

2015

Essays on the forecasts of ending stocks

Jinzhi Xiao
Iowa State University

Follow this and additional works at: <https://lib.dr.iastate.edu/etd>



Part of the [Economics Commons](#)

Recommended Citation

Xiao, Jinzhi, "Essays on the forecasts of ending stocks" (2015). *Graduate Theses and Dissertations*. 14895.
<https://lib.dr.iastate.edu/etd/14895>

This Dissertation is brought to you for free and open access by the Iowa State University Capstones, Theses and Dissertations at Iowa State University Digital Repository. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Essays on the forecasts of ending stocks

by

Jinzhi Xiao

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Major: Economics

Program of Study Committee:
Sergio H. Lence, Co-Major Professor
Chad E. Hart, Co-Major Professor
Dermot J. Hayes
David A. Hennessy
John R. Schroeter

Iowa State University

Ames, Iowa

2015

Copyright © Jinzhi Xiao, 2015. All rights reserved.

DEDICATION

This dissertation is dedicated to my family. My special gratitude goes to my parents for their unwavering support and encouragement. I would also like to dedicate this work to my many friends who have kept me going over the years.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	vi
ACKNOWLEDGMENTS	vii
CHAPTER 1 GENERAL INTRODUCTION.....	1
CHAPTER 2 USDA FORECASTS OF CROP ENDING STOCKS: HOW GOOD ARE THEY	3
2.1 Abstract	3
2.2 Introduction.....	3
2.3 The Model.....	8
2.3.1 Background Models	8
2.3.2 Proposed Model	12
2.4 Data	17
2.5 Estimation Methods	19
2.6 Results and Discussion	23
2.7 Conclusions.....	28
2.8 References.....	30
2.9 Appendix: Conditional Posterior Distributions for Model Parameters in the Gibbs Sampler	32
CHAPTER 3 THE ENDING STOCKS FORECASTS FROM PRIVATE ANALYSTS: AN EXAMINATION OF EFFICIENCY	37
3.1 Abstract	37
3.2 Introduction.....	37
3.3 The Model.....	43
3.4 Data	47
3.4.1 Data Sources and Structure.....	47
3.4.2 Creating a Representative Analyst.....	49
3.4.3 Individual Analysts' Forecasts.....	51
3.4.4 Descriptive Statistics.....	51
3.5 Estimation Methods	52
3.5.1 Estimation Procedures for Representative Analysts' Forecasts.....	52
3.5.2 Integration of Multiple Imputation and MCMC Estimations	55
3.6 Results and Discussion	57
3.6.1 Representative Analysts and USDA Forecasts	57

3.6.2	Individual Analysts' Forecasts.....	62
3.6.3	Discussion.....	63
3.7	Conclusions.....	66
3.8	References.....	68
3.9	Appendix: Conditional Posterior Distributions for Model Parameters in the Gibbs Sampler.....	70
CHAPTER 4	DO ANALYSTS FORECAST THE ENDING STOCKS OR THE USDA FORECASTS?.....	85
4.1	Abstract.....	85
4.2	Introduction.....	85
4.3	The Model.....	91
4.3.1	The Structure of Forecasts.....	91
4.3.2	Proposed Model.....	93
4.4	Data.....	100
4.5	Estimation Methods.....	102
4.6	Results and Discussion.....	105
4.6.1	Corn.....	105
4.6.2	Soybeans.....	107
4.6.3	Discussion: Forecasting Behaviors.....	108
4.7	Conclusions.....	109
4.8	References.....	109
4.9	Appendix: Conditional Posterior Distributions for Model Parameters in the Gibbs Sampler.....	111

LIST OF FIGURES

	Page
Figure 2.1 Standard deviations of USDA forecast revisions, by month.....	34
Figure 3.1 Monthly standard deviations of forecast revisions for corn	72
Figure 3.2 Monthly standard deviations of forecast revisions for soybeans	72
Figure 3.3 Monthly standard deviations of forecast revisions for wheat.....	73
Figure 3.4 The integrated procedure of multiple imputations and MCMC estimations	74
Figure 3.5 Estimated coefficients for individual analysts' forecasts for corn ending stocks	75
Figure 3.6 Estimated coefficients for individual analysts' forecasts for soybeans ending stocks.....	76
Figure 3.7 Estimated coefficients for individual analysts' forecasts for w ending stocks.....	77
Figure 3.8 Forecast timing for corn/soybean ending stocks	78
Figure 3.9 Forecast timing for wheat ending stocks	78
Figure 4.1 Timeline of USDA and analysts' forecasts of ending stocks for corn and soybeans	113

LIST OF TABLES

	Page
Table 2.1 Descriptive statistics for USDA forecast revisions	34
Table 2.2 Means and Standard Deviations for the estimates of USDA ending stocks forecasts, 1985/86 – 2013/14.....	35
Table 2.3 Medians and Credible Intervals for the estimates of USDA ending stocks forecasts, 1985/86 – 2013/14.....	36
Table 3.1 List of the selected frequent analysts	79
Table 3.2 Descriptive statistics for USDA and analysts' forecast revisions for corn, soybeans and wheat.....	80
Table 3.3 Parameter estimates for the ending stocks forecasts for corn, 2004/05 – 2013/14.....	81
Table 3.4 Parameter estimates for the ending stocks forecasts for soybeans, 2004/05 – 2013/14	82
Table 3.5 Parameter estimates for the ending stocks forecasts for wheat, 2004/05 – 2013/14	83
Table 3.6 Median and Credible Intervals of the estimated coefficients for ending stocks forecasts, 2004/05 – 2013/14	84
Table 4.1 Descriptive statistics for USDA and analysts' forecast revisions for corn, soybeans and wheat.....	114
Table 4.2 Means and standard deviations for the estimates of USDA and analysts ending stocks forecasts, 2004/05 – 2013/14	115
Table 4.3 Medians and 95% credible intervals for the estimates of USDA and analysts ending stocks forecasts, 2004/05 – 2013/14	116

ACKNOWLEDGMENTS

I would like to express my deepest gratitude to my co-major professors: Dr. Sergio Lence and Dr. Chad Hart. I would like to thank them for their incredibly valuable guidance and unwavering support throughout this research. The discussions with them have always been interesting and insightful. I am also indebted to them for their patience and encouragement during the hard times.

I am grateful to my committee members, Dr. Dermot Hayes, Dr. David Hennessy and Dr. John Schroeter, for their valuable criticisms and comments to this work. In addition, I would also like to thank the department faculty and graduate students who have offered me help over my graduate school life.

CHAPTER I

GENERAL INTRODUCTION

Ending stocks play an important role in decision making by market participants and policy makers. The provision of accurate forecasts of ending stocks is critical as it can timely reflect the market situation and reduce the uncertainty faced by decision makers. Over the years, the USDA and private analysts have been providing ending stocks forecasts. However, few studies have addressed USDA forecasts, and researchers have not investigated analysts' forecasts so far.

This dissertation focuses on analyzing the ending stocks forecasts issued by these two sources. It contains three essays which gradually delve into the USDA and private analysts' forecasting behaviors. The first essay advances existing models and analyze USDA forecasts. The proposed model focuses on forecast revisions and retains the link between forecasts and the ending stocks. The essay also introduces an error covariance structure specifically for ending stocks forecasts. The model is estimated using Bayesian Markov Chain Monte Carlo methods. Results show that USDA forecasts are inefficient.

Given this finding, the model is applied to private analysts' forecasts to find whether analysts can provide "better" forecasts. This part of analysis is performed in the second essay. Unlike USDA forecasts, analysts' forecasts are often incomplete. Thus, Essay Two first investigates analysts' forecasts as a group by combining them to create complete forecast data, and then proposes a method of integrating multiple imputation and MCMC estimations to analyze individual analysts' forecasts. Results show that analysts, as a group, are inefficient in making ending stocks forecasts. Besides, forecasting behavior vary across individual analysts.

Essay Two also finds that analysts and the USDA have similar behavior in forecasting corn and soybeans ending stocks. Thus, it is possible that their forecasts affect each other. Hence, essay three proposes a method to analyze the forecasts from these two sources together, utilizing the overlooked the information in previous studies. Results show that analysts actually forecast the USDA forecasts instead of the ending stocks.

The proposed models and methods are designed for general purposes. Thus, they can be easily extended by including additional features, as well as can be applied to other fixed-events.

CHAPTER 2

USDA FORECASTS OF CROP ENDING STOCKS: HOW GOOD ARE THEY?

2.1 Abstract

This paper evaluates forecasts of U.S. ending stocks for corn, soybeans and wheat issued by the U.S. Department of Agriculture (USDA) for the marketing years 1985/86 through 2013/14. The proposed efficiency tests focus on forecast revisions, with emphasis on the link between forecasts and final ending stocks. Forecast errors are decomposed into the sum of monthly unforecastable shocks and USDA's own idiosyncratic errors. The error covariance matrix contains both heteroscedasticity and auto-correlations. Results suggest that the USDA forecasts are inefficient, providing strong evidence that the USDA is conservative in forecasting the ending stocks. Shocks are found to be heteroscedastic, and idiosyncratic errors not negligible.

2.2 Introduction

The U.S. Department of Agriculture (USDA) provides different forecasts of supply and demand for major agricultural commodities in its monthly World Agricultural Supply and Demand Estimates (WASDE) reports. These forecasts include various balance sheet components for each crop, such as beginning stocks, imports, production, domestic use, exports, and ending stocks. It can be argued that the public provision of this information is valuable for market participants and enhances the overall functioning of agricultural markets. The WASDE forecasts not only provide the commodity's fundamental conditions for the private sector to make decisions, but also provide an important basis for relevant government policies (Allen 1994). Researchers have found that farmers, agribusinesses, government agencies and other market

participants place substantial value on WASDE forecasts, and adjust their market behavior accordingly. (*e.g.*, Bauer and Orazem 1994, Garcia *et al.* 1997, Isengildina-Massa *et al.* 2008a, 2008b, Adjemian 2012)

Ending stock is a measure of carryover of a commodity which enters the supply side of the market in the following marketing year. It reflects the imbalance of supply and demand of the commodity of interest. Many studies in agricultural forecasts have analyzed the accuracy and efficiency of USDA price and production forecasts. (*e.g.*, Irwin *et al.* 1994, Sanders and Manfredo 2002, 2003, Isengildina *et al.* 2004, 2006) In contrast, little attention has been paid to the ending stocks forecasts. To the best of our knowledge, only Botto *et al.* (2006) and Isengildina-Massa *et al.* (2013) have included ending stocks forecasts in their analyses. Botto *et al.* (2006) used a frequentist approach to investigate the accuracy of USDA ending stocks forecasts and estimated the trends in the forecast accuracy over the marketing years 1980/81 through 2003/04. They find a significant downward trend in the variance of forecast errors when forecast horizon shortens. They also find that almost all balance sheet categories are significant in explaining errors in ending stocks forecasts. Isengildina-Massa *et al.* (2013) analyzed how WASDE forecast errors are affected by selected behavioral and macroeconomic factors over the marketing years 1987/88 through 2009/10. They found strong evidence of inefficiency for both types of factors for the ending stocks forecasts.

USDA ending stocks forecasts are fixed-event forecasts, because they are made for a specific target (ending stocks), but have different forecast horizons. Previous research on fixed-event forecasts has often examined macroeconomic variables such as inflation rate, interest rate, and real and nominal GDP growth rates. (*e.g.*, Clements 1995, 1997, Romer and Romer 2000, Harvey *et al.* 2001, Clements *et al.* 2007) The models in the literature can be classified into two

main categories, namely, those based on Nordhaus (1987) and the ones following Davies and Lahiri (1995, 1999).

The models based on Nordhaus (1987) focus on forecast revisions. Nordhaus (1987) introduced a weak efficiency test which only uses information on past forecasts because the forecast history is always available to the public. The test consists of assessing whether changes in forecasts are affected by past forecast changes. Nordhaus (1987) applied the test to several macroeconomic, energy consumption, and oil price forecasts. He found significant autocorrelations in the revisions of these forecasts. Isengildina *et al.* (2006) extended Nordhaus' test and applied it to evaluate the USDA crop production forecasts. They found positive autocorrelations in forecast revisions. However, assumption of *i.i.d.* errors in the model may not be realistic. On one hand, the size of revisions may vary in different months, because the arrival of new information may have seasonal patterns. For example, in the early months of the forecasting cycle for crop ending stocks, revisions are likely to be larger relative to those in later months because of the uncertainty in crop production. On the other hand, errors can also be correlated if one takes into account that forecasters often correct their own errors in previous forecasts. Thus, it is interesting to analyze forecasters' behavior with a generalized error covariance structure.

Unlike the Nordhaus model, the framework advocated by Davies and Lahiri (1995, 1999) directly focuses on the forecast errors. They decompose the forecast errors into the sum of unforecastable shocks and the forecaster's own idiosyncratic errors. Their framework provides a way to explain the reason why forecasts made at a date closer to the target event tend to be more precise. Specifically, the fact that early forecasts typically have large variances can be explained by the stack of unforecastable shocks. As the forecasting horizons shorten, unforecastable shocks

are gradually revealed, so that less uncertainty is remained in forecasting. This approach is in line with studies of forecasts in other areas of fixed-events. (*e.g.*, Egelkraut *et al.* 2003, for crop production forecasts) Based on the Davies and Lahiri framework, Clements *et al.* (2007) first analyzed forecast revisions by differencing the forecast errors. In this way, they avoided the possible problem in the original Davies and Lahiri model that the dependent variables could be correlated with the errors. However, they only investigated the relationship between non-adjacent forecast revisions as endogeneity would occur if adjacent forecast revisions are used. Therefore they did not consider the impact of the most updated information. Besides, they did not fully estimate the error covariance matrix. Instead, they simplified the estimations by only considering the diagonal elements or restricting the idiosyncratic errors to be zero.

Based on the model proposed by Clements *et al.* (2007), the present study develops an estimation framework for examining the efficiency of fixed-event forecasts. We revisit the Nordhaus (1987) and Davies and Lahiri (1995, 1999) models and investigate forecast revisions by emphasizing the link between the forecasts and the forecast target, which is not included in the original Nordhaus test. Our framework also decomposes the errors into the sum of unforecastable shocks and USDA's own errors. Specifically, we take into account the forecaster's correction of their own errors. If such corrections occur and forecasts are efficient, then adjacent forecast revisions must be negatively correlated. Thus, the results from the original Nordhaus test are biased if forecasters in fact do correct their own errors.

The present study also makes the following three contributions to the literature. First, it introduced an error covariance matrix which is unique for crop ending stocks forecasts. Unlike other fixed event forecasts, there is a large number of forecast observations for the ending stock of a marketing year. Thus, it might not be wise to ignore the possible complex error structure

implied by these forecasts. By focusing solely on ending stocks forecasts, we can step further by building a specific error covariance structure for them and hence do a deeper investigation. Besides, some studies on fixed-event forecasts typically run ordinary least squares (OLS) first and then calculate the elements in the error covariance matrix. In contrast, the proposed error structure has reduced number of parameters to be estimated. This is especially useful for ending stocks forecasts which have many more forecast observations within a forecasting cycle. We also allow conditional heteroscedasticity in the error covariance matrix. On one hand, it reduces restrictions and fits the data better as it represents the size of shocks. On the other hand, it improves the estimates of the variances of the coefficients, increasing the credibility of the significance test. The proposed error covariance structure can be easily extended to other types of fixed-event forecasts.

Secondly, the test in our proposed framework is performed by estimating a system of equations instead of a single equation as in previous literature. The consideration of forecasters' correction of errors introduces autocorrelations and generates endogeneity if adjacent forecast revisions are used as explanatory variables. To address this issue, we treat the first forecast revision for each marketing year as exogenous, and all of the subsequent revisions as endogenous. The later revisions are formed by the joint effects of the first revisions, unforecastable shocks, and USDA's own errors. The revisions occurred within a marketing year are viewed as in the same panel. The efficiency test is then performed by estimating a system containing such panels from different marketing years. This method allows us to perform efficiency tests by only considering forecast revisions, which eliminates the problems that arise when the forecast errors are used directly.

Thirdly, the econometric analysis in our study is performed by means of a Bayesian Markov Chain Monte Carlo (MCMC) approach. Unlike frequentist approaches, Bayesian methods yield the full posterior distributions for the parameters of interest, which is especially useful when estimating parameters, such as error variances, are highly likely to be skewed. The method allows us to estimate the regression coefficients and the error covariance matrix in one iteration step.

The remainder of the study is organized as follows: Section 2.3 reviews the background models for analyzing the fixed-event forecasts and introduces the advocated model for evaluating crop ending stocks forecasts. Section 2.4 introduces the data and provides descriptive statistics. Section 2.5 presents the empirical MCMC methods employed for the estimation. Section 2.6 discusses the results, and the final section provides concluding remarks.

2.3 The Model

The present study evaluates USDA crop ending stocks forecasts by testing for bias and efficiency. The null hypothesis is as follows:

H_0 : USDA forecasts are unbiased and efficient forecasts of ending stocks.

A rejection of H_0 indicates that USDA forecasts are inefficient and can be improved upon by using existing information.

2.3.1 Background Models

Empirical studies on testing forecast bias and efficiency are typically based on Mincer and Zarnowitz (1969). To express their model in our notation, let S_t represent the realization of the ending stocks of a given commodity at the end of marketing year t . Let n be the forecast

horizon (*i.e.*, the number of months between the time the forecast is made and the final ending stocks) and $U_{t,n}$ be the USDA n -month-ahead forecast of the ending stocks S_t . The bias of the forecasts can be tested by fitting the Mincer and Zarnowitz (1969) regression:

$$S_t = a + bU_{t,n} + error_{t,n} \quad (2.1)$$

Under the null hypothesis $H_0: (a, b) = (0, 1)$, USDA forecasts of ending stocks are unbiased.

A preferred specification is obtained by imposing $\beta = 1$ in regression (2.1), which yields

$$S_t - U_{t,n} = a + error_{t,n} \quad (2.2)$$

The difference $S_t - U_{t,n}$ represents the error of the n -month-ahead USDA forecast. Regression (2.2) is widely used because it is more intuitive and does not require the forecast $U_{t,n}$ to be uncorrelated with the residual in regression (2.1).

The forecasts of ending stocks are fixed-event forecasts. Thus, the regression errors in (2.2) are inherently correlated with each other (*i.e.* $error_{t,n}$ and $error_{t,m>n}$ are correlated because they overlap over period n). If the covariance structure of the regression errors does not incorporate this correlation, estimates of the variances of the parameters could be biased, thereby undermining the credibility of inferences based on the significance tests.

The research based on Nordhaus (1987) suggests focusing on forecast revisions instead. Nordhaus points out that it is difficult to test strong efficiency in the form of rational expectations, because it is impossible to incorporate into the test all of the information available at the time the forecasts are issued. However, the information on past forecasts is always available to the public. Nordhaus thus introduced a weak efficiency test which is solely based on past forecasts. A forecast is said to be weakly efficient if both the current forecast error and

forecast revision are independent of all past forecast revisions. According to this, the following test can be performed to test for weak efficiency:

$$U_{t,n-1} - U_{t,n} = \gamma_t(U_{t,n} - U_{t,n+1}) + \zeta_{t,n}, \text{ for each } t \quad (2.3)$$

where $\zeta_{t,n}$ follows a normal distribution with fixed variance. For each year, the test is performed by pooling over all forecast revisions in that year. A γ_t significantly different from zero means rejection of weak efficiency, implying that forecasts can be improved upon by using information from past forecasts.

Isengildina *et al.* (2006) modified the Nordhaus model by pooling over all forecast revisions for a certain month instead of a certain year to reduce the number of regressions to be estimated. Isengildina *et al.* (2013) further extend it by including a bias term and additional public information which results in the following model:

$$U_{t,n-1} - U_{tn} = a_n + \gamma(U_{tn} - U_{t,n+1}) + b_n C_{t,n} + \zeta_{t,n}, \text{ for each } n \quad (2.4)$$

where a_n is the forecast revision bias, and $C_{t,n}$ is an explanatory variable representing publicly available information.

Unlike the Nordhaus model, Davies and Lahiri (1995, 1999) focus on forecast errors directly. They suggest that there are two types of errors for fixed-event forecasts. The first type consists of unforecastable shocks within the forecasting cycle. These shocks typically arise from elements which cannot be controlled by the forecaster, such as changes in economic structure, market conditions, or deviations of benchmark assumptions.¹ The second type is the forecaster's idiosyncratic errors, which stem from the forecaster's subjective views and/or his private model.

¹ For agricultural forecasts, one example is related to weather. The realizations of weather often deviate from the assumptions used in models in a forecasting cycle. The impact on forecasts can thus be explained as shocks which cannot be controlled by forecasters.

In the present notation, the Davies and Lahiri decomposition of the error term for a single forecaster² can be decomposed as

$$error_{t,n} = \lambda_{t,n} + \varepsilon_{t,n} \quad (2.5)$$

where $\lambda_{t,n}$ represents the unforecastable shock for forecast horizon n and marketing year t , and $\varepsilon_{t,n} \sim i.i.d. N(0, \sigma^2)$ is the idiosyncratic error. The shock term $\lambda_{t,n}$ can be further decomposed as the sum of *i.i.d.* monthly shocks:

$$\lambda_{t,n} = \sum_{j=0}^{n-1} k_{t,j} \quad (2.6)$$

where $k_{t,j} \sim i.i.d. N(0, \sigma_j^2)$ can be interpreted as monthly shocks. The idea underlying the decomposition of (2.6) is that a forecast made at a date closer to the target event tends to be more precise; hence it should have a smaller forecast error variance. Note that the structure in (2.6) implies that the unforecastable shocks $\lambda_{t,n}$ are correlated within each marketing year t . This is because of the overlaps of forecast horizons. Therefore, the bias test developed from the Davies and Lahiri framework can be written as

$$S_t - U_{t,n} = a_n + \sum_{j=1}^n k_{t,j} + \varepsilon_{t,n} \quad (2.7)$$

where a_n represents the forecast bias.

Based on the Davies and Lahiri framework, Clements *et al.* (2007) proposed analyzing forecast revisions by differencing the forecast errors. In the present notation, their test can be written as

$$U_{t,n-1} - U_{t,n} = a_n - a_{n-1} + \omega_{t,n} \quad (2.8)$$

² The Davies and Lahiri model is developed for a three-dimensional analysis of panel data. The notation representing the forecasters is omitted because we only consider a single forecaster – the USDA.

where $\omega_{t,n} \equiv k_{t,n} + \varepsilon_{t,n} - \varepsilon_{t,n-1}$, and $k_{t,n}$ is assumed to be homoscedastic. This model provides a thought to investigate the forecast revisions in the Davies and Lahiri framework. However, Clements *et al.* (2007) did not consider adjacent forecast revisions, which represent the impacts of most updated information. Besides, they simplified the estimation of the covariance matrix. Their estimation did not account for the negative correlations generated by the $\varepsilon_{t,n}$'s. They also proposed a simplification by restricting $\varepsilon_{t,n}$'s to be zero. Thus the error covariance matrix is not fully estimated and the estimated parameters do not reveal the characteristics of the forecasts. Thus, it is necessary to develop a method to obtain a full estimation of the matrix.

2.3.2 Proposed Model

Based on Clements *et al.* (2007), we develop an efficiency test which combines the characteristics of the Nordhaus (1987) and Davies and Lahiri (1995, 1999) models. We first decompose the forecast bias a_n as the sum of monthly biases so that individual bias is allowed as in Clements *et al.* (2007):

$$a_n = \sum_{j=1}^n \alpha_j \quad (2.9)$$

where α_j represents the bias of the n -month-ahead forecast revision. In this way, regression (2.7) becomes:

$$S_t - U_{tn} = \sum_{j=1}^n \alpha_j + \sum_{j=1}^n k_{tj} + \varepsilon_{tn} \quad (2.10)$$

Given the structure of regression (2.10), we apply first differencing as in Clements *et al.* (2007) to obtain a system of equations, where only forecast revisions of consecutive months appear:

$$\begin{cases} S_t - U_{t,1} &= \alpha_1 + k_{t,1} + \varepsilon_{t,1} \\ U_{t,1} - U_{t,2} &= \alpha_2 + k_{t,2} - \varepsilon_{t,1} + \varepsilon_{t,2} \\ &\vdots \\ U_{t,N-1} - U_{t,N} &= \alpha_N + k_{t,N} - \varepsilon_{t,N-1} + \varepsilon_{t,N} \end{cases} \quad (2.11)$$

In (2.11), N is the maximum forecast horizon for a marketing year, $\varepsilon_{t,n} \sim i.i.d. N(0, \sigma^2)$ are idiosyncratic errors, and $k_{t,j} \sim i.i.d. N(0, \sigma_j^2)$ are monthly shocks with different variances for different forecast horizons. The assumption of heteroscedastic shocks is reasonable for fixed-event forecasts, especially for ending stocks forecasts whose forecast horizon is long. For example, larger variances can be expected in early revisions of ending stocks forecasts because of the uncertainty from productions.³ In addition, seasonality in consumption, trade, and production patterns for many crops means that the arrival of new information is most likely to vary from month to month, again making it more realistic to assume heteroscedasticity in shocks.

Autocorrelations also exist in the residuals of system (2.11). They stem from the idiosyncratic errors and can be interpreted as forecasters' corrections of their own errors. For example, suppose forecasters make no idiosyncratic errors in a particular year τ except for misinterpreting a piece of information when issuing their n th forecast, causing it to unduly underestimate the ending stocks. That is, $\varepsilon_{\tau,m \neq n} = 0$ and $\varepsilon_{\tau,n} > 0$. Then the revision for the n th horizon ($U_{t,n} - U_{t,n+1}$) will be smaller (by $-\varepsilon_{\tau,n}$) than it should be, and it will be followed by an $(n - 1)$ th revision greater (by $\varepsilon_{\tau,n}$) than it would have been otherwise. The Nordhaus model does not build in this feature due to their strong *i.i.d.* assumptions on the idiosyncratic errors.

The efficiency test based on (2.11) consists of fitting

³ For example, shocks to production output, such as weather conditions, can be substantial.

$$\begin{cases} S_t - U_{t,1} &= \alpha_1 + \beta_1 X_{t,1} + k_{t,1} + \varepsilon_{t,1} \\ U_{t,1} - U_{t,2} &= \alpha_2 + \beta_2 X_{t,2} + k_{t,2} - \varepsilon_{t,1} + \varepsilon_{t,2} \\ &\vdots \\ U_{t,N-1} - U_{t,N} &= \alpha_N + \beta_N X_{t,N} + k_{t,N} - \varepsilon_{t,N-1} + \varepsilon_{t,N} \end{cases} \quad (2.12)$$

where $X_{t,n}$ can be a variable or a vector, representing one or more explanatory variables known at the time the forecast is made. The null hypothesis $H_0: \alpha_n = \beta_n = 0$ for all n indicates that $U_{t,n}$ is an efficient forecast of S_t , in the sense that forecast errors cannot be predicted using available information. As pointed out in the literature (*e.g.*, Nordhaus 1987), it is impossible to include all past information in $X_{t,n}$. Therefore, we follow Nordhaus (1987) and construct a test for weak efficiency by letting the past forecast revision be the explanatory variable, *i.e.*, $X_{t,n} = U_{t,n} - U_{t,n+1}$.

Our model builds upon the Nordhaus test by further identifying the unforecastable shocks and the forecaster's own errors. Given the proposed covariance structure, system (2.12) can no longer be estimated by OLS, because the explanatory variable $U_{t,n} - U_{t,n+1}$ is negatively correlated with the idiosyncratic error $\varepsilon_{t,n}$.⁴ Thus, estimate the equations as a system by treating the forecast revisions within the same marketing year as a panel.

Due to the limited size of the data about ending stocks forecasts, estimating $\alpha_1, \dots, \alpha_N$ and β_1, \dots, β_N separately results in too many parameter estimates relative to the number of observations. Hence, we impose the restrictions $\alpha_1 = \dots = \alpha_N = \alpha$ and $\beta_1 = \dots = \beta_N = \beta$, which lead to the following system of equations used for estimation purposes:

$$\begin{cases} S_t - U_{t,1} &= \alpha + \beta(U_{t,1} - U_{t,2}) + k_{t,1} + \varepsilon_{t,1} \\ U_{t,1} - U_{t,2} &= \alpha + \beta(U_{t,2} - U_{t,3}) + k_{t,2} - \varepsilon_{t,1} + \varepsilon_{t,2} \\ &\vdots \\ U_{t,N-2} - U_{t,N-1} &= \alpha + \beta(U_{t,N-1} - U_{t,N}) + k_{t,N-1} - \varepsilon_{t,N-2} + \varepsilon_{t,N-1} \end{cases} \quad (2.13)$$

⁴ This negative correlation leads to the OLS estimates of the slope coefficients (β) biased toward zero, as the problem is analogous to the well-known "attenuation" caused by measurement errors in the explanatory variables.

The last equation in (2.12) is discarded because the first forecast of each marketing year can only be used to calculate the explanatory variable. Therefore system (2.13) comprises only $N - 1$ equations. The system of equations offers an overall assessment of bias and efficiency of USDA forecast performance. Succinctly, letting $U_{t,0} = S_t$ and $\varepsilon_{t,0} = 0$, system (2.13) can be rewritten as

$$U_{t,n-1} - U_{t,n} = \alpha + \beta(U_{t,n} - U_{t,n+1}) + k_{t,n-1} - \varepsilon_{t,n-1} + \varepsilon_{t,n} \quad (2.14)$$

for $n = 1, \dots, N - 1$.

System (2.14) can be viewed as an improvement on both streams of the literature discussed earlier. It reduces the restrictions on the error covariance matrix by allowing for both regression error heteroscedasticity and autocorrelations. At the same time, it introduces an error covariance structure to estimate a minimal number of covariance parameters.

To see the covariance structure, consider system (2.14) for a single marketing year. The data is sorted by forecast horizon $n = 1, \dots, N$. If we assume that both $\varepsilon_{t,j}$'s and $k_{t,j}$'s are *i. i. d.*, the covariance matrix is

$$B = \begin{bmatrix} \sigma_1^2 + \sigma^2 & -\sigma^2 & 0 & \dots & 0 \\ -\sigma^2 & \sigma_2^2 + 2\sigma^2 & -\sigma^2 & & \vdots \\ 0 & -\sigma^2 & \ddots & -\sigma^2 & 0 \\ \vdots & & -\sigma^2 & \sigma_{N-2}^2 + 2\sigma^2 & -\sigma^2 \\ 0 & \dots & 0 & -\sigma^2 & \sigma_{N-1}^2 + 2\sigma^2 \end{bmatrix}_{(N-1) \times (N-1)} \quad (2.15)$$

It can be seen that the difference in total error variances are contributed by differences in variances of monthly shocks. Also there are correlations between adjacent regression equations.

Regarding to ending stocks, the maximum forecast horizon is 17 months for corn and soybeans, 14 months for wheat. As the maximum horizon is larger than 12 month, there are instances where ending stocks forecasts for two consecutive crop years are issued

simultaneously. Since shocks for consecutive crop years are likely to be correlated (*e.g.*, a negative demand shock will result in higher ending stocks for both the current and the following marketing year), we also estimate system (2.14) using the alternative covariance matrix (2.16) instead of (2.15):

$$\bar{B}_{(N-1) \times (N-1)} = \begin{bmatrix} \sigma_1^2 + \sigma^2 & -\sigma^2 & 0 & \dots & \sigma_1^2 & 0 & 0 & 0 \\ -\sigma^2 & \sigma_2^2 + 2\sigma^2 & -\sigma^2 & & & \sigma_2^2 & 0 & 0 \\ 0 & -\sigma^2 & \ddots & & & & \sigma_3^2 & 0 \\ \vdots & & & \ddots & & & & \sigma_4^2 \\ \sigma_1^2 & & & & \bar{\sigma}_{N-4}^2 + 2\sigma^2 & & & \vdots \\ 0 & \sigma_2^2 & & & & \bar{\sigma}_{N-3}^2 + 2\sigma^2 & -\sigma^2 & 0 \\ 0 & 0 & \sigma_3^2 & & & -\sigma^2 & \bar{\sigma}_{N-2}^2 + 2\sigma^2 & -\sigma^2 \\ 0 & 0 & 0 & \sigma_4^2 & \dots & 0 & -\sigma^2 & \bar{\sigma}_{N-1}^2 + \sigma^2 \end{bmatrix} \quad (2.16)$$

where $\bar{\sigma}_{n+12}^2 = \sigma_{n+12}^2 + \sigma_n^2$ for $1 < n < N - 12$. The assumption underlying covariance matrix (2.16) is that next year's shock $\bar{k}_{t+1,n+12}$ consists of this year's shock $k_{t,n}$ plus orthogonal shock $k_{t+1,n+12}$.⁵ That is,

$$\bar{k}_{t+1,n+12} = k_{t,n} + k_{t+1,n+12} \quad (2.17)$$

In this way, $Cov(\bar{k}_{t+1,n+12}, k_{t,n}) = \sigma_n^2$.

The structure of the data is characterized by two dimensions, namely, forecast horizon n and marketing year t . While the order of the equations does not affect the estimation, it is interesting to see the structure of the full covariance matrix. Without loss of generality, we can sort the data first by forecast horizon and then by marketing year.⁶ Thus, in the case of covariance matrix (2.16), the full $T(N - 1) \times T(N - 1)$ covariance matrix of the error term can be expressed as

⁵ Note that shocks with horizon greater than 12 are always for the next marketing year.

⁶ Instead, we can sort the data first by marketing year then by forecast horizon. The error covariance matrix will be different, but the estimation results won't change.

$$\bar{\Sigma} = \begin{bmatrix} \bar{B} & 0 & \dots & 0 \\ 0 & \bar{B} & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & \bar{B} \end{bmatrix}_{T \times T} \quad (2.18)$$

2.4 Data

The model is applied to data on U.S. ending stocks and their corresponding USDA monthly forecasts for three major agricultural commodities - corn, soybeans and wheat. U.S. ending stocks are obtained from the Grain Stocks Report released by the National Agricultural Statistics Service (NASS). The report is issued quarterly, typically in early January, and at the end of March, June and September. Specifically, ending stocks data for corn and soybeans are retrieved from the September report (the first report after the ending of the U.S. marketing year for these two commodities), whereas ending stocks data for wheat are retrieved from the June report.⁷ For each commodity, ending stocks from the marketing years 1984/85 through 2013/14, a total of 29 marketing years, are used to fit the model.

The USDA monthly forecasts are retrieved from the WASDE reports. The U.S. marketing year for corn and soybean starts in September and ends in August of the following calendar year. For each marketing year, the first USDA forecast for corn and soybeans is released in May, before the marketing year begins. The last forecast is released in September, after the marketing year ends and before the release of the ending stock of that marketing year. The U.S. marketing year for wheat is different for corn and soybeans, as it starts in June and ends in May of the following calendar year. However, the first USDA forecast for wheat ending stocks is also released in May (together with the first forecast for corn and soybeans), and the last

⁷ In rare cases there have been revisions of the ending stocks in the Grain Stocks Report, but they have been typically quite small. In this situation, we use the finalized ending stocks in later reports.

forecast is released in June of the following calendar year. Thus, for each marketing year there are 17 forecasts for corn and soybeans, and 14 forecasts for wheat.

Since ending stock values are strictly positive and with a distribution skewed to the right, all the data are transformed into logarithms. In this way, forecast revisions represent percentage changes instead of changes in levels. As discussed in the previous section, the structure of the model prevents us from using the entire forecast dataset to compute the dependent variable, as the first forecast of each marketing year can only be used to calculate the explanatory variable which represents the forecast revision of previous month. Thus the dependent variables consist of forecast revisions with maximum forecast horizons of 16 for corn and soybean, and 13 for wheat. In summary, the corn and soybean regressions comprise 464 ($= 29 \text{ years} \times 16 \text{ forecasts}$) observations, whereas the regressions for wheat have 377 ($= 29 \text{ years} \times 13 \text{ forecasts}$) observations.

Table 2.1 shows the descriptive statistics for the USDA forecast revisions for all three commodities.⁸ The means of forecast revisions for corn and soybeans are slightly negative, at -0.81% and -1.59% respectively. The mean for wheat is slightly positive, at 0.13%. The medians for corn and wheat are zero, whereas for soybeans it is slightly negative (-0.41%). For the standard deviations, they are much larger for corn and soybeans (for which it exceeds 10%) than for wheat (6.45%). The range of revisions is largest for soybeans, from -48.84% to 75.69%. Revisions for corn range from -64.78% to 40.98%. The range for wheat is smallest, from -22.23% to 23.02%.

⁸ The summary statistics contain all forecast revisions which enter as the dependent variables, hence do not contain the earliest forecast revisions, which are only used as explanatory variables.

Figure 2.1 depicts the monthly standard deviations of these forecast revisions, displayed in order of diminishing forecast horizons. For all three commodities, the standard deviations exhibit a decreasing trend as the forecast horizons shorten. The revisions for early months tend to be larger than revisions for other months. Also, the magnitude of the final revision is typically larger than the ones of the preceding revisions. The largest final revisions occur for soybeans, and the smallest ones are observed for wheat. In addition, monthly standard deviations are generally greater for corn and soybeans than for wheat. It is clear from Figure 2.1 that the standard deviations of forecast revisions vary by forecast horizon, which implies that it is important to build error heteroscedasticity into the model.

2.5 Estimation Methods

The proposed model is estimated using Bayesian Markov Chain Monte Carlo (MCMC) methods. The method greatly facilitates dealing with both heteroscedasticity and autocorrelation. Another advantage of the Bayesian approach is that it yields full posterior distributions for the parameters of interest. It is particularly useful when researchers try to characterize the property of parameters with a skewed posterior, such as error variances. This section outlines the joint posterior distributions of the parameters of the model, the choice of priors for the parameters, and the steps in the MCMC iterations.

To simplify the notations, the proposed regression system (2.14) is rewritten as:

$$y_{t,n} = \mathbf{x}_{t,n}\boldsymbol{\beta} + k_{t,n} - \varepsilon_{t,n-1} + \varepsilon_{t,n} \quad (2.19)$$

where $y_{t,n} \equiv U_{t,n-1} - U_{t,n}$, $\mathbf{x}_{t,n} \equiv [1 \ U_{t,n} - U_{t,n+1}]$, $\boldsymbol{\beta} \equiv [\alpha \ \beta]'$. The matrix form of each panel is

$$\mathbf{y}_t = \mathbf{x}_t \boldsymbol{\beta} + \mathbf{w} \mathbf{k}_t + \mathbf{p} \boldsymbol{\varepsilon}_t \quad (2.20)$$

where $\mathbf{y}_t = [y_{t,1}, \dots, y_{t,N-1}]'$, $\mathbf{x}_t \equiv [\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,N-1}]'$, $\mathbf{k}_t = [k_{t,1}, \dots, k_{t,N-1}]'$,⁹ and $\boldsymbol{\varepsilon}_t = [\varepsilon_{t,1}, \dots, \varepsilon_{t,N-1}]'$. \mathbf{w} is a matrix indicating the existence of elements in \mathbf{k}_t in each equation.¹⁰ And \mathbf{p} is a matrix indicating the existence of elements in $\boldsymbol{\varepsilon}_t$ in each equation. Combining the panels for all t 's, the full system can then be written as

$$\mathbf{Y} = \mathbf{X} \boldsymbol{\beta} + \mathbf{W} \mathbf{K} + \mathbf{P} \mathbf{E} \quad (2.21)$$

Since there is a common intercept α , for identification purposes one of the shock terms must be restricted to zero. Hence, without loss of generality, we set $k_{T,1}$ to be zero, *i.e.*, we assume that the final one-month forecast error only contains the forecaster's own idiosyncratic error.

Let $\Lambda = \{\boldsymbol{\beta}, \{\sigma_n^2\}_{n=1}^{N-1}, \sigma^2\}$ denote the set of parameters of the proposed model. The joint posterior density of Λ is the following:

$$p(\Lambda) = \Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \sigma^2 \boldsymbol{\Omega}) \prod_{jt} \Phi(k_{t,j} | \sigma_j^2) * \Phi(\boldsymbol{\beta} | \mathbf{M}, \mathbf{V}) p(\sigma^2) \prod_{j=1}^{N-1} p(\sigma_j^2) \quad (2.22)$$

where $\Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \sigma^2 \boldsymbol{\Omega})$ is the distribution of \mathbf{Y} , which is multivariate normal. $\sigma^2 \boldsymbol{\Omega}$ is the covariance matrix generated by the idiosyncratic errors. $\Phi(\boldsymbol{\beta} | \mathbf{M}, \mathbf{V})$ is the prior distribution of $\boldsymbol{\beta}$, which is multivariate normal with mean \mathbf{M} and variance \mathbf{V} .

In order to obtain the likelihood function and the full posterior, we derive a Gibbs Sampler based on generalized least squares (GLS). We cannot directly obtain the distributions for each $y_{t,n}$ because of the autocorrelations in the error terms. Thus the likelihood can only be

⁹ The use of *i.i.d.* shock components instead of actual shocks can greatly simplify the estimation process. Note that the actual shocks can always be reconstructed from the *i.i.d.* shock components using formula (2.17).

¹⁰ Note that \mathbf{w} is not an identity matrix. Based on formula (2.17), there are two 1's in the rows representing forecasts with horizon greater than 12.

expressed as a multivariate normal with covariance $\sigma^2 \mathbf{\Omega}$. To calculate the conditional posterior density for σ^2 , first note that the relationship between $\sigma^2 \mathbf{\Omega}$ and the idiosyncratic error vector \mathbf{E} is

$$\sigma^2 \mathbf{\Omega} = E[\mathbf{PE}(\mathbf{PE})'] = E[\mathbf{PEE}'\mathbf{P}'] = \mathbf{PE}[\mathbf{EE}']\mathbf{P}' = \sigma^2 \mathbf{PP}' \quad (2.23)$$

Also, the location of the idiosyncratic error in each observation is uniquely determined by the indicator matrix \mathbf{P} , which is a known matrix. Thus the conditional posterior density for σ^2 can be calculated based on premultiplying system (2.21) by \mathbf{P}^{-1} :

$$\mathbf{P}^{-1}\mathbf{Y} = \mathbf{P}^{-1}\mathbf{X}\boldsymbol{\beta} + \mathbf{P}^{-1}\mathbf{W}\mathbf{K} + \mathbf{E} \quad (2.24)$$

Or

$$\tilde{\mathbf{Y}} = \tilde{\mathbf{X}}\boldsymbol{\beta} + \tilde{\mathbf{W}}\mathbf{K} + \mathbf{E} \quad (2.25)$$

where $\tilde{\mathbf{Y}} = \mathbf{P}^{-1}\mathbf{Y}$, $\tilde{\mathbf{X}} = \mathbf{P}^{-1}\mathbf{X}$, $\tilde{\mathbf{W}} = \mathbf{P}^{-1}\mathbf{W}$.

Conditionally conjugate priors are adopted in the present analysis. In particular, the priors chosen for the parameters are:

$$\begin{aligned} \boldsymbol{\beta} &\sim N(\mathbf{M}, \mathbf{V}) \\ \varepsilon_{t,n} &\sim N(0, \sigma^2) \\ k_{t,n} &\sim N(0, \sigma_n^2) \\ \sigma, \sigma_n &\sim \text{Uniform}(0, \infty) \end{aligned} \quad (2.26)$$

for $n = 1, \dots, N - 1$ and $t = 1, \dots, T$. The prior distribution for the coefficient vector $\boldsymbol{\beta}$ is multivariate normal with mean $\mathbf{M} = [0 \ 0]'$ and covariance matrix $\mathbf{V} = 1000\mathbf{I}_{2 \times 2}$, where $\mathbf{I}_{2 \times 2}$ is a 2×2 identity matrix. The prior mean of $\boldsymbol{\beta}$ is chosen to be consistent with the null hypothesis of efficiency. The scale of the variance matrix \mathbf{V} is chosen to be large, so that the prior is non-informative. In this way, the draws of $\boldsymbol{\beta}$ are diffuse and widely spread around the mean zero. The uniform prior for the standard deviation parameters is chosen following Gelman (2006). This

prior is non-informative and can be viewed as a limit of the half- t family distributions, which is conditionally conjugate to the extent of more general folded-noncentral- t distributions.¹¹ The error covariance matrix we proposed is flexible in the sense that we let the estimation process determine the value of the correlations. If these correlations do exist, the estimated σ^2 will be of comparable value of the variances of $k_{t,n}$. Otherwise, if the forecasters' own errors are tiny compared to the outside shocks, the estimation process will push the error variance estimate to the variance of $k_{t,n}$, making σ^2 close to zero. The conditional posterior distributions for the parameters of the model are outlined in the Appendix.

The MCMC iteration steps for the model can be summarized as follows:

Step 1: Set up initial values for each parameter in the set Λ , as well as $\mathbf{K}^{(0)}$ and $\mathbf{E}^{(0)}$.

Step 2: Given $\{k_{t,n}^{(i)}, \sigma^{2(i)}\}$, draw $\boldsymbol{\beta}^{(i+1)}$ from a multivariate normal distribution.

Step 3: Given $\{\boldsymbol{\beta}^{(i+1)}, \sigma^{2(i)}, \{\sigma_n^2\}^{(i)}, \mathbf{K}_{-k_{t,n}}^{(i)}\}$, sequentially draw $k_{t,n}^{(i+1)}$ from a normal distribution for each $t = 1, \dots, T$ and $n = 1, \dots, N - 1$.

Step 4: Given $\{\boldsymbol{\beta}^{(i+1)}, \mathbf{K}^{(i+1)}\}$, update $\mathbf{E}^{(i+1)}$, and draw $\sigma^{2(i+1)}$ from an inverse gamma distribution.

Step 5: Given $\mathbf{K}^{(i+1)}$, sequentially draw $\sigma_n^{2(i+1)}$ from an inverse gamma distribution for each $n = 1, \dots, N - 1$.

Step 6: Set $i = i + 1$.

¹¹ The uniform prior serves a better role than the weakly-informative inverse gamma (ϵ, ϵ) prior distribution for the variance, where ϵ is a positive value close to zero. When using a gamma (ϵ, ϵ) instead of the uniform prior to estimate the model, inferences are found to be very sensitive to the choice of ϵ , because the value of estimated standard deviation parameters are quite small. In this way, the results from applying inverse gamma (ϵ, ϵ) priors are less likely to be non-informative in the present study if ϵ is not small enough relative to the estimated variances. Importantly, similar results are found when employing inverse gamma (ϵ, ϵ) priors with $\epsilon = 0.01$ and 0.001 . Moreover, there are no significant changes in the results when other forms of non-informative priors (e.g., $\sigma, \sigma_n \sim U(0, 1000)$, or $\sigma^2, \sigma_n^2 \sim U(0, \infty)$) are applied.

Step 7: Repeat Step 2 until the maximum iteration is reached.

For each commodity, the Gibbs Sampler is run for three Markov Chains for 100,000 iterations each. The first half of each chain is discarded as a burn-in period. Gelman and Rubin (1992) test is then applied to check the convergence of the remaining part of the chains. The Gelman and Rubin test statistic compares the variances of both within the chains and between the chains. Values of the statistics close to 1 indicate convergence.

2.6 Results and Discussion

Estimation results for the USDA forecasts for the marketing years 1984/85 through 2013/14 are summarized in Table 2.2 and 2.3.¹² Gelman and Rubin (1992) test statistics are below 1.03 for all parameters for all three commodities, which strongly suggests convergence of the Markov Chains. Table 2.2 displays the means and standard deviations for the estimated parameters, including the intercept, slope, and the standard deviations of the idiosyncratic error and unforecastable shocks. Panel B reports the medians and 95% credible intervals of the corresponding parameters. The sequence of standard errors for the unforecastable shocks is displayed in the order of increasing forecast horizons. Specifically, σ_n is the standard deviation of the shock corresponding to the n -month forecast horizon. For example, for corn and soybeans, σ_{16} is the standard error of shocks with a 16-month horizon, *i.e.*, between June and July for forecasts targeting the ending stocks of the following marketing year. σ_1 is the standard deviation of the shock corresponding to the final forecast revision, *i.e.*, the final forecast error. For wheat, because there are only 13 shocks, σ_{13} is the standard error of the earliest shocks, which is also

¹² Estimation results for the constant and slope remain basically unchanged if (2.15) is used as the component in the error covariance matrix instead. The estimated variances of shocks, however, are different due to the correlation assumption on simultaneous forecasts.

between June and July for forecasts targeting the ending stocks of the following marketing year. It is worth noting that the estimated parameters can be compared among all three commodities, because forecast revisions are all measured in natural logarithms (*i.e.*, percentage values).

The point estimate of the intercept α represents the bias of the USDA forecast revisions. For corn, it is positive and significant at the 5% level, suggesting that USDA forecast revisions have a tendency to be upwardly biased. The point estimate of $\alpha=0.5\%$ indicates that the USDA revises its forecasts up by about 0.5% each month.¹³ When adding up the forecast revisions, it shows that the USDA has a tendency to underestimate the ending stocks. For example, although the estimate of α is small, it can be inferred that $U_{t,16}$, the forecast made in June for the ending stock of the following marketing year tends to underestimate the ending stock by an average of 8% ($= 0.5\% \times 16$), which is no longer negligible. In other words, the positive and significant estimate of α shows that USDA forecasts of corn ending stocks are inefficient.

The point estimate of α for soybeans is -1.37%, the absolute value of which is more than twice as large as that for corn. The estimate is negative and significant, showing that on average the USDA adjusts its forecast down for about 1.37% each month. The finding also indicates that USDA has a tendency to overestimate the ending stocks. Specifically, it can be inferred that $U_{t,16}$, the forecast made in June for the ending stock of the following marketing year, tends to overestimate the ending stock by an average of 21.92% ($=1.37\% \times 16$). It contributes to the inefficiency of the USDA soybean forecasts, but in an opposite way from forecasts of corn ending stocks. For wheat, the estimate is -0.11%. However, both the 90% and 95% credible intervals contain zero. Therefore we cannot reject the null hypothesis that USDA wheat forecasts revisions are conditionally unbiased.

¹³ Recall that forecast revisions are defined by $U_{t,n-1} - U_{t,n}$.

It is interesting to find opposite signs for the estimates of corn and soybean. One possible explanation of this finding is that corn and soybean are close substitutes in production. The increase in the acreage of planting one crop will typically result in decrease in the acreage of the other, as the total land use for these two crops are stable over a short period. Therefore, if the demand conditions are unchanged, the ending stocks will be affected in a similar way as productions. The USDA forecasts reflect this relationship. Thus, a positive bias in the USDA forecast revisions for corn is accompanied by a negative bias in its soybean forecasts.

Coefficient β measures the association between two adjacent USDA forecast revisions, accounting for the endogeneity of past revisions. Note that it is different from the coefficient in Nordhaus (1987) which assumes *i. i. d.* residuals. For all three commodities, the point estimates of β are positive and significant at the 5% level, indicating inefficient USDA forecasts. The estimate for corn is 18.37% on average, meaning that if USDA adjusts its forecast up by 1% in the past month, its forecast will also be revised up by roughly 0.18% in the current month, given other conditions fixed. The estimate for wheat is 18.84%, which is similar in size to that for corn. The estimate for soybeans is 51.37%, which is largest and almost triple the estimates for corn and wheat.

Our findings show that the USDA is conservative in adjusting its ending stocks forecasts. In other words, the most recent USDA forecast does not fully represent the arrival of new information. It can also be argued that the USDA is more conservative to forecast ending stocks for soybean compared to corn and wheat, because the estimated association coefficient β for soybean is the largest. The results are consistent with previous research that government agency have a tendency to smooth their forecasts. (*e.g.* Isengildina *et al*, 2006).

An interesting question to ask is: why is the USDA conservative in making ending stocks forecasts? Isengildina *et al.* (2006) summarized several reasons that could explain why government agencies smooth their forecasts. These reasons include predicting based on weather conditions, forming a weighted average forecast based on earlier forecasts and current estimates, strategic behavior, *etc.* They claimed that for USDA crop production forecasts, the smoothing comes from conservation of farm operators' assessment and bias in using information. Ending stocks forecasts, however, are quite different from production forecasts. While production forecasts are generated based on surveys and satellite images, ending stocks forecasts are combinations of various forecasts of both demand and supply, which are inheritably more subjective. Vogel and Bange (1999) stated, "Throughout the growing season and afterwards, estimates are compared with new information on production and utilization, and historical revisions are made as necessary." Therefore, the USDA may include past forecasts in forming new ending stocks forecasts, as the ending stocks forecasts require much more subjective analysis on the demand side of the balance sheet. In this way, a new forecast can possibly be a weighted average of earlier forecasts and current estimates. This method of averaging is widely used to generate forecasts in industry.

Table 2.2 also displays the standard deviations of the monthly shocks. It can be seen that they tend to decrease as the forecast horizon shortens. The shocks are typically large for the first seven months of a forecasting cycle. This is because the production plans haven't been fully revealed, adding another layer of uncertainty to the ending stocks forecasts. Later forecast revisions are mainly attributed to the demand side only, and hence shocks are typically smaller. For corn, the standard deviations of monthly shocks range from 1.56% to 24.58%. Large shocks are expected to arrive in revisions in pre-marketing year June/July (18.63%), pre-marketing year

July/August (24.58%), September/October (14.59%), and December/January (13.38%). Note that the December/January revision corresponds to the release of the final production forecast of that marketing year, and signals the end of the role of domestic production in USDA ending stocks forecasts. Shocks in February/March (1.71%) and May/June (1.56%) are the smallest and lower than the standard deviations of USDA idiosyncratic errors (1.84%).

The standard deviations of monthly shocks for soybeans range from 1.93% to 19.87%. Compared to corn, the range is smaller and the lower bound is higher. Large shocks are expected in revisions in all of the first five months of the forecasting cycle from June to November. The September/October shock is largest at 19.87%, while the other four early shocks are at around 13%. Note that September/October is also the time of the first revision of soybean production forecasts. Contrary to the case of corn, the December/January shock (4.58%) is not large compared to other early shocks. Interestingly, there are seven shocks which have standard deviations smaller than that of the USDA idiosyncratic errors (4.65%).

For wheat, the standard deviations of shocks have a much narrower range, which is 3.17% to 9.84%. Large shocks are expected in revisions occurring in June/July (9.11%), July/August (9.84%), and September/October (8.59%). The standard deviations of shocks are all larger than those of the USDA idiosyncratic errors.

There is a jump in the standard deviations of final forecast revisions as it measures the difference between analysts' final forecasts and the ending stocks, which is different in nature with other forecast revisions. The estimate is largest for soybeans (16.52%) and smallest for wheat (7.08%). For corn, the estimate is 8.65%. One explanation of this finding is that there might be unexpected large demand changes during the final month of the marketing year. Another reason is that the models used by the USDA may not take into account some important

information which can last for as long as the full forecasting cycle. Hence, when the final stocks are released, the information which is ignored will come out in a sudden, which is captured as shocks in the proposed model, resulting in typically large final revisions.

The estimates of σ represent the standard deviations of USDA idiosyncratic errors. The point estimates are significantly greater than zero for all three commodities. The estimate for corn is 1.84% on average. The estimate for wheat is 1.31%, which is smallest. The estimate for soybeans is quite larger, at 4.65%. The estimated standard deviations of the idiosyncratic errors are larger than some estimates of the standard deviations of monthly shocks for corn and soybeans. The results validate our assumption on the existence of idiosyncratic errors, and they are not negligible compared to the unforecastable shocks.

2.7 Conclusions

We develop a framework to investigate the efficiency of USDA crop ending stocks forecasts based on the works of Clements *et al.* (2007), Nordhaus (1987) and Davies and Lahiri (1995, 1999). The proposed model analyzes adjacent forecast revisions with emphasis on the link between forecasts and the forecast target. The residuals are decomposed as the sum of monthly unforecastable shocks and USDA's own idiosyncratic errors. The postulated error covariance matrix then exhibits heteroscedasticity (due to the unforecastable shocks), as well as autocorrelation (due to the idiosyncratic errors).

We apply our estimation framework to USDA ending stocks forecasts for three major agricultural commodities - corn, soybeans and wheat. A total of 29 marketing years, from 1985/86 to 2013/14 are investigated. Estimation is conducted by means of a Bayesian Markov Chain Monte Carlo (MCMC) approach. This method allows us to estimate the coefficients and

the error covariance matrix in the same iteration. The MCMC method also allows the parameters to vary freely, so that the estimation results can be used to validate the postulated structure of the residual covariance matrix.

Results show that USDA forecasts are inefficient for all three commodities. Forecast revisions for corn and soybean are biased: the USDA has a tendency to underestimate the ending stocks for corn and overestimate the ending stocks for soybeans. We cannot reject the null hypothesis that USDA forecasts for wheat are unbiased. The slope coefficients for three commodities are all positive and significant, providing strong evidence against efficiency. The significantly positive slope estimates suggest that the USDA is conservative in adjusting its forecasts, and it might put a positive weight on its past forecasts. We also find that the unforecastable shocks are heteroscedastic. Shocks corresponding to early forecast revisions are typically large. The overall precision of shocks is smallest for soybeans and largest for wheat. Moreover, the USDA's own idiosyncratic errors are not negligible compared to the unforecastable shocks. Especially for corn and soybeans, the variances of the idiosyncratic errors can be larger than the shocks for certain months.

The estimation framework can be easily extended by adding other explanatory variables, such as macroeconomic variables or weather variables. Also, one can investigate the asymmetric impacts by decomposing the forecast revisions into positive and negative parts. Moreover, the framework can be easily applied to analyze other fixed-event forecasts as long as one can specify the correct error covariance matrix.

Given that USDA forecasts are found to be inefficient, an interesting question worth exploring is whether there are any forecasters who can provide better forecasts. In recent years, private analysts have started to provide their own ending stocks forecasts. For crop productions,

Garcia *et al.* (1997) found a decline in the informational value of USDA forecasts. However, the comparisons between USDA and analysts' forecasts have never been addressed before. Thus, it would be worth investigating whether analysts are efficient in forecasting ending stocks, and whether the analysts' forecasts can improve upon USDA forecasts.

2.8 References

- Adjemian, M.K. 2012. Quantifying the WASDE Announcement Effect. *American Journal of Agricultural Economics*. 94(1): 238-256.
- Allen, P.G. 1994. Economic Forecasting in Agriculture. *International Journal of Forecasting* 10:81-135.
- Bauer, R.F., and P.F. Orazem 1994. The Rationality and Price Effects of U.S. Department of Agriculture Forecasts of Oranges. *Journal of finance* 49(2): 681-685.
- Botto, A.C., O. Isengildina, S.H. Irwin, D.L. Good. 2006. Accuracy Trends and Sources of Forecast Errors in WASDE Balance Sheet Categories for Corn and Soybeans. Selected Paper prepared for presentation at the *American Agricultural Economics Association Annual Meeting*, Long Beach, California, July 23-26, 2006.
- Clements, M.P. 1995. Rationality and the Role of Judgement in Macroeconomic Forecasting. *The Economic Journal*. 105(429): 410-420.
- Clements, M.P. 1997. Evaluating the Rationality of Fixed-event Forecasts. *Journal of Forecasting*. 16:225-239.
- Clements, M.P., F. Joutz, and H.O. Stekler. 2007. An Evaluation of the Forecasts of the Federal Reserve: a Pooled Approach. *Journal of Applied Econometrics*. 22: 121-136.
- Davies, A., and K.Lahiri. 1995. A New Framework for Testing Rationality and Measuring Aggregate Shocks Using Panel Data. *Journal of Econometrics* 68: 205-227.
- Davies, A., and K.Lahiri. 1999. Re-examining the Rational Expectations Hypothesis Using Panel Data on Multi-period Forecasts. *Analysis of Panels and Limited Dependent Variable Models*, 226-254. Cambridge University Press.
- Egelkraut, T.M., P. Garcia, S.H. Irwin, and D.L. Good. 2003. An Evaluation of Crop Forecast Accuracy for Corn and Soybeans: USDA and Private Information Agencies. *Journal of Agricultural and Applied Economics* 35:79-95.

- Garcia, P., S. Irwin, R. Leuthold, and L. Yang. 1997. The Value of Public Information in Commodity Futures Markets. *Journal of Economic Behavior and Organization* 32(4): 559-570.
- Gelman, A. 2006. Prior Distributions for Variance Parameters in Hierarchical Models (Comment on the Article by Browne and Draper). *Bayesian Analysis* 1(3):515-534.
- Gelman A. and D.B. Rubin. 1992. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*. 7(4):457-472.
- Harvey, D.I., Leybourne, S.J., and Newbold, P. 2001. Analysis of a Panel of UK Macroeconomic Forecasts. *Econometrics Journal*. 4: 37-55.
- Irwin, S.H., M.E. Gerlow, and T.R. Liu. 1994. The Forecasting Performance of Livestock Futures Prices: A Comparison to USDA Expert Predictions. *Journal of Futures Markets*. 14:861-875.
- Isengildina, O., S.H. Irwin, and D.L. Good. 2004. Evaluation of USDA Interval Forecasts of Corn and Soybean Prices. *American Journal of Agricultural Economics*. 86:990-1004.
- Isengildina O., S.H. Irwin, and D.L. Good. 2006. Are Revisions to USDA crop Production Forecasts smoothed. *American Journal of Agricultural Economics*. 88(4): 1091-1104.
- Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008a. The Impact of Situation and Outlook Information in Corn and Soybean Futures Markets: Evidence from WASDE Reports. *Journal of Agricultural and Applied Economics* 40(1):89-103.
- Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008b. Impact of WASDE Reports on Implied Volatility in Corn and Soybean Markets. *Agribusiness* 24(4): 473-490.
- Isengildina-Massa, O., B. Karali, and S.H. Irwin. 2013. When do USDA forecasters make mistakes? *Applied Economics* 45(36): 5086–5103.
- Mincer, J. and V. Zarnowitz. 1969. The Evaluation of Economic Forecasts, in: J. Mincer (ed.) *Economic Forecasts and Expectations*. New York: National Bureau of Economic Research.
- Nordhaus , W.D. 1987. Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics* 69:667-674.
- Romer, C.D., and D.H. Romer (2000). Federal Reserve Information and the Behavior of Interest rates, *American Economic Review*. 90(3): 429-457.
- Sanders, D.R., and M.R. Manfredo. 2002. USDA Production Forecasts for Pork, Beef, and Broilers: An Evaluation. *Journal of Agricultural and Resource Economics*. 27:114-128.

Sanders, D.R., and M.R. Manfredo. 2003. USDA Livestock Price Forecasts: A comprehensive evaluation. *Journal of Agricultural and Resource Economics*, 28: 316–334.

Vogel, F.A., and G.A. Bange. 1999. Understanding USDA Crop Forecasts. *USDA miscellaneous publication* No. 1554.

2.9 Appendix

Conditional Posterior Distributions for Model Parameters in the Gibbs Sampler

The proposed model consists of system (2.21) and priors (2.25):

$$Y = X\beta + WK + P\Sigma$$

$$\beta \sim N(M, V)$$

$$\varepsilon_{t,n} \sim N(0, \sigma^2) \quad (2.A.1)$$

$$k_{t,n} \sim N(0, \sigma_n^2)$$

$$\sigma, \sigma_n \sim \text{Uniform}(0, \infty)$$

for $n = 1, \dots, N - 1, t = 1, \dots, T$. Let $\Omega \equiv P'P$. Given $\{\beta, K, \sigma^2, \Omega\}$, the dependent variable $y_{t,n}$ follows a multivariate normal distribution:

$$Y | \beta, K, \sigma^2, \Omega \sim N(X\beta + WK, \sigma^2 \Omega) \quad (2.A.2)$$

and the likelihood is $\Phi(Y | \beta, K, \sigma^2, \Omega)$. The posterior density of the set of model parameters is given by

$$p(\Lambda) = \Phi(Y | \beta, K, \sigma^2, \Omega) \prod_{jt} \Phi(k_{t,j} | \sigma_j^2) * \Phi(\beta | M, V) p(\sigma^2) \prod_{j=1}^{N-1} p(\sigma_j^2) \quad (2.A.3)$$

The conditional posterior density for β is

$$p(\beta | \Lambda \setminus \beta) = \Phi(Y | \beta, K, \sigma^2, \Omega) * \Phi(\beta | M, V) \quad (2.A.4)$$

Hence:

$$\begin{aligned} \boldsymbol{\beta}|\Lambda \backslash \boldsymbol{\beta} \sim N((\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X}/\sigma^2 + \mathbf{V}^{-1})^{-1}(\mathbf{X}'\boldsymbol{\Omega}^{-1}(\mathbf{Y} - \mathbf{W}\mathbf{K})/\sigma^2 \\ + \mathbf{V}^{-1}\mathbf{M}), (\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X}/\sigma^2 + \mathbf{V}^{-1})^{-1}) \end{aligned} \quad (2.A.5)$$

The conditional posterior density of $k_{t,n}$, $t = 1, \dots, T$, $n = 1, \dots, N - 1$ is

$$p(k_{t,n}|\Lambda \backslash k_{t,n}) = \Phi(\mathbf{Y}|\boldsymbol{\beta}, \mathbf{K}, \sigma^2\boldsymbol{\Omega}) * \Phi(k_{t,n}|\sigma_n^2) \quad (2.A.6)$$

Therefore,

$$\begin{aligned} k_{t,n}|\Lambda \backslash k_{t,n} \\ \sim N\left(\frac{\mathbf{W}'_{k_{t,n}}\boldsymbol{\Omega}^{-1}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}_{-k_{t,n}}\mathbf{K}_{-k_{t,n}})/\sigma^2}{\mathbf{W}'_{k_{t,n}}\boldsymbol{\Omega}^{-1}\mathbf{W}_{k_{t,n}}/\sigma^2 + 1/\sigma_n^2}, \frac{1}{\mathbf{W}'_{k_{t,n}}\boldsymbol{\Omega}^{-1}\mathbf{W}_{k_{t,n}}/\sigma^2 + 1/\sigma_n^2}\right) \end{aligned} \quad (2.A.7)$$

where $\mathbf{W}_{k_{t,n}}$ is the column of \mathbf{W} which indicates the monthly shock $k_{t,n}$, and $\mathbf{W}_{-k_{t,n}}$, $\mathbf{K}_{-k_{t,n}}$ are matrices with the column indicating $k_{t,n}$ deleted from \mathbf{W} , \mathbf{K} , respectively.

The conditional posterior density of σ_n^2 , $n = 1, \dots, N - 1$ is

$$p(\sigma_n^2|\Lambda \backslash \sigma_n^2) = \prod_{j=1}^N \Phi(k_{t,n}|\sigma_n^2) * p(\sigma_n^2) \quad (2.A.8)$$

Thus

$$\sigma_n^2|\Lambda \backslash \sigma_n^2 \sim IG((T - 1)/2, \sum_{t=1}^T k_{t,n}^2/2) \quad (2.A.9)$$

Finally, the conditional posterior of σ^2 is

$$p(\sigma^2|\Lambda \backslash \sigma^2) = \Phi(\mathbf{Y}|\boldsymbol{\beta}, \mathbf{K}, \sigma^2\boldsymbol{\Omega}) * p(\sigma^2) \quad (2.A.10)$$

so that

$$\sigma^2|\Lambda \backslash \sigma^2 \sim IG((TN - 1)/2, (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}\mathbf{K})'\boldsymbol{\Omega}^{-1}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}\mathbf{K})/2) \quad (2.A.11)$$




Figure 2.1. Standard deviations of USDA forecast revisions, by month.

Table 2.1. Descriptive statistics for USDA forecast revisions.

	Mean	Median	St. Dev.	Min	Max
Corn	-0.0081	0.0000	0.1073	-0.6478	0.4098
Soybeans	-0.0159	-0.0041	0.1153	-0.4884	0.7569
Wheat	0.0013	0.0000	0.0645	-0.2223	0.2302

Note: summary statistics are displayed in logarithms.

Table 2.2. Means and Standard Deviations for the estimates of USDA ending stocks forecasts, 1985/86 – 2013/14.

Parameter	Corn		Soybean		Wheat	
	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)
Coefficient						
α (intercept)	0.0050	(0.0023)**	-0.0137	(0.0028)**	-0.0011	(0.0028)
β (slope)	0.1837	(0.0382)**	0.5137	(0.0478)**	0.1884	(0.0694)**
Shock						
σ_1	0.0865	(0.0129)	0.1652	(0.0248)	0.0708	(0.0106)
σ_2	0.0319	(0.0079)	0.0755	(0.0138)	0.0344	(0.0078)
σ_3	0.0399	(0.0083)	0.0418	(0.0162)	0.0317	(0.0081)
σ_4	0.0709	(0.0109)	0.0418	(0.0143)	0.0448	(0.0083)
σ_5	0.0156	(0.0089)	0.0193	(0.0136)	0.0309	(0.0085)
σ_6	0.0521	(0.0090)	0.0633	(0.0152)	0.0422	(0.0080)
σ_7	0.0845	(0.0127)	0.0393	(0.0163)	0.0628	(0.0105)
σ_8	0.0171	(0.0087)	0.0195	(0.0136)	0.0535	(0.0090)
σ_9	0.0523	(0.0089)	0.0423	(0.0157)	0.0323	(0.0086)
σ_{10}	0.1338	(0.0194)	0.0458	(0.0200)	0.0859	(0.0128)
σ_{11}	0.0563	(0.0097)	0.0542	(0.0176)	0.0685	(0.0108)
σ_{12}	0.0766	(0.0121)	0.1365	(0.0229)	0.0984	(0.0144)
σ_{13}	0.1459	(0.0211)	0.1987	(0.0314)	0.0911	(0.0144)
σ_{14}	0.0870	(0.0150)	0.1325	(0.0246)		
σ_{15}	0.2458	(0.0351)	0.1321	(0.0238)		
σ_{16}	0.1863	(0.0267)	0.1214	(0.0228)		
Idiosyncratic Err.						
σ	0.0184	(0.0027)	0.0465	(0.0032)	0.0131	(0.0062)

Note: (*) and (**) denote parameter estimates significant at 10% and 5%, respectively. The standard errors for the idiosyncratic error and shocks are all significant at 5% level, and hence the indicators are omitted.

Table 2.3. Medians and Credible Intervals for the estimates of USDA ending stocks forecasts, 1985/86 – 2013/14

Parameter	2.5%	Corn Median	97.5%	2.5%	Soybean Median	97.5%	2.5%	Wheat Median	97.5%
Coefficient									
α (intercept)	0.0009	0.0049	0.0098	-0.0191	-0.0137	-0.0082	-0.0066	-0.0011	0.0046
β (slope)	0.1063	0.1848	0.2561	0.4212	0.5139	0.6060	0.0700	0.1822	0.3367
Shock									
σ_1	0.0654	0.0851	0.1157	0.1244	0.1625	0.2216	0.0534	0.0696	0.0949
σ_2	0.0164	0.0316	0.0484	0.0516	0.0743	0.1058	0.0178	0.0346	0.0494
σ_3	0.0249	0.0395	0.0577	0.0072	0.0423	0.0732	0.0129	0.0321	0.0464
σ_4	0.0529	0.0698	0.0958	0.0118	0.0420	0.0698	0.0291	0.0444	0.0621
σ_5	0.0007	0.0160	0.0328	0.0008	0.0170	0.0500	0.0108	0.0316	0.0462
σ_6	0.0364	0.0513	0.0718	0.0355	0.0626	0.0955	0.0271	0.0419	0.0586
σ_7	0.0635	0.0832	0.1131	0.0061	0.0397	0.0712	0.0437	0.0621	0.0854
σ_8	0.0013	0.0175	0.0337	0.0008	0.0176	0.0498	0.0374	0.0529	0.0730
σ_9	0.0370	0.0515	0.0720	0.0095	0.0428	0.0730	0.0103	0.0331	0.0475
σ_{10}	0.1017	0.1317	0.1776	0.0043	0.0467	0.0842	0.0643	0.0847	0.1146
σ_{11}	0.0396	0.0555	0.0775	0.0189	0.0541	0.0892	0.0497	0.0676	0.0922
σ_{12}	0.0564	0.0754	0.1037	0.0981	0.1343	0.1877	0.0746	0.0969	0.1309
σ_{13}	0.1110	0.1437	0.1932	0.1461	0.1957	0.2693	0.0659	0.0900	0.1226
σ_{14}	0.0619	0.0855	0.1204	0.0902	0.1304	0.1863			
σ_{15}	0.1879	0.2421	0.3253	0.0916	0.1299	0.1847			
σ_{16}	0.1425	0.1835	0.2464	0.0816	0.1197	0.1715			
Idiosyncratic Err.									
σ	0.0130	0.0185	0.0237	0.0407	0.0464	0.0531	0.0038	0.0125	0.0260

CHAPTER 3

THE ENDING STOCKS FORECASTS FROM PRIVATE ANALYSTS: AN EXAMINATION OF EFFICIENCY

3.1 Abstract

This paper examines the efficiency of private analysts' ending stocks forecasts for corn, soybeans and wheat for marketing years 2004/05 through 2013/14. The model proposed in Chapter 2, which focuses on adjacent forecast revisions while retaining the link between the forecasts and the ending stocks, is applied. This paper also investigates individual analysts' forecasts after filling the missing data points using multiple imputations. Results show that private analysts, as a group, are inefficient in making ending stocks forecasts for all three commodities. Analysts are also found to have similar forecasting behavior as the USDA for corn and soybeans. Moreover, estimations show that forecasting behavior varies across individual analysts.

3.2 Introduction

Ending stocks are a measure of the unutilized quantity of the commodity at the end of a marketing year. Entering into the market in the following marketing year, ending stocks can serve as reserve against unexpected supply and demand shocks in the commodity market. They are major indicators of the supply of the commodity, and have great impact on the price volatility and market participants' decision makings. As the United States is a major exporting country for agricultural commodities like corn, soybeans and wheat, the impact of stocks changes is much larger than for other countries in the global market. Therefore, the public has paid close attention

to the U.S. ending stocks level. The provision of accurate forecasts of ending stocks becomes extremely important because it can timely reflect the market situation and reduce the uncertainty faced by decision makers. In other words, it can significantly reduce the market risk and enhance the overall functioning of the commodity market.

For several decades, the U.S. Department of Agriculture (USDA) has provided ending stocks forecasts for major agricultural commodities in its monthly World Agricultural Supply and Demand Estimates (WASDE) reports. As released by the government agency, the forecasts in the WASDE reports are highly credited for their integrity and incorporation of comprehensive information. WASDE forecasts have been paid substantial attention by policy makers and various market participants, including farmers and agribusinesses. Researchers also find that decisions and behaviors of these participants have been adjusted following the update of the WASDE forecasts (*e.g.*, Bauer and Orazem 1994, Garcia *et al.* 1997, Isengildina-Massa *et al.* 2008a, 2008b, Adjemian 2012).

Recently, a new group of forecasters, private analysts, has arisen due to the reduced difficulty and cost of acquiring related information. A greater number of analysts has started to provide crop ending stocks forecasts over the past few years. Although private analysts have the same forecast target as the USDA, their forecasts are different. Specifically, their sources of information need not overlap with or be as comprehensive as those from the USDA. For example, crop-production related surveys conducted by the private sector may not cover as many farmers as those done by the USDA. This could possibly be due to costs or time issues. Also, those farmers who participate in the surveys from the private sector may differ from those who participate in USDA surveys. Secondly, although private analysts and the USDA may share some common information such as satellite maps or macroeconomic conditions, their

interpretation and inference might not be the same because they are performed by different analysts. Thirdly, it is possible that the objective of private analysts is distinct from that of the USDA. This could happen because some private analysts themselves are actively involved in the market and even pursue profit by directly trading the commodity. Thus, one would expect their forecasts to be more or less affected by their own trading objectives and strategies. On the other hand, USDA forecasts are widely considered to be objective because of the policy of the government agency and the goal for enhancing the functioning of agricultural markets as a public service.

Many researchers have investigated the effect of the rise of the private forecasts. For example, French *et al.* (1989) find that relatively large differences between the consensus analysts' forecasts and the forecasts from government agencies often lead to market volatility. Studies also find that analysts' start to compete with government agencies. In the case of crop output forecasts, Garcia *et al.* (1997) found a decline in the informational value of USDA forecasts as the private sector started to provide its own forecasts. Egelkraut *et al.* (2003) found that private agencies can compete with the USDA in making forecasts for several specific months. Further, Fortenberry and Sumner (1993) found that markets no longer react to USDA production forecasts, whereas McKenzie (2008) arrived at the opposite conclusion that futures prices continue to react to the USDA reports. Given the aforementioned findings, it is interesting to study the performance of private analysts as competitors to the USDA, especially in forecasts of crop ending stocks, which have been overlooked by the literature.

U.S. ending stocks forecasts are fixed-event forecasts because they are made for a specific target (the ending stocks), but have different forecast horizons. Previous research on fixed-event forecasts has often examined macroeconomic variables such as inflation rate, interest

rate, and GDP growth rates. (*e.g.*, Clements 1995, 1997, Romer and Romer 2000, Harvey *et al.* 2001, Clements *et al.* 2007) There are also some studies which focus on agricultural fixed-event forecasts such as USDA production forecasts (Sanders and Manfredo, 2002; Isengildina *et al.* 2006) and USDA ending stocks forecasts (Botto *et al.*, 2006; Isengildina-Massa *et al.*, 2013).

Previous models regarding efficiency of fixed-event forecasts can generally be categorized into two main strands. The first strand, which is based on Nordhaus (1987), builds on forecast revisions. Nordhaus (1987) introduced a weak efficiency test which only uses information on past forecasts because the forecast history is always available to public. The test consists of identifying the relationship of adjacent forecast revisions. However, the model's assumption of *i.i.d.* errors need not to be realistic, as it cuts off the inherent relationship between forecasts and the forecast target. In contrast, the second strand, which is advocated by Davies and Lahiri (1995, 1999), directly focuses on forecast errors. They decompose the forecast errors as the sum of unforecastable shocks and the forecaster's own idiosyncratic errors. Based on their work, Clements *et al.* (2007) analyzed forecast revisions by differencing the forecast errors. In this way, they can avoid possible problems in the original Davies and Lahiri model, which are generated by the correlations between the dependent variables and the residuals. However, for estimation purposes, they did not use the postulated theoretical structure of the covariance matrix. Instead, they adopted a simplified version of it.

The proposed the model in Chapter 2 further extends the Clements model by investigating adjacent forecast revisions with emphasis on the link between the forecasts and the forecast target. The model in Chapter 2 combined the Nordhaus (1987) and Davies and Lahiri (1995, 1999) approaches and developed a comprehensive test to analyze forecast efficiency. In particular, the model builds an error covariance matrix which takes into account both the

heteroscedasticity of the unforecastable shocks and the autocorrelations generated by the forecaster's correction of their own errors. The model implies that the existence of forecaster's own idiosyncratic errors will lead to negative correlations between adjacent forecast revisions if forecasts are efficient. This correlation structure serves as the link between the forecasts and the forecast target – the ending stocks. The model was applied to USDA ending stocks forecasts of corn, soybeans and wheat for marketing years 1985/86 – 2013/14. Results show that the USDA forecasts are inefficient and there is strong evidence that the USDA is conservative in making ending stocks forecasts.

The present study applies the model in Chapter 2 to explore whether private analysts' forecasts of ending stocks are efficient. The analysis focuses on three major agricultural commodities, namely, corn, soybeans and wheat. A total of 10 marketing years, 2004/05 – 2013/14, are investigated. As ending stocks forecasts from private analysts have similar forecasting schedules as those from the USDA, the error covariance matrix developed in Chapter 2 can also be used in describing the error structure of private analysts' forecasts. We form four representative analysts with different combinations of the analysts' forecast data and analyze the efficiency of their forecasts.

Private analysts' forecasts, like other data in social sciences, are often incomplete. This problem of missing data needs to be addressed in order to study the forecasts from individual analysts. Rubin (1976) proposes that there exist several types of missingness. The ideal type is missing completely as random (MCAR). Under MCAR, the data can be viewed as a subsample of the complete dataset. However, this assumption is too strong for most of the real-world datasets. The missing data can also be ignored under a less restrictive assumption: missing at random (MAR). Past researchers developed methods like listwise deletion and pairwise deletion

to discard the observations or variables which are incomplete. But these methods typically yield biased estimates even under MAR (*e.g.*, Arbuckle 1996, Wothke 2000). Besides, sometimes it results in insufficient data for researchers to use. Thus, recent studies have started to focus on data imputation in order to utilize the information in the incomplete observations which would otherwise have been discarded. Early imputation methods are typically single imputations, including hot deck imputations, mean imputations, regression imputations, *etc.* However, these methods ignore the uncertainty generated by the missing data. Hence, Rubin (1987) proposed a method of averaging the estimates from multiple imputed datasets. The method introduced certain degrees of randomness to the imputations to account for the uncertainty in predicting the missing data.

The present study also examines the forecasting behavior of individual analysts, the data of which are incomplete. We use a multiple imputation method, developed by Honaker and King (2010), to generate complete datasets for analysis. The method is designed specifically for time-series cross-section datasets, taking into account for various characteristics such as smooth time trends, correlations over time and space, *etc.* The imputed values are thus more accurate compared to previous imputation methods. We then develop an approach to use the imputed datasets in our estimation procedures.

The estimation is performed by a Bayesian Markov Chain Monte Carlo (MCMC) method. The method is particularly useful in analyzing the proposed model structure because it allows us to estimate the coefficients and the error covariance matrix in one iteration step. We adopt conjugate priors for estimations. These priors contain no subjective information so that the data dominates the estimation processes. In this way, the method developed can also serve as a justification of the model proposed structure. Moreover, the MCMC method allows us to

examine the full posterior distributions for the parameters of interest, especially for the variances of the errors on which the distributional assumptions are not normal.

The rest of the paper is organized as follows. Section 3.3 introduces the model for investigating the efficiency of private analysts' ending stocks forecasts. Section 3.4 describes the data and explains the proposed way of forming representative analysts. Section 3.5 introduces the MCMC method employed for the estimation and the approach used to integrate the imputed individual analysts' datasets into the estimation procedures. Section 3.6 describes the results and implications, and the final section concludes.

3.3 The Model

Chapter 2 introduced a model of evaluating fixed-event forecasts by focusing on forecast revisions instead of forecast errors, while retaining the link between the forecasts and the target event. The model is developed to deal with the issues generated by the endogeneity after the introduction of complex error covariance structures.

To apply the model to the private analysts' forecasts, let S_t be the realization of the ending stock of a given commodity at the end of marketing year t . Let n be the forecast horizon and $V_{t,n}$ be the representative analyst's n -month ahead forecast of the ending stock S_t . Based on Davies and Lahiri (1995, 1999), forecast errors can be evaluated by the following regression:

$$S_t - V_{t,n} = n\alpha + \sum_{j=1}^n k_{t,j} + \varepsilon_{t,n} \quad (3.1)$$

where α represents the bias of the forecast revisions, $k_{t,j} \sim i.i.d. N(0, \sigma_j^2)$ represents the shock of month j of the forecasting cycle for marketing year t , and $\varepsilon_{t,n} \sim i.i.d. N(0, \sigma^2)$ is the representative analyst's own idiosyncratic error. Thus $n\alpha$ is the full bias of the forecast error of

$V_{t,n}$, and $\sum_{j=1}^n k_{t,j}$ can be interpreted as the aggregate shock between forecast month n and the time of the revelation of the ending stock.

Given the structure of (3.1), Chapter 2 suggests applying first differencing as in Clements *et al.* (2007) to consider the forecast revisions instead of the forecast errors. Forecast revisions are more convenient for the econometric analysis than forecast errors, because the later contain information from later periods, which is not closely related to the information available when the current forecasts are made. Thus these pieces of future information can be safely excluded from the model. The link between the forecasts and the ending stocks, however, is still retained by the idiosyncratic errors. First-order differencing of (3.1) gives

$$\begin{cases} S_t - V_{t,1} &= \alpha + k_{t,1} & +\varepsilon_{t,1} \\ V_{t,1} - V_{t,2} &= \alpha + k_{t,2} & -\varepsilon_{t,1} + \varepsilon_{t,2} \\ &\vdots \\ V_{t,N-1} - V_{t,N} &= \alpha + k_{t,N} & -\varepsilon_{t,N-1} + \varepsilon_{t,N} \end{cases} \quad (3.2)$$

where N is the maximum forecast horizon for a marketing year.

The efficiency test based on (3.2) consists of fitting

$$\begin{cases} S_t - V_{t,1} &= \alpha + \beta X_{t,1} + k_{t,1} & +\varepsilon_{t,1} \\ V_{t,1} - V_{t,2} &= \alpha + \beta X_{t,2} + k_{t,2} & -\varepsilon_{t,1} + \varepsilon_{t,2} \\ &\vdots \\ V_{t,N-1} - V_{t,N} &= \alpha + \beta X_{t,N} + k_{t,N} & -\varepsilon_{t,N-1} + \varepsilon_{t,N} \end{cases} \quad (3.3)$$

where $X_{t,n}$'s represent explanatory variables which contain available information at the time the forecast $V_{t,n}$ is made. Efficiency can then be examined by testing the null hypothesis $H_0: \alpha = \beta = 0$. The representative analyst's forecasts V are said to be efficient under the null hypothesis, meaning that the revisions cannot be predicted using available information.

One reasonable candidate for $X_{t,n}$ is the representative analyst's most recent forecast revision $V_{t,n} - V_{t,n+1}$, because it represents the most up-to-date information and is always available. Substituting $X_{t,n}$ for $V_{t,n} - V_{t,n+1}$, regression (3.3) becomes

$$\begin{cases} S_t - V_{t,1} &= \alpha + \beta(V_{t,1} - V_{t,2}) + k_{t,1} & +\varepsilon_{t,1} \\ V_{t,1} - V_{t,2} &= \alpha + \beta(V_{t,2} - V_{t,3}) + k_{t,2} & -\varepsilon_{t,1} + \varepsilon_{t,2} \\ \vdots & \\ V_{t,N-2} - V_{t,N-1} &= \alpha + \beta(V_{t,N-1} - V_{t,N}) + k_{t,N} & -\varepsilon_{t,N-1} + \varepsilon_{t,N} \end{cases} \quad (3.4)$$

The last equation in (3.3) is dropped as there are no past revisions associated to the earliest forecast revision. However, the inclusion of most recent forecast revisions as explanatory variables brings out an endogeneity problem if the regression is estimated equation by equation, because the forecast revisions appear both as dependent variables and explanatory variables. Specifically, there exist correlations between the revisions and the idiosyncratic errors. To address the problem, Chapter 2 proposes estimating the parameters by recognizing the equations as a system. The earliest forecast revisions are treated as exogenous, while the remaining revisions are then treated as endogenous.

The model also introduces a covariance matrix with a minimal number of parameters to be estimated. Without loss of generality, the data can be sorted by increasing forecast horizon $n = 1, \dots, N$. Take corn for example. Based on Chapter 2, the error covariance structure of a typical marketing year can be characterized as

$$B_{(N-1) \times (N-1)} = \begin{bmatrix} \sigma_1^2 + \sigma^2 & -\sigma^2 & 0 & \dots & 0 & 0 & 0 & 0 \\ -\sigma^2 & \sigma_2^2 + 2\sigma^2 & -\sigma^2 & & & 0 & 0 & 0 \\ 0 & -\sigma^2 & \ddots & & & & 0 & 0 \\ \vdots & & & \ddots & & & & 0 \\ 0 & & & & \sigma_{N-4}^2 + 2\sigma^2 & & & \vdots \\ 0 & 0 & & & & \sigma_{N-3}^2 + 2\sigma^2 & -\sigma^2 & 0 \\ 0 & 0 & 0 & & & -\sigma^2 & \sigma_{N-2}^2 + 2\sigma^2 & -\sigma^2 \\ 0 & 0 & 0 & 0 & \dots & 0 & -\sigma^2 & \sigma_{N-1}^2 + \sigma^2 \end{bmatrix} \quad (3.5)$$

The diagonal elements represent the total error variances. It can be seen that the differences are contributed by the heteroscedastic shocks. The subdiagonal elements, as well as the superdiagonal, are the negative of the variances of the idiosyncratic errors. They represent the autocorrelations between adjacent equations, which also serve as the link between forecasts and the forecast target - the ending stocks.

Moreover, there exist correlations between forecasts made at the same month but for different marketing years. Thus the *i. i. d.* assumption imposed on the monthly shocks associated with these forecasts might no longer be realistic. Instead, the monthly shocks with horizons greater than 12 are redefined as $\bar{k}_{t+1,n+12}$, where

$$\bar{k}_{t+1,n+12} = k_{t,n} + k_{t+1,n+12} \text{ for } 1 \leq n \leq N - 12 \quad (3.6)$$

Thus, $\bar{k}_{t+1,n+12}$ contains the shock in the current marketing year $k_{t,n}$ and the shock $k_{t+1,n+12}$ which solely affects the ending stock in the next marketing year. In this way, the covariance between $\bar{k}_{t+1,n+12}$ and $k_{t,n}$ is σ_n^2 . Let $\bar{\sigma}_{n+12}^2$ be the variance of forecasts with horizon greater than 12.

Then

$$\bar{\sigma}_{n+12}^2 = \sigma_{n+12}^2 + \sigma_n^2 \quad (3.7)$$

For corn and soybean forecasts, the error covariance matrix then becomes

$$\begin{aligned}
& \bar{B}_{(N-1) \times (N-1)} \\
& = \begin{bmatrix} \sigma_1^2 + \sigma^2 & -\sigma^2 & 0 & \dots & \sigma_1^2 & 0 & 0 & 0 \\ -\sigma^2 & \sigma_2^2 + 2\sigma^2 & -\sigma^2 & & & \sigma_2^2 & 0 & 0 \\ 0 & -\sigma^2 & \ddots & & & & \sigma_3^2 & 0 \\ \vdots & & & \ddots & & & & \sigma_4^2 \\ \sigma_1^2 & & & & \bar{\sigma}_{N-4}^2 + 2\sigma^2 & & & \vdots \\ 0 & \sigma_2^2 & & & & \bar{\sigma}_{N-3}^2 + 2\sigma^2 & -\sigma^2 & 0 \\ 0 & 0 & \sigma_3^2 & & & -\sigma^2 & \bar{\sigma}_{N-2}^2 + 2\sigma^2 & -\sigma^2 \\ 0 & 0 & 0 & \sigma_4^2 & \dots & 0 & -\sigma^2 & \bar{\sigma}_{N-1}^2 + \sigma^2 \end{bmatrix} \quad (3.8)
\end{aligned}$$

To see the structure of the full error covariance matrix, note that without loss of generality, we can sort the data first by marketing year $t = 1, \dots, T$, and then by forecast horizon $n = 1, \dots, N$. Then the covariance matrix of the error terms can be expressed as a $T(N-1) \times T(N-1)$ block diagonal matrix:

$$\Sigma = \begin{bmatrix} \bar{B} & & & \\ & \bar{B} & & \\ & & \bar{B} & \\ & & & \bar{B} \end{bmatrix}_{T \times T} \quad (3.9)$$

3.4 Data

3.4.1 Data Sources and Structure

The model is applied to data on U.S. ending stocks and their corresponding private analysts' monthly forecasts, for three major agricultural commodities – corn, soybeans and wheat. A total of 10 marketing years, from 2004/05 through 2013/14, are investigated in the present study. The forecasts from the USDA within the same period are also analyzed as comparisons.

U.S. ending stocks are obtained from the Grain Stocks Report released by the National Agricultural Statistics Services (NASS). The Grain Stocks Report is published quarterly,

typically in early January, and late March, June, September. As the U.S. marketing year for corn and soybeans starts in September and ends in August of the following calendar year, we retrieve the final ending stocks from the September report – the first report after the ending of the U.S. marketing year. For wheat, the U.S. marketing year starts in June and ends in May of the following calendar year. Thus, the wheat ending stocks are obtained from the June report.

The USDA monthly forecasts are retrieved from the WASDE reports. They are used as benchmarks to which private analysts' forecasts are compared. For each marketing year, the first USDA forecast for corn and soybeans is released in May before the marketing year begins. The last forecast is released in early September, after the ending of the marketing year but before the release of the ending stock. Thus there are 17 forecasts in a forecasting cycle for corn and soybeans. For wheat, there are 14 forecasts instead, with the first one released in May, together with that for corn and soybeans, and the last released in June of the following calendar year.

The private analysts' forecast data have the same format as the USDA data. Analysts' forecasts obtained from the monthly Survey of U.S. Grain and Soybeans Carryout Forecasts conducted by Dow Jones Commodities Services. It is worth noting that the surveys are typically released a few days before the release of the WASDE reports. Hence, there could be differences between private analysts' forecasts and their USDA counterpart, as the information used to generate them might not be the same.

As discussed in the previous section, the first forecast of each marketing year can only be used to construct an explanatory variable (*i.e.*, for the second revision). Thus for corn and soybeans, the dependent variable consists of only 16 forecast revisions, which yields 160 (=16 forecasts \times 10 marketing years) observations in the regression system. For wheat, the dependent

variable consists of 13 revisions and the regression system has 130 (=16 forecasts \times 10 marketing years) observations.

3.4.2 Creating a Representative Analyst

Private analysts' forecasts are often incomplete, as they are not required by policies or laws to provide consecutive forecasts every month. Private analysts' forecasts sometimes are missing due to various reasons. For example, an analyst may fail to respond to the survey on time. On the other hand, there might also be some issues in collecting analysts' responses to the survey, resulting in missing data points for those analysts.

When examining the private analysts' forecasts, we find another reason that caused the data points to appear missing: business changes. A private analyst may stop providing forecasts if the company ceases its business in research or trading in commodities and their derivatives, or shifts its business focus to other areas not related to forecasting. Also, due to the rapid development in the agribusiness industry, many new forecasters emerged over the past decade. These forecasters typically have missing data points in early periods of the marketing years examined in this study.

We investigated the background of all private forecasters in the Survey of U.S. Grain and Soybeans Carryout Forecasts. We combine forecasts from several analysts due to business merges, who were identified as separate forecasters in the original data. Take the ABN Amro / Fortis case for example. Most parts of business of ABN Amro, including commodity services, were taken over by Fortis in 2007. Later Fortis began providing its own ending stocks forecasts. We also find in the data that the time when ABN Amro stopped forecasting and Fortis started forecasting matched perfectly. Thus, we are able to combine these two forecasters into a single

one. Similar cases applied to Alaron, which was acquired by PFGBest in 2009, and A.G. Edwards & Sons, which was acquired by Wachovia Securities in 2007. The final dataset, therefore, contains 54 analysts who provide at least one forecast during the 10 marketing years investigated.¹

It is also worthwhile to look at the forecasts which are provided by analysts who have few missing observations. We define frequent forecasters as those analysts who provide forecasts for more than 60% of the time.² These analysts are typically large companies which have been in the industry for a long time, for example, Allendale, Citigroup, US Commodities, *etc.* As these businesses may have more resources to collect data and perform analysis than other forecasters, it is interesting to see whether their forecasts have different patterns and if these differences exist, how they differ. For corn and soybeans, we have identified 10 analysts as frequent forecasters. For wheat, 9 analysts have been viewed as frequent forecasters. The names and number of forecasts of these analysts are reported in Table 3.1.

The present model utilizes the following four combinations of analysts' data as the representatives of analysts' forecasts:

AA: The average of analysts' forecasts

MA: The median of analysts' forecasts

AFA: The average of selected frequent analysts' forecasts

MFA: The median of selected frequent analysts' forecasts

¹ There is a special case for which we are able to find the date when the two businesses merged, but the two firms have overlapping forecasts. In this case, we still treat the two businesses as separate analysts because there could be some unidentified reasons resulting in such overlap.

² The threshold of 60% is selected not only for the number of available observations, but also based on the missing patterns. For example, an analyst who reports slightly less than 60% but did not forecast in the first 4 marketing years is not considered as a frequent forecaster. In the present study, we try to select analysts whose forecasts span the whole 10 marketing years.

We choose the average of analysts' forecasts (AA) as a representative because it is widely accepted by the public as a benchmark of consensus analysts' forecasts. The median of analysts' forecasts (MA) is chosen because the number of analysts who offer forecasts varies for each month. Hence, the median is more likely to be more stable if there exist outliers.³ The same rationales apply to the mean and median of the selected frequent analysts' forecasts. As comparisons, the USDA forecasts are also estimated using the same model, as they are viewed as the "official" forecasts by the public.

3.4.3 Individual Analysts' Forecasts

The present study also investigates the efficiency of individual analysts' forecasts. We apply the estimations to the forecasts from the selected frequent analysts, that is, 10 analysts for corn and soybeans, and 9 analysts for wheat. The data from other analysts' are not analyzed because of the high level of missingness, which would otherwise introduce too much uncertainty to the estimations. We then use a method of multiple imputations to fill the missing data points of these analysts' forecasts. The details and the estimations steps are discussed together in Section 3.4.

3.4.4 Descriptive Statistics

Table 3.2 shows the descriptive statistics for the four representative analysts' forecast revisions and the USDA counterpart for all three commodities.^{4 5} The overall means are not significantly different from zero for all three commodities. The overall medians are also zero or

³ In other words, the median of analysts' forecasts can be more credible as consensus if some analysts made forecasts with large deviations from the majority.

⁴ The earliest forecast revisions are also included in calculating the descriptive statistics.

⁵ The statistics for the individual analysts' forecasts are omitted due to space limits.

close to zero. The standard deviations are much larger for corn and soybeans than for wheat. For corn, the five sets of data show that forecast revisions range from approximately -60% to 40%. The AA and AFA data have both lower upperbounds and lower lowerbounds. The MA and MFA data have both higher upperbounds and higher lowerbounds. For soybeans, the USDA forecast revisions range from -38.30% to 75.69%. The range for the analysts' forecasts are quite different. The AA and MA data have both higher upperbounds and higher lowerbounds, whereas the AFA and MFA data have lower upperbounds and lower lowerbounds. For wheat, the USDA data have smaller range, at -16.48% to 18.63%. The ranges for the four representative analysts' datasets are quite similar and larger, at around -20% to 24%.

Figure 3.1 depicts the monthly standard deviations of the forecast revisions. The sequences are displayed in the order of diminishing forecast horizons. It can be seen from the figure that for all three commodities, the standard deviations exhibit a general decreasing trend as the forecast horizons shorten. Besides, there is a jump in the standard deviations for the final revisions for all forecasts data. It can also be seen that standard deviations of forecast revisions are greater for corn and soybeans stocks than for wheat stocks.

3.5 Estimation Methods

3.5.1 Estimation Procedures for Representative Analysts' Forecasts

The model is estimated using a Bayesian Markov Chain Monte Carlo (MCMC) method, as it greatly facilitates dealing with an error structure characterized by both heteroscedasticity and autocorrelation. In this section, we first outline the joint posterior distributions of the parameters, and then introduce the choice of prior for each parameter. Finally, we list the steps in the MCMC estimation procedures.

To simplify the notation, we rewrite the proposed regression system as:

$$y_{t,n} = \mathbf{x}_{t,n}\boldsymbol{\beta} + k_{t,n-1} - \varepsilon_{t,n-1} + \varepsilon_{t,n} \quad (3.10)$$

where $y_{t,n} \equiv V_{t,n-1} - V_{t,n}$, $\mathbf{x}_{t,n} \equiv [1 \ V_{t,n} - V_{t,n+1}]$, $\boldsymbol{\beta} \equiv [\alpha \ \beta]'$. Then the matrix form of panels of regressions for each marketing year is,

$$\mathbf{y}_t = \mathbf{x}_t\boldsymbol{\beta} + \mathbf{w}\mathbf{k}_t + \mathbf{p}\boldsymbol{\varepsilon}_t \quad (3.11)$$

where $\mathbf{y}_t = [y_{t,1}, \dots, y_{t,N-1}]'$, $\mathbf{x}_t \equiv [\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,N-1}]'$, $\mathbf{k}_t = [k_{t,1}, \dots, k_{t,N-1}]'$, and $\boldsymbol{\varepsilon}_t = [\boldsymbol{\varepsilon}_{t,1}, \dots, \boldsymbol{\varepsilon}_{t,N-1}]'$. \mathbf{w} is a matrix indicating the existence of elements in \mathbf{k}_t in each corresponding equation. And \mathbf{p} is a matrix indicating the existence of elements in $\boldsymbol{\varepsilon}_t$ in each corresponding equation. The full system can then be written as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{W}\mathbf{K} + \mathbf{P}\mathbf{E} \quad (3.12)$$

where each character represents the vector containing the same character with subscripts $t = 1, \dots, T$. For identification purposes, $k_{T,1}$ is set to be zero, as there is a common intercept α to be estimated. In other words, we assume that the final forecast error in marketing year T only contains the representative analyst's own idiosyncratic error.

Let $\Lambda = \{\boldsymbol{\beta}, \{\sigma_n^2\}_{n=1}^{N-1}, \sigma^2\}$ be the set of the parameters of the proposed model. The joint posterior density of Λ is

$$p(\Lambda) = \Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \sigma^2 \boldsymbol{\Omega}) \prod_{jt} \Phi(k_{t,j} | \sigma_j^2) * p(\boldsymbol{\beta}) p(\sigma^2) \prod_{j=1}^{N-1} p(\sigma_j^2) \quad (3.13)$$

$\Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \sigma^2 \boldsymbol{\Omega})$ is the distribution of \mathbf{Y} , which is multivariate normal. $\sigma^2 \boldsymbol{\Omega}$ is the covariance matrix generated by the idiosyncratic errors. $p(\boldsymbol{\beta})$ is the prior distribution of $\boldsymbol{\beta}$. $p(\sigma^2)$ is the prior distribution of the representative analyst's idiosyncratic errors. $p(\sigma_j^2)$ is the prior distribution of the shocks of month j .

To calculate the posterior distribution for each parameter, a Gibbs Sampler is derived based on general least squares. We use conditionally conjugate priors for the estimations.

Specifically, the priors chosen for the parameters are:

$$\begin{aligned}
 \boldsymbol{\beta} &\sim N(\boldsymbol{M}, \boldsymbol{V}) \\
 \varepsilon_{t,n} &\sim N(0, \sigma^2) \\
 k_{t,n} &\sim N(0, \sigma_n^2) \\
 \sigma, \sigma_n &\sim \text{Uniform}(0, \infty)
 \end{aligned} \tag{3.14}$$

for $n = 1, \dots, N - 1, t = 1, \dots, T$. The prior distribution for the coefficient vector $\boldsymbol{\beta}$ is multivariate normal with mean $\boldsymbol{M} = \mathbf{0}_2$ and covariance matrix $\boldsymbol{V} = 1000\boldsymbol{I}_{2 \times 2}$, where $\mathbf{0}_2$ is a 2×1 vector of zeros and $\boldsymbol{I}_{2 \times 2}$ is a 2×2 identity matrix. Thus the posterior of $\boldsymbol{\beta}$ also follows a multivariate normal distribution. The prior mean of $\boldsymbol{\beta}$ is chosen to be consistent of the null hypothesis of efficiency. The scale of variance is chosen to be large so that the prior is non-informative. In this way, the draws of $\boldsymbol{\beta}$ will be diffused and widely spread around the mean zero. The uniform prior for the standard deviation parameters is chosen following Gelman (2006). This prior is non-informative and can be viewed as a limit of the half- t family distributions, which is conditionally conjugate to the extent of more general folded-noncentral- t distributions. The posterior of the standard deviation parameters follows an inverse gamma distribution, a family of folded-noncentral- t distributions. The details of the conditional posterior distributions for the parameters of the model are outlined in the Appendix.

The MCMC iteration steps for the model can be summarized as follows:

Step 1: Set up initial values for each parameter in the set Λ , as well as $\boldsymbol{K}^{(0)}$ and $\boldsymbol{E}^{(0)}$.

Step 2: Given $\{k_{t,n}^{(i)}, \sigma^{2(i)}\}$, draw $\boldsymbol{\beta}^{(i+1)}$ from a multivariate normal distribution.

Step 3: Given $\{\boldsymbol{\beta}^{(i+1)}, \sigma^{2(i)}, \{\sigma_n^2\}^{(i)}, \mathbf{K}_{-k_{t,n}}^{(i)}\}$, sequentially draw $k_{t,n}^{(i+1)}$ from a normal distribution for each $t = 1, \dots, T$ and $n = 1, \dots, N - 1$.

Step 4: Given $\{\boldsymbol{\beta}^{(i+1)}, \mathbf{K}^{(i+1)}\}$, update $\mathbf{E}^{(i+1)}$, and draw $\sigma^{2(i+1)}$ from an inverse gamma distribution.

Step 5: Given $\mathbf{K}^{(i+1)}$, sequentially draw $\sigma_n^{2(i+1)}$ from an inverse gamma distribution for each $n = 1, \dots, N - 1$.

Step 6: Set $i = i + 1$.

Step 7: Repeat Step 2 until the maximum iteration is reached.

For each dataset, the Gibbs Sampler is run for three Markov Chains for at most 80,000 iterations each.⁶ The first half of each chain is discarded as a burn-in period. Gelman and Rubin (1992) tests are then applied to check the convergence of the remaining part of the chains. The Gelman and Rubin test statistic compares the variances of both within the chains and between the chains. Values of the statistics close to 1 indicate convergence.

3.5.2 Integration of Multiple Imputation and MCMC Estimations

Because of the missing values, additional steps need to be included in order to utilize the individual analysts' forecast data. For the selected frequent analysts' forecasts, we use multiple imputations to fill the missing data points, thus creating a list of full datasets. The imputation is performed using the Amelia II package in R. The Amelia II package is developed by Honaker *et al.* (2014), with the imputation method based on Honaker and King (2010). The method is designed specifically for imputing time-series cross-section datasets, and it contains features

⁶ The lengths of the Markov Chains are different for different datasets. This is because for some datasets the chains converge quickly, so that it is not necessary to run additional iterations.

which are not considered in standard multiple imputation methods, such as smooth time trend, correlations over time and space, *etc.*

The method proposed by Honaker and King (2010) assumes that the complete data follow a multivariate normal distribution. However, this condition is commonly not satisfied for social science data. Thus data transformations are necessary to obtain more accurate imputation results. For the crop ending stocks forecasts, we make the following transformations so that the transformed data follow a distribution closer to multivariate normal:

$$\text{Transformed Forecast} = \ln \frac{\text{Original Forecast}}{\text{Ending Stock}} \quad (3.15)$$

The transformed forecasts are thus the percentage deviations from the ending stocks.

The imputation method also assumes that the missing data is MAR. To make this assumption more plausible, we include two additional variables to increase the predictive power of the imputations.⁷ In particular, we include the average forecasts of all analysts, which serve as the information from other analysts' forecasts. We also include the USDA forecasts, which represent the information from another source – the government agency. Both included variables are complete, so that the package can fully utilize this additional information.

Based on Rubin (1987), only a few imputations are sufficient enough to generate valid and accurate imputed values. Typically, researchers choose 3 to 10 imputations for their datasets. In the present study, as there are at most 40% missing observations of the selected frequent analysts' forecasts, we generate 10 imputed datasets so that pooling estimates can be more stable. To aggregate the estimated results from the multiple imputed datasets, Rubin (1987) suggested a simple averaging and calculate the standard errors based on within-imputation variance and

⁷ The MAR assumption cannot be tested in this study as the test requires information from the missing part of the data (Schafer and Graham 2002). Thus we can only try to make the assumption more plausible.

between-imputation variance. However, it cannot be directly applied to the Bayesian MCMC simulations employed in the present study. Hence, we propose a method to directly use the simulated Markov Chains to compute the estimated mean values and the credible intervals.

Our integrated method of imputation and estimation is described as follows. To analyze an individual analyst's forecasts, we first transform the original data for imputations. Then, the imputed data are transformed back so that they can be used in the estimation process. We then run 3 Markov Chains for each of the 10 imputed datasets, using the iteration steps introduced in the previous section. The 3 chains have different starting values for each parameter, but have the same sets of starting values for each imputed dataset. The Gibbs Sampler is run for 40,000 iterations for each Markov Chain. The first half of each chain is then discarded as a burn-in period. We then combine the remaining half of the chains from all the imputed datasets with same starting values, creating 3 aggregated chains with 200,000 ($= 40,000 / 2 \text{ observations} \times 10 \text{ datasets}$) iterations each. Finally, we apply Gelman and Rubin (1992) test statistics to check the convergence of the aggregated chains. A graphical illustration of the process is depicted in Figure 3.4.

3.6 Results and Discussion

3.6.1 Representative Analysts and USDA Forecasts

Tables 3.3 – 3.6 summarize the estimation results of the parameters for the five sets of data. The means and standard deviations for the estimated coefficients, shocks and idiosyncratic errors are reported in Table 3.3 (corn), Table 3.4 (soybeans), and Table 3.5 (wheat). The sequence of standard deviations for the shocks is displayed in order of increasing forecast horizons. Table 3.6 reports the medians and 95% credible intervals of the constants and slopes

for the five sets of data for all three commodities. The Gelman and Rubin (1992) test statistics are below 1.03 for all parameters in all datasets, strongly suggesting convergence of the Markov Chains.

Corn

The point estimate of the intercept α represents the bias of the forecast revisions. For the AA, AM, FAA data, the estimates are all positive but insignificant. Thus we cannot reject the null hypothesis that the representative analysts' forecast revisions are unbiased. The estimates of α for the FAM data are positive and significant at 10% level, indicating that the FAM representative analyst slightly revises its forecasts up each month, and hence the revisions are weakly biased.

Coefficient β measures the association between two adjacent forecast revisions given the error covariance structure. The estimates for the slope β for the four sets of data are all positive and significant at 5% level. The estimated β is at around 16% for the AA and AM data. It is slightly lower, at 15.21% on average, for the FAA data. The estimate is much higher for the FAM data, at 20.29% on average. The results show that if the representative analyst adjusts its forecast up by 1% in the past month, its forecast will also be revised by 0.15% - 0.2%⁸ in the current month. Thus the findings show that analysts are conservative in adjusting their ending stocks forecasts for corn. In other words, the most recent forecast does not necessarily fully represent the arrival of new information. Instead, it can be viewed as a weighted average of the new information and the previous forecast.

⁸ The exact number depends on the dataset used.

The magnitudes of shocks are similar for the four representative analysts. Specifically, the results show that analysts typically exhibit relatively large shocks in the first two months. Besides, there is a jump in the magnitude of shock in the final month. The estimated standard deviations of the idiosyncratic errors are similar for the AA, AM, FAA data, at around 2.5% on average. It is larger, at 3.19% on average, for the FAM data. It is also worthwhile to note that the estimated standard deviations of the idiosyncratic errors are larger than those of the shocks of August/September (σ_2), providing evidence that the idiosyncratic errors are not negligible.

Soybeans

The estimates of the intercept α are quite similar for the four representative analysts. In particular, the estimates are all negative, at around -1.5%, and significant at the 5% level. This bias indicates that the representative analysts constantly revise their forecasts down by roughly 1.5% per month on average. Thus the analysts have a tendency to overestimate the soybeans ending stocks.

The point estimates of β are around 50% for all four analysts samples. They are all significant at the 5% level, serving as strong evidence that analysts' forecasts are inefficient. In other words, it can be argued that analysts' typically retain about 50% of the informational value of their own previous forecasts. As this magnitude is much higher than that for corn, analysts are said to be more conservative in adjusting their forecasts.

The estimates of the standard deviations of shocks are similar across the four analyst samples. The ranges are from roughly 4% to 31%. It can be found that shocks are expected to be large during September/October (σ_{13}) and during the final revisions. They are observed to be larger for soybeans than for corn. The estimates of the standard deviations of the idiosyncratic

errors are slightly more than 7% for all four analyst samples. They are also larger than those for corn forecasts.

Wheat

Results are quite different for analysts' forecasts of wheat ending stocks. The estimates of α are not significantly different from zero for the AA, AM, FAM data. Thus there isn't enough evidence that these representative analysts' forecasts are biased. However, for the FAA data, the estimate of α is -0.94% on average and significant at 5% level, indicating that the forecasts are biased downward. That is to say, the FAA representative analyst tends to overestimate the ending stocks, adjusting its forecasts down by about 1% each month.

The estimates of the slope β are all positive for the four datasets. For the AA and AM data, the point estimates are 21.16% and 24.63% respectively, both significantly greater than zero at 5% level. The findings indicate that analysts, as a group, are conservative in forecasting ending stocks. Thus their forecasts are inefficient. For the FAA and FAM data, although the estimates are positive, they are not significantly different from zero. Thus there is not enough evidence that the representative analysts are conservative in making forecasts.

The estimates of the standard deviations of shocks are similar across the four samples. The ranges are roughly 1% to 11%. Large shocks typically come in June/July (σ_{13}), July/August (σ_{12}), September/October (σ_{10}), and the final revisions. However, the sizes are much smaller compared to those for corn and soybeans. The estimate of the standard deviations of idiosyncratic errors is smallest for the AA data, at 0.89% on average, and largest for the FAM data, at 2% on average. The estimates for the AM and FAA data are similar, at around 1.3% on average.

Comparisons with USDA forecasts

The estimates for the USDA forecasts from marketing years 2004/05 to 2013/14 are also displayed in tables 3.3 – 3.6 for comparisons. For corn, the estimates of the constants and slopes are very close between the analysts' data and the USDA data. Specifically, we don't find enough evidence that USDA forecasts are biased at 5% significance level. The USDA forecasts are inefficient due to positive and significant association between adjacent forecasts revisions. There are some differences in the estimated standard deviations of the shocks, but in general they follow the same patterns. The estimated standard deviations of the USDA idiosyncratic errors are smaller, at 2.06% on average. In other words, USDA are more precise in making ending stocks forecasts than the private analysts.

Similar conclusions apply to the USDA and analysts' soybeans forecasts. In particular, the estimates show that USDA forecasts are also biased downward. The USDA tends to be conservative in adjusting its forecasts, as the estimated slope coefficient is at 56.32% on average, which is positive and significant at the 5% level. The estimates for the standard deviations of the shocks are generally close for the two types of forecasters. And the estimated standard deviation of idiosyncratic errors of the USDA forecasts is at 6.36% on average, about 1% smaller than analysts' forecasts.

Results are substantially different for wheat forecasts. There is not enough evidence that USDA forecasts are inefficient. However, the representative analysts' forecasts in the AA, AM, FAA data are all found to be inefficient. Although efficiency cannot be rejected for the FAM representative analysts' forecasts, there exist some differences in the average of the estimated coefficients and parameters in the error covariance matrix.

3.6.2 Individual Analysts' Forecasts

Figure 3.5 – 3.7 show graphical results of the estimated intercepts (α) and slopes (β) for the selected frequent analysts' forecasts, as well as the 95% credible intervals of the estimates.⁹ The Gelman and Rubin test statistics are all below 1.03, strongly suggesting the convergence of the aggregated Markov Chains.

For corn forecasts, it can be seen that the estimated intercepts are not significantly different from zero for any of the 10 analysts' forecasts. Thus, we cannot find bias in the individual analysts' forecasts. The situation is different for the estimated slopes, as 2 of them are positive and significant at the 5% level, indicating inefficiencies for the forecasts from these two analysts. Results are different for soybeans forecasts. Estimated intercepts for 9 analysts are found negative and significant at the 5% level. We also find that the estimated slopes for 7 analysts are positive and significant at the 5% level, and one slope is positive and significant at the 10% level. For wheat forecasts, there is only one analyst who has negative and significant intercept. The estimated slopes for the 9 analysts are all insignificant.

Results show that there exists diversity in analysts' forecasting behaviors. This diversity can be attributed to a number of factors, such as the data sources and the forecasting models. It can also come from the subjective opinions of the analysts. We can divide these analysts into two distinct groups based on how close their forecasting patterns are to that of the USDA. Take corn for example. There are 8 analysts whose forecasts are found efficient. The remaining 2 analysts are found inefficient in forecasting the ending stocks. However, these 2 analysts have similar forecast behaviors as the USDA. For soybeans, there are 7 analysts whose forecasting behaviors

⁹ The estimated parameters for the error covariance matrix – the standard deviations of unforecastable shocks (σ_n^2) and the idiosyncratic errors (σ^2) - are omitted to save space.

are close to the USDA. And for wheat, the number is 8. This finding shows that some analysts are able to capture the forecasting behavior of the USDA, because their forecasts are typically released before the USDA does. More subjectively, it is possible that these analysts have incorporated the USDA forecasts into their forecasting models. Thus, they may try to make forecasts as close as possible to the upcoming USDA forecasts. These arguments, however, cannot be applied to the analysts who are found to be efficient in forecasting the ending stocks but are different from the USDA. However, the forecasts from these analysts can be considered “good” forecasts because they fully reveal the new information, and the forecasts cannot be predicted from their own history.

3.6.3 Discussion

Occurrence of shocks

As discussed in the previous section, for all three commodities, the estimated standard deviations of shocks exhibit differences between private analysts’ and USDA forecasts. A natural question arises is that: as shocks are from the outside environments and are common information to both groups of forecasters, why do the differences exist in their estimated magnitudes? The answers to this question can be separated into two types of explanations: the objective and the subjective.

Objective explanations come from the time schedule of these forecasts. Note that USDA typically publishes its forecasts during the 9th – 14th of each month. Private analysts’ forecasts are finalized and collected several days before the release of the USDA forecasts. Graphical illustrations of the forecast timings are presented in Figure 3.8 for corn and soybeans and Figure 3.9 for wheat. The time interval between a consensus analysts’ forecast and its corresponding

USDA forecast is almost one week, which is not negligible compared to the time interval between two own forecasts. Thus, if there happen to be some disturbances during the interval between analysts' and USDA forecasts, these disturbances can be observed differently by the analysts and the USDA. Specifically, for the private analysts, the disturbances would be allocated to the next shock because analysts have already made their forecasts. In contrast, the disturbances for the USDA are allocated to the current shock, as the USDA has not finalized its forecasts yet. In this way, the different allocations of such disturbances result in different observations of shocks. Thus, the estimated standard deviations of these shocks could be different for private analysts and the USDA.

Subjective explanations come from the measurements of the shocks. We can assume, without loss of generality, that private analysts and the USDA do not necessarily have the same measurements when facing the same shocks. There are two reasons which can possibly lead to this result. First, the data sources can be different for the two groups of forecasters. For example, in early months where production has not been finalized, the surveys conducted by the USDA and private analysts may cover different subjects. As the shock may affect these farmers differently, the observance of the shocks will also be different. Besides, the survey sample size matters as well, because the calculated population characteristics are highly related to the sample characteristics. Thus, the estimated shocks will be different for the two groups of forecasters if different survey sample sizes have been employed. Secondly, this finding can also result from different models or mathematical methods applied by these forecasters. Examples include, but also not limited to, quantifications of deviations of weather benchmarks, and estimations of a sudden demand shock. It is worthwhile to note that these disagreements about shocks are not the

forecasters' own idiosyncratic errors. In the present study, the latter are separated based on the forecaster's corrections of such errors.

Forecasting Behaviors

Results show that the USDA is inefficient in forecasting the ending stocks for corn and soybeans. But for wheat, there isn't enough evidence to reject the null hypothesis of efficiency for USDA forecasts. The results also show that private analysts, as a group, have similar forecasting behaviors as the USDA for corn and soybeans, and distinct forecasting behaviors for wheat. These results are interesting when we also take into account the findings in Chapter 2 that the USDA is conservative in making ending stocks forecasts for these three commodities during a larger time span of 29 marketing years. Different from other fixed-event forecasts, ending stocks forecasts are inherently more subjective because they are combinations of various components in both the demand and supply side. Thus it might not be enough to provide a good forecast by only considering the deterministic components, *i.e.*, the market conditions or other outside information. Instead, forecasters will likely add additional autoregressive terms to partake some weights in their forecasting process. For example, past forecasts can serve as prior means to current forecasts. This helps "stabilize" the forecasts, resulting in the positive estimated slopes in our study. As Litterman (1986) pointed out in his research, the inclusion of autoregression can provide forecasts which are "as accurate, on average, as those used by the best known commercial forecasting services". This method has also been advanced in the past few decades and is widely adopted in generating forecasts. It is also interesting to look at the wheat ending stocks forecasts. The USDA might have "improved" its forecasts, so that although the estimated slope is positive, it is no longer significant for the more recent 10 years data, while

in Chapter 2, the estimate is positive and significant for 29 years data. However, it is possible that private analysts, as a group, did not detect the change in USDA forecasting behaviors.

Therefore, analysts ending stocks for wheat, like for corn and soybeans, are still conservative.

3.7 Conclusions

We applied the model in Chapter 2 to investigate the efficiency of private analysts' ending stocks forecasts for three major agricultural commodities: corn, soybeans and wheat. Private analysts' forecast data for marketing years 2004/05 – 2013/14 are analyzed and the USDA forecasts of the same period are included for comparisons. The model incorporates the most recent revisions as explanatory variables. Besides, it builds an error covariance matrix to distinguish unforecastable outside shocks and the forecaster's own idiosyncratic errors. The assumption on the monthly shocks is in line with the fact that forecast errors decrease as forecast horizons shorten. The introduction of forecasters' corrections of their own errors serves as the link between forecasts and the forecast targets – the ending stocks.

Private analysts' forecast data are grouped into four representatives: the average of analysts' forecasts, the median of analysts' forecasts, the average of selected frequent analysts' forecasts, and the median of selected frequent analysts' forecasts. Estimation results show that private analysts' forecasts, as a group, are inefficient for all three commodities. In particular, there is strong evidence that they are conservative in forecasting. The same results also apply to the selected frequent analysts' forecasts for corn and soybeans. However, results are different for wheat as the estimated slopes are not significantly different from zero. Besides, results indicate that there are differences in the magnitudes of unforecastable shocks. The idiosyncratic errors of

private analysts' forecasts are not negligible because the estimated standard deviations are significantly greater than zero, and larger than some of the outside shocks.

Estimations of the USDA forecasts show that the USDA is conservative in making ending stocks forecasts for corn and soybeans, indicating inefficiency in its forecasts. The estimated intercepts and slopes are similar to those of the private analysts for these two commodities. However, we cannot find enough evidence that USDA wheat forecasts are inefficient.

We also apply our model to the individual analysts' forecasts. We use a multiple imputation method to fill the missing data points in the selected frequent analysts' forecasts. Then we develop a method to use the imputed data for the MCMC estimations. Results show that there exists diversity across individual analysts' forecasts. The analysts can be basically categorized into two groups, with one group exhibiting forecasting behavior similar to the USDA.

An interesting finding in the present study is the similarity of the forecasting behaviors of the representative analysts and the USDA for corn and soybeans. A natural question to ask is: why are their forecasts so close to each other? We cannot deny that the private analysts and the USDA may share a substantial amount of information. Remember that, as discussed at the beginning, private analysts are market participants whose objective in the market might be different from the USDA. Thus, given that analysts release their forecasts first, it is possible that they try to make forecasts as close as possible to the USDA forecasts. This guess, however, is currently unexplored for private analysts' forecasts, and is to be further investigated in future research.

3.8 References

- Arbuckle, J.L. 1996 Full Information Estimation in the Presence of Incomplete Data. In Marcoulides, G.A. and Schumacker, R.E. [Eds.] *Advanced structural equation modeling: Issues and Techniques*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Adjemian, M.K. 2012. Quantifying the WASDE Announcement Effect. *American Journal of Agricultural Economics* 94(1): 238-256.
- Bauer, R.F., and P.F. Orazem 1994. The Rationality and Price Effects of U.S. Department of Agriculture Forecasts of Oranges. *Journal of finance* 49(2): 681-685.
- Botto, A.C., O. Isengildina, S.H. Irwin, D.L. Good. 2006. Accuracy Trends and Sources of Forecast Errors in WASDE Balance Sheet Categories for Corn and Soybeans. Selected Paper prepared for presentation at the *American Agricultural Economics Association Annual Meeting, Long Beach, California, July 23-26, 2006*.
- Clements, M.P. 1995. Rationality and the Role of Judgement in Macroeconomic Forecasting. *The Economic Journal*. 105(429): 410-420.
- Clements, M.P. 1997. Evaluating the Rationality of Fixed-event Forecasts. *Journal of Forecasting*. 16:225-239.
- Clements, M.P., F. Joutz, and H.O. Stekler. 2007. An Evaluation of the Forecasts of the Federal Reserve: a Pooled Approach. *Journal of Applied Econometrics*. 22: 121-136.
- Davies, A., and K.Lahiri. 1995. A New Framework for Testing Rationality and Measuring Aggregate Shocks Using Panel Data. *Journal of Econometrics* 68: 205-227.
- Davies, A., and K.Lahiri. 1999. Re-examining the Rational Expectations Hypothesis Using Panel Data on Multi-period Forecasts. *Analysis of Panels and Limited Dependent Variable Models*, 226-254. Cambridge University Press.
- Egelkraut, T.M., P. Garcia, S.H. Irwin, and D.L. Good. 2003. An Evaluation of Crop Forecast Accuracy for Corn and Soybeans: USDA and Private Information Agencies. *Journal of Agricultural and Applied Economics* 35:79-95.
- Fortenberry, T.R., and D.A. Sumner. 1993. The Effects of USDA Reports in Futures and Options Markets. *Journal of Futures Markets* 13:157-173.
- French, K.R., R. Leftwich and N. Uhrig (1989). The Effects of Scheduled Announcements on Futures Markets. Working Paper 373, *Center for Research in Security Prices, Graduate School of Business, University of Chicago*.
- Garcia, P., S. Irwin, R. Leuthold, and L. Yang. 1997. The Value of Public Information in Commodity Futures Markets. *Journal of Economic Behavior and Organization* 32(4): 559-570.

- Gelman, A. 2006. Prior Distributions for Variance Parameters in Hierarchical Models (Comment on the Article by Browne and Draper). *Bayesian Analysis* 1(3):515-534.
- Gelman A. and D.B. Rubin. 1992. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*. 7(4):457-472.
- Harvey, D.I., Leybourne, S.J., and Newbold, P. 2001. Analysis of a Panel of UK Macroeconomic Forecasts. *Econometrics Journal*. 4: 37-55.
- Honaker, J. and G. King. 2010. What to Do about Missing Values in Time-Series Cross-Section Data. *American Journal of Political Science* 54(2): 561–581.
- Honaker, J., G. King, and M. Blackwell. 2014. AMELIA II: A Program for Missing Data. <https://cran.r-project.org/web/packages/Amelia/vignettes/amelia.pdf>.
- Isengildina O., S.H. Irwin, and D.L. Good. 2006. Are Revisions to USDA crop Production Forecasts smoothed. *American Journal of Agricultural Economics*. 88(4): 1091-1104.
- Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008a. The Impact of Situation and Outlook Information in Corn and Soybeans Futures Markets: Evidence from WASDE Reports. *Journal of Agricultural and Applied Economics* 40(1):89-103.
- Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008b. Impact of WASDE Reports on Implied Volatility in Corn and Soybeans Markets. *Agribusiness* 24(4): 473-490.
- Isengildina-Massa, O., B. Karali, and S.H. Irwin. 2013. When do USDA forecasters make mistakes? *Applied Economics* 45(36): 5086–5103.
- Litterman, R.B. 1986. Forecasting With Bayesian Vector Autoregressions - Five Years of Experience. *Journal of Business & Economic Statistics* 4(1):25-38.
- 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.
- Nordhaus , W.D. 1987. Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics* 69:667-674.
- Romer, C.D., and D.H. Romer. 2000. Federal Reserve Information and the Behavior of Interest rates, *American Economic Review*. 90(3): 429-457.
- Rubin, D.B. 1976. Inference and Missing Data. *Biometrika* 63(3): 581-592.
- Rubin, D.B. *Multiple Imputation for Nonresponse in Surveys*. New York: John Wiley.

- Sanders, D.R., and M.R. Manfredo. 2002. USDA Production Forecasts for Pork, Beef, and Broilers: An Evaluation. *Journal of Agricultural and Resource Economics*. 27:114-128.
- Schafer, J.L., and Graham, J.W. 2002. Missing Data: Our View of the State of the Art. *Psychological Methods* 7: 147-177.
- Wothke, W. 2000. Longitudinal and Multigroup Modeling with Missing Data. In Little, T.D., Schnabel, K.U., and Baumert, J. [Eds.]. *Modeling Longitudinal and Multilevel Data: Practical Issues, Applied Approaches, and Specific Examples*. Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers. 219-240, 269-281.

3.9 Appendix

Conditional Posterior Distributions for Model Parameters in the Gibbs Sampler

The proposed model consists of regression (3.4) and priors (3.14):

$$Y = X\beta + WK + P\Sigma$$

$$\beta \sim N(M, V)$$

$$\varepsilon_{t,n} \sim N(0, \sigma^2) \quad (3.A.1)$$

$$k_{t,n} \sim N(0, \sigma_n^2)$$

$$\sigma, \sigma_n \sim \text{Uniform}(0, \infty)$$

for $n = 1, \dots, N, t = 1, \dots, T$. Let $\Omega \equiv P'P$. Given $\{\beta, K, \sigma^2, \Omega\}$, the dependent variable $y_{t,n}$

follows a multivariate normal distribution:

$$Y | \beta, K, \sigma^2 \Omega \sim N(X\beta + WK, \sigma^2 \Omega) \quad (3.A.2)$$

and the likelihood is $\Phi(Y | \beta, K, \sigma^2 \Omega)$. The posterior density of the set of model parameters is given by

$$(\Lambda) = \Phi(Y | \beta, K, \sigma^2 \Omega) \prod_{jt} \Phi(k_{t,j} | \sigma_j^2) * p(\beta) p(\sigma^2) \prod_{j=1}^{N-1} p(\sigma_j^2) \quad (3.A.3)$$

The conditional posterior density for β is

$$p(\boldsymbol{\beta}|\Lambda\backslash\boldsymbol{\beta}) = \Phi(Y|\boldsymbol{\beta}, \mathbf{K}, \sigma^2\boldsymbol{\Omega}) * \Phi(\boldsymbol{\beta}|\mathbf{M}, \mathbf{V}) \quad (3.A.4)$$

Hence:

$$\begin{aligned} \boldsymbol{\beta}|\Lambda\backslash\boldsymbol{\beta} \sim N((\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X}/\sigma^2 + \mathbf{V}^{-1})^{-1}(\mathbf{X}'\boldsymbol{\Omega}^{-1}(\mathbf{Y} - \mathbf{W}\mathbf{K})/\sigma^2 \\ + \mathbf{V}^{-1}\mathbf{M}), (\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X}/\sigma^2 + \mathbf{V}^{-1})^{-1}) \end{aligned} \quad (3.A.5)$$

The conditional posterior density of $k_{t,n}$, $t = 1, \dots, T$, $n = 1, \dots, N$ is

$$p(k_{t,n}|\Lambda\backslash k_{t,n}) = \Phi(Y|\boldsymbol{\beta}, \mathbf{K}, \sigma^2\boldsymbol{\Omega}) * \Phi(k_{t,n}|\sigma_n^2) \quad (3.A.6)$$

Therefore,

$$k_{t,n}|\Lambda \backslash k_{t,n} \sim N\left(\frac{\mathbf{W}'_{k_{t,n}}\boldsymbol{\Omega}^{-1}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}_{-k_{t,n}}\mathbf{K}_{-k_{t,n}})/\sigma^2}{\mathbf{W}'_{k_{t,n}}\boldsymbol{\Omega}^{-1}\mathbf{W}_{k_{t,n}}/\sigma^2 + 1/\sigma_n^2}, \frac{1}{\mathbf{W}'_{k_{t,n}}\boldsymbol{\Omega}^{-1}\mathbf{W}_{k_{t,n}}/\sigma^2 + 1/\sigma_n^2}\right) \quad (3.A.7)$$

where $\mathbf{W}_{k_{t,n}}$ is the column of \mathbf{W} which indicates the monthly shock $k_{t,n}$, and $\mathbf{W}_{-k_{t,n}}$, $\mathbf{K}_{-k_{t,n}}$ are matrices with the column indicating $k_{t,n}$ deleted from \mathbf{W} , \mathbf{K} , respectively.

The conditional posterior density of σ_n^2 , $n = 1, \dots, N$ is

$$p(\sigma_n^2|\Lambda\backslash\sigma_n^2) = \prod_{j=1}^N \Phi(k_{t,n}|\sigma_n^2) * p(\sigma_n^2) \quad (3.A.8)$$

Thus

$$\sigma_n^2|\Lambda\backslash\sigma_n^2 \sim IG((T-1)/2, \sum_{t=1}^T k_{t,n}^2/2) \quad (3.A.9)$$

Finally, the conditional posterior of σ^2 is

$$p(\sigma^2|\Lambda\backslash\sigma^2) = \Phi(Y|\boldsymbol{\beta}, \mathbf{K}, \sigma^2\boldsymbol{\Omega}) * p(\sigma^2) \quad (3.A.10)$$

so that

$$\sigma^2|\Lambda\backslash\sigma^2 \sim IG((TN-1)/2, (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}\mathbf{K})'\boldsymbol{\Omega}^{-1}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}\mathbf{K})/2) \quad (3.A.11)$$




Figure 3.1. Monthly standard deviations of forecast revisions for corn




Figure 3.2. Monthly standard deviations of forecast revisions for soybeans




Figure 3.3. Monthly standard deviations of forecast revisions for wheat

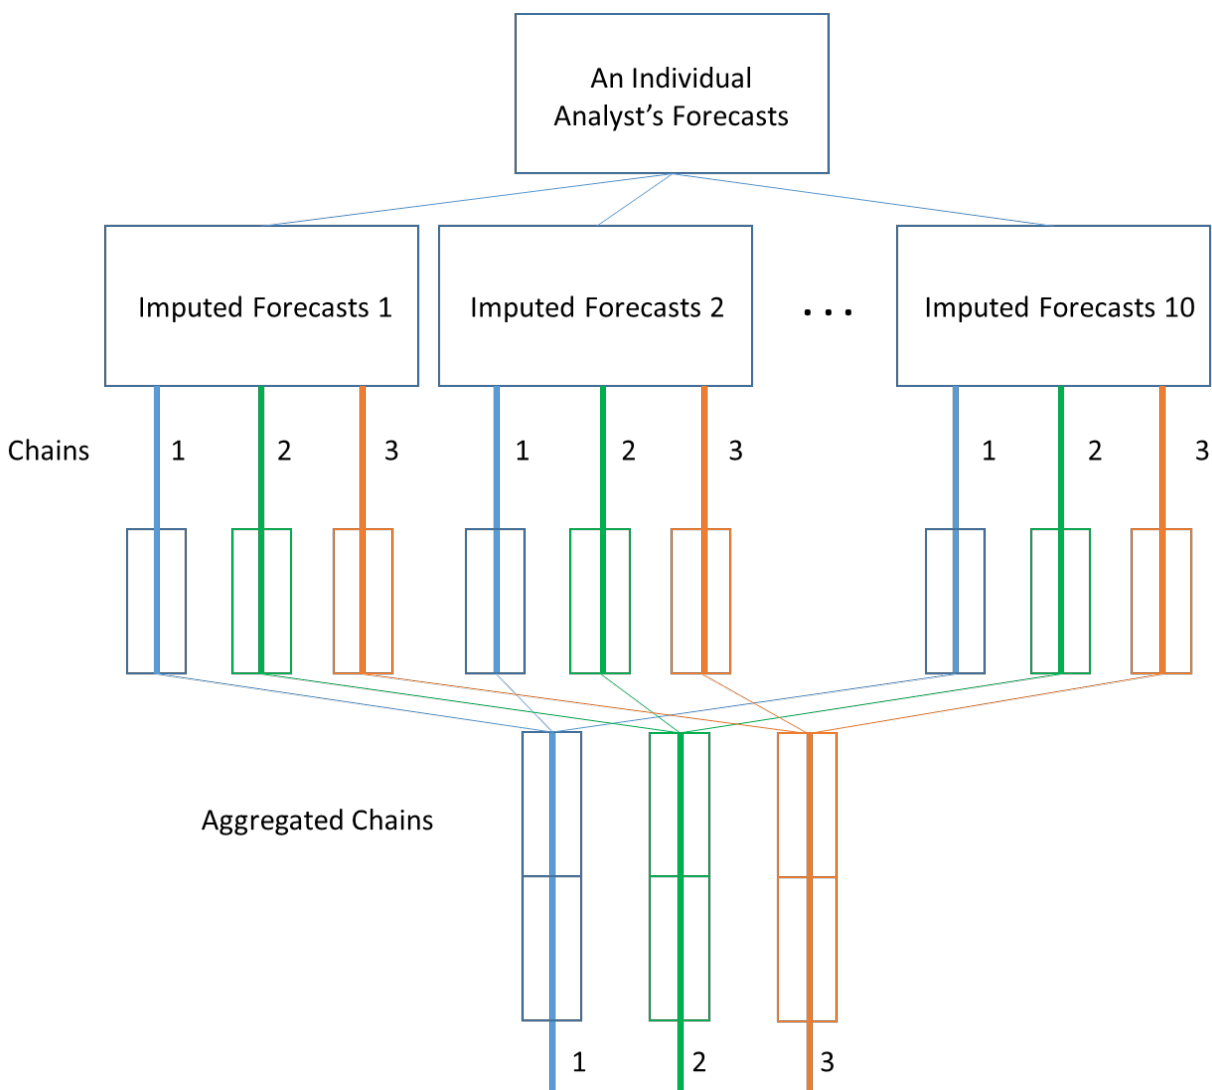


Figure 3.4. The integrated procedure of multiple imputations and MCMC estimations

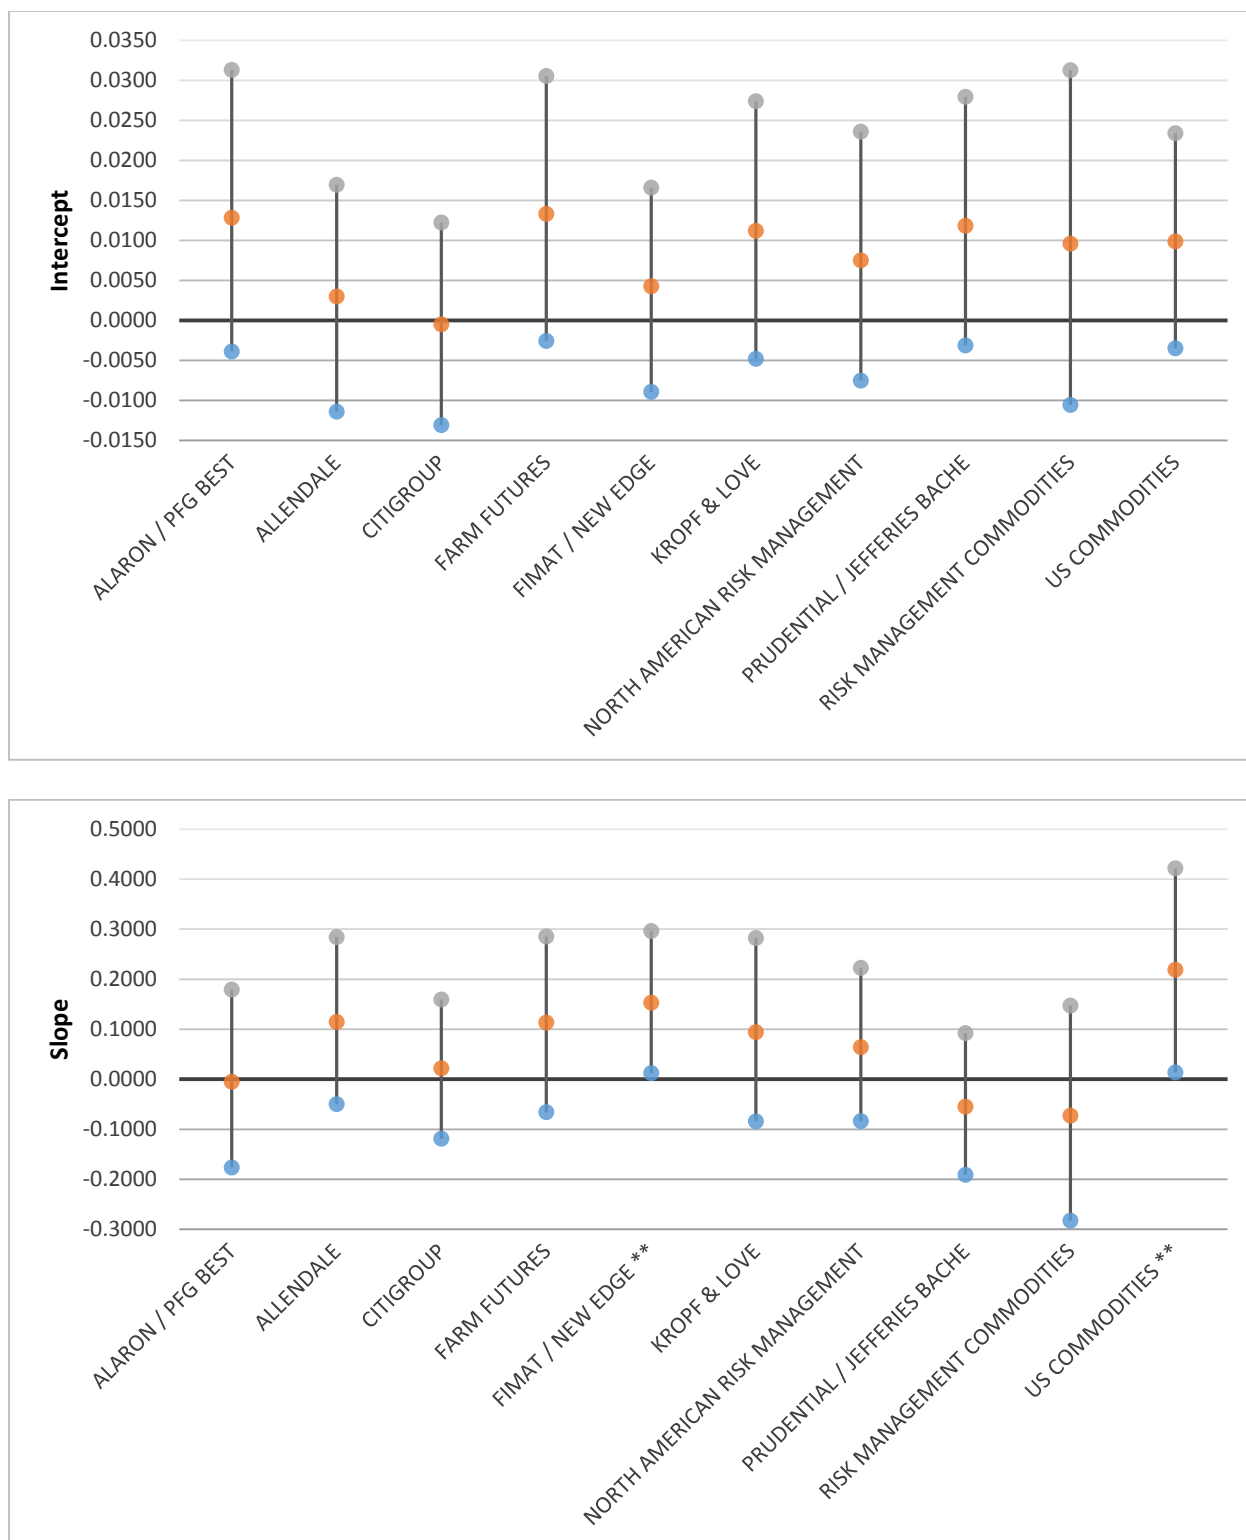


Figure 3.5. Estimated coefficients for individual analysts' forecasts for corn ending stocks

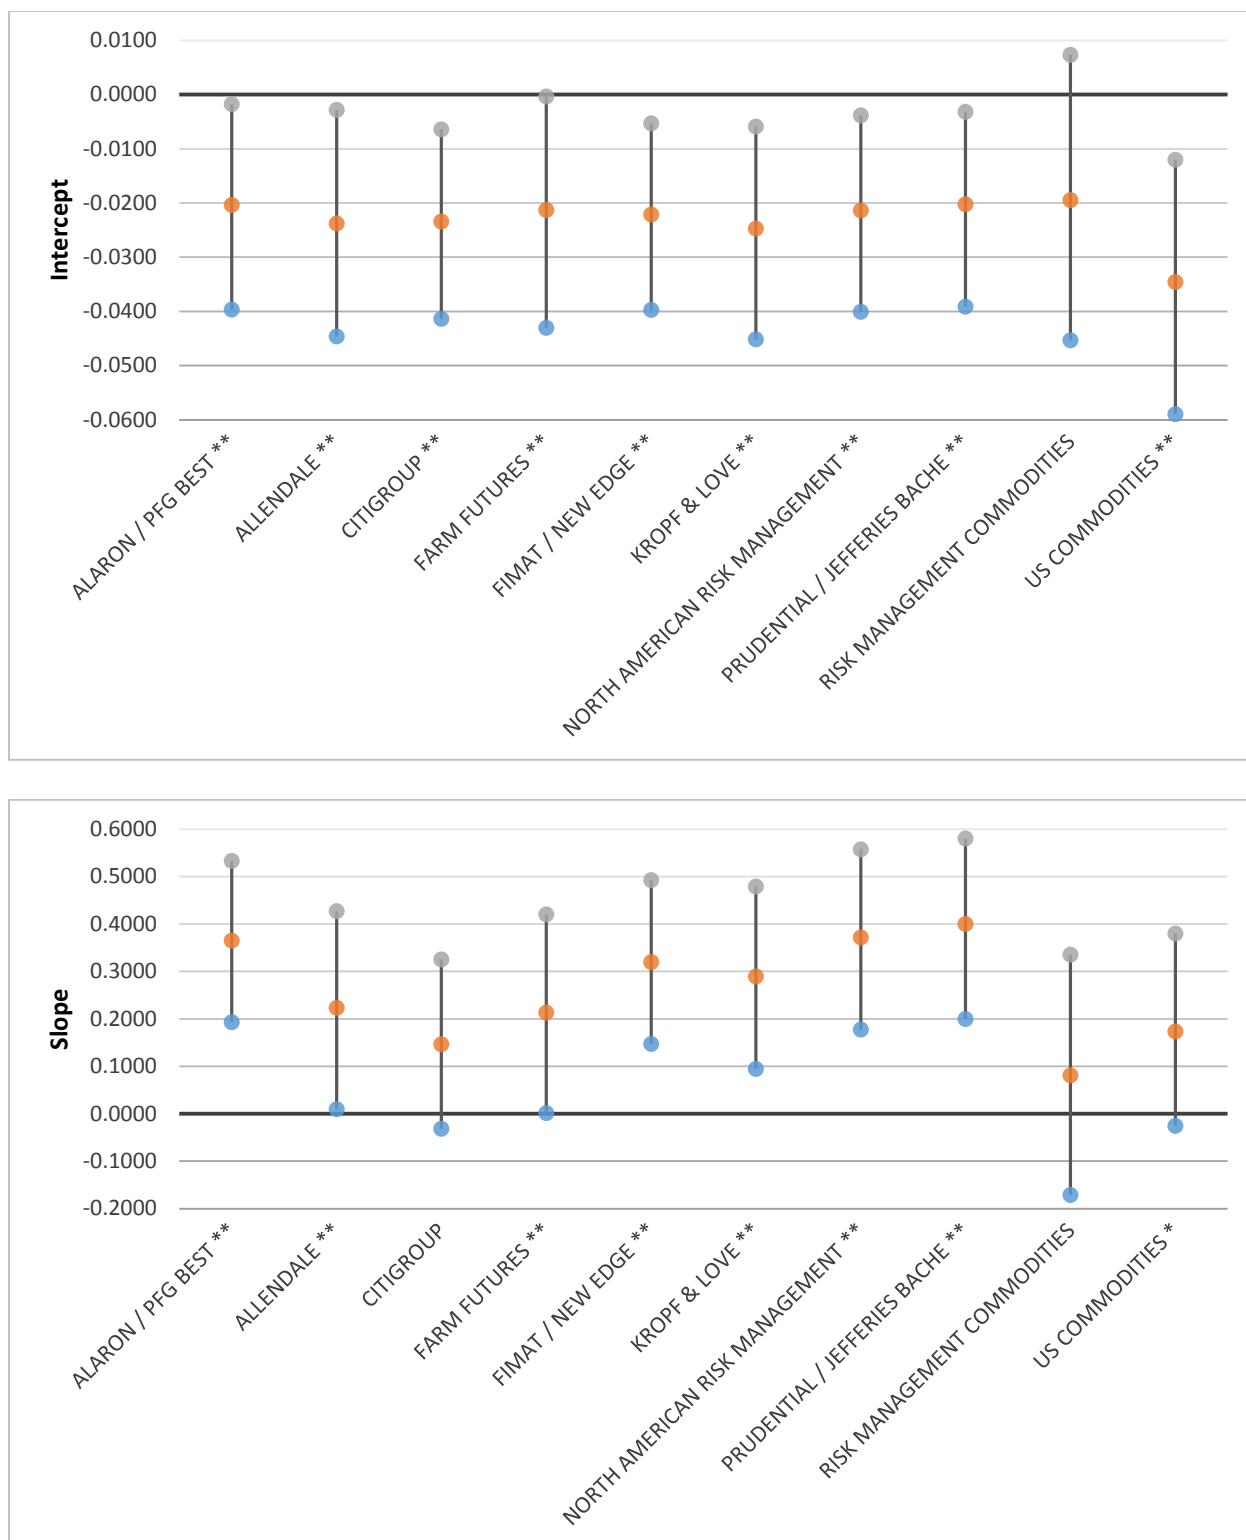


Figure 3.6. Estimated coefficients for individual analysts' forecasts for soybeans ending stocks

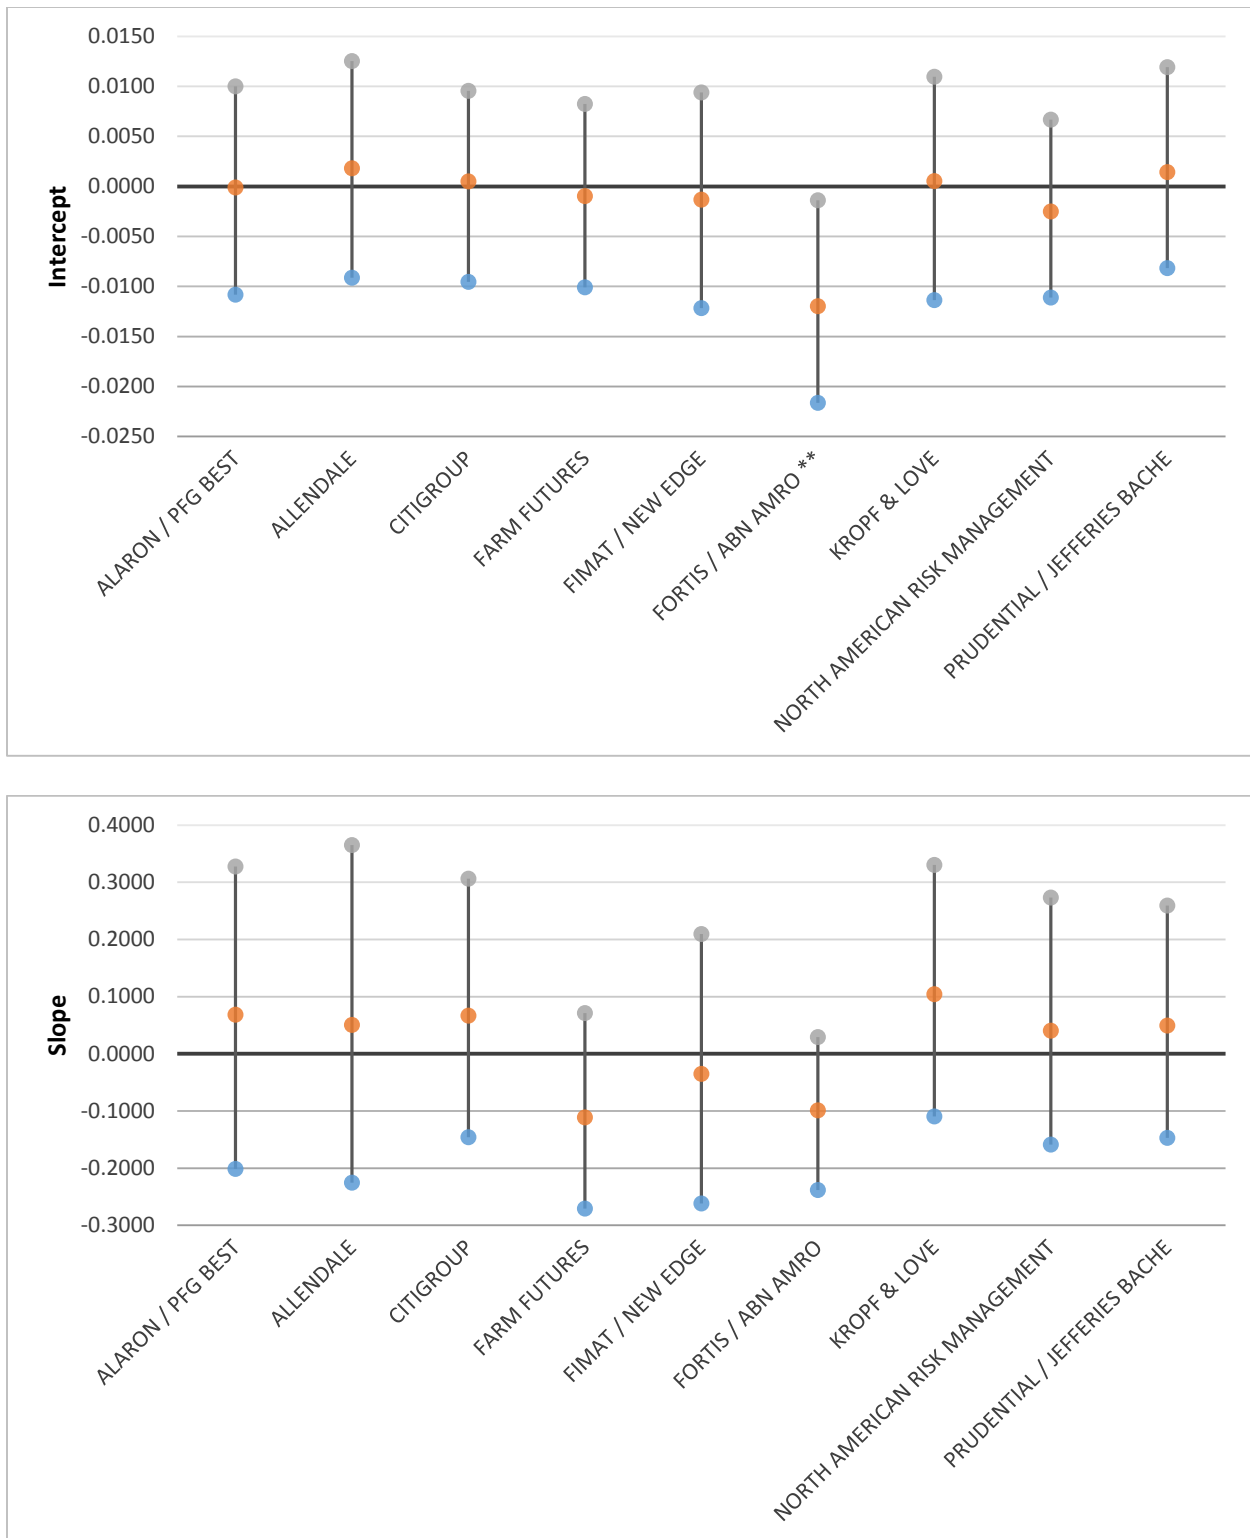


Figure 3.7. Estimated coefficients for individual analysts' forecasts for wheat ending stocks




Figure 3.8. Forecast timing for corn/soybean ending stocks




Figure 3.9. Forecast timing for wheat ending stocks

Table 3.1. List of the selected frequent analysts

Analyst by Name	Corn		Soybeans		Wheat	
	# of obs.	% of obs.	# of obs.	% of obs.	# of obs.	% of obs.
Alaron / PFG Best	126	74.12%	126	74.12%	107	76.43%
Allendale	161	94.71%	161	94.71%	134	95.71%
Citigroup	163	95.88%	164	96.47%	133	95.00%
Farm Futures	115	67.65%	115	67.65%	91	65.00%
Fimat / New Edge	167	98.24%	167	98.24%	137	97.86%
Fortis / ABN Amro	-	-	-	-	86	64.13%
Kropf & Love	130	76.47%	132	77.65%	110	78.57%
North American Risk Management	128	75.29%	128	75.29%	106	75.71%
Prudential / Jefferies Bache	168	98.82%	161	94.71%	136	97.14%
Risk Management Commodities	119	70.00%	118	69.41%	-	-
US Commodities	150	88.24%	149	87.65%	-	-

Note: The total number of observations for corn and soybeans is 170. The total number of observations for wheat is 140. The frequent analysts are defined as those who have more than 60% of observations.

Table 3.2. Descriptive statistics for USDA and analysts' forecast revisions for corn, soybeans and wheat

Commodity	Datasets	Mean	Median	St. Dev.	Min	Max
Corn	Analysts Avg.	0.0080	0.0034	0.1271	-0.6372	0.3805
	Analysts Med.	0.0080	0.0066	0.1294	-0.5967	0.4260
	Selected Freq. Ana. Avg.	0.0083	0.0148	0.1281	-0.6706	0.3783
	Selected Freq. Ana. Med.	0.0087	0.0079	0.1293	-0.6190	0.4191
	USDA	0.0097	0.0000	0.1249	-0.5988	0.4098
Soybeans	Analysts Avg.	-0.0132	-0.0123	0.1388	-0.4953	0.5821
	Analysts Med.	-0.0117	-0.0082	0.1405	-0.4533	0.6747
	Selected Freq. Ana. Avg.	-0.0128	-0.0206	0.1407	-0.6024	0.5308
	Selected Freq. Ana. Med.	-0.0118	-0.0151	0.1445	-0.6133	0.5457
	USDA	-0.0132	0.0000	0.1387	-0.3830	0.7569
Wheat	Analysts Avg.	0.0021	-0.0007	0.0564	-0.2140	0.2409
	Analysts Med.	0.0015	0.0000	0.0571	-0.1989	0.2471
	Selected Freq. Ana. Avg.	0.0019	-0.0013	0.0592	-0.2392	0.2429
	Selected Freq. Ana. Med.	0.0019	0.0000	0.0590	-0.2189	0.2471
	USDA	0.0023	0.0000	0.0590	-0.1648	0.1863

Note: summary statistics are displayed in logarithmic values.

Table 3.3. Parameter estimates for the ending stocks forecasts for corn, 2004/05 – 2013/14

Parameter	USDA		Analysts Avg.		Analysts Med.		Selected Ana. Avg.		Selected Ana. Med.	
	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)
Coefficient										
α (intercept)	0.0041	(0.0051)	0.0073	(0.0052)	0.0075	(0.0048)	0.0088	(0.0056)	0.0096	(0.0055)*
β (slope)	0.1675	(0.0549)**	0.1646	(0.0640)**	0.1601	(0.0672)**	0.1521	(0.0684)**	0.2029	(0.0760)**
Idiosyncratic Err.										
σ	0.0206	(0.0043)	0.0257	(0.0058)	0.0267	(0.0055)	0.0246	(0.0057)	0.0319	(0.0068)
Shock										
σ_1	0.1637	(0.0494)	0.1169	(0.0356)	0.1163	(0.0354)	0.1163	(0.0356)	0.1193	(0.0367)
σ_2	0.0559	(0.0194)	0.0173	(0.0127)	0.0189	(0.0137)	0.0133	(0.0109)	0.0219	(0.0157)
σ_3	0.0464	(0.0188)	0.0438	(0.0205)	0.0426	(0.0204)	0.0358	(0.0195)	0.0354	(0.0212)
σ_4	0.1146	(0.0330)	0.1188	(0.0352)	0.1189	(0.0351)	0.1099	(0.0327)	0.1039	(0.0333)
σ_5	0.0213	(0.0139)	0.0752	(0.0265)	0.0750	(0.0265)	0.0833	(0.0277)	0.0826	(0.0302)
σ_6	0.0843	(0.0254)	0.0703	(0.0258)	0.0880	(0.0296)	0.0684	(0.0249)	0.0735	(0.0297)
σ_7	0.1052	(0.0303)	0.1134	(0.0342)	0.1380	(0.0405)	0.1038	(0.0317)	0.1036	(0.0344)
σ_8	0.0135	(0.0104)	0.0437	(0.0216)	0.0442	(0.0218)	0.0579	(0.0231)	0.0536	(0.0264)
σ_9	0.0509	(0.0195)	0.1302	(0.0386)	0.1312	(0.0394)	0.1392	(0.0406)	0.1364	(0.0415)
σ_{10}	0.1396	(0.0399)	0.0893	(0.0282)	0.1028	(0.0314)	0.0865	(0.0276)	0.0909	(0.0307)
σ_{11}	0.1003	(0.0294)	0.0413	(0.0207)	0.0343	(0.0201)	0.0572	(0.0229)	0.0511	(0.0258)
σ_{12}	0.0496	(0.0194)	0.1617	(0.0459)	0.1513	(0.0435)	0.1805	(0.0509)	0.1560	(0.0459)
σ_{13}	0.2163	(0.0597)	0.1430	(0.0417)	0.1360	(0.0401)	0.1407	(0.0408)	0.1440	(0.0436)
σ_{14}	0.0984	(0.0328)	0.1044	(0.0329)	0.1092	(0.0341)	0.1173	(0.0352)	0.1182	(0.0378)
σ_{15}	0.2724	(0.0770)	0.2965	(0.0836)	0.2776	(0.0784)	0.3074	(0.0855)	0.2909	(0.0822)
σ_{16}	0.2713	(0.0756)	0.2529	(0.0709)	0.2835	(0.0801)	0.2327	(0.0653)	0.2641	(0.0755)

Note: (*) and (**) denote parameter estimates significant at 10% and 5%, respectively. The standard errors for the idiosyncratic error and shocks are all significant at 5% level, and hence the indicators are omitted.

Table 3.4. Parameter estimates for the ending stocks forecasts for soybeans, 2004/05 – 2013/14

Parameter	USDA		Analysts Avg.		Analysts Med.		Selected Ana. Avg.		Selected Ana. Med.	
	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)
Coefficient										
α (intercept)	-0.0176	(0.0060)**	-0.0156	(0.0067)**	-0.0155	(0.0068)**	-0.0155	(0.0064)**	-0.0157	(0.0067)**
β (slope)	0.5632	(0.0675)**	0.5125	(0.0747)**	0.5011	(0.0792)**	0.5014	(0.0746)**	0.5231	(0.0782)**
Idiosyncratic Err.										
σ	0.0636	(0.0065)	0.0735	(0.0069)	0.0773	(0.0073)	0.0713	(0.0068)	0.0760	(0.0072)
Shock										
σ_1	0.2383	(0.0740)	0.2372	(0.0744)	0.2290	(0.0727)	0.2255	(0.0717)	0.2345	(0.0742)
σ_2	0.1331	(0.0415)	0.1451	(0.0458)	0.1511	(0.0481)	0.1371	(0.0441)	0.1467	(0.0473)
σ_3	0.0401	(0.0289)	0.0413	(0.0303)	0.0473	(0.0329)	0.0481	(0.0324)	0.0612	(0.0392)
σ_4	0.0429	(0.0291)	0.0451	(0.0312)	0.0514	(0.0335)	0.0405	(0.0296)	0.0414	(0.0309)
σ_5	0.0298	(0.0246)	0.0535	(0.0381)	0.0543	(0.0396)	0.0447	(0.0344)	0.0503	(0.0379)
σ_6	0.1139	(0.0451)	0.0956	(0.0484)	0.0883	(0.0487)	0.1058	(0.0467)	0.0822	(0.0473)
σ_7	0.0578	(0.0366)	0.0627	(0.0400)	0.0635	(0.0412)	0.0540	(0.0377)	0.0629	(0.0405)
σ_8	0.0341	(0.0268)	0.0367	(0.0289)	0.0354	(0.0278)	0.0432	(0.0319)	0.0412	(0.0314)
σ_9	0.0435	(0.0302)	0.0401	(0.0310)	0.0403	(0.0315)	0.0403	(0.0308)	0.0413	(0.0318)
σ_{10}	0.0471	(0.0342)	0.0450	(0.0343)	0.0526	(0.0374)	0.0438	(0.0339)	0.0518	(0.0370)
σ_{11}	0.0662	(0.0399)	0.1117	(0.0523)	0.1068	(0.0557)	0.1142	(0.0524)	0.0980	(0.0529)
σ_{12}	0.1756	(0.0586)	0.1330	(0.0602)	0.1275	(0.0623)	0.1635	(0.0599)	0.1600	(0.0626)
σ_{13}	0.3215	(0.0937)	0.3260	(0.0965)	0.3450	(0.1026)	0.3092	(0.0924)	0.3051	(0.0935)
σ_{14}	0.1273	(0.0609)	0.1641	(0.0679)	0.1583	(0.0700)	0.1949	(0.0708)	0.1937	(0.0747)
σ_{15}	0.1256	(0.0531)	0.0664	(0.0454)	0.0850	(0.0549)	0.0732	(0.0501)	0.1040	(0.0601)
σ_{16}	0.1316	(0.0563)	0.0839	(0.0521)	0.0665	(0.0482)	0.0736	(0.0487)	0.0678	(0.0488)

Note: (*) and (**) denote parameter estimates significant at 10% and 5%, respectively. The standard errors for the idiosyncratic error and shocks are all significant at 5% level, and hence the indicators are omitted.

Table 3.5. Parameter estimates for the ending stocks forecasts for wheat, 2004/05 – 2013/14

Parameter	USDA		Analysts Avg.		Analysts Med.		Selected Ana. Avg.		Selected Ana. Med.	
	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)
Coefficient										
α (intercept)	-0.0025	(0.0039)	-0.0007	(0.0033)	0.0014	(0.0043)	-0.0094	(0.0036)**	-0.0034	(0.0050)
β (slope)	0.0913	(0.0861)	0.2116	(0.0915)**	0.2463	(0.1049)**	0.1092	(0.0981)	0.1948	(0.1312)
Idiosyncratic Err.										
σ	0.0103	(0.0060)	0.0089	(0.0052)	0.0137	(0.0075)	0.0134	(0.0044)	0.0200	(0.0063)
Shock										
σ_1	0.0916	(0.0275)	0.1066	(0.0318)	0.1064	(0.0319)	0.1102	(0.0330)	0.1055	(0.0321)
σ_2	0.0335	(0.0116)	0.0114	(0.0068)	0.0140	(0.0091)	0.0099	(0.0071)	0.0125	(0.0097)
σ_3	0.0165	(0.0089)	0.0381	(0.0130)	0.0424	(0.0174)	0.0229	(0.0114)	0.0345	(0.0173)
σ_4	0.0505	(0.0153)	0.0421	(0.0132)	0.0375	(0.0141)	0.0453	(0.0148)	0.0360	(0.0165)
σ_5	0.0531	(0.0163)	0.0394	(0.0141)	0.0339	(0.0159)	0.0263	(0.0123)	0.0296	(0.0160)
σ_6	0.0437	(0.0142)	0.0564	(0.0170)	0.0475	(0.0176)	0.0731	(0.0220)	0.0549	(0.0220)
σ_7	0.0586	(0.0184)	0.0333	(0.0113)	0.0346	(0.0142)	0.0393	(0.0144)	0.0312	(0.0168)
σ_8	0.0548	(0.0168)	0.0303	(0.0109)	0.0363	(0.0142)	0.0583	(0.0182)	0.0387	(0.0175)
σ_9	0.0292	(0.0126)	0.0679	(0.0193)	0.0595	(0.0186)	0.0695	(0.0210)	0.0610	(0.0220)
σ_{10}	0.1069	(0.0304)	0.1043	(0.0290)	0.1033	(0.0298)	0.1073	(0.0305)	0.1020	(0.0304)
σ_{11}	0.0618	(0.0182)	0.0586	(0.0172)	0.0534	(0.0179)	0.0755	(0.0219)	0.0697	(0.0232)
σ_{12}	0.0942	(0.0263)	0.0803	(0.0228)	0.0759	(0.0234)	0.0679	(0.0201)	0.0726	(0.0232)
σ_{13}	0.0954	(0.0273)	0.1115	(0.0314)	0.1273	(0.0361)	0.1130	(0.0327)	0.1203	(0.0355)

Note: (*) and (**) denote parameter estimates significant at 10% and 5%, respectively. The standard errors for the idiosyncratic error and shocks are all significant at 5% level, and hence the indicators are omitted.

Table 3.6. Median and Credible Intervals of the estimated coefficients for ending stocks forecasts, 2004/05 – 2013/14

Parameter		α			β		
		2.5%	50%	97.5%	2.5%	50%	97.5%
Corn	AA	-0.0033	0.0074	0.0173	0.0454	0.1624	0.2962
	AM	-0.0020	0.0076	0.0164	0.0338	0.1578	0.3040
	FAA	-0.0018	0.0086	0.0215	0.0203	0.1518	0.2867
	FAM	-0.0012	0.0097	0.0209	0.0571	0.2020	0.3601
	USDA	-0.0054	0.0036	0.0158	0.0640	0.1658	0.2776
Soybeans	AA	-0.0293	-0.0155	-0.0028	0.3653	0.5135	0.6580
	AM	-0.0287	-0.0155	-0.0028	0.3418	0.5027	0.6532
	FAA	-0.0284	-0.0155	-0.0031	0.3608	0.5108	0.6530
	FAM	-0.0296	-0.0154	-0.0031	0.3622	0.5250	0.6724
	USDA	-0.0298	-0.0174	-0.0067	0.4285	0.5638	0.6963
Wheat	AA	-0.0072	-0.0007	0.0058	0.0357	0.2126	0.3848
	AM	-0.0069	0.0014	0.0086	0.0678	0.2408	0.4722
	FAA	-0.0161	-0.0097	-0.0020	-0.0788	0.1086	0.3100
	FAM	-0.0155	-0.0033	0.0058	-0.0666	0.1954	0.4533
	USDA	-0.0097	-0.0028	0.0057	-0.0673	0.0855	0.2729

CHAPTER 4

DO ANALYSTS FORECAST THE ENDING STOCKS OR THE USDA FORECASTS?

4.1 Abstract

Previous researchers typically analyze the USDA and private analysts' ending stocks forecasts separately, ignoring the interactions between these two types of forecasters. The present study recognizes the alternating forecast structure and builds a model to integrate the forecasts from these two sources. The model assumes that the forecasts are made from a synthetic forecaster who is comprised of the USDA and private analysts. A system which focuses on forecast revisions is then proposed for analysis. Results show that for corn, the USDA and analysts are forecasting each other, but their forecasts are both inefficient. For soybeans, the USDA is targeting the ending stocks, and private analysts are efficiently forecasting the USDA forecasts.

4.2 Introduction

The ending stocks of an agricultural commodity are recognized as a major indicator of the relative balance of supply and demand of the commodity in the market. They play important roles not only in policy makers' decision making, but also in market participants' planning for the upcoming marketing year. Thus, it is critical to provide the public with timely and accurate forecasts of ending stocks, reducing the uncertainty faced by the decision makers, and hence reducing market volatility.

The preparation of ending stocks forecasts requires sizable efforts and costs in data collection and analysis of both the supply and demand sides of the market. Thus historically,

ending stocks forecasts have been offered as a public service and provided by the U.S. Department of Agricultural (USDA), a federal executive department. The USDA has long been releasing crop ending stocks forecasts in its monthly World Agricultural Supply and Demand Estimates (WASDE) reports. The agency prepares these forecasts under the policy requirements and the goal of enhancing the overall functioning of the agricultural markets. The USDA forecasts are thus highly credited for its integrity, objectivity, and the incorporation of the most comprehensive information. Researchers have found that market participants place substantial value on the WASDE reports and adjust their market behavior accordingly (*e.g.*, Isengildina-Massa *et al.* 2008a, 2008b, Adjemian 2012).

Besides the USDA, the private sector has also issued its own ending stocks forecasts. The recent rise in the number of private forecasts can be attributed to the reduced difficulty and cost of acquiring related information, as well as advances in technologies and analytical methods. More and more private analysts have joined the group of ending stocks forecasters over the past decade. Their forecasts also contain important information and are closely followed by the public. However, the private forecasts may not have exactly the same characteristics as the USDA forecasts. Specifically, their sources of information need not overlap or be as comprehensive as the ones from the USDA. For example, for the factors from the supply side, surveys conducted by the private sector may have different samples from the ones conducted by the USDA. In addition, the use of the same pieces of information may be different as well. This is because the analysis is performed by different analysts who can make different interpretations and inferences using various analytical methods. Moreover, private forecasts are not regulated by the policies which the USDA must follow. Some private analysts themselves are actively involved in the market and pursue profits by directly trading the corresponding commodities.

Thus it would not be surprising if their forecasts were affected by their own trading objectives and strategies.

Crop ending stocks forecasts are fixed-event forecasts because there exist multiple forecasts for a specific event – the ending stock of a marketing year. As the private forecasts are typically published several days ahead of the USDA forecasts, it cannot be ruled out that the USDA and private analysts affect each other in forecasting the stocks. Thus, under the existing framework, it is difficult to investigate the behavior of one forecaster given the existence of its competitor, because the competitor's forecasts are not included in the analysis.

In addition, given that analysts' forecasts are published ahead of the USDA forecasts, one natural question arises: do analysts in fact forecast the ending stocks or the upcoming USDA forecasts? It is interesting to identify the true target of the analysts' forecasts. On one side, analysts can choose the upcoming USDA forecasts as the short-term target, because successful forecasts allow market participants to gain an advantage, thus mitigating the risk and even obtaining immediate profits. On the other side, analysts can choose to forecast the ending stocks if their focus is on the long term. However, as researchers typically cannot obtain the specific models that analysts employed, inferences about analysts' forecasting behavior can only be made from limited public information, or the historical forecasts that analysts have made.

Similar arguments can be applied to the USDA forecasts. The USDA may take into account the analysts' forecasts, as these forecasts reflect expectations of market participants. In other fixed-events, research has found that relatively large differences between the consensus analysts' forecasts and the government forecasts often lead to market volatility (*e.g.*, French *et al.* 1989). There is also a large number of studies addressing the announcement effect of the government forecasts in various areas (*e.g.*, Adjemian 2012). This announcement effect

introduces more uncertainty along the path toward the revealing the forecasting target. Thus, it is possible that the USDA will incorporate it as a factor in generating its forecasts, and also consider the forecasts from the private sector.

The above question is typically overlooked in studies for other fixed-event forecasts. The reason is that either the forecast horizons are short, or the number of forecasts in a forecasting cycle is small. If the forecast horizon is short, it is irrelevant to forecast the government forecasts or forecast the outcome, as the differences will be quite small. If the number of forecasts in a forecasting cycle is small, there is limited information to identify the true forecast target. For example, studies addressing Federal Reserve's macroeconomic projections such as nominal / real GDP growth, inflation, or unemployment typically focus on 4 or 5 quarterly forecasts. In agricultural events, only 5 monthly forecasts are used in studies on corn and soybeans production forecasts.

But the situation is different for crop ending stocks forecasts for corn and soybeans. Over the past three decades, the USDA has issued a total of 17 monthly forecasts for each marketing year's ending stocks for these two commodities. The forecast cycle is long, covering from May to September of the following calendar year, and the number of forecasts within the forecasting cycle is much more than for other fixed-events. Thus, it is more likely that the two targets – the upcoming USDA forecasts and the final ending stocks – can differ from each other, making it more possible for us to separate and investigate the private analysts' forecasting behavior.

Chapter 3 finds that both the USDA and the consensus analysts' forecasts are inefficient during marketing years 2004/05 – 2013/14. In particular, both the USDA and the consensus analysts are conservative in adjusting their forecasts, and the estimated parameters are quite similar. This finding of similar forecasting patterns shows that there is possibility that one

forecaster is forecasting the other. Thus, it is interesting to identify the true relationship between the USDA forecasts, the analysts' counterpart, and the final ending stocks.

To the best of our knowledge, this problem has seldom been addressed in the past literature in fixed-event forecasts, as most researchers analyze private forecasts and government forecasts separately. Take for example the agricultural price and production forecasts. Garcia *et al.* (1997) treated analysts' forecasts as competing forecasts of the USDA forecasts, thus implicitly assuming that private analysts, similar to the USDA, are directly forecasting the target outcome (the production of the commodity). They find a decline in the informational value of USDA forecasts, which is consistent with the rise in the provision of private forecasts. Egelkraut *et al.* (2003) evaluate the accuracy of USDA and analysts' forecasts based on the same implicit assumption. On the other hand, McKenzie (2008) conducts his research under the assumption that private forecasts are unbiased estimates of the government forecasts. As for the ending stocks forecasts, which very few researchers have been addressed, it is necessary to specify the forecasting behavior of the two groups of forecasters for future research.

Chapter 2 introduces an estimation framework that combines the two main strands in fixed-event forecast studies. The first strand is introduced by Nordhaus (1987), which focuses on forecast revisions. The second strand is introduced by Davies and Lahiri (1995, 1999), which directly addresses forecast errors and suggests error decompositions. Clements *et al.* (2007) first suggest analyzing forecast revisions by differencing the forecast errors. Chapter 2 further extends the Clements model by investigating adjacent forecast revisions, while retaining the link between the forecasts and the forecast target. Chapter 2 introduces an error covariance matrix which takes into account both the heteroscedasticity of the shocks and the autocorrelations generated by forecaster's corrections of their own errors.

Based on the model in Chapter 2, the present study proposes a framework to analyze the efficiency of USDA and analysts' forecasts in a single system. The underlying assumption is that the USDA and private analysts' forecasts affect each other along the path toward the revealing of the ending stocks. Specifically, we recognize the alternating forecast structure of ending stocks and further divide the time interval between own forecast revisions into two intervals with different meanings. This division of time intervals allows us to examine each forecaster in more detail, and analyze the role played by the competing forecaster. Besides, we can focus on utilizing the necessary and most updated information only. We examine whether one forecaster is forecasting the other and then determine the interactions between them. The method is applied to the ending stocks forecasts for corn and soybeans from marketing years 2004/05 to 2013/14.

We develop a Bayesian Markov Chain Monte Carlo (MCMC) method to estimate the parameters in the system. In the MCMC method, all parameters are estimated in one iteration step. Conjugate priors are adopted for the estimation. The chosen priors are also uninformative, so that the data dominates the estimation processes. The method developed can also serve as a justification of the proposed modelling structure.

The rest of this study is organized as follows: Section 4.3 first introduces the alternating forecast structure of USDA and analysts' ending stocks forecasts. Then it advances the model in Chapter 2 to estimate the efficiency of forecasts from these two sources in a single system. Section 4.4 outlines the source and descriptive statistics of the data. Section 4.5 introduces the empirical methods employed in the estimation. In Section 4.6, we describe the estimation results and perform further analysis. The final section provides concluding remarks.

4.3 The Model

4.3.1 The Structure of Forecasts

The discussion in the introduction shows that the treatment of multiple forecasts as competitors may generate biased results by ignoring the information behind the true forecast structure. The present study, however, adds another dimension to the multiple forecasts analysis by considering their forecast timing. In this way, two forecasts in different forecasting series which are made “during the same month” are no longer treated as released at the same time. As a result, this originally overlooked information allows us to further investigate the relationship between these series of forecasts.

This section specifies the structure of ending stocks forecasts for corn and soybeans. A typical marketing year for corn and soybeans begins on September 1st, and ends on August 31st of the next calendar year. The ending stock of the past marketing year is then released at the end of September. Within a forecasting cycle, both the USDA and the analysts provide 17 monthly forecasts of the ending stocks. The first forecast is published in May before the marketing year starts, whereas the last forecast is released in September after the marketing year ends and before the ending stock report is finalized. Thus there exist two forecasts for the ending stocks of different marketing years in the months from May through September.

The USDA forecasts are released between the 9th and 12th day of each month. Instead, private analysts typically finish making their forecasts several days before the publishing of the USDA counterparts. The survey of private analysts’ forecasts are then collected and released to the public. The average time between the release of the survey and the USDA forecasts is five days.

Since private analysts' forecasts are released several days ahead of the USDA forecasts, it is reasonable to assume that the information content in these two groups of forecasts is different. In fact, forecasters can only take advantage of the information which is available before they finalize their forecasts. As the USDA releases its forecasts several days later, it can include additional information which private analysts, who have already forecasted, cannot obtain or utilize until the following month's forecasts. This specification of analysts' and USDA forecast structure further divides the time intervals, as well as the information contents, into smaller partitions. Thus the information contents can be allocated to either analysts' updates of most recent USDA forecasts or USDA updates of most recent analyst's forecasts. As a result, we can investigate the relationship between these two group of forecasters based on the partitioned information.

To illustrate the forecast structure more clearly, let S_t be the ending stock of marketing year t for either commodity, $U_{t,n}$ be the USDA n -month-ahead forecast of S_t , and $V_{t,n}$ be the representative analyst's n -month-ahead forecast of S_t . Figure 4.1 depicts the timeline of the USDA and representative analyst's forecasts. It can be seen that the time interval of a forecaster's own revision (either $U_{t,n} - U_{t,n+1}$ or $V_{t,n} - V_{t,n+1}$) are further divided as the combination of the following two segments:

Type A: The time interval between the USDA forecast issued in the previous month and the analyst's forecast issued in the current month (e.g., $V_{t,n} - U_{t,n+1}$).

Type B: The time interval between the analyst's forecast and the subsequent USDA forecast issued in the same month (e.g., $U_{t,n} - V_{t,n}$).

The Type A interval typically lasts slightly more than three weeks. It includes most of the period following the USDA forecasts for the previous month. It covers the main influx of new

information that is used to update the previous forecasts. On the other hand, the Type B interval only lasts for slightly less than a week. It refers to the time between two forecasts issued in the same month. Thus, the time interval of the USDA own forecast revision ($U_{t,n} - U_{t,n+1}$) can be viewed as covering a Type A segment, the analyst's update of past month USDA forecast ($V_{t,n} - U_{t,n+1}$), and a Type B segment, the USDA update of the most recent analyst's forecast ($U_{t,n} - V_{t,n}$). Similarly, the time interval of the analyst's own forecast revision ($V_{t,n} - V_{t,n+1}$) can be decomposed into a Type B segment, the USDA update of the analyst's forecast ($U_{t,n+1} - V_{t,n+1}$), followed by a Type A segment, the analyst's update of the past month USDA forecast $V_{t,n} - U_{t,n+1}$. This specification of alternating forecast structure adds another dimension to the forecasting data, enabling us to investigate the relationship between these two groups of forecasters.

4.3.2 Proposed Model

This section describes the model to test the efficiency of both USDA and private analysts' forecasts in the context of alternating forecast structure. Previous research on fixed-event forecasts typically treat government and private forecasts as competitors. Tests were performed separately for each group of forecasts with null hypotheses like the following:

H_0 : Government forecasts are unbiased and efficient forecasts of the fixed event.

H_0 : Private analyst's forecasts are unbiased and efficient forecasts of the fixed event.

Based on the alternating forecast structure of the ending stocks, we propose testing the following two hypotheses:

H_0 : USDA forecasts are unbiased and efficient forecasts of ending stocks, given the analysts' forecasts.

H_0 : Analysts' forecasts are unbiased and efficient forecasts of ending stocks given the USDA forecasts.

In this way, the analysis on the forecasting behavior of one forecaster is no longer isolated, because the forecasts from the competitor play a role in the evaluations. As discussed in the previous section, the advantage of including forecasts from the competitor is that they contain most up-to-date information. Information within past own forecasts are thus no longer the newest after the inclusion of the most recent forecasts from the competitor.

The bias and efficiency of the forecasts can be tested by means of transformations of the Mincer and Zarnowitz (1969) regression:

$$S_t - U_{t,n} = a_U + \beta_U X_U + error_{U,t,n} \quad (4.1)$$

$$S_t - V_{t,n} = a_V + \beta_V X_V + error_{V,t,n} \quad (4.2)$$

The null hypothesis $H_0: a_U = \beta_U = 0$ indicates that the USDA forecasts are unbiased and efficient. Similar results apply to the analysts' forecasts. However, it may not be wise to directly consider forecast errors ($S_t - U_{t,n}, S_t - V_{t,n}$), as they are related to future forecasts, which are not directly affected by current information. A natural substitute for the target of the ending stocks (S_t) is the next month own forecast, *i.e.*, $U_{t,n+1}$ for the USDA and $V_{t,n+1}$ for the analysts. Chapter 2 introduced an efficiency test based on these alternative forecasting targets. The proposed efficiency test in Chapter 2 can be expressed as the following systems of equations:

$$U_{t,n-1} - U_{t,n} = a_U + \beta_U (U_{t,n} - U_{t,n+1}) + k_{U,t,n-1} - \varepsilon_{U,t,n-1} + \varepsilon_{U,t,n} \text{ for USDA} \quad (4.3)$$

$$V_{t,n-1} - V_{t,n} = a_V + \beta_V (V_{t,n} - V_{t,n+1}) + k_{V,t,n-1} - \varepsilon_{V,t,n-1} + \varepsilon_{V,t,n} \text{ for analysts} \quad (4.4)$$

where k and ε are the Davies and Lahiri (1995, 1999) decompositions of the error term in regressions (4.1) and (4.2), representing unforecastable shocks and idiosyncratic errors, respectively. The most recent own forecast revisions are the most obvious candidates to include as past information, because the own forecast history is always available. The null hypothesis $H_0: (\alpha, \beta) = (0, 0)$ implies that the forecasts are efficient.

The incorporation of the competing forecasts and the recognition of the alternating forecast structure can be used to further decompose the information structure in Chapter 2. As discussed in the previous section, there exists an analyst's forecast $V_{t,n-1}$ between the two consecutive USDA forecasts $U_{t,n}$ and $U_{t,n-1}$. Thus the information in the period between $U_{t,n}$ and $U_{t,n-1}$ can be further divided into two pieces: Type A information between $U_{t,n}$ and $V_{t,n-1}$, and Type B information between $V_{t,n-1}$ and $U_{t,n-1}$. Similarly, the information between the analyst's revisions $V_{t,n}$ and $V_{t,n-1}$ are the sum of Type A information between $U_{t,n}$ and $V_{t,n-1}$, and Type B information between $V_{t,n}$ and $U_{t,n}$.

Given this fact, we can apply first differencing with the considerations of both groups of forecasters. For notational purposes, we can treat that the forecasts of $V_{t,n}$ and $U_{t,n}$'s are made from a synthetic forecaster W . Let $W_t = [U_{t,1}, V_{t,1}, U_{t,2}, V_{t,2}, \dots, U_{t,N}, V_{t,N}]'$, a vector of $V_{t,n}$ and $U_{t,n}$'s. Let m be the subscript of W_t , representing the m th element of W_t . We now follow the steps in Chapter 2 to develop the estimation framework.

The bias test following Davies and Lahiri (1995, 1999) can be written as

$$S_t - W_{t,m} = \sum_{j=1}^m \alpha_j + \sum_{j=1}^m k_{t,j} + \varepsilon_{t,m} \quad (4.5)$$

where α_j is the monthly bias coefficient, representing either α_U or α_V . $k_{t,j} \sim i.i.d. N(0, \sigma_j^2)$ is the monthly unforecastable shock, and $\varepsilon_{t,m} \sim i.i.d. N(0, \sigma^2)$ is the forecaster's idiosyncratic error.

Denote M as the maximum forecast horizon of the synthetic forecaster. The value of M is then twice of the maximum forecast horizon of either the USDA or the analyst. First differencing (4.5) gives the following system of equations:

$$\begin{cases} S_t - W_{t,1} = \alpha_1 + k_{t,1} + \varepsilon_{t,1} \\ W_{t,1} - W_{t,2} = \alpha_2 + k_{t,2} - \varepsilon_{t,1} + \varepsilon_{t,2} \\ \vdots \\ W_{t,M-1} - W_{t,M} = \alpha_M + k_{t,M} - \varepsilon_{t,M-1} + \varepsilon_{t,M} \end{cases} \quad (4.6)$$

The efficiency test based on (4.6) consists of fitting

$$\begin{cases} S_t - W_{t,1} = \alpha_1 + \beta_1 X_{t,1} + k_{t,1} + \varepsilon_{t,1} \\ W_{t,1} - W_{t,2} = \alpha_2 + \beta_2 X_{t,2} + k_{t,2} - \varepsilon_{t,1} + \varepsilon_{t,2} \\ \vdots \\ W_{t,M-1} - W_{t,M} = \alpha_M + \beta_M X_{t,M} + k_{t,M} - \varepsilon_{t,M-1} + \varepsilon_{t,M} \end{cases} \quad (4.7)$$

where $X_{t,m}$ represents one or more explanatory variables known when the forecast $W_{t,m}$ is made.

The null hypothesis $H_0: \alpha_n = \beta_n = 0$ for all n indicates efficiency of the group of forecasts W .

In order to estimate the specific parameters for the USDA and the representative analyst, the elements of W_t need to be reverted back to the original forecasts. Thus system (4.7) becomes

$$\begin{cases} S_t - U_{t,1} = \alpha_1 + \beta_1 X_{t,1} + k_{t,1} + \varepsilon_{t,1} \\ U_{t,1} - V_{t,1} = \alpha_2 + \beta_2 X_{t,2} + k_{t,2} - \varepsilon_{t,1} + \varepsilon_{t,2} \\ \vdots \\ V_{t,N-1} - U_{t,N} = \alpha_{M-1} + \beta_{M-1} X_{t,M-1} + k_{t,M-1} - \varepsilon_{t,M-2} + \varepsilon_{t,M-1} \\ U_{t,N} - V_{t,N} = \alpha_M + \beta_M X_{t,M} + k_{t,M} - \varepsilon_{t,M-1} + \varepsilon_{t,M} \end{cases} \quad (4.8)$$

Each equation in system (4.8) can be classified into one of the previously proposed two categories based on the dependent variables. That is, the revisions in the dependent variables are either Type A revisions $V_{t,n-1} - U_{t,n}$ or Type B revisions $U_{t,n} - V_{t,n}$. Additional restrictions are imposed on the parameters to reflect the characteristics of revisions by their types. Specifically,

we restrict α and β 's to be the same for each category. Besides, the variances of the idiosyncratic errors can also be separated. In notations, the following restrictions are applied:

$$\begin{aligned}\alpha_1 &= \alpha_3 = \dots = \alpha_{M-1} = \alpha_A, \alpha_2 = \alpha_4 = \dots = \alpha_M = \alpha_B \\ \beta_1 &= \beta_3 = \dots = \beta_{M-1} = \beta_A, \beta_2 = \beta_4 = \dots = \beta_M = \beta_B \\ \varepsilon_1, \varepsilon_3, \dots, \varepsilon_{M-1} &\sim N(0, \sigma_A^2), \varepsilon_2, \varepsilon_4, \dots, \varepsilon_M \sim N(0, \sigma_B^2)\end{aligned}\tag{4.9}$$

Further changing the subscripts of X , k and ε 's for corresponding types of revisions in a similar way, we have

$$\left\{ \begin{array}{l} S_t - U_{t,1} = \alpha_A + \beta_A X_{t,1,A} + k_{t,1,A} + \varepsilon_{t,1,A} \\ U_{t,1} - V_{t,1} = \alpha_B + \beta_B X_{t,1,B} + k_{t,1,B} - \varepsilon_{t,1,A} + \varepsilon_{t,1,B} \\ \vdots \\ V_{t,N-1} - U_{t,N} = \alpha_A + \beta_A X_{t,N,A} + k_{t,N,A} - \varepsilon_{t,N-1,B} + \varepsilon_{t,N,A} \\ U_{t,N} - V_{t,N} = \alpha_B + \beta_B X_{t,N,B} + k_{t,N,B} - \varepsilon_{t,N,A} + \varepsilon_{t,N,B} \end{array} \right. \tag{4.10}$$

System (4.10) can be interpreted as a joint estimation of efficiency of USDA and the representative analyst's forecasts. Type A equations represent the test of whether USDA forecasts are efficient given the analyst's forecasts, whereas Type B equations test whether analyst's forecasts are efficient given USDA forecasts.

The explanatory variables $X_{t,n,A}$ and $X_{t,n,B}$ can be any variable which is exogenous or predetermined. One possible candidate is the lag of the dependent variable. Given the alternating forecast structure, the first two lags of the dependent variables can be included because they represent different types of information. The first lag represents the most recent forecast revision of the alternative type. Specifically, if $V_{t,n-1} - U_{t,n}$ is the dependent variable, $U_{t,n} - V_{t,n}$ is the first lag. If $U_{t,n} - V_{t,n}$ is the dependent variable, $V_{t,n} - U_{t,n+1}$ entered as one of the explanatory variables. The second candidate for explanatory variable is the most recent own forecast revision. For example, if $V_{t,n-1} - U_{t,n}$ is the dependent variable, $V_{t,n} - U_{t,n+1}$ is the most recent

own forecast revision and is included as an explanatory variable. Similarly, $V_{t,n} - U_{t,n+1}$ is included if $U_{t,n} - V_{t,n}$ is the dependent variable. The final two observations in (4.10) are dropped from the dependent variables because the corresponding past forecast revisions do not exist.

Thus, system (4.10) becomes:

$$\left\{ \begin{array}{l} S_t - U_{t,1} = \alpha_A + \beta_{1A}(U_{t,1} - V_{t,1}) + \beta_{2A}(V_{t,1} - U_{t,2}) + k_{t,1,A} + \varepsilon_{t,1,A} \\ U_{t,1} - V_{t,1} = \alpha_B + \beta_{1B}(V_{t,1} - U_{t,2}) + \beta_{2B}(U_{t,2} - V_{t,2}) + k_{t,1,B} - \varepsilon_{t,1,A} + \varepsilon_{t,1,B} \\ \vdots \\ V_{t,N-2} - U_{t,N-1} = \alpha_A + \beta_{1A}(U_{t,N-1} - V_{t,N-1}) + \beta_{2A}(V_{t,N-1} - U_{t,N-2}) + k_{t,N-1,A} - \varepsilon_{t,N-2,B} + \varepsilon_{t,N-1,A} \\ U_{t,N-1} - V_{t,N-1} = \alpha_B + \beta_{1B}(V_{t,N-1} - U_{t,N-2}) + \beta_{2B}(U_{t,N-2} - V_{t,N-2}) + k_{t,N-1,B} - \varepsilon_{t,N-1,A} + \varepsilon_{t,N-1,B} \end{array} \right. \quad (4.11)$$

To interpret the structure and parameters, take for example the Type A equations:

$$V_{t,n-1} - U_{t,n} = \alpha_A + \beta_{1A}(U_{t,n} - V_{t,n}) + \beta_{2A}(V_{t,n} - U_{t,n-1}) + k_{t,n,A} - \varepsilon_{t,n-1,B} + \varepsilon_{t,n,A} \quad (4.12)$$

The dependent variable can be viewed as the decomposed form of the forecast error of $U_{t,n}$, *i.e.*,

$S_t - U_{t,n}$, with contents related to future information being discarded. β_{1A} is the coefficient

representing the impact of the most recent forecast revision of the alternative category. In

addition, it is worthwhile to note that the explanatory variable $U_{t,n} - V_{t,n}$ can be viewed as the

difference between the forecast errors of $V_{t,n}$ and $U_{t,n}$, *i.e.* $(S_t - V_{t,n}) - (S_t - U_{t,n})$. Thus, β_{1A}

can be interpreted using arguments similar to the forecast encompassing test by Granger and

Newbold (1973, 1986). In this way, $\beta_{1A} = 0$ implies that $U_{t,n}$ encompasses $V_{t,n}$ in the sense that

there is no statistically significant increase in expected squared error loss if $V_{t,n}$ is excluded. In

other words, $U_{t,n}$ carries all the information in $V_{t,n}$ in predicting $V_{t,n-1}$. Therefore, one can reach

a conclusion that $U_{t,n}$ is a forecast of $V_{t,n-1}$ if the estimated β_{1A} is not significantly different

from zero. β_{2A} is the coefficient indicating the impact of the most recent own forecast revision. It

measures the behavior of the forecast revisions based on the forecaster's own forecast history

given the competing forecasts. $\beta_{2A} = 0$ implies that forecast revisions cannot be inferred from their own past revisions. Similar arguments can be made for Type B equations.

The joint estimation also introduces an error covariance matrix which is typically ignored if the two types of regressions are performed separately. We impose limited restrictions on the error covariance, while retaining the link between forecasts and the ending stocks. Specifically, the error covariance matrix allows for both heteroskedasticity in the shocks and autocorrelations generated by the forecaster's own errors. The variances of unforecastable shocks are assumed to be the same for each forecast horizon, but are different for each category, so that $k_{t,n,A} \sim i.i.d. N(0, \sigma_{n,A}^2)$ and $k_{t,n,B} \sim i.i.d. N(0, \sigma_{n,B}^2)$. The variance of the forecaster's idiosyncratic errors are also assumed to be the same within each category. Based on the rationale in Chapter 2, the covariance matrix for a typical marketing year is as follows:¹

$$\bar{B}_{2(N-1) \times 2(N-1)} = \begin{bmatrix} \bar{\sigma}_{N-1,B}^2 + 2\sigma_B^2 & -\sigma_A^2 & 0 & \dots & \sigma_{4,B}^2 & 0 & 0 & 0 \\ -\sigma_A^2 & \bar{\sigma}_{N-1,A}^2 + 2\sigma_A^2 & -\sigma_B^2 & & \sigma_{4,A}^2 & 0 & 0 & \\ 0 & -\sigma_B^2 & \bar{\sigma}_{N-2,B}^2 + 2\sigma_B^2 & -\sigma_A^2 & & \sigma_{1,B}^2 & 0 & \\ \vdots & & -\sigma_A^2 & \bar{\sigma}_{N-2,A}^2 + 2\sigma_A^2 & & & \sigma_{1,A}^2 & \\ \sigma_{4,B}^2 & & & & \ddots & & \vdots & \\ 0 & \sigma_{4,A}^2 & & & & \ddots & -\sigma^2 & 0 \\ \vdots & 0 & \sigma_{1,B}^2 & & & -\sigma^2 & \sigma_{1,B}^2 + 2\sigma_B^2 & -\sigma_A^2 \\ 0 & 0 & 0 & \sigma_{1,A}^2 & \dots & 0 & -\sigma_A^2 & \sigma_{1,A}^2 + \sigma_A^2 \end{bmatrix} \quad (4.13)$$

where $\bar{\sigma}_n^2 = \sigma_n^2 + \sigma_{n-12}^2$ for $n > 12$. The diagonal elements represent the variances of the residuals. They are decomposed into the sum of the variances of the monthly shocks and the idiosyncratic errors. The sub-diagonal elements of the matrix represent the covariance of errors of two adjacent equations, which is determined by the variance of the forecaster's idiosyncratic errors. The interpretations for $\bar{\sigma}_n^2$ come from the assumption in Chapter 2: for $n > 12$, an

¹ The data is sorted by forecast horizon $n = N, \dots, 1$.

additional variance is included because the shocks which appear in these months can affect the two forecasts made in the same month but are for different marketing years.

The structure of the data is characterized by two dimensions, namely, forecast horizon n and marketing year t . While the order of the equations does not affect the estimation, it is interesting to see the structure of the full covariance matrix. Without loss of generality we can sort the data first by marketing year and then by forecast horizon. Then the full covariance matrix Σ is block diagonal with (4.13) as the blocks:

$$\Sigma = \begin{bmatrix} \bar{B} & 0 & \dots & 0 \\ 0 & \bar{B} & \dots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \dots & \bar{B} \end{bmatrix}_{T \times T} \quad (4.14)$$

4.4 Data

The data consist of U.S. ending stocks and their corresponding USDA and private analysts' monthly forecasts, for two major agricultural commodities – corn and soybeans. The marketing years used for the analysis are from 2004/05 through 2013/14, a total of 10 marketing years.

U.S. ending stocks are obtained from the Grain Stocks Reports released by the National Agricultural Statistical Services (NASS). The report is published quarterly. As the marketing year for corn and soybeans ends in August, the ending stocks data are retrieved from the September report. The USDA monthly forecasts are obtained from the WASDE reports. For the two commodities, there are 17 forecasts published by the USDA in a forecasting cycle, with the first forecast released in May before the marketing year begins and the last one issued in September after the marketing year ends. The private analysts' forecasts data have the same

format as the USDA counterpart. The data are obtained from the monthly Surveys of U.S. Grain and Soybeans Carryout Forecasts, which are conducted by Dow Jones Commodities Services. Therefore, based on the alternating forecast structure proposed in the previous section, we have 34 forecast revisions for each marketing year. Among these forecast revisions, half of them are Type A revisions (*i.e.*, the USDA forecast revisions of analyst forecasts), whereas the other half are Type B revisions (*i.e.*, the representative analyst's forecast revisions of the past month's USDA forecast).

The present study uses the average and median of analysts' forecasts as representatives of analysts' forecasts. We include 54 analysts who have provided at least one forecast during the marketing years of 2004/05 through 2013/14. The main reason underlying this choice is that in Chapter 3 the forecasts of these two representative analysts, as well as the USDA forecasts, are all found inefficient. Thus, it is possible to separate the forecast targets: the upcoming competing forecasts or the final ending stocks. Another reason is that the public generally focuses more on the analysts as a group, instead of a single analyst. So we use the statistics with all analysts included, such as the most recognized average of forecasts. The median of analysts' forecasts is included because the number of analysts who make forecasts varies for each month, and sometimes the median can be more credible as the analysts' consensus forecasts.

Descriptive statistics for each type of the USDA and representative analyst's revisions, as well as own forecast revisions, are depicted in Table 4.1. For all datasets, the means and medians of both Type A and Type B revisions are close to zero. The standard deviations of Type A and Type B revisions are smaller than those of the two own forecast revisions. This is because they are subintervals of the later. Besides, the standard deviations of Type A revisions are found to be larger than those of Type B revisions, as the Type A intervals are larger than Type B intervals.

For the same reason, the range of Type A revisions are observed larger than that of Type B revisions.

4.5 Estimation Methods

The proposed system of equations is estimated using Bayesian Markov Chain Monte Carlo (MCMC) methods. Bayesian MCMC methods offer a convenient way to estimate a model with the advocated error covariance structure. Besides, the proposed MCMC methods allow the data to play an essential role in validating the proposed structure. For example, if the proposed autocorrelations in the error covariance matrix do not exist, the MCMC simulations would generate results with variances of idiosyncratic errors close to zero, and much smaller compared to the variances of shocks.

To explain the estimation methods, we can write the system (4.11) as a collection of the Type A and Type B equations:

$$\begin{cases} V_{t,n-1} - U_{t,n} = \alpha_A + \beta_{1A}(U_{t,n} - V_{t,n}) + \beta_{2A}(V_{t,n} - U_{t,n-1}) + k_{t,n,A} - \varepsilon_{t,n-1,B} + \varepsilon_{t,n,A} \\ U_{t,n} - V_{t,n} = \alpha_B + \beta_{1B}(V_{t,n} - U_{t,n-1}) + \beta_{2B}(U_{t,n-1} - V_{t,n-1}) + k_{t,n,B} - \varepsilon_{t,n,A} + \varepsilon_{t,n,B} \end{cases} \quad (4.15)$$

for $n = 1, \dots, N - 1$ and $t = 1, \dots, T$. In matrix form,

$$\mathbf{y}_t = \mathbf{x}_t \boldsymbol{\beta} + \mathbf{w} \mathbf{k}_t + \mathbf{p} \boldsymbol{\varepsilon}_t \quad (4.16)$$

$$\text{where } \mathbf{y}_t = \begin{bmatrix} S_t - U_{t,1} \\ U_{t,1} - V_{t,1} \\ \vdots \\ V_{t,N-2} - U_{t,N-1} \\ U_{t,N-1} - V_{t,N-1} \end{bmatrix}, \mathbf{x}_t = \begin{bmatrix} 1 & U_{t,1} - V_{t,1} & V_{t,1} - U_{t,2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & V_{t,1} - U_{t,2} & U_{t,2} - V_{t,2} \\ & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & U_{t,N-1} - V_{t,N-1} & V_{t,N-1} - U_{t,N-2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & V_{t,N-1} - U_{t,N-2} & U_{t,N-2} - V_{t,N-2} \end{bmatrix},$$

$$\boldsymbol{\beta} = \begin{bmatrix} \alpha_A \\ \beta_{1A} \\ \beta_{2A} \\ \alpha_B \\ \beta_{1B} \\ \beta_{2B} \end{bmatrix}, \mathbf{k}_t = \begin{bmatrix} k_{t,1,A} \\ k_{t,1,B} \\ \vdots \\ k_{t,N-1,A} \\ k_{t,N-1,B} \end{bmatrix}, \boldsymbol{\varepsilon}_t = \begin{bmatrix} \varepsilon_{t,1,A} \\ \varepsilon_{t,1,B} \\ \vdots \\ \varepsilon_{t,N-1,A} \\ \varepsilon_{t,N-1,B} \end{bmatrix}. \mathbf{w} \text{ is a matrix indicating the existence of elements in}$$

\mathbf{k}_t in each equation. \mathbf{p} is a matrix indicating the existence of elements in $\boldsymbol{\varepsilon}_t$ in each equation. \mathbf{w} and \mathbf{p} do not vary with t . The full system can be further written as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{W}\mathbf{K} + \mathbf{P}\boldsymbol{\varepsilon} \quad (4.17)$$

where each character represents the vector containing the same character with subscripts $t = 1, \dots, T$. For identification purposes, we set $k_{T,1,A}$ and $k_{T,1,B}$ to be zero.

Our estimation is performed by Gibbs Sampling. We use conditionally conjugate priors for each parameter and derive the corresponding posterior distributions. Let $\boldsymbol{\Lambda} = \{\boldsymbol{\beta}, \{\sigma_{n,A}^2\}_{n=1}^N, \{\sigma_{n,B}^2\}_{n=1}^N, \sigma_A^2, \sigma_B^2\}$. $\boldsymbol{\Lambda}$ is the set of the parameters to be estimated in the model. The joint posterior density of $\boldsymbol{\Lambda}$ is

$$\begin{aligned} p(\boldsymbol{\Lambda}) &= \Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega}) \prod_{jt} \Phi(k_{t,j,A} | \sigma_{j,A}^2) \prod_{jt} \Phi(k_{t,j,B} | \sigma_{j,B}^2) \\ &\quad * \Phi(\boldsymbol{\beta} | \mathbf{M}, \mathbf{V}) p(\sigma_A^2) p(\sigma_B^2) \prod_{j=1}^N p(\sigma_{j,A}^2) \prod_{j=1}^N p(\sigma_{j,B}^2) \end{aligned} \quad (4.18)$$

where $\Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega})$ is a multivariate normal distribution with mean $\mathbf{X}\boldsymbol{\beta} + \mathbf{W}\mathbf{K}$ and variance $\boldsymbol{\Omega}$.

The conditionally conjugate priors chosen for the parameters are:

$$\begin{aligned} \boldsymbol{\beta} &\sim N(\mathbf{M}, \mathbf{V}) \\ \varepsilon_{t,n,A} &\sim N(0, \sigma_A^2), \varepsilon_{t,n,B} \sim N(0, \sigma_B^2) \\ k_{t,n,A} &\sim N(0, \sigma_{n,A}^2), k_{t,n,B} \sim N(0, \sigma_{n,B}^2) \\ \sigma_A, \sigma_B, \sigma_{n,A}, \sigma_{n,B} &\sim \text{Uniform}(0, \infty) \end{aligned} \quad (4.19)$$

for $n = 1, \dots, N$, and $t = 1, \dots, T$. The prior distribution for the coefficient vector $\boldsymbol{\beta}$ is multivariate normal with mean $\mathbf{M} = \mathbf{0}_6$ and covariance matrix $\mathbf{V} = 1000\mathbf{I}_{6 \times 6}$, where $\mathbf{0}_6$ is a 6×1 vector of zeros and $\mathbf{I}_{6 \times 6}$ is a 6×6 identity matrix. The prior mean of $\boldsymbol{\beta}$ is chosen to be zero so that we cannot reject the null hypothesis of efficiency in default. The scale of the variance is chosen to be large so that the priors are non-informative. The draws of $\boldsymbol{\beta}$ will thus be diffused and widely spread around the mean zero. The improper uniform prior for the standard deviation parameters is chosen following Gelman (2006). This prior is non-informative and can be viewed as a limit of the half- t family distributions, which is conditionally conjugate to the extent of more general folded-noncentral- t distributions. The conditional posterior distribution for each parameter is outlined in the Appendix.

The MCMC iteration steps for the proposed model are as follows:

Step 1: For each chain, set up initial values for each parameter in the set $\mathbf{\Lambda}$, as well as for $\mathbf{K}^{(0)}$ and $\mathbf{E}^{(0)}$.

Step 2: Given $\{\mathbf{K}^{(i)}, \sigma_A^{2(i)}, \sigma_B^{2(i)}\}$, draw $\boldsymbol{\beta}^{(i)}$ from its posterior, which is a multivariate normal distribution.

Step 3: Given $\{\boldsymbol{\beta}^{(i+1)}, \sigma_A^{2(i)}, \sigma_B^{2(i)}, \{\sigma_{n,A}^2\}^{(i)}, \{\sigma_{n,B}^2\}^{(i)}, \mathbf{K}_{-k_{t,n,m}}^{(i)}\}$, sequentially draw $k_{t,n,m}^{(i+1)}$ from its posterior, which is a normal distribution, for each $t = 1, \dots, T$, $n = 1, \dots, N - 1$, and type $m = A, B$.

Step 4: Given $\{\boldsymbol{\beta}^{(i+1)}, \mathbf{K}^{(i+1)}\}$, update $\mathbf{E}^{(i+1)}$. Then draw $\sigma_m^{2(i+1)}$ from its posterior, which is an inverse gamma distribution, for each type $m = A, B$.

Step 5: Given $\mathbf{K}^{(i+1)}$, sequentially draw $\sigma_{n,m}^{2(i+1)}$ from its posterior, an inverse gamma distribution, for each $n = 1, \dots, N - 1$ and type $m = A, B$.

Step 6: Set $i = i + 1$.

Step 7: Repeat Steps 2-6 until the maximum iteration is reached.

For each dataset, we run three Markov Chains with different starting values. Each chain is proceeded for 200,000 iterations. The first 100,000 iterations of each chain are discarded as burn-in period. Gelman and Rubin (1992) tests are then applied to check the convergence of the remaining part of the chains. The test statistic compares the variances of both within the chains and between the chains. The convergence is indicated by the values of the statistic which are close to 1.

4.6 Results and Discussion

Estimation results are summarized in Tables 4.2 and 4.3. Table 4.2 reports the means and standard deviations for the estimated coefficients and idiosyncratic errors for each combination of USDA and the representative analyst's forecasts, *i.e.*, USDA *vs.* the average or median of the representative analyst's forecasts, for corn and soybeans. The range of the standard errors of the unforecastable shocks is also reported. Table 4.3 reports the medians and 95% credible intervals for the intercepts and the slopes. Gelman and Rubin (1992) test statistics are below 1.1 for all parameters of all four datasets, suggesting convergence of the Markov Chains.

4.6.1 Corn

The point estimate of the intercept α_A represents the bias of the USDA forecasts. The estimates are positive but insignificant for the two dataset, suggesting that we cannot rule out that USDA forecasts are unbiased. The point estimate of the intercept α_B represents the bias of the

representative analyst's forecasts. The estimates are also positive but insignificant, thus we cannot find evidence that the representative analyst's forecasts are biased.

The point estimate of the slope β_{1A} measures the relationship between the USDA forecasts and their immediate preceding forecast revisions. The estimates are 3.6% with analysts' average forecasts as the representative and 4.31% with analysts' median forecasts as the representative. The estimates are positive but insignificant, suggesting that there is no significant impact from the immediate preceding forecast revisions. Similarly, the point estimate of the slope β_{1B} measures the counterpart for the representative analyst's forecasts. The estimates are negative, at -6.31% with analysts' average forecasts as the representative and -8.42% with analysts' median forecasts as the representative. However, the estimates are also insignificant.

The point estimates of β_{2A} and β_{2B} measure the efficiency with respect to own past forecast revisions. The estimate of β_{2A} is 19.24% with analysts' average forecasts as the representative. The estimate is significant at 5% level. That is to say, if USDA adjusts its forecasts up by 1% in the past month, its forecasts will also be revised up by roughly 0.19% on average. The estimate of β_{2A} is 14.49% with analysts' median forecasts as the representative. The estimate is smaller, but also significant at 5% level. The estimates of β_{2B} are 16.31% and 15.01% for the two dataset respectively. Estimates are positive and also significant.

For the dataset with analysts' average forecasts as the representative, the estimates of the standard errors of Type A unforecastable shocks range from 0.73% to 26.29%. For Type B unforecastable shocks, the range is 1.82% - 13.06%. The upper limit of the Type A range is larger, which is consistent with the fact that Type A time intervals are larger than Type B time intervals. The estimates of the standard errors of idiosyncratic error are 1.64% and 0.56% on

average, due to the same argument. Similar results are found for the dataset with analysts' median forecasts as the representative.

4.6.2 Soybeans

The estimates of parameters for soybeans exhibit different patterns. the point estimates of α_{1A} , β_{1A} and β_{2A} are all significantly different from zero. Specifically, the estimates of α_A are -2.41% and -2.39%, respectively, for the two dataset, indicating that USDA has a tendency to overestimate the ending stocks. The estimates of β_{1A} are 63.59% and 66.39% respectively. The estimates of β_{2A} are 41.65% and 37.46% respectively. The point estimates of α_B , β_{1B} and β_{2B} , however, are different from the USDA counterpart. Although the estimates are all positive for both datasets, they are insignificantly different from zero. This finding shows that there is not enough evidence to reject the efficiency of the representative analyst's forecasts.

For the error covariance parameters, the estimated standard errors range from 1.61% to 29.89% (1.6% - 31.14%) for Type A shocks and from 2.13% to 15.01% (1.54% - 14.96%) for Type B shocks. The estimated standard errors of the idiosyncratic errors are 5.75% (5.93%) for Type A and 0.94% for Type B. Similar to the case in corn, the upper limits for the Type A shocks are larger, and the size of the Type A idiosyncratic errors are larger. It is also interesting to find that all these error covariance parameters are larger than the respective ones for corn. The results thus indicates that the shocks in soybeans ending stocks are larger than those in corn ending stocks, and both USDA and analysts have lower precision in forecasting the soybeans ending stocks.

4.6.3 Discussion: Forecasting Behaviors

The forecasting behavior for both the USDA and the analysts can be inferred from the estimation results. For corn, because the estimate for β_{1A} is not significantly different from zero, similar to the arguments in Granger and Newbold (1973, 1986) we can conclude that the USDA forecast $U_{t,n}$ encompasses the analysts' forecast $V_{t,n}$ in forecasting $V_{t,n-1}$. In other words, $U_{t,n}$ has all the information in $V_{t,n}$ and can substitute the later in forecasting $V_{t,n-1}$. Thus we can claim that the USDA is actually forecasting the upcoming analysts' forecasts. Similar arguments can be applied to the representative analyst's forecasts, as β_{1B} is found to be insignificant. The forecasts made by the USDA and analysts, however, are still inefficient with respect to their own forecast revision history because of the significant positive estimates of β_{2A} and β_{1B} .

For soybeans, the forecast encompassing argument for the USDA is rejected due to the significantly positive estimate of β_{1A} . That is to say, the USDA forecast $U_{t,n}$ does not contain all the information in $V_{t,n}$, thus cannot be viewed as the forecast of $V_{t,n-1}$. Hence it can be argue that the USDA is not forecasting the upcoming analyst's forecasts. Besides, the USDA forecasts are inefficient because of the significant estimates of α_{1A} , β_{1A} and β_{2A} . In contrast, for the representative analyst's forecasts, the forecast encompassing argument cannot be rejected. Thus the representative analyst's forecasts can be viewed as the forecasts of the upcoming USDA forecasts. Moreover, there is not enough evidence that the analyst's forecasts are inefficient, because the estimated α_B , β_{1B} and β_{2B} are not significantly different from zero.

It is interesting to find that, for both commodities, the analysts are actually forecasting the USDA forecasts. Thus, we can argue that the private analysts focus more on the short-term targets. As discussed in the introduction, this could be possibly because of the analysts' role in the market. Successful forecasts of the upcoming USDA forecasts can help the analysts mitigate

the risk or even gain profit from trading. This is more attractive to the private sector as they are not required by policy to provide accurate and objective forecasts to the public.

4.7 Conclusions

We developed a model to jointly investigate the efficiency of the USDA and analysts' ending stocks forecasts. The model recognizes the alternating forecast structure and performs the estimations of USDA and the representative analyst's forecasts in a single system. In this way, we take into account the interaction of these two forecasters which have been overlooked in previous studies. The model is applied to USDA and analysts' forecasts of ending stocks for corn and soybeans from marketing years 2004/05 to 2013/14. A Bayesian MCMC method is developed for estimations. Results show that for corn, the USDA and analysts are forecasting each other, but their forecasts are inefficient. For soybeans, the USDA is targeting the ending stocks, and private analysts are efficiently forecasting the USDA forecasts. Thus, it can be concluded that private analysts focus more on the short-term. Hence, future models can directly build on the assumption that the forecasting target of private analysts is the USDA forecasts instead of the ending stocks of corn and soybean.

4.8 References

- Adjemian, M.K. 2012. Quantifying the WASDE Announcement Effect. *American Journal of Agricultural Economics* 94(1): 238-256.
- Clements, M.P., F. Joutz, and H.O. Stekler. 2007. An Evaluation of the Forecasts of the Federal Reserve: a Pooled Approach. *Journal of Applied Econometrics*. 22: 121-136.
- Davies, A., and K.Lahiri. 1995. A New Framework for Testing Rationality and Measuring Aggregate Shocks Using Panel Data. *Journal of Econometrics* 68: 205-227.

- Davies, A., and K.Lahiri. 1999. Re-examining the Rational Expectations Hypothesis Using Panel Data on Multi-period Forecasts. *Analysis of Panels and Limited Dependent Variable Models*, 226-254. Cambridge University Press.
- Egelkraut, T.M., P. Garcia, S.H. Irwin, and D.L. Good. 2003. An Evaluation of Crop Forecast Accuracy for Corn and Soybeans: USDA and Private Information Agencies. *Journal of Agricultural and Applied Economics* 35:79-95.
- French, K.R., R. Leftwich and N. Uhrig (1989). The Effects of Scheduled Announcements on Futures Markets. Working Paper 373, *Center for Research in Security Prices, Graduate School of Business, University of Chicago*.
- Garcia, P., S. Irwin, R. Leuthold, and L. Yang. 1997. The Value of Public Information in Commodity Futures Markets. *Journal of Economic Behavior and Organization* 32(4): 559-570.
- Granger, C.W.J., and P. Newbold. 1973. Some Comments on the Evaluation of Economic Forecasts. *Applied Economics* 5:35-47.
- Granger, C.W.J., and P. Newbold. 1986. *Forecasting Economic Time Series* (2nd ed.), Orlando, FL: Academic Press.
- Gelman, A. 2006. Prior Distributions for Variance Parameters in Hierarchical Models (Comment on the Article by Browne and Draper). *Bayesian Analysis* 1(3):515-534.
- Gelman A. and D.B. Rubin. 1992. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*. 7(4):457-472.
- Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008a. The Impact of Situation and Outlook Information in Corn and Soybeans Futures Markets: Evidence from WASDE Reports. *Journal of Agricultural and Applied Economics* 40(1):89-103.
- Isengildina-Massa, O., S.H. Irwin, D.L. Good, and J.K. Gomez. 2008b. Impact of WASDE Reports on Implied Volatility in Corn and Soybeans Markets. *Agribusiness* 24(4): 473-490.
- 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.
- Mincer, J. and V. Zarnowitz. 1969. The Evaluation of Economic Forecasts, in: J. Mincer (ed.) *Economic Forecasts and Expectations*. New York: National Bureau of Economic Research.
- Nordhaus , W.D. 1987. Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics* 69:667-674.

4.9 Appendix

Conditional Posterior Distributions for Model Parameters in the Gibbs Sampler

The proposed model consists of system (4.15) and priors (4.19):

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{W}\mathbf{K} + \mathbf{P}\boldsymbol{\Sigma}$$

$$\boldsymbol{\beta} \sim N(\mathbf{M}, \mathbf{V})$$

$$\varepsilon_{t,n,A} \sim N(0, \sigma_A^2), \varepsilon_{t,n,B} \sim N(0, \sigma_B^2) \quad (4.A.1)$$

$$k_{t,n,A} \sim N(0, \sigma_{n,A}^2), k_{t,n,B} \sim N(0, \sigma_{n,B}^2)$$

$$\sigma_A, \sigma_B, \sigma_{n,A}, \sigma_{n,B} \sim \text{Uniform}(0, \infty)$$

for $n = 1, \dots, N, t = 1, \dots, T$. Let the subscript m be either A or B . Let $\boldsymbol{\Omega} \equiv \mathbf{P}\boldsymbol{\Sigma}(\mathbf{P}\boldsymbol{\Sigma})'$. Given $\{\boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega}\}$, the dependent variable \mathbf{Y} follows a multivariate normal distribution:

$$\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega} \sim N(\mathbf{X}\boldsymbol{\beta} + \mathbf{W}\mathbf{K}, \boldsymbol{\Omega}) \quad (4.A.2)$$

and the likelihood is $\Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega})$. The posterior density of the set of model parameters is given by

$$\begin{aligned} p(\Lambda) = & \Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega}) \prod_{jt} \Phi(k_{t,j,A} | \sigma_{j,A}^2) \prod_{jt} \Phi(k_{t,j,B} | \sigma_{j,B}^2) \\ & * \Phi(\boldsymbol{\beta} | \mathbf{M}, \mathbf{V}) p(\sigma_A^2) p(\sigma_B^2) \prod_{j=1}^N p(\sigma_{j,A}^2) \prod_{j=1}^N p(\sigma_{j,B}^2) \end{aligned} \quad (4.A.3)$$

The conditional posterior density for $\boldsymbol{\beta}$ is

$$p(\boldsymbol{\beta} | \Lambda \setminus \boldsymbol{\beta}) = \Phi(\mathbf{Y} | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega}) * \Phi(\boldsymbol{\beta} | \mathbf{M}, \mathbf{V}) \quad (4.A.4)$$

Hence:

$$\begin{aligned} \boldsymbol{\beta} | \Lambda \setminus \boldsymbol{\beta} \sim & N((\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X} + \mathbf{V}^{-1})^{-1}(\mathbf{X}'\boldsymbol{\Omega}^{-1}(\mathbf{Y} - \mathbf{W}\mathbf{K}) \\ & + \mathbf{V}^{-1}\mathbf{M}), (\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X} + \mathbf{V}^{-1})^{-1}) \end{aligned} \quad (4.A.5)$$

The conditional posterior density of $k_{t,n,m}, t = 1, \dots, T, n = 1, \dots, N, m = A, B$ is

$$p(k_{t,n,m} | \Lambda \setminus k_{t,n,m}) = \Phi(Y | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega}) * \Phi(k_{t,n,m} | \sigma_{n,m}^2) \quad (4.A.6)$$

Therefore,

$$k_{t,n,m} | \Lambda \setminus k_{t,n,m} \sim N\left(\frac{\mathbf{W}'_{k_{t,n,m}} \boldsymbol{\Omega}^{-1} (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}_{-k_{t,n,m}} \mathbf{K}_{-k_{t,n,m}})}{\mathbf{W}'_{k_{t,n,m}} \boldsymbol{\Omega}^{-1} \mathbf{W}_{k_{t,n,m}} + 1/\sigma_{n,m}^2}, \frac{1}{\mathbf{W}'_{k_{t,n,m}} \boldsymbol{\Omega}^{-1} \mathbf{W}_{k_{t,n,m}} + 1/\sigma_{n,m}^2}\right) \quad (4.A.7)$$

where $\mathbf{W}_{k_{t,n,m}}$ is the column of \mathbf{W} which indicates the monthly shock $k_{t,n,m}$, and $\mathbf{W}_{-k_{t,n,m}}$,

$\mathbf{K}_{-k_{t,n,m}}$ are matrices with the column indicating $k_{t,n,m}$ deleted from \mathbf{W}, \mathbf{K} , respectively.

The conditional posterior density of $\sigma_{n,m}^2, n = 1, \dots, N, m = A, B$, is

$$p(\sigma_{n,m}^2 | \Lambda \setminus \sigma_{n,m}^2) = \prod_{j=1}^N \Phi(k_{t,n,m} | \sigma_{n,m}^2) * p(\sigma_{n,m}^2) \quad (4.A.8)$$

Thus

$$\sigma_{n,m}^2 | \Lambda \setminus \sigma_{n,m}^2 \sim IG((T-1)/2, \sum_{t=1}^T k_{t,n,m}^2 / 2) \quad (4.A.9)$$

Finally, the conditional posterior of $\sigma_m^2, m = A, B$, is

$$p(\sigma_m^2 | \Lambda \setminus \sigma_m^2) = \Phi(Y | \boldsymbol{\beta}, \mathbf{K}, \boldsymbol{\Omega}) * p(\sigma_m^2) \quad (4.A.10)$$

so that

$$\sigma^2 | \Lambda \setminus \sigma^2 \sim IG((TN-1)/2, (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}\mathbf{K})' (\mathbf{P}_m \mathbf{P}_m')^{-1} (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{W}\mathbf{K}) / 2) \quad (4.A.11)$$

where $\mathbf{P}_m = \mathbf{I}_m \mathbf{P}$. \mathbf{I}_m is an indicator vector which indicate the locations of Type m idiosyncratic errors in $\boldsymbol{\Sigma}$. In particular, $\mathbf{I}_A = [1 \ 0 \ 1 \ 0 \ \dots \ 1 \ 0]'$, $\mathbf{I}_B = [0 \ 1 \ 0 \ 1 \ \dots \ 0 \ 1]'$.




Figure 4.1. Timeline of USDA and analysts' forecasts of ending stocks for corn and soybeans

Table 4.1. Descriptive statistics for USDA and analysts forecast revisions for corn and soybeans

Data	Mean	Median	St. Dev.	Min	Max
<hr/>					
Corn, Average					
Type A Revisions	0.0078	-0.0019	0.1111	-0.5966	0.3676
Type B Revisions	0.0002	-0.0011	0.0702	-0.2413	0.1840
USDA own Revisions	0.0097	0.0000	0.1249	-0.5988	0.4098
Analysts own Revisions	0.0080	0.0034	0.1271	-0.6372	0.3805
<hr/>					
Corn, Median					
Type A Revisions	0.0084	0.0000	0.1111	-0.5626	0.3746
Type B Revisions	-0.0004	0.0000	0.0712	-0.2468	0.2011
USDA own Revisions	0.0097	0.0000	0.1249	-0.5988	0.4098
Analysts own Revisions	0.0080	0.0066	0.1294	-0.5967	0.4260
<hr/>					
Soybeans, Average					
Type A Revisions	-0.0068	-0.0192	0.1220	-0.3724	0.6496
Type B Revisions	-0.0064	-0.0029	0.0740	-0.2776	0.1912
USDA own Revisions	-0.0132	0.0000	0.1387	-0.3830	0.7569
Analysts own Revisions	-0.0132	-0.0123	0.1388	-0.4953	0.5821
<hr/>					
Soybeans, Median					
Type A Revisions	-0.0070	-0.0040	0.1193	-0.4212	0.6799
Type B Revisions	-0.0047	0.0000	0.0757	-0.2744	0.1719
USDA own Revisions	-0.0132	0.0000	0.1387	-0.3830	0.7569
Analysts own Revisions	-0.0117	-0.0082	0.1405	-0.4533	0.6747

Note: summary statistics are displayed in logarithms.

Table 4.2. Means and standard deviations for the estimates of USDA and analysts ending stocks forecasts, 2004/05 – 2013/14

Parameter	Corn				Soybeans			
	USDA vs. Ana. Avg.		USDA vs. Ana. Med.		USDA vs. Ana. Avg.		USDA vs. Ana. Med.	
	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)	Mean	(St. dev.)
Coefficient								
α_A	0.0013	(0.0026)	0.0008	(0.0021)	-0.0242	(0.0066)**	-0.0239	(0.0067)**
β_{1A}	0.0360	(0.0435)	0.0431	(0.0347)	0.6359	(0.1346)**	0.6639	(0.1031)**
β_{2A}	0.1924	(0.0498)**	0.1449	(0.0389)**	0.4165	(0.0765)**	0.3746	(0.0831)**
α_B	0.0010	(0.0043)	0.0014	(0.0040)	0.0071	(0.0065)	0.0067	(0.0070)
β_{1B}	-0.0631	(0.0691)	-0.0842	(0.0678)	0.0618	(0.0694)	0.1143	(0.0820)
β_{2B}	0.1631	(0.0568)**	0.1501	(0.0548)**	0.0682	(0.0582)	0.0610	(0.0579)
Shocks (range)								
σ_{kA}	0.0073 – 0.2629		0.0049 – 0.2791		0.0161 – 0.2989		0.0160 – 0.3114	
σ_{kB}	0.0182 – 0.1306		0.0183 – 0.1318		0.0213 – 0.1501		0.0154 – 0.1496	
Idiosyncratic Err.								
σ_A	0.0164	(0.0026)	0.0135	(0.0020)	0.0575	(0.0053)	0.0593	(0.0055)
σ_B	0.0056	(0.0033)	0.0033	(0.0021)	0.0094	(0.0054)	0.0094	(0.0050)

Note: (*) and (**) denote parameter estimates significant at 10% and 5%, respectively. The standard errors for the idiosyncratic errors and shocks are all significant at 5% level, and hence the indicators are omitted.

Table 4.3. Medians and 95% credible intervals for the estimates of USDA and analysts ending stocks forecasts, 2004/05 – 2013/14

Parameter	USDA vs. Ana. Avg.			USDA vs. Ana. Med.		
	2.5%	Median	97.5%	2.5%	Median	97.5%
Corn						
Coefficient						
α_A	-0.0041	0.0013	0.0062	-0.0033	0.0008	0.0049
β_{1A}	-0.0469	0.0359	0.1262	-0.0243	0.0430	0.1092
β_{2A}	0.0998	0.1916	0.2967	0.0713	0.1446	0.2222
α_B	-0.0074	0.0010	0.0093	-0.0057	0.0014	0.0090
β_{1B}	-0.1915	-0.0643	0.0735	-0.2103	-0.0810	0.0505
β_{2B}	0.0495	0.1621	0.2872	0.0458	0.1478	0.2597
Idiosyncratic Err.						
σ_A	0.0117	0.0164	0.0216	0.0100	0.0133	0.0180
σ_B	0.0004	0.0054	0.0126	0.0002	0.0030	0.0081
Soybeans						
Coefficient						
α_A	-0.0374	-0.0241	-0.0113	-0.0367	-0.0240	-0.0105
β_{1A}	0.3708	0.6363	0.9192	0.4471	0.6736	0.8476
β_{2A}	0.2638	0.4168	0.5651	0.2132	0.3742	0.5365
α_B	-0.0057	0.0070	0.0201	-0.0072	0.0068	0.0201
β_{1B}	-0.0677	0.0594	0.2020	-0.0430	0.1134	0.2812
β_{2B}	-0.0418	0.0681	0.1851	-0.0538	0.0604	0.1748
Idiosyncratic Err.						
σ_A	0.0480	0.0571	0.0686	0.0495	0.0590	0.0708
σ_B	0.0007	0.0091	0.0204	0.0009	0.0093	0.0196