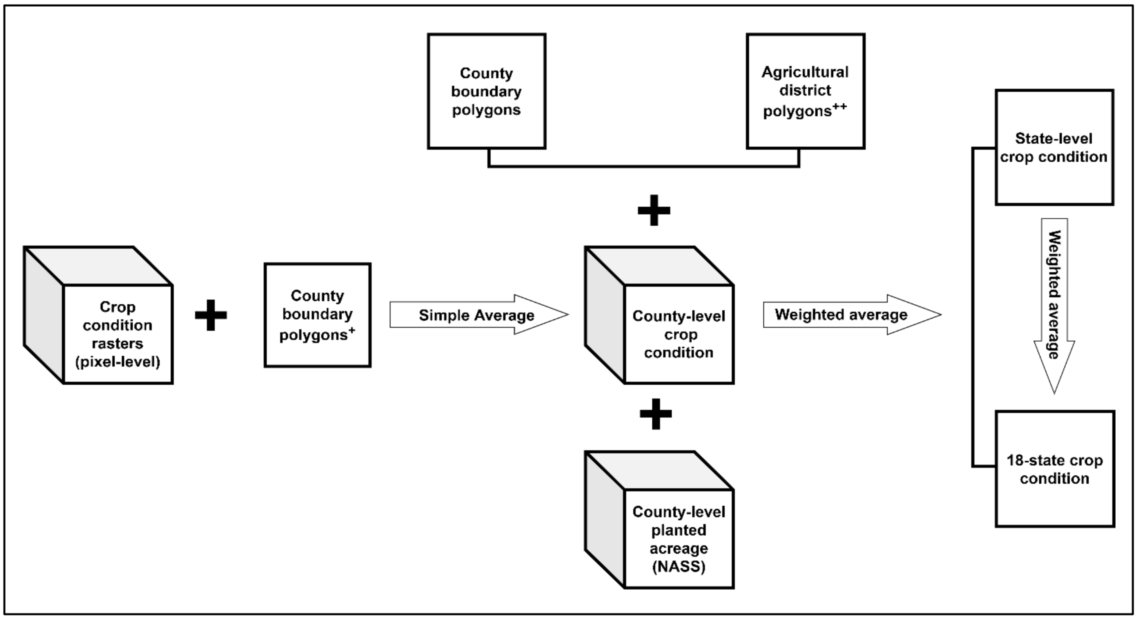

Figure 4.5. Aggregation Procedure from Pixel-level to 18-state Crop Condition

Note:

+ County boundary shapefiles are obtained from the official website of the United States Census Bureau: https://www.census.gov.html.

++ Agricultural district shapefiles for the entire USA are obtained from Esri ArcGIS Online Platform: https://www.esri.com/en-us/arcgis/products/arcgis-open-data. The agricultural district boundaries are necessary to identify which counties are included in the “Other (combined) counties” of each district. Only then, we determine the weighted crop condition of these counties, and their respective contribution to the overall crop condition of the state and 18-state production area of each crop. For a given agricultural district, the exact acreage planted in each of the “Other (combined) counties” is not available, but their sum is available. Using this sum, we first calculate the acreage share of the whole combined group. Then, we use their respective shares in the total number of available crop condition pixels of the group to split further the group acreage share into county shares. Without combining the county boundary files with the agricultural district boundaries file, it is not possible to locate the counties within these groups (all named the same in NASS database: “Other (combined) counties”) to their right position on the map, neither to correctly determine the number of active crop condition pixels per week within each of them.

A few exceptions are a couple of large differences in 2018 for corn, and medium-sized differences starting late 2020 for both crops. There is no obvious explanation for the former, but it is unlikely due to our aggregation procedure – since the procedure closely replicates the original reports for the most part.[84] The latter is similar for both crops and can be explained by the fact that, starting from 2020, NASS Quick Stats stopped reporting the acreage planted in "Other (combined) counties" within each agricultural district. As noted, without this information, it is not possible to account for the contribution of those counties to the state- and national-level crop conditions. As it turns out, this missing information underestimates the aggregated 18-state condition of both crops in that year by about two percent on average. This result is not surprising, provided that the "Other (combined) counties" altogether make up a considerable acreage share of 18 production states. But if NASS still takes them into account to produce the state-level crop condition while the information is no longer available to the public, then the forecast errors stemming from their absence are likely incorporated into the true market surprises.

Consistent with the plots, the summary statistics of g_{cc}, o_{cc} and $ln\ (g_{cc}/o\_cc)$ for corn (panel a) and soybean (panel b) of Table 4.1 indeed confirm that the two datasets are in general very close when converted to a unified format. The average discrepancy between the two datasets is less than two percent (*resp.* 0.5 percent) for corn (*resp.* soybeans). Nevertheless, for consistency between the predicted condition (pixel-wise) and the actual condition, in what follows we stick to the gridded condition when calculating the surprise proxy. That way, potential discrepancies are neutralized and do not bias our surprise estimates.

[84] These are weeks 32, 41, and 42 in 2018. The first week is associated with missing data for soybeans as well. Hence, we suspect that some systematic issue might occur in the production of the gridded dataset for that specific week.

Figure 4.6. Comparison between the Gridded Condition and Original Condition Dataset during 2015-2021

Table 4.1. Summary Statistics – gridded vs. Original Crop Condition Data during 2015-2021

	Mean	Median	SD	Min	Max	Jarque-Bera test statistics†	No. non-missing Obs
A. Corn							
g_{cc}	3.6917	3.7366	0.1424	3.3837	3.9133	**10.27*** **	153
o_{cc}	3.6977	3.7300	0.1279	3.4700	3.9200	**10.37*** **	153
$ln\ (g_{cc}/o_{cc})$	-0.018	0.019	0.0165	-0.1081	0.0206	**3.23e+3*** **	153
B. Soybeans							
g_{cc}	3.6317	3.6382	0.1382	3.3819	3.8614	**9.62*** **	135
o_{cc}	3.6487	3.6400	0.1230	3.4200	3.8600	**8.40** **	135
$ln\ (g_{cc}/o_{cc})$	-0.0048	-0.0034	0.0086	-0.0343	0.0088	**56.80*** **	135

Statistical significance code: *** 0.01 **0.05 *0.10

Note:Table 4.1 provides summary statistics for corn condition (panel a) and soybean (panel b) condition of 18 production states according to CPCRs during 2015-2021. For each crop, the condition index values are calculated with two datasets: the gridded, "county-level" dataset (g_{cc}) and the original, "state-level" dataset (o_{cc}). The statistics of the log-difference between the two datasets ($ln\ (g_{cc}/o_{cc}$) are also provided.

† Jarque-Bera test for the null hypothesis that the sample comes from a normal distribution with an unknow mean and variance.

4.5.2 *ML workflow in practice*

We choose Extreme Gradient Boosting (XGB) for our empirical analysis. This tree-based ML algorithm has gained its popularity in recent time for its flexibility and predictive performance (see Roznik, Mishra, and Boyd (2023) and the references cited therein). XGB does not require hand-crafted feature engineering and thus saves us from having to make arbitrary assumptions about the functional form. In the same vein, pre-train feature selection is not necessary, which allows the model to freely exploit different random sets of weather predictors at a higher temporal frequency than the reported condition (*i.e.*, daily weather *vs.* weekly CPCRs). Our own preliminary experiments also show that XGB is much more computationally efficient compared to other methods such as Random Forests or Neural Networks, which makes it especially attractive for our large datasets. The computation is parallelized and distributed on the University of Illinois' high-performance computing cluster.

As emphasized earlier, we do not aim for the highest accuracy, so to give our models the freedom to search for the most suitable surprise proxy in any range of predictive accuracy. Thus, we define a hyperparameter space that is fairly broad in terms of prediction variability, from poorly to highly accurate. Our hyperparameter space contains 12,800 combinations, which are the cartesian products of six important XGB hyperparameters whose value ranges spread as widely as possible. Details of these hyperparameters and of the resulting distributions of predicted crop condition are presented in Appendix 4.A.2.

As previously discussed, it is important to allow the possibility of updating the model regularly. However, changing the configuration too frequently (*e.g.*, on a weekly basis) is practically costly, both in terms of computational resource and of degrees of freedom (when evaluating correlations with commodity price returns). In our particular case, due to the seasonal nature of crop production, it is plausible that the development of a full crop cycle is required to evaluate model performance. Hence, we allow for the predicting model to be updated after each crop year. We follow the procedure in Figure 4.3, in particular:

1 – For each crop year, we train 12,800 model candidates on all data available prior to the crop season, using all hyperparameter combinations. Then, we use each of them to predict the crop condition throughout the year. Within each year, we also include the condition and progress of the previous week to the set of predictors. This is to ensure that all previous relevant information up to the weekend before each report is included in all models.

2 – For each week, the pixel-wise predicted condition is aggregated to 18-state predicted condition using the procedure described in Section 4.5.1, denoted g_{pc_t}. The standardized

surprise proxy, x_t is defined as the log-difference between the aggregated gridded condition, g_{cc_t}, and this predicted value[85]:

$$x_t = ln\left(\frac{g_{cc_t}}{g_{pc_t}}\right) \qquad (4.18).$$

3 – To evaluate these 12,800 $\mathrm{x_t}$ series, we regress the post-event close-to-open returns of the new-crop future contracts on each of them to obtain

$$r_t = \hat{b}_0 + \hat{b}_1 x_t + \hat{v}_t \qquad (4.19).$$

Then, following the reasoning in Section 4.4.3, we select the x_t series that yields the highest R^2 as the best surprise proxy for that year.[86]

4 – We repeat the procedure for each year separately until 2021. That way, we can estimate the surprise for 6 years in total, from 2016 to 2021.

4.5.3 *Model outcome*

Our position is that adopting the best predictions from some given model as a proxy for market expectations is not justifiable. However, in the absence of an accurate prediction, this hypothesis cannot be tested. Thus, in the first place, we must prove that a highly accurate prediction can be generated by our model. Only then can we reliably test its explanatory power to price movements against another candidate.

Figure 4.7 and Figure 4.8 summarize the prediction distributions of our ML pipeline for corn and soybean conditions, respectively. Weekly actual crop condition indices are plotted in green, solid lines. For each year, each model candidate produces a series of crop condition prediction spanning the crop season – which typically includes 20-22 weeks. Hence, for each reported week, we obtain a sample of 12,800 predictions of crop condition. The sample medians of the distributions are plotted in blue dashed lines. The areas shaded in light blue mark the range where 95 percent of model predictions fall in

[85] The log-difference surprises are conventionally used (*e.g.,* Garcia *et al.* (1997), Karali, Irwin, and Isengildina-Massa (2019), Cao and Robe (2022)) to the interpretation of the slope coefficient from different units of measurement and scales of the surprises. It does not alter the results conceptualized in Section 4.3.

[86] We also calculate the Bayesian Information Criterion (BIC) for every regression. Results are consistent with the R^2 and is available upon request.

(*i.e.*, 47.5 percent above and 47.5 percent below the median).[87] Among these 12,800 series, the series with the smallest root mean aggregated square error (RMWSE)[88] – "best predictions" – are plotted in orange, solid line. The series that explain the most variations in the post-event new-crop future price returns (*i.e.*, yield the highest R^2 for Equation 19) are plotted in red, solid lines[89]. In the same manner, in Figure 4.9 (corn) and Figure 4.10 (soybeans) we plot the distributions of log-difference prediction errors, *i.e.*, the surprise proxies x_t.

Consistently across all years and commodities, our best predictions are very close to the actual conditions reported by the USDA, indicating that our model design can predict the condition of a crop with very high accuracy. The sample medians of predictions follow closely in terms of predictive accuracy. In some cases (in 2016 for both corn and soybeans and 2017 for soybeans), the median series and the best series even coincide. This minimal difference suggests that the median can be a promising candidate when predictive accuracy is the main focus. This said, it has a slight tendency of overpredicting in some years for corn, compared to the best forecast series. The strong predictive performance of our models is also apparent from the distribution of predictions as a whole. Since the number of predictions falling within the 95% area is distributed evenly above and below the median, the narrower area above the median lines in in Figure 4.7 (corn) and Figure 4.8 (soybeans) indicates that the majority of model predictions concentrates on the highest tier of crop condition index range. Because the actual figures (in green) are distributed mostly in this upper range, as confirmed by the summary statistics in Table 4.1, it follows that the likelihood of getting close to the true value of the condition using our hyperparameter space is relatively high. Comparing the two commodities, our best models produce slightly better predictions for corn than for soybean conditions. This difference can be attributed to the fact that corn production in the USA is not only relatively larger than soybean in terms of acreages, but is also characterized by longer growing periods (USDA 2017).[90] The pixel-week dataset for corn is thus significantly

[87] The mean prediction plots are available in Figure 4.11 and Figure 4.12 of Appendix 4.A.2. As it shows, the mean and the median predictions are very similar, both in magnitude and in pattern.

[88] In general, this is a modified version of the root mean square error (RMSE) criterion used in conventional ML applications. For each week, squared forecast errors at pixel level are aggregated to 18-state condition using the procedure described in Figure 4.5. Then, we take the average of all weeks in the year and finally take the square root of it to obtain the RMWSE.

[89] The plots of the R^2 against RMWSE of the surprise candidates are provided in Appendix 4.A.2.

[90] According to NASS Quick Stats database, in the period 2015-2021, the average planted acreage of corn across 18 production main production states are approximately 81 million acres per year. For soybeans, it is around 78 million acreages per year.

larger for corn, which boosts predictive performance.[91] The reverse patterns of the log-difference errors in Figure 4.9 (corn) and Figure 4.10 (soybeans) are consistent with the information conveyed from Figure 4.7 (corn) and Figure 4.8 (soybeans).

Table 4.2 provides the summary statistics of these series continuously from 2016 through 2021. We also compute the test statistics for the Jarque-Bera test of normality for both the prediction and the log-difference error series. Similar to Table 4.1, the test consistently rejects the normality of all series – except the best surprise proxy for corn condition. Thus, in the last columns, when performing paired sample test of each series against the best market expectation and best surprise proxies, we report the nonparametric Wilcoxon signed rank test statistics. Consistent with our observations from the graphs, Table 4.2 confirms that for the entire period from 2016 to 2021, the best models produced unbiased and highly accurate forecasts for both crops.[92] On average, the log-difference errors of the best prediction series are not significantly different from zero. The median forecasts of both crops exhibit some upward bias but with small magnitudes (less than 2 percent for both crops), as can be observed in the prediction plots (Figure 4.7 for corn and Figure 4.8 for soybeans) and log-difference error plots (Figure 4.9 for corn and Figure 4.10 for soybeans).

[91] The average number of pixel-week observations in corn dataset is about 953.2 thousand per year. For soybeans, it is about 522.5 thousand pixel-week per year.

[92] For each crop, we also compute Wilcoxon signed rank test statistic for the null hypothesis that the best prediction and the actual condition comes from the same population. The test fails to reject the null for both crops.

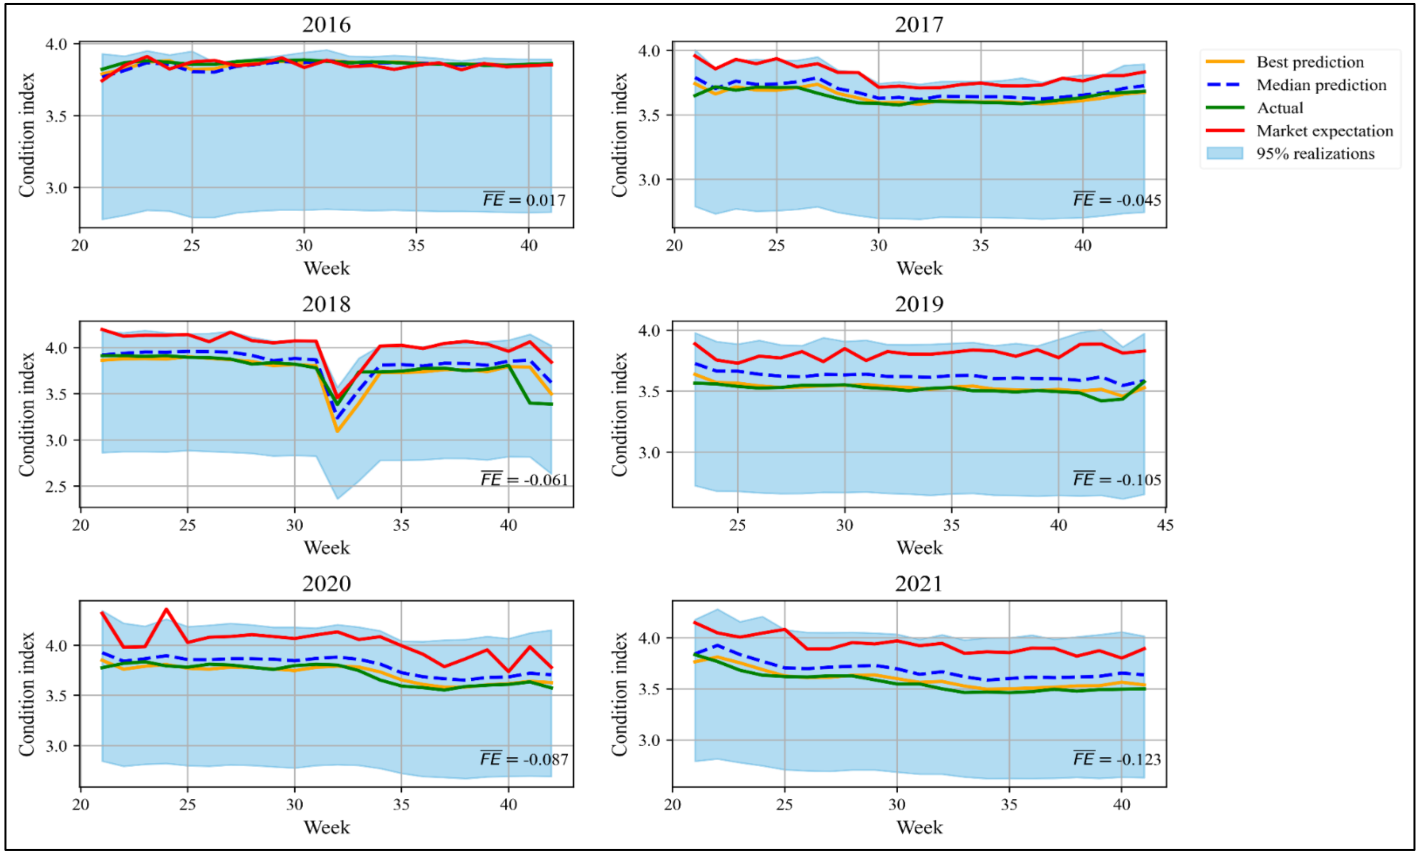

Figure 4.7. Model Predictions for Crop Condition during 2016-2021 – Corn

Note:

- Best prediction: the outcome of the model candidate which produces the smallest root mean weighted square error (RMWSE), averaged across all weeks of the year.
- Median prediction: the sample median of 12,800 predicted crop condition values that are produced by 12,800 model candidates. Each of them is an 18-state aggregate using the procedure in Figure 4.5.
- Actual: 18-state aggregated crop condition from NASS gridded dataset, which is denoted as g_{cc} in Section 4.5.1.
- Market expectation: the outcome of the model candidate which produces the surprise proxy that best explains post-even market returns in that year. That is, among all candidates, the regression in equation (4.19) yields the highest R^2 with this surprise proxy.
- 95% realization: from 2.5-percentile to 97.5-percentile of 12,800 predicted crop condition values as explained above. In other words, 47.5 percent of predictions above the median prediction and 47.5 percent of predictions below the median prediction fall within this range.
- $\overline{FE}$ is the yearly simple average of the median prediction.
- Figure 4.8 below is produced analogously for soybeans. The disrupted gap on the subplot 2018 is due to the missing dataset in week 32 for soybeans, as discussed in Section 4.5.1. Since we need the condition and progress of the previous week as predictors, the prediction for week 33 cannot be made neither.

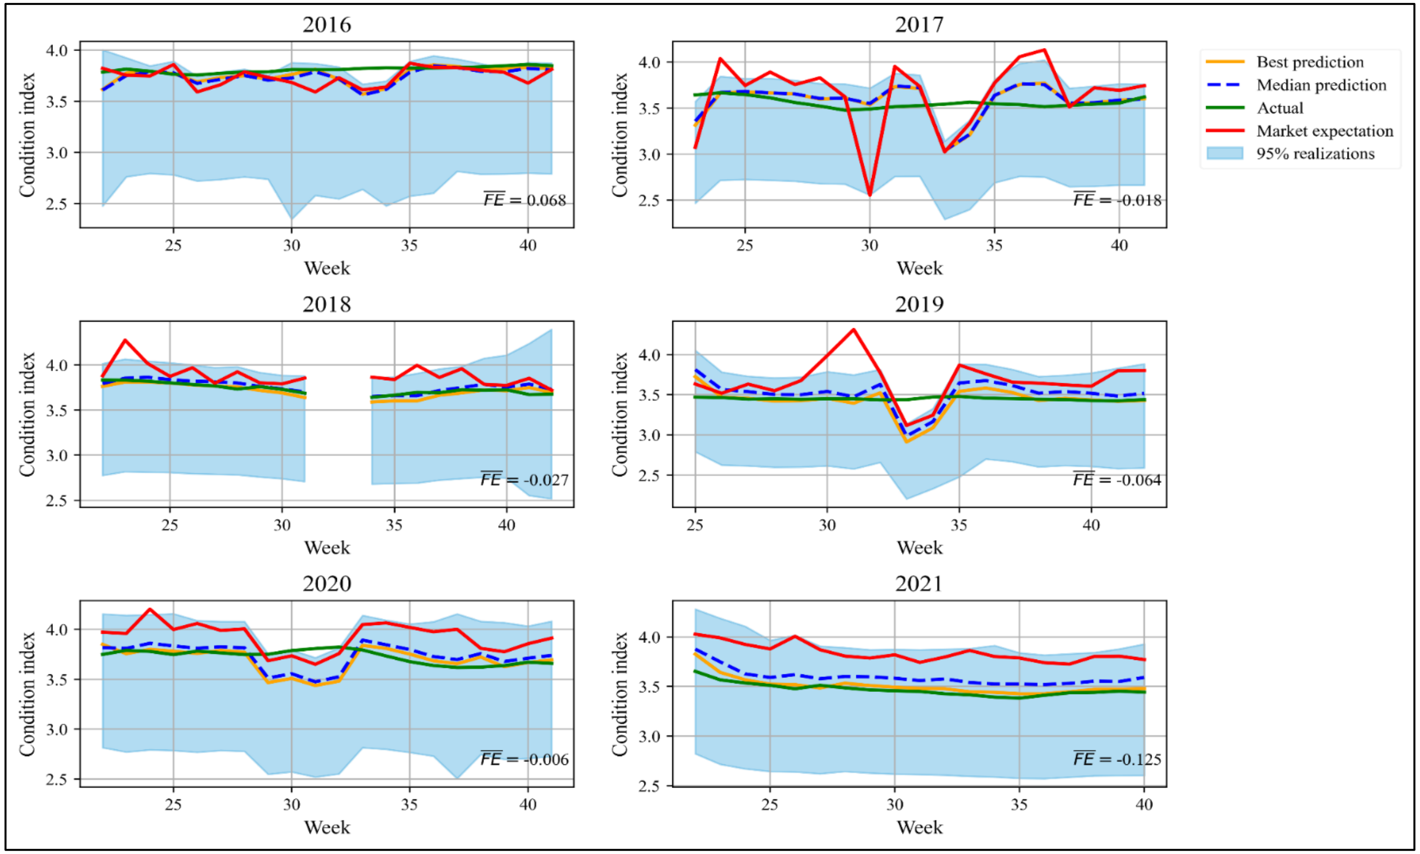

Figure 4.8. Model predictions for Crop Condition during 2016-2021 – Soybeans

<u>Note</u>: See notes below Figure 4.7

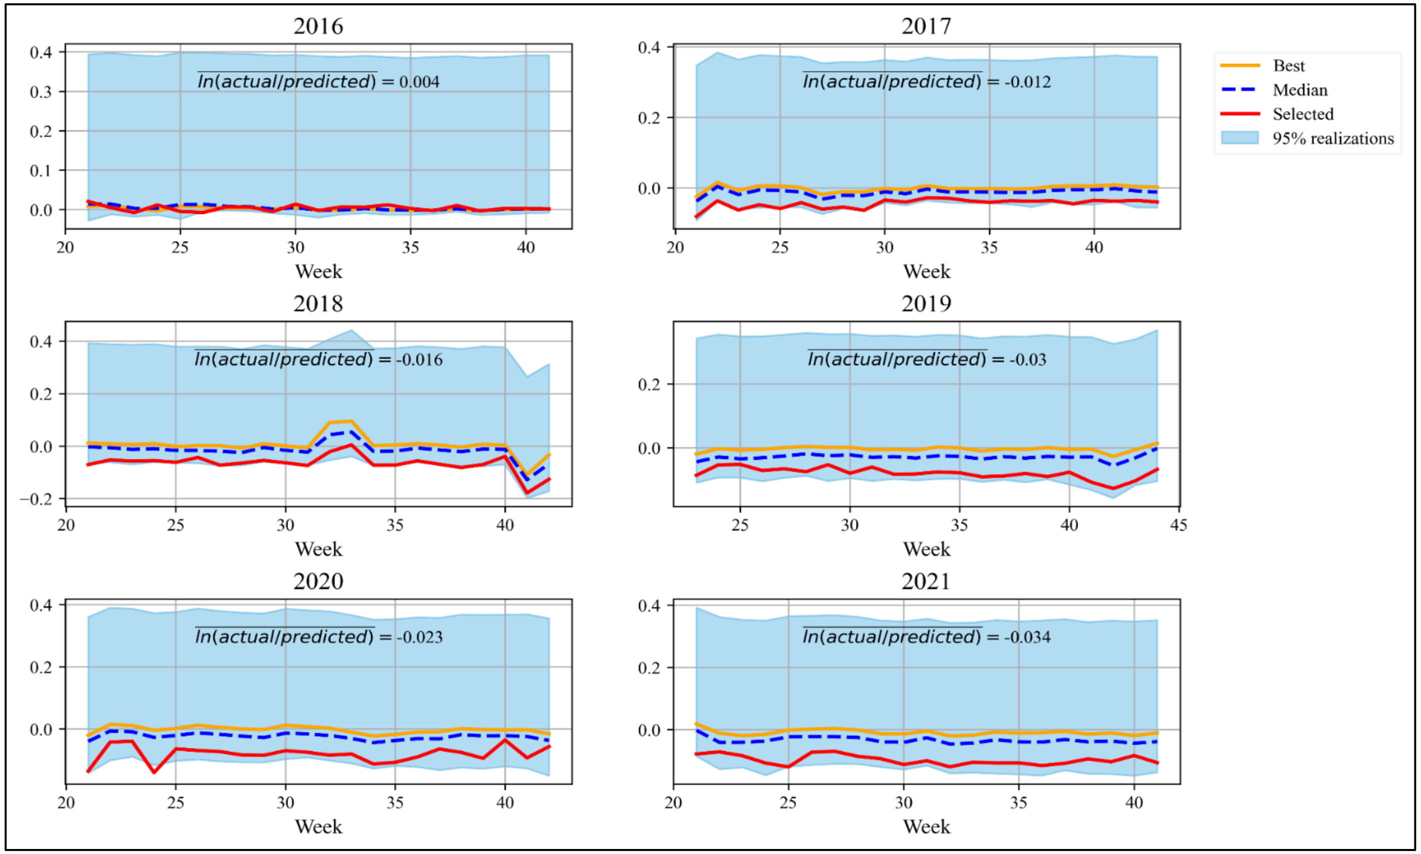

Figure 4.9. Model prediction errors in log-difference during 2016-2021 – Corn

Note:

- Best: the log-difference error of the model candidate which produces the smallest root mean weighted square error (RMWSE), averaged across all weeks of the year.
- Median: the sample median of 12,800 aggregate log-difference error of 12,800 model candidates. The actual and predicted conditions used to calculate this log-difference error is aggregated over 18 production state using the procedure in Figure 4.5.
- Selected: the surprise proxy that best explains post-even market returns in that year. That is, among all candidates, the regression in equation (4.19) yields the highest R^2 with this surprise proxy.
- 95% realization: from 2.5-percentile to 97.5-percentile of 12,800 log-difference error values as explained above. In other words, 47.5 percent of x_t above the median prediction and 47.5 percent of x_t below the median error fall within this range.
- $\overline{ln\ (actual/predicted)}$ is the yearly simple average of the median log-difference error.
- Figure 4.10 below is produced analogously for soybeans. The disrupted gap on the subplot 2018 is due to the missing dataset in week 32 for soybeans, as discussed in Section 4.5.1. Since we need the condition and progress of the previous week as predictors, the prediction for week 33 cannot be made neither.

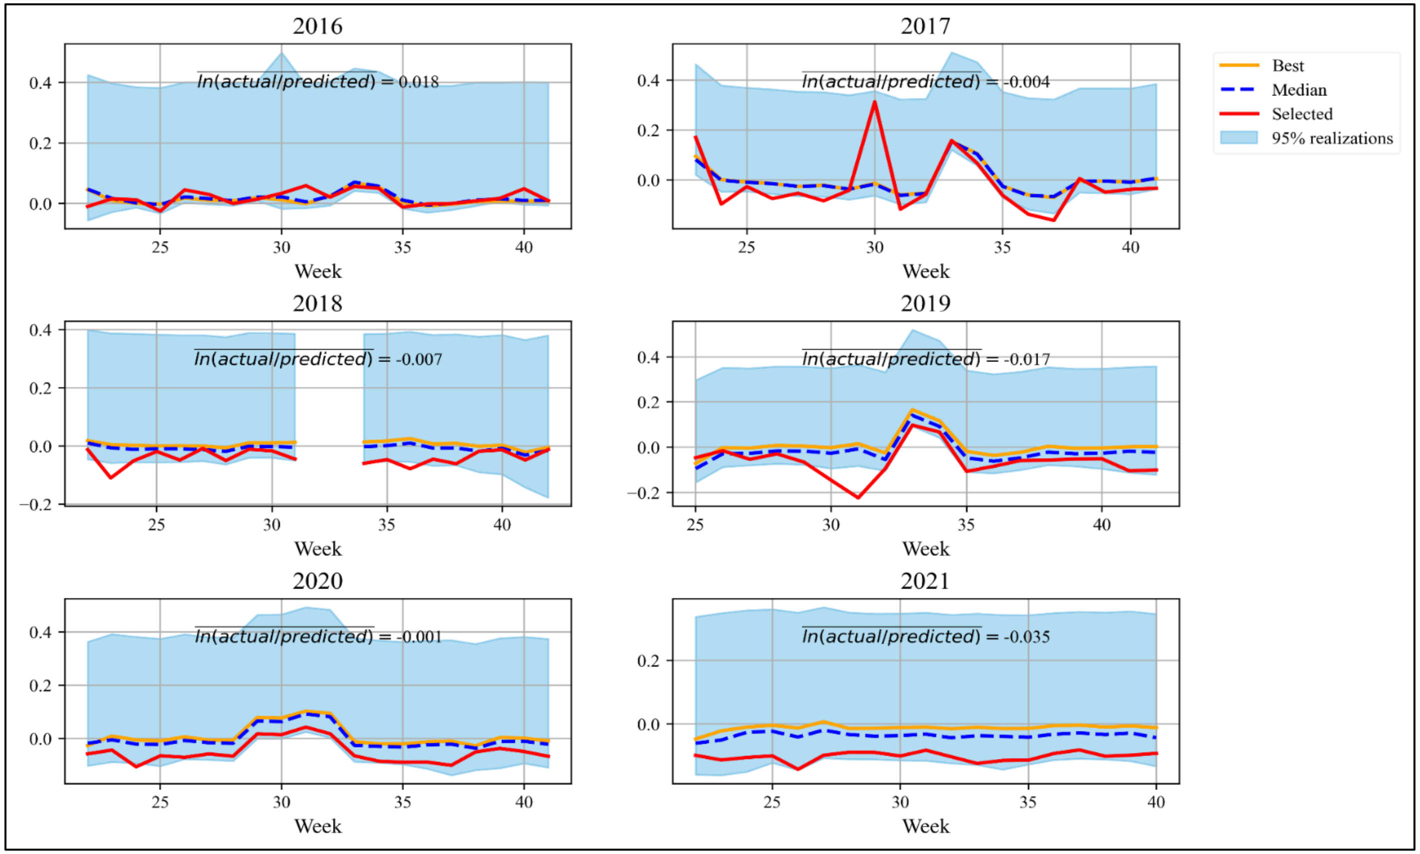

Figure 4.10. Model prediction errors in log-difference during 2016-2021 – Soybeans

Note: See notes below Figure 4.9.

Table 4.2. Summary Statistics – Predicted Crop Condition and Errors during 2016-2021

	Mean	**Median**	**SD**	**Min**	**Max**	**Jarque-Bera test statistics**†	**Wilcoxon test statistics**††	**No. non-missing Obs**
A. Corn								
Actual condition§	3.6774	3.6528	0.1489	3.3837	3.9133	**9.12*** **	3677	131
Best prediction	3.6823	3.6659	0.1416	3.0935	3.9017	**6.15** **	N/A	131
Median prediction	3.7450	3.7269	0.1224	3.2392	3.9603	**4.43* **	**223*** **	131
Best market expectation proxy	3.9064	3.8731	0.1397	3.4570	4.3601	**6.87** **	**141*** **	131
Log-difference errors of best prediction§§	-0.001	-0.001	0.018	-0.109	0.096	**2284.6*** **	3650	131
Median log-difference errors	-0.018	**-0.019*** **	0.020	-0.128	0.054	**250.05*** **	**141*** **	131
Best surprise proxy	-0.061	**-0.066*** **	0.039	-0.178	0.021	0.648	**192*** **	131

Statistical significance code: *** 0.01 ** 0.05 * 0.10

Table 4.2 (cont.). Summary Statistics –predicted crop condition and errors during 2016-2021

	Mean	Median	SD	Min	Max	Jarque-Bera test statistics†	Wilcoxon test statistics††	No. non-missing Obs
B. Soybeans								
Actual condition§	3.6292	3.6433	0.1491	3.3819	3.8614	**11.15*** **	3089	115
Best prediction	3.6150	3.6415	0.1739	2.9120	3.8569	**48.48*** **	N/A	115
Median prediction	3.6566	3.6747	0.1584	2.9860	3.8924	**114.69*** **	**553*** **	115
Best market expectation proxy	3.7869	3.8014	0.2337	2.5538	4.3123	**328.72*** **	**492*** **	115
Log-difference errors of best prediction§§	0.004	-0.002	0.037	-0.071	0.166	**198.65*** **	3067	115
Median log-difference errors	-0.007	**-0.016*** **	0.040	-0.094	0.153	**112.83*** **	**555*** **	115
Best surprise proxy	-0.041	**-0.051*** **	0.070	-0.223	0.312	**177.19*** **	**498*** **	115

Note: Table 4.2 provides summary statistics for corn (panel a) and soybean (panel b) condition, their predicted conditions and log-difference errors. Statistics are reported for the aggregate of 18 production states during 2016-2021.

† Jarque-Bera test for the null hypothesis that the sample comes from a normal distribution with an unknow mean and variance.

†† For each condition series that is not the best prediction, Wilcoxon signed rank test statistics is reported for the null hypothesis that it comes from the same population with the best predictions. Likewise, for each log-difference error series that is not the errors of the best prediction series, the null hypothesis is that it comes from the same population with the best prediction's log-different errors. For the best predictions' log different error series, the null is that it comes from a population with zero median. For the best prediction series, the test is not applicable since by construction, crop condition index (and consequently its best prediction) takes a minimum value of 1.

§ Actual condition, best prediction, median prediction and best market expectation proxy are defined as in Figure 4.7.

§§ Log-difference error of best prediction, median log-difference errors and best surprise proxy are defined as Figure 4.9.

Most importantly, Figures 4.7 to 4.10 and Table 4.2 convincingly prove that the best predictions are far different from the best market expectation proxies. This holds true for almost every year and for both crops, with the exception of 2016. In Figure 4.7 (corn) and Figure 4.8 (soybeans), not only are the model-selected proxies for market expectations (in red) clearly distant from the best prediction and the median series, but they also oscillate more strongly than all other series on the graphs. Occasionally, they spike outside the zone of 95-percent predictions. The Wilcoxon test statistics in Table 4.2 consistently reject the null hypothesis that these returns-derived surprise series are drawn from the same distribution with the best forecasts' errors, neither the median forecasts' errors. The test also rejects the null that the median of the surprise proxy is not significantly different from zero.

According to Table 4.2, the post-event price movements during 2016-2021 reveal that the CPCRs surprise the market by 6 percent below what had been expected for corn condition, and 5 percent for soybean condition – using the median statistic. This finding suggests that the market as a whole have a general tendency of overestimating the corn and soybean condition during this period. In extreme cases, the CPCRs even report soybean conditions 22 percent below market expectations (for week 31 in 2019) and corn conditions 18 percent below market expectations (for week 41 in 2018).

The 2019 growing season is considered a historically bad soybean season with late planting, flooding, and below-average temperature (USDA 2020). Since that year's soybean crop's progress and condition kept falling below historical records, they must have been more difficult to predict. Thus, the strong swings in expectations seen on the plot are plausible.

Nevertheless, in the case of corn, this underestimation should be taken with caution due to the unreasonable discrepancy detected in Section 4.5.1. Because the corn condition reported in the gridded dataset is about 9 percent below its counterpart in the traditional dataset (consistent with a 10 percent difference between the best prediction and the actual condition of the same week, as can be observed on the plot), it is possible that the true surprise is about half of the figure reported above. Even so, it is still the largest negative surprise for corn, and does not alter the median surprise reported above.

While the CPCRs do not bring much better-than-expected news regarding corn condition, our analysis shows that for soybean the announced condition can be 31 percent better than expected by the market (for week 30 in 2017). This large positive surprise is consistent with the fact that soybean condition had been consistently dropping since the beginning of the planting season until that point (as showed clearly in Figure 4.6), which reasonably caused a generally pessimistic outlook among observers (Irwin 2017). Hence, when soybean condition finally reverses that week, the resulting positive surprise is credibly large.

Finally, it is interesting to note that the derived surprises for both crops exhibit very similar patterns in 2021. According to price signals, the market consistently overestimates the condition of both crops by about 10 percent during that year. Our investigation in Section 4.5.1 reveals that the gridded condition is below the traditional figures for both crops in the same period only by 2 percent due to the unallocated pixels in "Other/combined counties". This is about the same size with the upward bias in the median predictions for that year. Even when we adjust for that bias, the reports still bring considerable negative surprises to the markets, as reflected in the price returns. There are two possible explanations for this phenomenon. First, as discussed in Section 4.5.1, most market participants might overestimate the crop condition because NASS considers some part of the planted acreage that is no longer publicly available in their state-level crop condition estimates. As both corn and soybean conditions deteriorate throughout 2021 crop season, it is not improbable that this unobservable crop portion pulls down the actual estimates, compared to market expectations. Since the original reports are produced at the state level and then interpolated to pixel-level for the gridded data, all the pixels within each state will inherit these lower-than-expected estimates – both the unallocated ones and the remaining ones. Thus, having removed the -2 percent of the excluded pixels, the other -8 percent difference from actual condition is likely the market surprise due to such unobservable information. The second possibility is that the reported crop condition throughout contradicts various yield forecasts that year. For example, all NASS yield surveys (in August, September, October, November, and end-of-season) had consistently reported near-record yield estimates for that year, both for corn and soybeans (USDA 2017). The final figures in the Crop Production 2021 Summary report indeed confirms that corn yields reach 177 bushels per acre (bpa) and soybean yields reach 51.4 bpa – which are the highest and the second highest yields since 1990, respectively. Clearly, it is reasonable for the market to expect better crop condition, provided such consistently high yield outlooks – especially provided that such negative correlation between condition and yield estimates rarely occurs (Bundy and Gensini 2022).[93]

[93] The authors document similar contradiction between corn condition and yield anomaly signals for 2020, 2017 and 2005 as well – but out of 35 years from 1986 to 2020. Thus, the overoptimistic pattern in corn condition expectations for 2020 and 2017 can be reasoned similarly. This explanation is also consistent with the situation of soybeans in some part of 2017, except those pessimistic plunges as discussed earlier.

4.5.4 *Explanatory power of the surprise proxies for the period 2016-2021*

Having derived the surprise series and shown that they can explain the situation of specific years, it is helpful to check their explanatory power to market returns for the entire period 2016-2021. With this step, we can achieve several objectives. First, as a robustness check, we can verify if the price impact of the derived surprises is consistent with what theory suggests. Second, conditioning on its theoretical validity, we can compare its explanatory power to price movements against the forecast errors of the best predictions, as well as the median predictions. Finally, we can draw conclusion about the market impact of it for the given period.

In Table 4.3, we report the regression results using Equation (4.19) for corn and soybeans in the entire period 2016-2021. Consistent with the yearly results, the best proxies for crop condition surprises yield the highest R^2. More importantly, the negative sign of the slope coefficients is consistent with the theory: a better-than-expected condition (*i.e.*, positive surprise) is price bearish because it implies that the crop size at the end of the season will be larger than what the market had been expected, and vice versa. The coefficients are highly significant for both corn and soybeans, supporting that the crop condition news do move the market as soon as the reports are released, though with small magnitude. The size of the reaction is similar for corn and soybeans, which is about -0.017. Thus, an average surprise of -6 percent for corn and -5 percent for soybeans (as presented in Table 4.2) are translated into approximately one percent (*resp.* 0.8 percent) increase in close-to-open returns on corn (*resp.* soybean) new-crop futures.

Table 4.3. Corn and Soybean New-crop Future Returns Explained by Different Surprise Proxies during 2016-2021

Dependent variable: $r_t = \ln P_{t,1} - \ln P_{t,0}$ †

	Corn			Soybeans		
	Best surprise proxy	Best predictions' errors	Median predictions' errors	Best surprise proxy	Best predictions' errors	Median predictions' errors
Constant	**-0.0012***** **(0.0004)**	-0.0002 (0.0003)	**-0.0007*** **(0.0004)**	-0.0003 (0.0003)	0.0005 (0.0003)	0.0002 (0.0003)
Log-difference surprises: x_t	**-0.0174**** **(0.0071)**	-0.0243 (0.0218)	-0.0282 (0.9175)	**-0.0167***** **(0.0052)**	**-0.019**** **(0.0082)**	**-0.0197***** **(0.007)**
Observations	131	131	131	115	115	115
R^2	0.0292	0.0117	0.0207	0.1285	0.0448	0.0582

Statistical significance code: *** 0.01 ** 0.05 * 0.10

Note: Table 4.3 reports the market impact estimation results for corn and soybean condition surprise during 2016-2021. Heteroskedasticity-consistent standard errors are reported in brackets.

† For both commodities, the dependent variable is the future returns of new-crop contracts (December for corn and November for soybeans). Similar to Lehecka (2014), we use the close-to-open returns from the event day to the next trading day as post-even market returns. This is because CPCRs are released after the closing time of the CME markets, as explained in Section 4.2.2.

4.6 Conclusion

One important function of public announcements about the demand and supply fundamentals of commodities is to guarantee information transparency, and thus to improve market efficiency. But how much new information they still bring to the market – or to put it differently, how much information in them had been anticipated by the market in advance – remains an unsettled question. The question becomes even more relevant in the era of big data analytics, when information is increasingly accessible and processible – in both volume and speed. In this paper, we propose a novel method to tackle this longstanding challenge with one such public report: the USDA Crop Progress and Condition reports.

Our theoretical framework extends semi-strong form Efficient Market Hypothesis with an important feature: to allow for the model that best approximates market expectations to be updated over time. To accommodate such great deal of flexibility in search of aggregated expectations, we argue that nonparametric ML can be an effective and powerful tool – as much as it can be exploited by individual market participants in prediction. Towards that end, we deploy the recent availability of the reports' gridded dataset with substantial improvements in spatial resolution. We develop an innovative ML modeling framework based on Extreme Gradient Boosting algorithm, which allows us to generate a large set of predicted corn and soybean conditions for the period 2016-2021. Our metric to evaluate the correlation between these candidates and the true market expectations is unique but intuitive: the ability to explain post-event price movements.

Despite the fact that highly accurate predictions can be generated within our models, post-release price movements point to an average CPCR surprise of 5-6 percent from what the market might have expected during 2016-2021, with occasional spikes up to 30 percent in absolute terms. There are a few over-pessimistic episodes observed in soybean condition expectations. But more often, both markets tend to hold more optimistic expectations about the crop rating than what actually comes out later in CPCRs. Given that crop condition rating only provides a transitory prospect of final crop size (as evident from the disagreement with yield forecasts) and is to be updated frequently throughout the crop season, the small market reaction to it is not surprising.

Our work contributes to the extant literature in many regards. Most importantly, we offer a new methodological approach to the problem of disentangling the unanticipated component of public announcements. We show that ML can be a good alternative to traditional parametric analyses for this complex setting, but only after we modify it properly according to the underlying theories. Though the relevant information set for market expectations varies with markets and report types, our ML modeling framework is applicable in various contexts. For example, by complementing the analyst forecast errors with our ML-based surprise proxies, we can shed more light on the puzzle of natural gas returns on the days of Energy Information Administration's storage announcements, as recently raised by Prokopczuk, Simen and Wichmann (2021). Given the excellent predictive performance of our predicting models, they can also be employed to several other research directions, including in-season price forecasting (Adjemian, Bruno, and Robe 2020), spatial variation of local bases (Bain and Fortenbery 2017) or crop insurance (Sherrick 2015). One interesting question is to which extent crop progress and condition predictions can help reduce the moral hazard in prevent planting claims in the early phase of crop season, as detected by Wu, Goodwin, and Coble (2020). In this particular context of the CPCRs, despite rapid developments in information technology, we reject the doubt that the market can anticipate crop condition well in advance and the CPCRs no longer provide new information to market participants, as questioned by Bain and Fortenbery (2017). Thus, not only market participants but also other user groups – to whom crop condition information is relevant, such as crop insurers and agronomists – should pay attention to the CPCRs.

For expositional purpose and consistent with previous literature, our study makes use of simple linear, bivariate specification for the returns-surprise equation. It can be argued that this choice is sufficient to evaluate the correlation with price movements among the surprise candidates, especially due to the distinct timing of the CPCR report releases. However, future research in different contexts should discover the possibility to incorporate more complex market reaction processes into the model selection phase. One example is the S-shape pattern of surprise-return relationship that has long been observed in stock price reactions to earning announcement surprises (Freeman and Tse 1992; Kinney, Burgstahler, and Martin 2002).

Acknowledgement

We are grateful to Don Fullerton, Francoise Robe, Jared Hutchins, Joe Janzen, Olga Isengildina-Massa and Scott Irwin for helpful discussions and suggestions. We thank conference participants at the 2022 and 2023 Conferences on Applied Commodity Price Analysis, Forecasting, and Market Risk Management (NCCC-134 in St. Louis, MO), the 2022 Meeting of the Agricultural and Applied Economics Association (AAEA in Anaheim, CA), and seminar participants at the U.S. Commodity Futures Trading Commission (CFTC), the USDA's National Agricultural Statistics Service (NASS) and the Universities of Illinois at Urbana Champaign, for helpful comments on an early draft. This paper was partly supported by a NIFA grant under Multistate Project Number S1072. Heckelei acknowledges partial funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2070 – 390732324). Gebrekidan acknowledges partial funding by the German Research Foundation under the CRC/Transregio 228: Future Rural Africa: Future-making and social-ecological transformation (Project number: 328966760). Cao gratefully acknowledges financial support from the Bonn International Graduate School at the University of Bonn's Agricultural Faculty. Robe also gratefully acknowledges the financial support received in his capacity as The Clearing Corporation Foundation Professor in Derivatives Trading at the University of Illinois.

While this paper was written, Robe also served as a part-time consulting senior economist at the CFTC; no CFTC resources or confidential information were used for this project. The views expressed in this article are those of the authors only and may not be attributed to the CFTC, the Commissioners, or other CFTC staff. The authors are solely responsible for any error or omission.

4.7 References

Adjemian, M. K. (2012). "Quantifying the WASDE Announcement Effect." *American Journal of Agricultural Economics*, *94*(1), 238–256. https://doi.org/10.1093/ajae/aar131

Adjemian, M. K., Bruno, V. G. and Robe, M. A. (2020). "Incorporating Uncertainty into USDA Commodity Price Forecasts." *American Journal of Agricultural Economics*, *102*(2), 696–712. https://doi.org/10.1002/ajae.12075

Bain, R. and Fortenbery, T. R. (2017). "Impact Of Crop Condition Reports on National and Local Wheat Markets". *Journal of Agricultural and Applied Economics*, *49*(1), 97–119.

Bartolini, L. (2008). Comment: "Noisy Macroeconomic Announcements, Monetary Policy, and Asset Prices" (Asset Prices and Monetary Policy). Chicago: University of Chicago Press.

Baur, R. F. and Orazem, P. F. (1994). "The Rationality and Price Effects of U.S. Department of Agriculture Forecasts of Oranges." *The Journal of Finance*, *49*(2), 681–695. https://doi.org/10.1111/j.1540-6261.1994.tb05157.x

Beguería, S. and Maneta, M. P. (2020). "Qualitative crop condition survey reveals spatiotemporal production patterns and allows early yield prediction." *Proceedings of the National Academy of Sciences*, *117*(31), 18317–18323.

Boudoukh, J., Richardson, M., Shen, Y. and Whitelaw, R. F. (2007). Do asset prices reflect fundamentals? Freshly squeezed evidence from the OJ market." *Journal of Financial Economics*, *83*(2), 397–412. https://doi.org/10.1016/j.jfineco.2004.06.017

Brunnermeier, M. K. (2005). "Information Leakage and Market Efficiency". *The Review of Financial Studies*, *18*(2), 417–457. https://doi.org/10.1093/rfs/hhi015

Bundy, L. R. and Gensini, V. A. (2022). "An assessment of USDA corn condition ratings across the U.S. Corn Belt." *Agronomy Journal*, *114*(1), 601–617. https://doi.org/10.1002/agj2.20973

Cao, A. N. Q. and Robe, M. A. (2022). "Market uncertainty and sentiment around USDA announcements." *Journal of Futures Markets*, *42*(2), 250–275.

Chiang, C.-H., Dai, W., Fan, J., Hong, H. and Jun, T. (2019). "Robust Measures of Earnings Surprises." *The Journal of Finance*, *74*(2), 943–983. https://doi.org/10.1111/jofi.12746

Dorfman, J. H. and Karali, B. (2015). “A Nonparametric Search for Information Effects from USDA Reports.” *Journal of Agricultural and Resource Economics*, *40*(1), 124–143. Retrieved from www.jstor.org/stable/44131280

Ederington, L. H. and Lee, J. H. (1993). “How Markets Process Information: News Releases and Volatility.” *The Journal of Finance*, *48*(4), 1161–1191. https://doi.org/10.1111/j.1540-6261.1993.tb04750.x

Fama, E. F., Fisher, L., Jensen, M. C. and Roll, R. (1969). “The adjustment of stock prices to new information.” *International Economic Review*, *10*(1), 1–21.

Freeman, R. N. and Tse, S. Y. (1992). “A Nonlinear Model of Security Price Responses to Unexpected Earnings.” *Journal of Accounting Research*, *30*(2), 185. https://doi.org/10.2307/2491123

Gao, F., Anderson, M.C., Zhang, X., Yang, Z., Alfieri, J.G., Kustas, W. P., Mueller, R., Johnson, D.M., Prueger, J. H. (2017). “Toward mapping crop progress at field scales through fusion of Landsat and MODIS imagery.” *Remote Sensing of Environment*, *188*, 9–25. https://doi.org/10.1016/j.rse.2016.11.004

Garcia, P., Irwin, S. H., Leuthold, R. M. and Yang, L. (1997). “The value of public information in commodity futures markets.” *Journal of Economic Behavior & Organization*, *32*(4), 559–570. https://doi.org/10.1016/S0167-2681(97)00013-9

Gousgounis, E. and Onur, E. (2018). “The effect of pit closure on futures trading”: *Journal of Commodity Markets*, *10*, 69-90. https://doi.org/10.1016/j.jcomm.2017.11.002

Hirshleifer, D., Lim, S. S. and Teoh, S. H. (2009). “Driven to Distraction: Extraneous Events and Underreaction to Earnings News.” *The Journal of Finance*, *64*(5), 2289–2325. https://doi.org/10.1111/j.1540-6261.2009.01501.x

Irwin, S. (2017). “Should We Be Surprised about the U.S. Average Yield of Corn and Soybeans in 2017?” *Farmdoc Daily*, *7*(201). Retrieved from https://farmdocdaily.illinois.edu/2017/11/should-we-be-surprised-about-us-average-yield.html

Irwin, S. and Hubbs, T. (2018). “Measuring the Accuracy of Forecasting Corn and Soybean Yield with Good and Excellent Crop Condition Ratings.” *Farmdoc Daily*, *8*.

Karali, B., Irwin, S. H. and Isengildina-Massa, O. (2019). “Supply Fundamentals and Grain Futures Price Movements.” *American Journal of Agricultural Economics*, *44*(3), 362. https://doi.org/10.1002/ajae.12012

Kinney, W., Burgstahler, D. and Martin, R. (2002). “Earnings Surprise “Materiality” as Measured by Stock Returns.” *Journal of Accounting Research*, *40*(5), 1297–1329. https://doi.org/10.1111/1475-679X.t01-1-00055

Kruse, J. and Smith, D. (1994). “Yield estimation throughout the growing season.”

Lehecka, G. V. (2014). “The value of USDA crop progress and condition information: Reactions of corn and soybean futures markets.” *Journal of Agricultural and Resource Economics*, 88–105.

Mourtzinis, S., Rattalino Edreira, J. I., Conley, S. P. and Grassini, P. (2017). “From grid to field: Assessing quality of gridded weather data for agricultural applications.” *European Journal of Agronomy*, *82*, 163–172. https://doi.org/10.1016/j.eja.2016.10.013

Oliveira, M., Torgo, L. and Santos Costa, V. (2021). “Evaluation Procedures for Forecasting with Spatiotemporal Data.” *Mathematics*, *9*(6), 691. https://doi.org/10.3390/math9060691

Prokopczuk, M., Simen, C. W. and Wichmann, R. (2021). “The Natural Gas Announcement Day Puzzle.” *The Energy Journal*, *42*(2). https://doi.org/10.5547/01956574.42.2.mpro

Rigobon, R. and Sack, B. (2008). 8. Noisy Macroeconomic Announcements, Monetary Policy, and Asset Prices: “Asset Prices and Monetary Policy.” In J. Y. Campbell (Ed.) (pp. 335–370). University of Chicago Press. https://doi.org/10.7208/9780226092126-011

Roznik, M., Mishra, A. K. and Boyd, M. S. (2023). “Using a machine learning approach and big data to augment WASDE forecasts: Empirical evidence from US corn yield.” *Journal of Forecasting.* Advance online publication. https://doi.org/10.1002/for.2956

Schlenker, W. and Roberts, M. J. (2009). “Nonlinear temperature effects indicate severe damages to U.S. Crop yields under climate change.” *Proceedings of the National Academy of Sciences of the United States of America*, *106*(37), 15594–15598. https://doi.org/10.1073/pnas.0906865106

Schnepf, R. (2006). “Price Determination in Agricultural Commodity Markets: A Primer.”

Sherrick, B. (2015). “Understanding the Implied Volatility (IV) Factor in Crop Insurance.” *Farmdoc Daily*, *5*.

Stevens, S. C. (1991). “Evidence for a weather persistence effect on the corn, wheat, and soybean growing season price dynamics.” *Journal of Futures Markets*, *11*(1), 81–88.

Storm, H., Baylis, K. and Heckelei, T. (2020). “Machine learning in agricultural and applied economics.” *European Review of Agricultural Economics*, *47*(3), 849–892.

UCAR (2022). “PRISM High-Resolution Spatial Climate Data for the United States: Max/min temp, dewpoint, precipitation | NCAR - Climate Data Guide.” Retrieved from https://climatedataguide.ucar.edu/climate-data/prism-high-resolution-spatial-climate-data-united-states-maxmin-temp-dewpoint

USDA, NASS (2017). NASS - Quick Stats. Retrieved from https://data.nal.usda.gov/dataset/nass-quick-stats

USDA, NASS (2020). Crop Production (January 2020). Retrieved from https://downloads.usda.library.cornell.edu/usda-esmis/files/tm70mv177/qn59qm27j/k3569m194/crop0120.pdf

USDA, NASS (2021). “Crop Progress and Condition Layers.” Retrieved from https://www.nass.usda.gov/Research_and_Science/Crop_Progress_Gridded_Layers/CropProgressDescription.pdf

Wardlow, B. D., Kastens, J. H. and Egbert, S. L. (2006). “Using USDA Crop Progress Data for the Evaluation of Greenup Onset Date Calculated from MODIS 250-Meter Data.” *Photogrammetric Engineering & Remote Sensing*, *72*(11), 1225–1234. https://doi.org/10.14358/pers.72.11.1225

Westcott, P. C. and Jewison, M. (2013). “Weather effects on expected corn and soybean yields.” *Washington DC: USDA Economic Research Service FDS-13g-01*.

Worrall, G., Rangarajan, A. and Judge, J. (2021). “Domain-Guided Machine Learning for Remotely Sensed In-Season Crop Growth Estimation.” *Remote Sensing*, *13*(22), 4605. https://doi.org/10.3390/rs13224605

Wu, S., Goodwin, B. K. and Coble, K. (2020). “Moral hazard and subsidized crop insurance.” *Agricultural Economics*, *51*(1), 131–142. https://doi.org/10.1111/agec.12545

Ying, J., Chen, Y. and Dorfman, J. H. (2019). “Flexible Tests for USDA Report Announcement Effects in Futures Markets.” *American Journal of Agricultural Economics*, *94*(1), 238. https://doi.org/10.1093/ajae/aaz013

4.8 Appendices

4.A.1 Predictors used to train XGB models for pixel-level crop condition prediction

Table 4.4. List of Models' Predictors

Category	Variable	Data type	Spatial resolution	Temporal resolution	Temporal lags	Total number of predictors
Temporal information	Year	Categorical	9x9 km	Weekly	current week	1
	Calendar week	Cyclical[94]	9x9 km	Weekly	current week	2
Spatial information	State	Categorical	9x9 km	weekly	current week	1
Previous CPCRs	Crop progress	continuous	9x9 km	weekly	one previous week	1
	Crop condition	Continuous	9x9 km	weekly	one previous week	1
Weather variables	Minimum temperature	Continuous	4x4 km	daily	current week and two previous weeks	21
	Maximum temperature	Continuous	4x4 km	daily	current week and two previous weeks	21
	Average temperature	Continuous	4x4 km	daily	current week and two previous weeks	21
	Precipitation	Continuous	4x4 km	daily	current week and two previous weeks	21
	Mean dew point temperature	Continuous	4x4 km	daily	current week and two previous weeks	21
	Minimum vapor pressor deficit	continuous	4x4 km	daily	current week and two previous weeks	21
	Maximum vapor pressure deficit	continuous	4x4 km	daily	current week and two previous weeks	21
Total						**153**

[94] cyclical transformation into sin and cosin numerical variables

4.A.2 XGB implementation: further notes

Table 4.5 provides an overview of the XGB hyperparameters used for our training process described in Section 4.4.3 and 4.5.2. Figure 4.11 (*resp.* Figure 4.12) then depict the full range of corn (*resp.* soybean) condition predictions generated by 12,800 model candidates corresponding to these 12,800 hyperparameter combinations. The two Figure are constructed analogously to the plots of selected 95 percent of realizations in Figure 4.7 and Figure 4.8, but here we also add the full range from lowest to highest prediction values for each week (*i.e.*, the yellow-shaded area below the 2.5-percentile and above the 97.5-percentile). As shown by both Figure 4.10 and Figure 4.11, our choice of hyperparameter spaces successfully serves the purpose of generating a wide range of predictions, from poorly underpredicting to extremely overpredicting. This dispersed distribution of predicting outcomes help reducing the risk of selecting bias surprise proxies due to the lack of outcomes distributed in some certain ranges of the distributions.

In addition, we provide the plots of the R^2 of the market impact regression using Equation (19) against the Root Mean Weighted Square Error (RMWSE) of our surprise proxy population in Figure 4.13 and Figure 4.14 for corn and soybeans, respectively. These plots strengthen our arguments in Section 4.4.3 and 4.5.3, that there is no obvious relationship between predictive accuracy and capability to explain post-release price movements – except the fact that for none of the year-crop combinations, the most accurate model is selected as the model which generates the most correlated errors with price returns. Neither is it true for the poorly predicting candidates.

Table 4.5. Hyperparameter Space Used for XGB

Hyperparameter	Data type	Value range/set	Total number of possible values
n_estimators: number of decision trees to be boosted; more trees tend to increase the precision and robustness of the prediction, but at the cost of time and computational resource.	integer	[20,1000]; increasing 20 stepwise	50
max_depth: the maximum depth of a tree. Deeper trees increase the complexity of the model and thus improve model fit, but at the risk of overfitting.	Integer	[3, 10]; increasing 1 stepwise	8
colsample_bytree: the fraction by which a subset of predictors is randomly sampled for each tree. A lower fraction increases the variability of features (*i.e.*, predictors) among the trees and thus leads to a less conservative model.	Fraction	[0.8]	1
subsample: the fraction of observations randomly sampled for each tree. A lower fraction increases the variability of the samples used among the trees and thus leads to a less conservative model.	Fraction	[0.5; 0.8]	2
learning_rate: step size shrinkage used in updating model weights to prevents overfitting. A lower value indicates a slower speed of updates.	Fraction	[0.01; 0.1; 0.5; 1]	4
Reg_lambdas: Ridge's type of regulation on weights. A higher value makes the model more conservative.	Fraction	[0.00001; 0.01; 0.1; 1]	4
		Cartesian product of all sets	**12,800**

Source: XGBoost documentation page, https://xgboost.readthedocs.io/en/stable/parameter.html

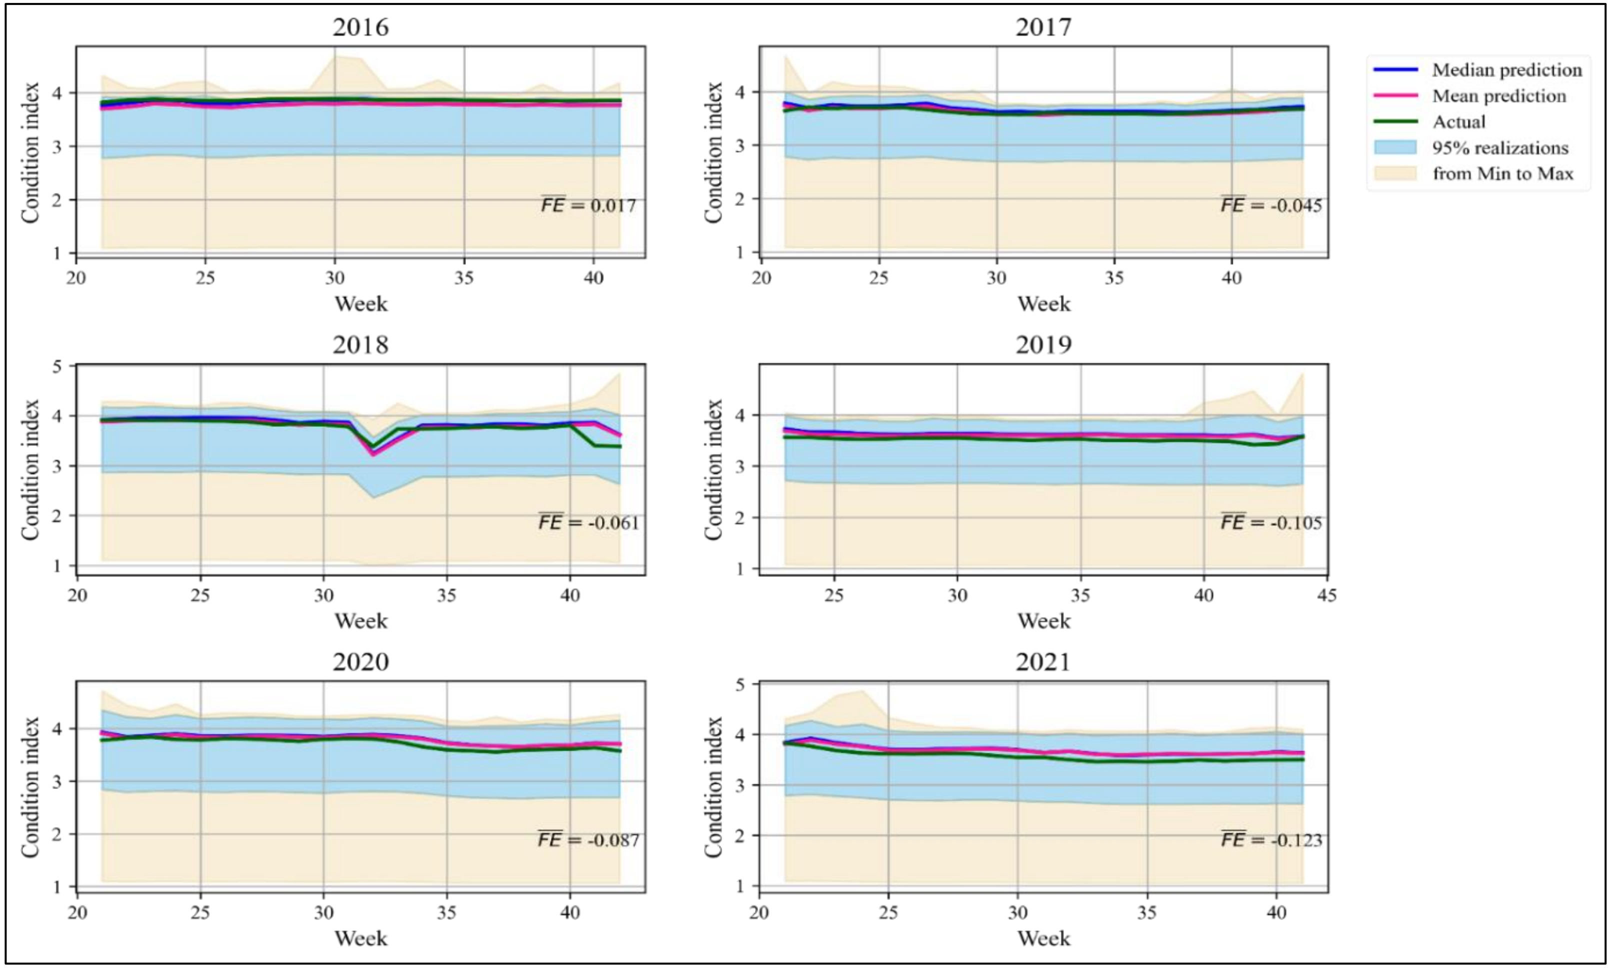

Figure 4.11. Full Distributions of Predicted Crop Condition Generated by 12,800 Model Candidates for Each Year during 2016-2021 – Corn

Note:

- Median prediction: the sample median of 12,800 predicted crop condition values that are produced by 12,800 model candidates. Each of them is an 18-state aggregate using the procedure in Figure 4.5.
- Mean prediction: the sample mean of 12,800 predicted crop condition values that are produced by 12,800 model - candidates. Each of them is an 18-state aggregate using the procedure in Figure 4.5.
- Actual: 18-state aggregated crop condition from NASS gridded dataset, which is denoted as g_{cc} in Section 4.5.1.
- 95% realization: from 2.5-percentile to 97.5-percentile of 12,800 predicted crop condition values as explained above. In other words, 47.5 percent of predictions above the median prediction and 47.5 percent of predictions below the median prediction fall within this range.
- from Min to Max: from the lowest predicted value to the highest predicted value of crop condition among 12,800 predicted condition series.
- $\overline{FE}$ is the yearly simple average of the median prediction.
- Figure 4.12 below is produced analogously for soybeans. The disrupted gap on the subplot 2018 is due to the missing dataset in week 32 for soybeans, as discussed in Section 4.5.1. Since we need the condition and progress of the previous week as predictors, the prediction for week 33 cannot be made neither.

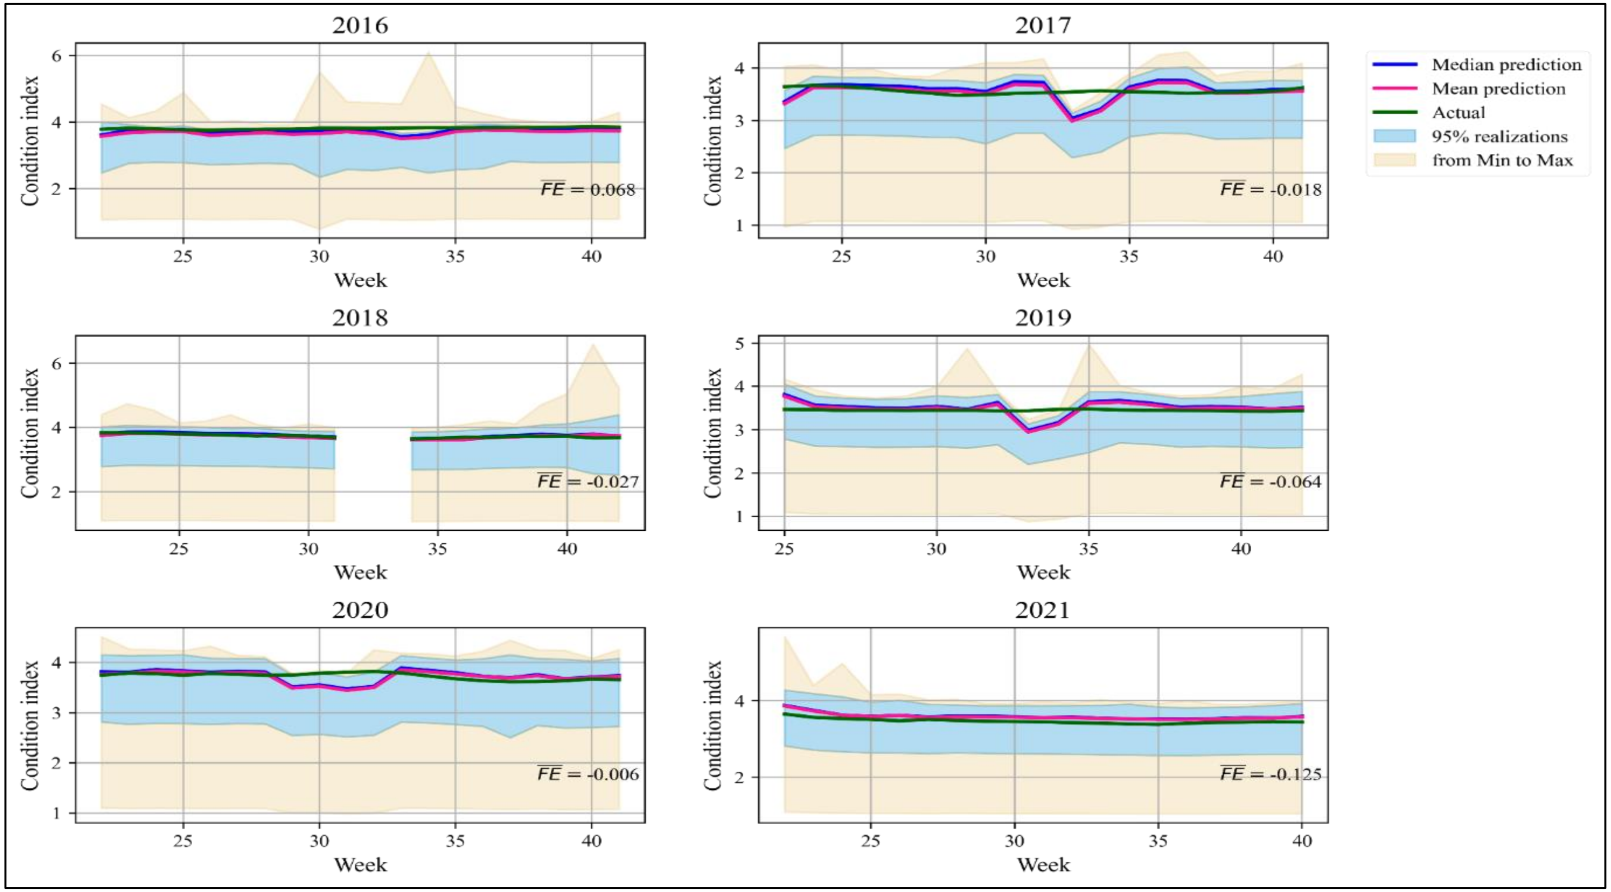

Figure 4.12. Full Distributions of Predicted Crop Condition Generated by 12,800 Model Candidates for Each Year during 2016-2021 – Soybeans

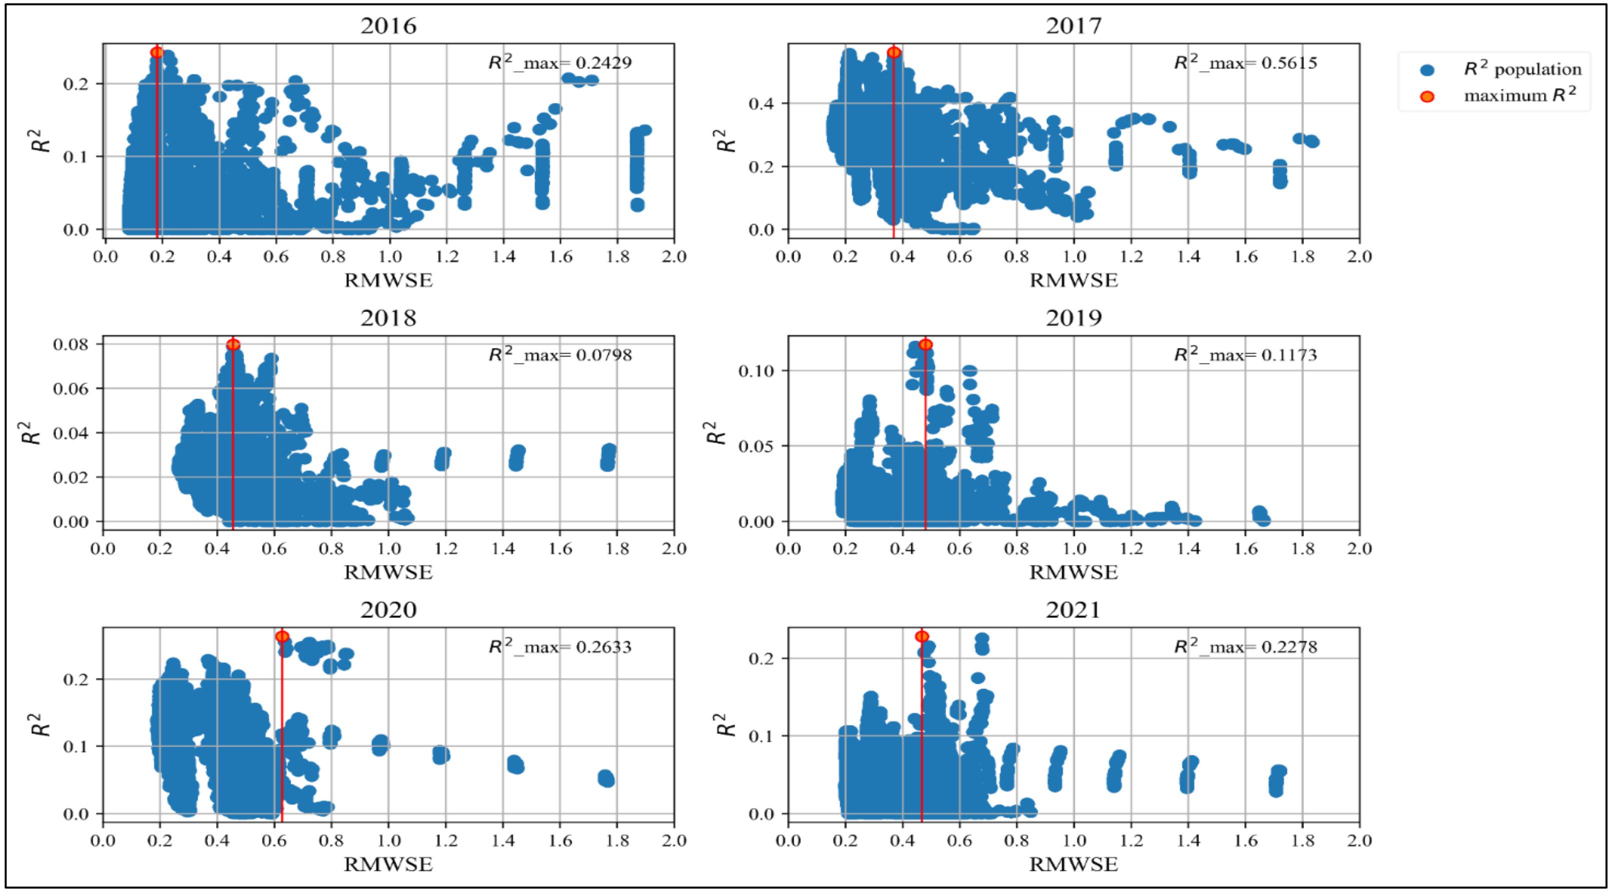

Figure 4.13. Post-event Returns Variation Explained by Models' Surprise Proxies during 2016-2021 – Corn

4.8 Appendices

Note:

- RMWSE: Root Mean Weighted Square Error. This is a modified version of the root mean square error (RMSE) criterion used in conventional ML applications. For each week, squared forecast errors at pixel level are aggregated to 18-state condition using the procedure described in Figure 4.5. Then, we take the average of all weeks in the year and finally take the square root of it to obtain the RMWSE
- As it appears on Figure 4.7, 95 percent of predictions are distributed within a range of ± 2 around the medians (in absolute terms of the continuous index's value range). Hence, for better visualization, in this Figure we only focus on those series with a maximum RMWSE value of 2. As it becomes clear on the graphs, all the series which highest R^2 have RMWSE below 2.
- Figure 4.14 below is produced analogously for soybeans.

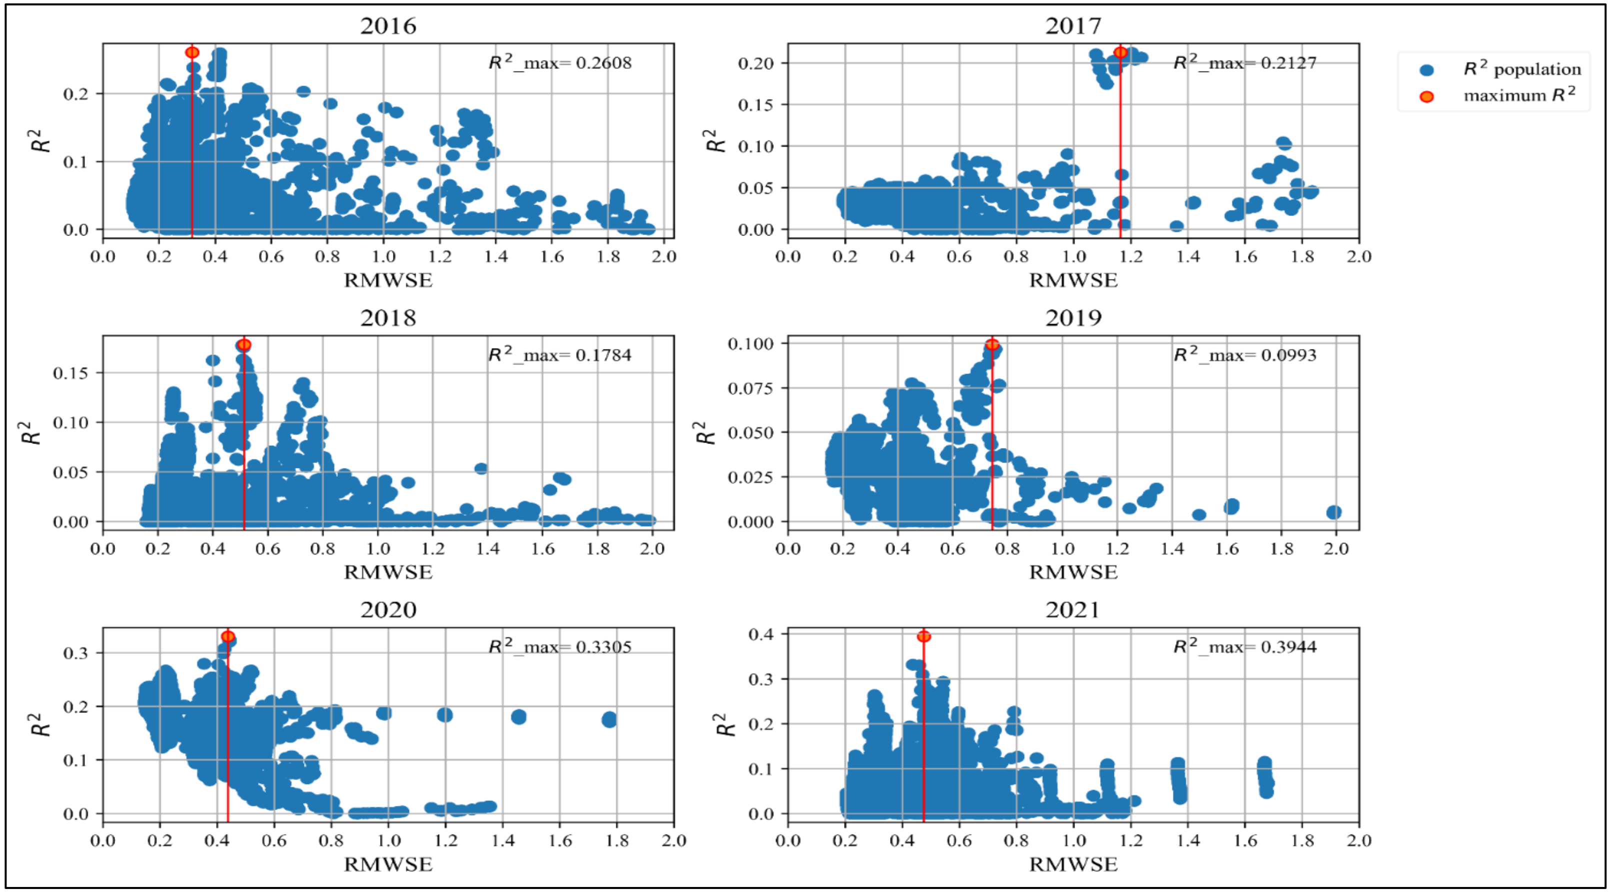

Figure 4.14. Post-event Returns Variation Explained by Models' Surprise Proxies during 2016-2021 – Soybeans