



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

**Should Farmers Follow the Recommendations of Market Advisory Services?
A Bayesian Approach to Performance Evaluation**

by

Silvina M. Cabrini, Scott H. Irwin and Darrel L. Good¹

*Paper presented at the NCR-134 Conference on Applied Commodity Price
Analysis, Forecasting and Market Risk Management
St. Louis, Missouri, April 17-18, 2006*

Copyright 2005 by Silvina M. Cabrini, Scott H. Irwin and Darrel L. Good. All rights reserved.
Readers may make verbatim copies of this document for non-commercial purposes by any
means, provided that this copyright notice appears on all such copies.

¹ Silvina M. Cabrini (cabrini@uiuc.edu) is a Graduate Research Assistant for the AgMAS Project in the Department of Agricultural and Consumer Economics at the University of Illinois at Urbana-Champaign. Scott H. Irwin is the Laurence J. Norton Professor of Agricultural Marketing and Darrel L. Good is a Professor in the Department of Agricultural and Consumer Economics at the University of Illinois at Urbana-Champaign.

Should Farmers Follow the Recommendations of Market Advisory Services?

A Hierarchical Bayesian Approach to Estimation of Expected Performance

Abstract

This paper employs a Bayesian hierarchical approach to estimate individual expected performance of market advisory programs in corn and soybeans. This estimation procedure is a conservative approach compared to traditional estimation, since it reduces estimation error in the expected gains from following top-performing advisory programs. Three versions of the model are estimated. The first combines information across the entire sample, while the second includes skeptical beliefs based on the efficient market hypothesis. The third divides programs into two groups based on the degree of activeness in marketing recommendations. Results indicate that even when skeptical beliefs are incorporated into the model a few programs in corn and several programs in soybeans appear to be better marketing alternative compared to a naïve strategy that mimics the market benchmark. More specifically, a skeptical farmer can expect to increase the price received for corn by 1% and the price received for soybeans by 5% following the single top-ranked program.

Key Words: Bayesian hierarchical models, corn, market advisory service, pricing performance, soybeans

Introduction

Agricultural market advisory services are popular with U.S. farmers (Patrick et al., 1998; Norvell and Lattz, 1999). For a subscription fee, these firms provide market analysis and pricing advice to farmers. In particular, they make recommendations on how to market crops using various instruments, including cash sales, forward, futures, and options contracts. Advisory services typically deliver reports with market information and marketing recommendations via daily email or web pages, with some offering multiple updates each day. Marketing recommendations are specific, indicating the portion of a crop that should be marketed, the marketing tool, and the timing of transactions. For example a service can recommend, “Buy May 2005 soybean puts today with a strike price of \$5.00/bu. for 50% of expected production.” Market advisory services conduct market research and employ fundamental and/or technical analysis to identify profitable marketing alternatives.

In 1994, the Agricultural Market Advisory Services (AgMAS) Project was initiated at the University of Illinois to evaluate the performance of agricultural market advisory services. The AgMAS Project has evaluated at least 23 advisory programs for 10 crop years. AgMAS subscribes to the services that are followed and records marketing recommendations on a real-time basis. The price that a crop farmer in central Illinois would receive for each crop by following the recommendations of each advisory service is computed and compared to external benchmarks. Empirical findings have been disseminated through various AgMAS research reports, providing valuable information to farmers selecting a market advisor. The most recent report presents the pricing performance in corn and soybeans for the 1995 to 2004 crop years (Irwin et al., 2006a). The average price obtained by exactly following the recommendations of market advisors is higher than the average price offered by the market for both crops, although

the price differences are not large. The average price difference between advisory prices and market benchmarks is 2¢/bu. to 5¢/bu. for corn and between 14¢/bu. to 16¢/bu. for soybeans. This research provides only weak evidence of advisory services as a group outperforming external benchmarks.

Previous studies on advisory services focus on the performance of services as group, without emphasizing expected performance of individual programs (Irwin et al., 2006a; 2006b). However, there is a wide range of pricing performance across advisory services. Moreover, there is some evidence of performance persistence at the extremes of advisory service performance rankings, which suggests that a subset of services may have more attractive expected performance than other services. Further evidence in this regard is provided by the results of the first chapter in this dissertation, which indicate that more “active” market advisory programs tend to have better pricing performance compared to less active programs. Therefore, farmers and other market participants are likely to be interested in the pricing performance of individual advisory programs.

Since numerous advisory programs are available to farmers, choosing between them requires computation of individual expected performance estimates. The simplest procedure to compute expected performance for a given advisory program is to average past performance observations for that program. However, this procedure requires estimation of numerous individual parameters based on a relatively small number of past observations. The most complete report on advisory service performance to date in corn and soybean markets provides data on 10 or fewer crop years for about 25 advisory services each year (Irwin et al., 2006a). In estimation problems like this, individual sample averages tend to over-fit the data and unusually high and low expected performance values may appear due to estimation error. This implies that the gains from following top performing programs are likely to be overestimated when traditional estimates are considered (e.g., Jorion, 1986; Michaud, 1989; Marcus, 1990; Grauer, 2002)²

An alternative estimation procedure is to combine the pricing information for all advisory services and compute pooled estimates of performance, as in previous research (Irwin et al., 2006a; 2006b). A pooled estimator is more reliable since it is based on a larger number of past observations. However, this approach imposes the potentially restrictive assumption that all services have identical expected pricing performance.

Neither traditional sample averages nor pooled estimates provide the most appropriate information for farmers considering contracting with advisory services. Instead, a model that combines the information for the group of advisory programs without assuming that all programs have equal expected performance is a more reasonable estimation procedure. Bayesian hierarchical models have this characteristic. A hierarchical model based on the normal distribution produces estimators that are weighted averages of separate and pooled estimates. This type of estimator is called a *shrinkage* estimator, since individual estimates are shrunk to common values. Shrinkage estimators have been applied to different estimation problems. For example, Efron and Morris (1975) discuss the application of shrinkage estimators to predict the

² Marcus (1990) analyses a problem related to the one studied in the current article. His paper focuses on determining whether a good performance history of the top-performing fund manager could be simple due to chance or it is more likely to be related to superior manager skills.

batting averages for baseball players and the incidence of toxoplasmosis for cities in El Salvador. Gelman et al. (2004) employ a Bayesian hierarchical model to estimate the effect of several coaching programs on SAT test scores. Allenby et al. (2005) employ a Bayesian hierarchical model to estimate consumer preferences. In the finance literature, several studies have employed shrinkage estimators to compute expected stock returns, and results indicate that these estimators outperform traditional sample estimates, in particular when the sample size is small (e.g., Jorion, 1986; Grauer, 2002).

The simplest application of Bayesian hierarchical models to individual performance estimation is to consider the sample of advisory programs as a whole, with separate performance estimates shrunk towards the overall pooled value. Alternatively, some farm decision-makers may be skeptical about the ability of advisory services to outperform the market, and therefore, unwilling to base expected performance exclusively on past performance observations. For instance, a farmer may be strongly influenced by the efficient market hypothesis and believe that corn and soybean markets generally are efficient. The corollary belief is that it is difficult if not impossible to enhance income based on the marketing recommendations of advisory services (Brorsen and Anderson 1994; Zulauf and Irwin, 1998; Tomek and Peterson, 2005). The views of a strong believer in the efficient market hypothesis can be incorporated in the Bayesian hierarchical model by adding a prior whose parameters imply that all services have an expected performance close to zero. A similar problem has been considered in the finance literature for the performance of mutual fund managers (e.g., Baks et al., 2001).

The two versions of the Bayesian hierarchical model discussed above assume there is no a priori reason to subdivide programs into groups with similar characteristics. However, this may not be the best assumption. For instance, consider the belief that advisors with superior information and performance tend to recommend more “active” marketing programs. This belief is consistent with the results presented in the first essay of this dissertation, where it was found that more active market advisory programs tend to outperform more conservative programs. Under this view, it is reasonable to group advisory programs based on the degree of activeness and estimate a Bayesian hierarchical model for each group.

The purpose of this essay is evaluate whether farmers can expect higher than average market prices by selling crops following the recommendations of individual market advisory programs in corn and soybeans. Hence, the focus of this study is the estimation of expected performance for individual programs. A Bayesian hierarchical approach is employed to estimate individual expected performance for advisory services tracked by the AgMAS Project in corn and soybeans over 1995 through 2004. Three versions of the model are estimated: 1) a version based exclusively on the sample data, 2) a version that includes skeptical beliefs about the performance of advisory services, and 3) a version that divides programs into two groups based on the degree of activeness in marketing recommendations. The posterior distribution of individual expected performance is computed by simulation. The numerical simulation uses the inverse cumulative distribution function method and sampling from normal distributions (Gelman et al., 2004). Bayesian point estimates and 90% confidence intervals are employed to identify subsets of programs that represent attractive marketing alternatives for farmers.

Data and Non-Bayesian Estimates

Data on corn and soybean net advisory prices for the 1995 to 2004 crop years are obtained from AgMAS project's records. All programs with two or more performance observations will be considered in this study.³ A complete list of programs tracked by AgMAS is presented in table 1.

Four different benchmarks are employed in previous AgMAS evaluations of pricing performance of market advisory services (Irwin et al., 2006a). Two market benchmarks represent the average price offered by the market, net of storage costs, over the indicated marketing window. The marketing window for the 24-month market benchmark starts in August of the year before harvest and finishes in August after harvest. The marketing window for the 20-month market benchmark starts in January of the harvest year and finishes in August after harvest. Two farmer benchmarks represent the average price received by farmers net of storage costs. The first farmer benchmark is the average price received by farmers as reported by the USDA. The second farmer benchmark is based on Illinois cash bid prices and the marketing weights reported by the USDA. A detailed description of the procedures employed to compute advisory prices and benchmarks is presented in Irwin et al. (2006a). In the current study results were computed using the four benchmarks; however a complete set of results is presented only for the 24-month market benchmark. The results computed using the other benchmarks are mentioned briefly in the discussion section and are available from the authors upon request.

The primary measure of advisory program performance is the difference between the price received by a farmer who markets grain following a program's recommendations and a given benchmark price:

$$(1) \quad y_{jt} = NAP_{jt} - BP_t$$

where NAP_{jt} is the net advisory price for program j in crop year t and BP_t is the benchmark price in crop year t . The traditional estimator for an advisory program's performance is simply the individual sample average:

$$(2) \quad \bar{y}_j = \frac{1}{T_j} \sum_{t=1}^{T_j} y_{jt}$$

where T_j is the number of past performance observations available for program j . Traditional estimators are commonly used because they are straightforward to compute and understand. However, they have the drawback that unusually high and low values may appear due to estimation error, in particular when the number of time series observations is low and the number of advisory programs is large.

Figure 1 shows traditional point estimates and 90% confidence intervals for the expected performance of advisory programs tracked by the AgMAS project. The values in figure 1 are obtained by estimating expected performance separately for each advisory program. Point estimates are sample averages (equation 2) and confidence intervals are computed using the standard errors for these averages. Panel A shows expected performance estimates for corn. Point estimates range from 30¢/bu. above the market benchmark to 19¢/bu. below. Slightly more than half of the programs have positive expected pricing performance, but only three have

³ A minimum of two observations is necessary to estimate the standard error of the separate estimate of expected performance.

an expected price that is significantly greater than the benchmark price ($\alpha = 0.10$). In addition, three programs have significantly negative performance in corn. Panel B shows expected performance estimates for soybeans. Point estimates range from 70¢/bu. above the benchmark to 29¢/bu. below. Nearly 80% of the programs have positive point estimates for expected performance and eight have significantly positive performance ($\alpha = 0.10$). None of the programs have significantly negative performance in soybeans.

An alternative approach for the estimation of expected pricing performance of market advisory services is to assume that there is not enough data to estimate individual performance, and therefore, information is pooled to obtain one estimate of expected performance for the group of advisory programs. The precision-weighted pooled estimate is:

$$(3) \quad \hat{y}^{pool} = \frac{\sum_{j=1}^N \frac{1}{\hat{\sigma}_{\bar{y}_j}^2} \bar{y}_j}{\sum_{j=1}^N \frac{1}{\hat{\sigma}_{\bar{y}_j}^2}}$$

where N is the number of advisory programs considered and $\hat{\sigma}_{\bar{y}_j}^2$ is the variance estimator of \bar{y}_j . Note that this pooled estimator is different from the simple average of individual expected performance across programs. The pooled estimator is a weighted average, where the weights are the inverse of the squared standard error of each estimate. A simple average would be reasonable under the assumption that individual estimates have the same error, or in other words that the standard deviation of performance is the same for all programs. However, as presented in figure 1, the data employed in this study suggest that standard deviation of the performance is different across programs, and hence a weighted pooled estimate is more appropriate.

Pooled estimates are a measure of performance for advisory programs as a group. Under the assumption that all advisory programs have the same expected performance, pooled estimates fully describe expected performance for the group of advisory programs considered. The last dot in each of the panels in figure 1 is the pooled estimate for expected performance. This value is 0.5¢/bu. for corn with a 90% confidence interval of -0.6¢/bu. to 1.7¢/bu. The pooled estimate for soybeans is 8.5¢/bu. with a 90% confidence interval of 5.8¢/bu. to 11.1¢/bu. These estimates imply that farmers should be indifferent between following an advisory program in corn and applying a naïve strategy of spreading sales along the marketing window, but farmers would prefer following recommendations from an advisory program in soybeans.

Bayesian Hierarchical Model

Separate individual or pooled estimates imply extreme assumptions about the expected performance of advisory programs. On one hand, separate estimates imply that expected performance is completely independent across programs. On the other hand, pooled estimates imply that all programs have the same expected performance. A situation between these two alternatives seems more reasonable and can be implemented by Bayesian hierarchical models. A hierarchical model based on the normal distribution produces *shrinkage* estimators that are weighted averages of individual and pooled estimates. For example, the shrinkage estimator for

the expected performance of advisory program j is a weighted average of the individual sample mean (\bar{y}_j) and the Bayesian pooled estimate $(\hat{\mu}^{pool})$:⁴

$$(4) \quad \hat{\theta}_j^{shrink} = (1 - w) \bar{y}_j + w \hat{\mu}^{pool}.$$

The coefficient w is defined as the *shrinkage intensity* since it indicates how much individual estimates are shrunk towards pooled values.

A simplified diagram of the structure of the normal hierarchical model for market advisory program performance is presented in figure 2. There are two levels of parameters in this model. The expected performance for each advisory program, $\theta = (\theta_1, \dots, \theta_N)$, is in the lower level and hyperparameters, (μ, τ) , that combine information for all programs in the sample are shown in the higher level. The general structure of a Bayesian hierarchical model includes a prior distribution for the parameters, $p(\theta)$, that can be decomposed into a conditional prior given the hyperparameters, $p(\theta|\mu, \tau)$, and the prior for the hyperparameters, $p(\mu, \tau)$, sometimes called the hyper-prior:

$$(5) \quad p(\theta) = p(\mu, \tau) p(\theta|\mu, \tau).$$

Then the related joint posterior distribution can be expressed as:

$$(6) \quad p(\mu, \tau, \theta|y) \propto p(\mu, \tau, \theta) p(y|\mu, \tau, \theta) = p(\mu, \tau, \theta) p(y|\theta)$$

where y is the sample information (data distribution). The last equality holds because the hyperparameters affect $p(y)$ only through the parameters θ . The key characteristic of this model is that individual performance parameters share a common prior. This prior distribution is not subjective or based on information that precedes data collection; instead it is constructed from the whole sample. In this context, not only data on price performance of a particular program is helpful in estimating the expected performance for that program, but also information from the rest of the programs contributes to the estimation. A detailed derivation of the probabilistic model for hierarchical models under normality is presented in Gelman et al. (2004). A description of the main points of the model employed in this study follows.

Performance for program j is assumed to be normally distributed with mean θ_j and variance v_j^2 . The simplifying assumption that variances are known is made such that:⁵

$$(7) \quad y_{jt} | \theta_j \sim N(\theta_j, v_j^2). \quad \text{distribution of } y_{jt}$$

⁴ The formula for the pooled estimate is presented below in equation (3.10). It has a similar structure to the formula for \hat{y}^{pool} (equation 3.3).

⁵ Although this assumption is not true in actual applications, it is commonly used as a good approximation in this type of estimation procedures. Traditional sample estimates for the variances are employed (Gelman, 2004).

As mentioned before, individual expected performance estimates share the same prior. Specifically, this prior is a normal distribution with mean μ and variance τ :

$$(8) \quad \theta_j | \mu, \tau \sim N(\mu, \tau) \quad \text{prior distribution of } \theta_j$$

where the parameter τ defines the prior uncertainty and, as explained below, determines the shrinkage intensity. Combining the sample likelihood derived from equation (7) with the prior distribution (equation 8), the posterior distribution of θ_j conditional on μ and τ is obtained:

$$(9) \quad \theta_j | \mu, \tau, y \sim N(\hat{\theta}_j, V_j) \quad \text{conditional posterior distribution of } \theta_j$$

$$\text{where } \hat{\theta}_j = \frac{\frac{1}{\sigma_{\bar{y}_j}^2} \bar{y}_j + \frac{1}{\tau^2} \mu}{\frac{1}{\sigma_{\bar{y}_j}^2} + \frac{1}{\tau^2}}, \quad V_j = \frac{1}{\frac{1}{\sigma_{\bar{y}_j}^2} + \frac{1}{\tau^2}}, \quad \text{and } \sigma_{\bar{y}_j}^2 = \frac{v_j^2}{T_j}.$$

The above equation shows that the posterior distribution for each program's expected performance is also normal with a mean equal to the weighted average of the sample mean for that program and the mean of the prior distribution. Note that the point estimate for θ_j is the shrinkage estimator, \hat{y}_j^{shrink} , presented in equation (4).⁶ Note also that the greater the variance of the sample mean $\sigma_{\bar{y}_j}^2$ the more the individual estimate is shrunk towards μ . Finally, the greater the prior uncertainty, measured by τ^2 , the lower the shrinkage intensity.

Up to this point, the posterior distribution of expected performance is defined in terms of the hyperparameters μ and τ . A full Bayesian treatment of hierarchical models includes the definition of a prior distribution for the hyperparameters. Following Gelman et al. (2004), an uninformative prior is employed here for $\hat{\mu}$. The use of this uninformative prior in hierarchical models is reasonable since the entire sample is employed to estimate μ and the total number of observations is large enough to justify relying only on the sample for the estimation of this parameter. The posterior distribution of μ conditional on τ is also normal with a mean equal to the Bayesian precision pooled estimate ($\hat{\mu}$)⁷:

$$(10) \quad \mu | \tau, y \sim N(\hat{\mu}, V_\mu) \quad \text{conditional posterior distribution of } \mu$$

⁶ The shrinkage coefficient, w , is equal to $\frac{1}{\tau^2} / \left(\frac{1}{\sigma_{\bar{y}_j}^2} + \frac{1}{\tau^2} \right)$.

⁷ Note that \hat{y}^{pool} in equation (3.3) is similar to $\hat{\mu}$, with the difference that the prior uncertainty, τ , is included in the computation of $\hat{\mu}$.

$$\text{where } \hat{\mu} = \frac{\sum_{j=1}^N \frac{1}{\sigma_{\bar{y}_j}^2 + \tau^2} \bar{y}_j}{\sum_{j=1}^N \frac{1}{\sigma_{\bar{y}_j}^2 + \tau^2}} \text{ and } V_{\mu} = \left[\sum_{j=1}^N \frac{1}{\sigma_{\bar{y}_j}^2 + \tau^2} \right]^{-1}$$

Finally, the posterior distribution of τ is:

$$(11) \quad p(\tau|y) \propto p(\tau) V_{\mu}^{1/2} \prod_{j=1}^N (\sigma_{\bar{y}_j}^2 + \tau^2)^{-1/2} \exp \left(-\frac{(\bar{y}_j - \hat{\mu})^2}{2(\sigma_{\bar{y}_j}^2 + \tau^2)} \right) \quad \text{posterior distribution of } \tau$$

An uninformative uniform prior distribution for τ is also assumed. According to Gelman (2004) this type of distribution performs well when the number of groups (advisory programs), is greater than two or three, as is the case in this study. Note that the distribution of τ depends on the dispersion of y_{jt} within programs ($\sigma_{\bar{y}_j}^2$) and the dispersion of \bar{y}_j across programs $\left[(\bar{y}_j - \hat{\mu})^2 \right]$. For high variability of y_{jt} within programs and low dispersion of \bar{y}_j across programs, small values of τ will be more likely and the optimal shrinking intensity will be higher. Also, the number of observations has an effect on the shrinkage intensity. Separate estimates for programs with less time-series observations, which are less reliable, have higher values of $\sigma_{\bar{y}_j}^2$ and will be shrunk more towards pooled estimates.

Skeptical Beliefs

Some farm decision-makers may be skeptical about the ability of advisory services to outperform the market, and therefore, unwilling to base expected performance estimates exclusively on past performance observations, as in the basic hierarchical model outlined in the previous section. Farmers may be strongly influenced by the efficient market hypothesis and believe that it is difficult if not impossible to enhance income based on the marketing recommendations of advisory services (Brorsen and Anderson 1994; Zulauf and Irwin, 1998; Tomek and Peterson, 2005). These views can be incorporated in the Bayesian hierarchical model by adding a prior whose parameters imply that all services have an expected performance close to zero. In this study a normal distribution is employed as a prior for skeptical beliefs (equation 12). While there are other ways to model skeptical beliefs, the normal distribution is chosen here to simplify the computation of the posterior distribution.

$$(12) \quad \theta_j \sim N(\mu_0, \tau_0^2) \quad \text{prior distribution of } \theta_j \text{ under skeptical view}$$

The mean (μ_0) and variance (τ_0^2) of the skeptical prior depend on the strength of the skeptical beliefs. The stronger the skeptical beliefs, the closer μ_0 is to zero and the smaller the values of τ_0^2 . At one extreme, both μ_0 and τ_0^2 equal to zero implies that the decision maker is absolutely sure that advisory programs will obtain an average price equal to the market benchmark price

and no data set would change this view. On the other hand, small positive values of μ_0 and large values of τ_0^2 imply that skeptical beliefs are not as strong.

The parameters of the prior distribution for a given view can be obtained from the answers to the following simple questions about advisory services performance:

Question 1: What is the most likely difference between the price obtained by following the recommendations of advisory programs and the market benchmark price?

Question 2: What is the probability that an advisory program outperforms the market benchmark, on average, for more than 5%?

The answers to these two questions should not depend on the data set employed in the estimation. The mean of the prior is set equal to the answer to the first question and the variance is determined based on the mean and the answer to the second question. In particular, the skeptical view employed in this study corresponds to a decision maker who believes that expected performance is most likely to be zero but there is a 1% probability that expected performance is more than 5% of the market benchmark price. The computation of the prior parameters based on the answer to the questions is described in section 5. This approach is similar to the one employed by Baks et al. (2001), where a prior distribution for mutual funds managers' ability to beat the market is obtained from the answer to similar questions.⁸

The posterior distribution of expected performance from the Bayesian hierarchical model with skeptical beliefs θ_j^{skip} combines the posterior distribution of the hierarchical model with the skeptical prior to obtain the following posterior distribution:

$$(13) \quad \theta_j^{skip} | \mu, \tau, \mu_0, \tau_0, y \sim N(\hat{\theta}_j^{skip}, V_j^{2\ skip})$$

$$\text{where } \hat{\theta}_j^{skip} = \frac{\frac{1}{V_j} \hat{\theta}_j + \frac{1}{\tau_0^2} \mu_0}{\frac{1}{V_j} + \frac{1}{\tau_0^2}} \text{ and } V_j^{2\ skip} = \frac{1}{\frac{1}{V_j} + \frac{1}{\tau_0^2}}$$

Pricing Performance and Degree of Activeness

A first two versions of the Bayesian hierarchical model combine performance information for all advisory programs in the sample. These versions of the model are consistent with the view that, a priori (before the data collection), there is no reason to believe that certain advisory programs are superior to others. An alternative belief is that advisors with superior information and performance tend to recommend more “active” marketing programs. This belief is consistent

⁸ Baks et al. use a more complex functional form for the prior, an asymmetric distribution with a lower bound and a right tail of a normal distribution, which leads to a more complex posterior distribution of expected performance. The idea in this study is that fund managers are expected to loose at most the transaction costs, on average. The authors argue that losses greater than this limit imply consistently trading on misinformation or the existence of behavioral biases.

with the results presented in the first essay of this dissertation, where it was found that more active market advisory programs tend to outperform more conservative programs.

To accommodate the belief that more active advisory programs have superior performance, the third version of the model places advisory programs into groups based on an activeness index. The activeness index measures the extent that recommended transactions imply speculation (bets) on future price movements. A detailed description of the computation of the activeness index, along with average index values for each program is presented in the second chapter of this dissertation. The hierarchical clustering method of complete linkage (Johnson and Wichern, 2002) is employed to form the groups. A hierarchical clustering method starts with as many clusters (activeness groups) as individual objects (programs). The most similar objects are grouped and then groups are merged gradually according to their similarities. This method is used to divide the programs into two groups. Ten of the programs were classified in a more active group and 24 in a more conservative group; the list of programs for each group is presented in the results section. In this case there is a hierarchical structure (figure 1) within each activeness group. Advisory programs in each activeness group share a common prior distribution (equation 11) that combines the information of the programs included in the group, then, separate estimates are shrunk towards pooled values computed from a group of programs with a similar degree of activeness.

Simulation Procedures

In the hierarchical model it is assumed that observed data (pricing performance observations) are normally distributed in the population with a different mean for each group (program), and the group means (average performance for each program) are also normality distributed. To assess whether these normality assumptions are supported by the data the Jarque-Bera normality test was applied to the pricing performance observations for each of the programs in corn and soybeans, as well as the distribution of average performance across programs. Normality was not rejected at the 5% significance level for any distribution of individual program pricing performance. Normality of the distribution of average performance across programs is rejected at the 5% level for corn and but not for soybeans. Although these results indicate some evidence of departures from normality, overall, the normal hierarchical model is a reasonable estimation alternative for the problem being evaluated in the current study.

The computation of the posterior distribution of expected performance for the first version of the model is accomplished via simulation in three steps. The first step is to use the sample information to compute the posterior distribution of τ , $p(\tau|y)$, and to simulate τ using the inverse cumulative density function method. The second step is to simulate μ by drawing from its conditional posterior normal distribution $p(\mu|\tau, y)$, given the simulated values for τ . Finally, the simulation of θ_j is accomplished by sampling from its conditional posterior normal distribution $p(\theta_j|\mu, \tau, y)$ given the simulated values for τ and μ . A detailed description of each step of the simulation procedure is presented below.

To begin, the average past performance for each program (\bar{y}_j) and the corresponding variance ($\hat{\sigma}_{\bar{y}_j}$) is computed, as these statistics are necessary information to compute the posterior

distribution of τ using equation (11). The computation of $p(\tau|y)$ is done for a grid of equally spaced values of τ within the specified range (Figure 3). Note that given the sample data, the most likely value for τ is \$0.03/bu. in the corn model and \$0.10/bu. in the soybeans model, and the dispersion is greater in the soybean distribution. Recall that τ is the measure of uncertainty of the common prior distribution, therefore, the differences in the distributions presented in figure 3 imply that expected performance estimates are likely to have higher shrinkage for corn than for soybeans.

The cumulative distribution of τ is computed based on the posterior probability distribution. A normalizing factor needs to be applied because the probability distribution presented in equation (11) is defined up to an unknown normalizing constant. The normalizing factor (K) is calculated by adding up all values of $p(\tau|y)$ for the range of τ considered. Then, the values of $p(\tau|y)$ are divided by the normalizing factor and the cumulative density function is computed by adding up the values of $p(\tau|y)/K$ for τ less than or equal to the given value. For example, the cumulative probability for $\tau = 0.056$ in corn is 95%. This cumulative density function is employed to simulate τ .

In the second step, μ is simulated by drawing from its conditional posterior normal distribution $p(\mu|\tau, y)$ (equation 10), given the simulated values for τ . Note that the mean and variance of the posterior distribution of μ can be computed based on the sample statistics \bar{y}_j and $\hat{\sigma}_{\bar{y}_j}$ and a given value of τ . For example, if the simulated value for τ in corn is 0.058, the mean and variance of the normal distribution in equation (10) are computed with $\tau = 0.058$.

In the final step, the θ_j are simulated by sampling from the relevant conditional posterior normal distribution $p(\theta_j|\mu, \tau, y)$ (equation 9) given the simulated values for τ and μ . Suppose that simulated values for τ and μ are 0.058 and 0.015, respectively, then the means and variances for the posterior distribution of expected performance for each program (equation 9) are computed with $\tau = 0.058$ and $\mu = 0.015$, and the θ_j are simulated from these distributions.

The impact of the Bayesian estimation procedure on expected performance estimates can be illustrated graphically. Recall that τ is the measure of uncertainty of the common prior distribution (equation 8). The greater τ is, the more the individual shrinkage estimates will be close to separate estimates. On the other hand, for low values of τ shrinkage estimates are more similar to pooled estimates. Figure 4 shows the relationship between the values for τ and the shrinkage intensity (w from equation 4) for a subset of 10 corn advisory programs. Since the purpose of this figure is to illustrate the relationship between τ and the shrinkage intensity, a small number of programs is employed to make the figure readable. At the left extreme of the figure τ equals zero and the shrinkage coefficient is one, which means that all individual estimates equal the pooled estimate. When τ is zero there is no uncertainty about the prior, which is equivalent to assuming that all programs have exactly the same performance. Moving to the right in the figure, the degree of uncertainty in the prior distribution increases and less weight is given to the pooled estimate and more weight to separate estimates. The figure shows

that for a given level of τ the shrinkage intensity is quite different across programs. Programs with high uncertainty in the individual estimates (large confidence intervals in figure 1) have higher shrinkage intensity. Large uncertainty for a particular program can be due to high variability in performance across years or a small number of available observations. For instance, programs #9 and #11, which have large confidence intervals in figure 1, also have large shrinkage intensity for a given value of τ compared to the rest of the programs. Data is available in all 10 crop years for program #9, but it has a large variability in performance across years. Data is available for program #11 in only two crop years.

Continuing the illustration, figure 5 plots the conditional posterior means of individual expected performance $(E(\theta_j|\tau, y))$ for each value of τ . At the left extreme of the figure, τ equals zero and all individual estimates are equal to the pooled estimated. At the right extreme of the figure, with a value for τ of 0.16, individual expected performance is quite different across programs, with the estimates close to the traditional separate estimates. It is also possible to see in this figure how shrinkage intensity varies across programs. For example, note that the lines for programs #9 and #13 cross each other. This occurs because the first program has a higher individual estimate but also higher shrinkage intensity compared to the second. By considering the information in figures 3 to 5 it is possible to say that, based on the sample information, the most reasonable estimates for corn pricing performance for these programs imply a substantial shrinkage towards the pooled value.

In the second version of the Bayesian model skeptical beliefs are added and individual estimates are shrunk towards zero. The degree of shrinkage towards zero depends on the prior uncertainty (τ_0^2 in equation 12) and the variance of individual parameters (V_j in equation 9). Recall that the subjective prior employed in this study corresponds to a decision maker who believes that expected performance is most likely to be zero with a 1% probability that the expected difference between the advisory price and the benchmark is more than 5% of the benchmark. The average benchmark price is \$2.28/bu. for corn and \$5.86/bu. for soybeans during the sample period considered. Therefore, this statement implies a mean of zero ($\mu_0 = 0$) and a standard deviation of 5¢/bu. for corn and 13¢/bu. for soybeans. Standard deviations are computed by first computing a value equal to 5% of the average benchmark price for corn and soybeans (11¢/bu. for corn, 29¢/bu. for soybeans) and then dividing this number by z :

$$(14) \quad \tau_0 = (0.05 * \overline{BP}) / z \quad \text{where } P(Z \geq z) = 0.01 \rightarrow z = 2.33$$

and Z has standard normal distribution. Similar to the Bayesian model without skeptical beliefs, the shrinkage intensity varies towards zero across programs depending on the uncertainty in individual estimates. The computation of the posterior distribution of θ_j^{kep} under skeptical beliefs is again simulated from the normal distribution given the simulated values of the other parameters $(\mu, \tau, \theta_j, V_j)$.

In the third version, the simulation procedure is applied to each of the two activeness groups separately. For both crops, the optimal shrinkage intensity is greater for the conservative group compared to the active group. The most likely values for τ are 0.10 and 0.02 for the active and conservative groups in corn, respectively, and 0.21 and 0.08 for the active and conservative groups in soybeans, respectively. The differences in shrinkage intensity are due to a

larger number of programs with less dispersion of average performance across programs in the conservative group compared to the active group.

Results and Discussion

Tables 2 and 3 present individual expected performance point estimates based on traditional and Bayesian models for corn and soybeans, respectively. Programs are ordered from highest to lowest based on traditional performance estimates (third column). The Bayesian point estimates are the median of the posterior distribution of the parameters. The fourth column presents the Bayesian point estimates for the hierarchical model based on the entire sample of programs. Expected performance estimates for most programs are strongly shrunk towards pooled values in most cases (the Bayesian pooled performance estimates, $\hat{\mu}$, in the model for the entire sample are 0.5¢/bu. and 13¢/bu. for corn and soybeans, respectively). Shrinkage intensity generally is higher for corn programs compared to soybean programs. For instance, *Agri-Mark* (program #12) has expected performance of 5¢/bu. and 38¢/bu. under traditional estimation for corn and soybeans, respectively, and corresponding Bayesian estimates of only 0.8¢/bu. and 17¢/bu. The shrinkage intensity in this case (w in equation 4) is 93% for corn and 85% for soybeans. In some cases traditional and Bayesian hierarchical estimates are quite similar. For example, *AgLine by Doane-cash only* (#7) has an expected performance of 14¢/bu. in soybeans under traditional estimation and 13¢/bu. under Bayesian estimation, with a shrinkage intensity of 50%. Recall that the lower the precision of the individual estimate, the greater the shrinkage intensity towards the pooled values, therefore programs with wide confidence intervals for the separate estimates (figure 1) have higher shrinkage intensity.

The fifth column in the tables presents expected performance point estimates for the Bayesian model under skeptical beliefs. Comparing the fourth and fifth columns in table 2 it is evident that the point estimates with and without skeptical beliefs are very similar for corn, with differences being smaller than 1¢/bu. in all cases. In other words, estimation results do not change much when skeptical views are added to the model. This occurs because the performance of corn advisory programs as a group matches the skeptical prior distribution reasonably well. That is, on average, expected performance of advisory programs is close to zero and therefore the skeptical prior is similar to the prior without skeptical beliefs. Table 3 shows that the effect of adding skeptical beliefs is stronger in the soybeans model. This is the case because performance of advisory programs as a group in soybeans is superior compared to performance in the corn market.

The last two columns in tables 2 and 3 presents the point estimates for the Bayesian hierarchical model by activeness groups. An “A” indicates that the program belongs to the most active group and a “C” that belongs to the most conservative group. Ten programs form the active group: *Ag Financial Strategies*, *Ag Review*, *AgResource*, *Agri-Mark*, *Brock (hedge)*, *Harris Weather/Elliott Advisory*, *Progressive Ag*, *Stewart-Peterson Advisory Reports*, *Top Farmer Intelligence* and *Utterback Marketing Services* (#2, #6, #9, #12, #20, #27, #32, #37, #39 and #40). The remaining programs belong to the conservative group.⁹ The Bayesian pooled

⁹ In the clustering procedure program #9 is classified in one group, programs #2, #6, #12, #20, #27, #32, #37, #39 and #40 in a second group and the rest of the programs in a third group. For the estimation of the Bayesian hierarchical model by groups program #9 and the second group of programs are combined in one group called “active”.

performance estimate, $\hat{\mu}$, is 3¢/bu. and 0.4¢/bu. for the active and conservative groups in corn, respectively, and 18¢/bu. and 12¢/bu. for the active and conservative groups in soybeans. Note that active programs have lower shrinkage intensity and most of the programs with high performance estimates belong to the active group, however, also two of the active programs are at the bottom of the ranking.

Note that there are some differences in the ordering of programs under the different estimation procedures. This is due to the differences in shrinkage intensity across programs. For example, *AgriVisor (aggressive cash)* (#13) is ranked 4th according to the traditional estimation, 1st for the hierarchical model with skeptical beliefs, and only 6th in the hierarchical model with groups based on activeness. Not surprisingly then, a decision maker will choose different programs depending on the beliefs that he/she is willing to incorporate in the model.

Figure 6 is the graphical representation of the point estimates for expected performance under the different estimation procedures. The dots above the lower gray line are programs expected to outperform the market benchmark by more than 1% (1% of the average benchmark price is \$0.02/bu. for corn and \$0.06/bu. for soybeans). The dots above the higher gray line are programs expected to outperform the market benchmark by more than 5% (\$0.11/bu. for corn and \$0.29/bu. for soybeans). This figure nicely illustrates shrinkage effects in the different Bayesian hierarchical models.

Based on traditional expected performance point estimates (panel A, figure 6), 18 out of the 35 corn programs have positive expected performance, 12 programs are expected to outperform the benchmark by more than 1% and four programs by more than 5%. In contrast, only one program (#13) is expected to outperform the benchmark by more than 1% and none are expected to outperform the benchmark by more than 5% based on the Bayesian hierarchical model for the entire sample and the model including skeptical beliefs. Six programs have expected performance greater than 1% (#6, #9, #12, #13, #32 and #40) and 2 programs (#9 and #40) greater than 5% according to the hierarchical Bayesian model by activeness groups. Most of the higher performing programs belong to the active group. Finally, if the whole posterior distribution of expected performance is considered, not just the point estimated presented in the figure, there is a 75% or more probability that expected performance is greater than zero for the programs that outperform the benchmark by more than 1% or 5% (with the exception of program #12).

The second panel in figure 6 shows that performance of advisory programs in the soybean market generally is superior to performance in the corn market. Based on traditional expected performance estimates, 27 out of the 34 soybean programs have positive expected performance, 25 programs are expected to outperform the benchmark by more than 1%, and 7 programs by more than 5%. Based on the Bayesian hierarchical model for the entire sample, 33 programs have positive expected performance, 30 have expected to outperform the benchmark by more than 1%, and one program (#29) is expected to outperform the benchmark by more than 5%. In the model including skeptical beliefs, 33 programs have positive expected performance, 27 programs are expected to outperform the benchmark by more than 1% and none are expected to outperform the benchmark by more than 5%. According to the hierarchical Bayesian model by activeness groups, 32 programs have positive expected performance, 30 programs have expected performance greater than 1% and two programs (#9 and #32) greater than 5%. Note that these latter two programs belong to the active group. Considering the whole posterior distribution of expected performance, there is a 75% or more probability that expected

performance is greater than zero for all programs that outperform the benchmark by more than 1% (with exception of program #6) and there is a 95% or more probability that expected performance is greater than zero for the programs that outperform the benchmark by more than 5% and.

Based on the results presented above, it is evident that the answer to the question of whether farmers should follow the advice of market advisory programs depends on the beliefs that the decision-maker is willing to include in the estimation model and the magnitude of expected pricing performance that he/she considers desirable. For instance, consider a decision maker who, based on market efficiency, has a skeptical prior and is willing to follow an advisory program only if it is expected to increase price received by more than 5%. This skeptical decision-maker would prefer to adopt a strategy that mimics the market benchmark rather than following any advisory program for both crops. Now consider a more “optimistic” decision-maker who is willing to group advisory programs by the degree of activeness for performance estimation and is interested in following programs with expected performance greater than 1%. Several programs in corn and most programs in soybeans are better marketing alternatives than the market benchmark for this decision-maker.

Farmers and other market participants naturally are also interested in the pricing performance of the top advisory program. To that end, the posterior distributions of the two highest ranked programs in corn and soybeans are shown in figures 7 and 8. These posterior distributions provide information on the upper bound of the benefits from following advisory programs. The top-ranked program in corn are *AgResource* (#9) and *AgriVisor (aggressive cash)* (#13), depending on the estimation model considered. *AgResource* is the top-ranked program based on traditional estimates and the Bayesian hierarchical model by activeness groups. Panel A of figure 7 shows the posterior distribution of expected pricing performance for this program under each version of the Bayesian model. The median of the distribution varies substantially across models: 2¢/bu., 1¢/bu., and 13¢/bu. in the models for the whole sample, with skeptical beliefs, and by activeness group, respectively. According to the posterior distributions, there is a 90% probability that expected performance is between -3¢/bu. and 9¢/bu. in the model for the whole sample, between -3¢/bu. and 6¢/bu. in the model with skeptical beliefs, and between -1¢/bu. and 32¢/bu. in the model by activeness groups.

The top-ranked program in corn based on the models for the whole sample and with skeptical beliefs is *AgriVisor (aggressive cash)* (#13). Panel B of figure 7 shows the posterior distribution of expected pricing performance for this program. The median of the distribution varies moderately across the models: 4¢/bu., 3¢/bu. and 2¢/bu. according to the models for the whole sample, with skeptical beliefs, and by activeness groups, respectively. Based on the Bayesian models, there is a 90% probability that expected performance for *AgriVisor (aggressive cash)* is between -0.02¢/bu. and 11¢/bu. in the model for the whole sample, between 0.2¢/bu. and 8¢/bu. in the model with skeptical beliefs, and between -0.5¢/bu. and 9¢/bu. in the model by activeness group. It is interesting to note that in all cases the 90% interval contains zero or values very close to zero, which indicates that even the top-ranked programs in corn have a substantial chance of not outperforming the market benchmark.

The top-ranked program in soybeans based on traditional estimates and the Bayesian hierarchical model by activeness groups is again *AgResource*, while the top-ranked program based on the models for the whole sample and with skeptical beliefs is *Northstar Commodity* (#29). Panels A and B of figure 8 show the expected distribution of pricing performance for

these two programs based on the different estimation models. Like corn, the median of *AgResource* distribution varies substantially across models: 27¢/bu., 17¢/bu., and 50¢/bu., in the models for the whole sample, with skeptical beliefs, and by activeness groups, respectively. The expected pricing performance for *Northstar Commodity* varies much less across the models: 30¢/bu., 24¢/bu., and 26¢/bu. For *AgResource*, there is a 90% probability that expected performance is between 10¢/bu. and 47¢/bu. in the model for the whole sample, between 5¢/bu. and 30¢/bu. in the model with skeptical beliefs, and between 21¢/bu. and 90¢/bu. in the model by activeness groups. In the case of *Northstar Commodity*, there is a 90% of probability that expected performance is between 18¢/bu. and 43¢/bu. in the model for the whole sample, between 14¢/bu. and 35¢/bu. for the model with skeptical beliefs, and between 13¢/bu. and 41¢/bu. for the model by activeness groups. The 90% intervals indicate that top-ranked programs in soybeans are highly likely to outperform the market benchmark.

The results presented in the last two paragraphs provide information on expected gains from following top-performing advisory programs in corn and soybeans. For a more complete evaluation of the magnitude of the upper bound on the benefits from following advisory programs, expected performance of the top-ranked programs can be expressed on a per-acre basis. Consider the most conservative approach, which is based on the Bayesian model with skeptical beliefs. Expected performance in corn is 3¢/bu. by following the top performing program's (#13) marketing recommendations. This value implies an expected annual gain of around \$4.5/acre for a farm with an average yield of 150 bu./acre. In soybeans, the highest expected performance (program #29) is 24¢/bu. Following the recommendations of this program, farmers can expect an annual gain of \$11/acre for a farm with an average yield of 47 bu./acre. Therefore, the combined gains for a 50/50 corn and soybeans farm would be \$8/acre. This value is small, but not trivial, considering that the average net income for an Illinois grain farm over 1995-2004 was \$61/acre (as reported by Lattz et al., 2005).

In the decision of whether to follow an advisory program farmers should also compare expected pricing performance with associated subscription costs. These fees are charged annually on a per-farm basis and represent small values when expressed on a per bushel for a medium-sized commercial farm. For the advisory programs in the sample, annual subscription ranges from \$100 to \$600 per farm. The subscription to the most expensive program represents a cost of only 0.8¢/bu. for a farmer growing 500 acres of corn with an average yield of 150bu./acre. However, there are other costs related to following advisory programs, such as the cost of implementing, monitoring, and managing the marketing strategies recommended by advisory programs. While these costs are difficult to measure, they may well be large enough to offset a considerable portion of expected benefits, if any, of following advisory programs.

The results presented so far in this essay are based entirely on comparisons of advisory programs to the 24-month market benchmark. As mentioned in the data section, AgMAS performance evaluations also consider three other benchmarks. Expected performance for individual programs was also estimated using the alternative benchmarks. In corn, the 24-market benchmark has, on average, the highest price among the benchmarks. Therefore, advisory programs have somewhat more attractive performance when compared to the rest of the benchmarks. For instance, when compared against the other three benchmarks, more than half of the corn programs have positive expected performance greater than 1% of the benchmark price and one program greater than 5% for the Bayesian model with skeptical beliefs (compare these values with the ones plotted in figure 6, panel A). In the case of soybeans, the 24- and 20-month

market benchmarks have similar average prices, while the farmer benchmarks have higher average prices. Therefore, the performance of advisory programs in soybeans is less attractive when compared to farmer benchmarks. Still, based on skeptical estimates, a few programs have expected performance greater than 1% of the average benchmark price for both farmer benchmarks in soybeans. Overall, the results for the other three benchmarks do not alter the basic conclusions reached using the 24-month market benchmark.

Summary and Conclusions

This paper employs a Bayesian hierarchical approach to estimate expected performance of market advisory programs in corn and soybeans. This estimation procedure is a conservative approach compared to traditional estimation, since it reduces estimation error in the expected gains from following top-performing advisory programs. Three versions of the model are estimated. The first combines information across the entire sample, while the second includes skeptical beliefs based on the efficient market hypothesis. The third divides programs into two groups based on the degree of activeness in marketing recommendations. The data consist of past observations of pricing performance for corn and soybean advisory programs for the 1995 to 2004 crop years and are obtained from AgMAS project records. The posterior distribution of individual expected performance is computed by simulation.

The Bayesian hierarchical model produces shrinkage estimators that are weighted averages of individual and pooled estimates and adding a skeptical prior shrinks performance estimates towards zero. The answer to the question of whether farmers should follow the advice of market advisory programs depends on the beliefs that the decision-maker is willing to include in the estimation model and the magnitude of expected pricing performance that he/she considers desirable. Results indicate that even when skeptical beliefs are incorporated into the model a few programs in corn and several programs in soybeans appear to be better marketing alternative compared to a naïve strategy that mimics the market benchmark. More specifically, a skeptical farmer can expect to increase the price received for corn by 1% and the price received for soybeans by 5% following the single top-ranked program. These values imply a combined expected annual gain for a 50/50 corn and soybeans farm of \$8/acre. Whether these gains would offset the cost of implementing, monitoring, and managing the recommended marketing strategies is still an open question.

While risk is not directly measured in this study, it should be noted that programs with higher performance variability are penalized in the Bayesian estimation model with higher shrinkage intensity, and therefore, become less attractive to farmers. A more comprehensive study of advisory services performance, including the measurement of the risk level of advisory programs is an interesting extension of the current study. The hierarchical Bayesian approach can be also applied to evaluate the benefits from following different combinations of marketing advisory programs, in a portfolio optimization context. In this case, the estimation of the covariance matrix for advisory pricing performance represents a challenge given the data availability restrictions and a hierarchical Bayesian model is an appropriate estimation procedure.

The Bayesian approach implemented in this study also provides an interesting framework for more general evaluations of grain marketing alternatives. For instance, a Bayesian hierarchical model with and without skeptical beliefs can be employed in the estimation of

expected gains of different combinations of cash and derivatives transactions, and marketing contracts offered by grain companies. These applications of Bayesian modeling in grain marketing represent interesting opportunities for further research.

References

- Allenby, G.M., P.E. Rossi, and R.E. McCulloch. "Hierarchical Bayes Models: A Practitioners Guide," Working paper, The Ohio State University, January 2005.
- Baks, K.P., A. Metrick, and J. Wachter. "Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation." *Journal of Finance* 56(2001):45-85.
- Brorsen, B.W., and K. B. Anderson. "Cash Wheat Marketing: Strategies for Real People." *Journal of Agribusiness* 12(1994):85-94.
- Efron, B., and C. Morris. "Data Analysis Using Stein's Estimator and its Generalization." *Journal of the American Statistical Association* 70(1975):311-319
- Gelman, A., J.B. Carlin, H.S. Stern, and D.B. Rubin. "Hierarchical Models." Ch 5 in *Bayesian Data Analysis*. New York: Chapman and Hall/CRC, 2004.
- Gelman, A. "Prior Distributions for Variance Parameters in Hierarchical Models" Working paper, Columbia University, 2004.
- Grauer, R.R. "On the Predictability of Stock Market Returns: Evidence from Industry-Rotation Strategies" Working paper, Simon Fraser University, 2002.
- Irwin, S.H., D.L. Good, J. Martines-Filho, and R. M. Batts. "The Pricing Performance of Market Advisory Services in Corn and Soybeans Over 1995-2004," AgMAS Project Research Report 2006-02, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, April 2006a.
- Irwin, S.H., D.L. Good, and J. Martines-Filho. "The Performance of Market Advisory Services in Corn and Soybeans." *American Journal of Agricultural Economics* 88(2006b):162-181.
- Johnson, R.A., and D.W. Wichern. "Clustering, Distance Methods and Ordination." Chapter 12 in *Applied Multivariate Statistical Analysis*. Prentice Hall. Upper Saddle River. New Jersey. 2002.
- Jorion, P. "Bayes-Stein Estimation for Portfolio Analysis." *Journal of Financial and Quantitative Analysis* 21(1986):279-292.
- Lattz, D.H., C.E. Cagley and D.D. Raab. *Summary of Illinois Farm Business Records for 2004*, Circular 1388-05, University of Illinois Extension, 2005.
- Marcus, A.J. "The Magellan Fund and Market Efficiency." *Journal of Portfolio Management* 17(1990):85-88.

- Michaud R.O. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal* (1989)45:31-42.
- Norvell, J.M., and D.H. Lattz. "Value-Added Crops, GPS Technology and Consultant Survey: Summary of a 1998 Survey to Illinois Farmers," Working paper, College of Agricultural, Consumer and Environmental Sciences, University of Illinois, 1999
- Patrick, G.F., W.N. Musser and D.T. Eckman. "Forward Marketing Practices and Attitudes of Large-Scale Midwestern Grain Farmers." *Review of Agricultural Economics* 20(1998):38-53.
- Tomek, W.G. and H.H. Peterson. "Implications of Commodity Price Behavior for Marketing Strategies." *American Journal of Agricultural Economics* 87(2005):1258-1264.
- Zulauf, C.R., and S.H. Irwin. "Market Efficiency and Marketing to Enhance Income of Crop Producers." *Review of Agricultural Economics* 20(1998):308-331.

Table 1. List of Market Advisory Programs Tracked by the AgMAS Project over the 1995-2004 Crop Years

ID	Market Advisory Program	Crop Years	ID	Market Advisory Program	Crop Years
1	Ag Alert for Ontario	(1996)	22	Co-Mark	(2000-2002)
2	Ag Financial Strategies	(2001-2004)	23	Freese-Notis	(1995-2004)
3	Ag Market Professional (cash only)	(2004)	24	Grain Field Marketing	(2001-2004)
4	Ag Market Professional (hedge)	(2004)	25	Grain Field Report	(1995)
5	Ag Profit by Hjort	(1995-1999)	26	Grain Marketing Plus	(2000-2001)
6	Ag Review	(1995-2004)	27	Harris Weather/Elliott Advisory	(1995-1996)
7	AgLine by Doane (cash only)	(1995-2004)	28	North American Ag	(1995)
8	AgLine by Doane (hedge)	(1996-2004)	29	Northstar Commodity	(2001-2004)
9	AgResource	(1995-2004)	30	Pro Farmer (cash only)	(1995-2004)
10	Agri-Edge (cash only)	(1995-1996)	31	Pro Farmer (hedge)	(1995-2004)
11	Agri-Edge (hedge)	(1995-1996)	32	Progressive Ag	(1996-2004)
12	Agri-Mark	(1995-2000)	33	Prosperous Farmer	(1995)
13	AgriVisor (aggressive cash)	(1995-2004)	34	Risk Management Group (cash only)	(1999-2004)
14	AgriVisor (aggressive hedge)	(1995-2004)	35	Risk Management Group (futures & options)	(1999-2004)
15	AgriVisor (basic cash)	(1995-2004)	36	Risk Management Group (options)	(1999-2004)
16	AgriVisor (basic hedge)	(1995-2004)	37	Stewart-Peterson Advisory Reports	(1995-2004)
17	Allendale (futures & options)	(1996-2004)	38	Stewart-Peterson Strictly Cash	(1995-1999)
18	Allendale (futures only)	(1995-2004)	39	Top Farmer Intelligence	(1995-2004)
19	Brock (cash only)	(1995-2004)	40	Utterback Marketing Services	(1997-2004)
20	Brock (hedge)	(1995-2004)	41	Zwicker Cycle Letter	(1995-1998)
21	Cash Grain	(1999)			

Notes: A crop year is a two-year marketing window from September of the year previous to harvest through August of the year after harvest. The Allendale (futures & options) program is offered only for corn.

Table 2. Traditional and Bayesian Estimates of Expected Pricing Performance for Market Advisory Program, Corn, 1995-2004 Crop Years

Program ID	Number of Observations	Traditional Separate Estimates	Bayesian Hierarchical Model	Bayesian Hierarchical Model with Skeptical Beliefs	Bayesian Hierarchical Model by Activeness Group	group ¹
		(\$/bu.)	(\$/bu.)	(\$/bu.)	(\$/bu.)	
9	10	0.298	0.015	0.011	0.129	A
11	3	0.242	0.015	0.010	0.008	C
40	8	0.174	0.021	0.018	0.121	A
13	10	0.127	0.040	0.033	0.024	C
6	10	0.091	0.018	0.014	0.073	A
32	9	0.077	0.011	0.010	0.060	A
14	10	0.057	0.015	0.012	0.009	C
7	10	0.051	0.020	0.018	0.015	C
12	6	0.047	0.008	0.006	0.035	A
8	9	0.043	0.021	0.017	0.014	C
15	10	0.027	0.012	0.010	0.008	C
16	10	0.023	0.010	0.009	0.006	C
39	10	0.022	0.009	0.008	0.022	A
34	6	0.013	0.008	0.007	0.007	C
35	6	0.013	0.006	0.005	0.005	C
24	4	0.012	0.010	0.009	0.006	C
41	4	0.008	0.006	0.005	0.004	C
36	6	0.003	0.005	0.005	0.003	C
19	10	-0.001	0.004	0.003	0.004	C
20	10	-0.003	0.006	0.004	0.011	A
29	4	-0.004	0.003	0.002	0.002	C
22	4	-0.009	-0.001	0.000	0.000	C
10	3	-0.010	0.005	0.005	0.002	C
21	2	-0.013	-0.004	-0.004	-0.002	C
23	10	-0.016	0.002	0.002	0.001	C
17	9	-0.018	0.001	0.001	0.000	C
38	6	-0.030	-0.004	-0.004	-0.004	C
18	10	-0.046	0.003	0.003	0.001	C
27	2	-0.056	0.006	0.005	0.021	A
37	10	-0.082	-0.008	-0.006	-0.056	A
30	10	-0.101	-0.013	-0.011	-0.007	C
31	10	-0.104	-0.019	-0.013	-0.012	C
26	3	-0.124	-0.001	0.000	-0.001	C
5	5	-0.131	-0.003	-0.004	-0.002	C
2	4	-0.187	-0.013	-0.008	-0.107	A

Note: The advisory programs' names are listed in table 1.

(1) "A" indicates that the program belongs to the active group and "C" to the conservative group.


Table 3. Traditional and Bayesian Estimates of Expected Pricing Performance for Market Advisory Programs, Soybeans, 1995-2004 Crop Years

Program ID	Number of Observations	Traditional Separate Estimates	Bayesian Hierarchical Model	Bayesian Hierarchical Model with Skeptical Beliefs	Bayesian Hierarchical Model by Activeness Group	
		(\$/bu.)	(\$/bu.)	(\$/bu.)	(\$/bu.)	group ¹
9	10	0.702	0.267	0.173	0.497	A
32	9	0.535	0.201	0.134	0.364	A
29	4	0.429	0.304	0.240	0.261	C
22	4	0.390	0.196	0.135	0.166	C
12	6	0.383	0.166	0.098	0.265	A
41	4	0.356	0.195	0.128	0.162	C
24	4	0.333	0.158	0.095	0.133	C
40	8	0.241	0.150	0.093	0.206	A
21	2	0.231	0.138	0.087	0.123	C
14	10	0.216	0.168	0.121	0.149	C
16	10	0.199	0.165	0.121	0.146	C
20	10	0.189	0.153	0.102	0.185	A
13	10	0.157	0.141	0.112	0.131	C
31	10	0.153	0.145	0.114	0.131	C
37	10	0.152	0.139	0.106	0.157	A
27	2	0.152	0.127	0.085	0.171	A
7	10	0.140	0.134	0.111	0.130	C
30	10	0.140	0.130	0.102	0.128	C
11	3	0.140	0.138	0.089	0.122	C
15	10	0.134	0.126	0.099	0.127	C
10	3	0.134	0.127	0.082	0.116	C
8	7	0.112	0.121	0.091	0.110	C
39	10	0.106	0.111	0.087	0.120	A
26	3	0.078	0.121	0.072	0.106	C
19	10	0.066	0.077	0.071	0.083	C
18	10	0.046	0.084	0.066	0.089	C
34	6	0.027	0.079	0.054	0.086	C
23	10	-0.006	0.018	0.020	0.028	C
38	6	-0.009	0.009	0.011	0.018	C
35	6	-0.009	0.069	0.051	0.087	C
36	6	-0.027	0.056	0.045	0.073	C
5	5	-0.065	0.073	0.060	0.088	C
2	4	-0.086	-0.044	-0.040	-0.069	A
6	10	-0.294	0.062	0.045	-0.020	A

Note: The advisory programs' names are listed in table 1.

(1) "A" indicates that the program belongs to the active group and "C" to the conservative group.

Panel A. Corn Pricing Performance



Panel B. Soybean Pricing Performance





Figure 1. Expected Performance of Market Advisory Programs, Traditional Separate Point Estimates and 90% Confidence Intervals


Note: Advisory programs' names are listed in table 1. The dots in the figures represent the point estimates and the lines the 90% confidence intervals for expected performance.



Note: θ_j is the expected performance for program j ; (μ, τ) are the parameters of the common prior distribution.

Figure 2. Diagram for the Structure of the Hierarchical Model for Advisory Programs' Expected Performance

Panel A. Corn Pricing Performance



Panel B. Soybean Pricing Performance


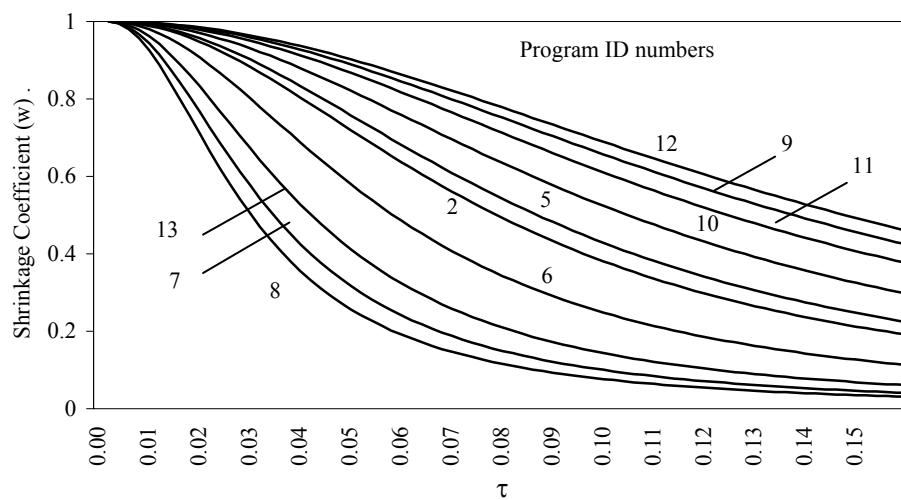


Figure 3. Marginal Posterior Density of τ (\$/bu.)



Note: The shrinkage coefficient is the weight for the pooled estimate in the shrinkage estimators.

Figure 4. Shrinkage Intensity vs. τ for Corn Performance Estimation


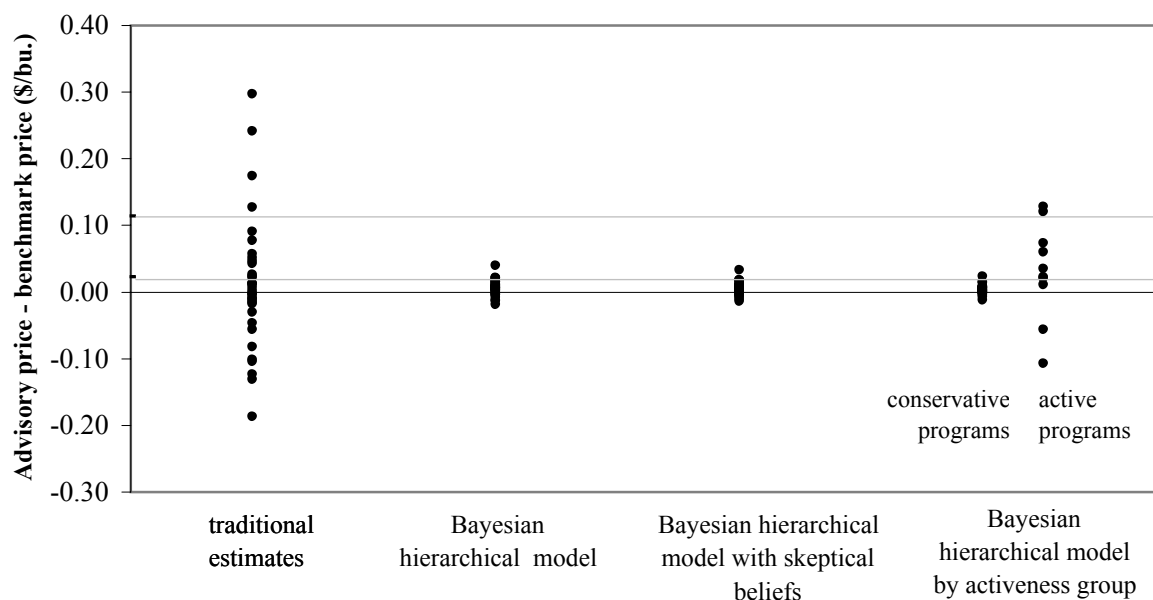



Figure 5. Posterior Point Estimates for Expected Corn Pricing Performance for Different Levels of Shrinkage Intensity

Panel A. Corn Pricing Performance



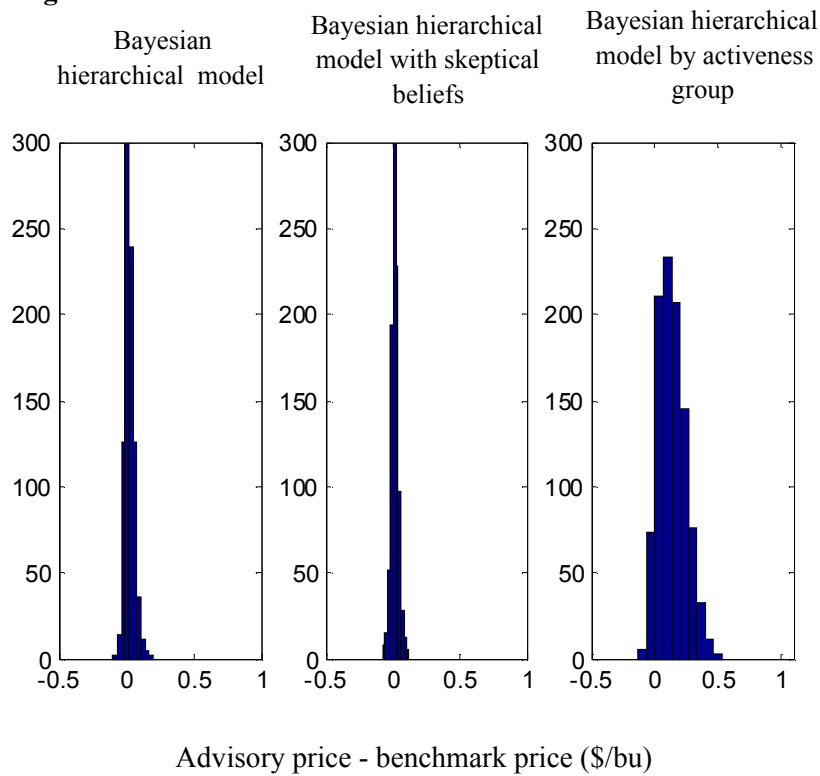
Panel B. Soybean Pricing Performance



Note: Back dots represent the point estimates for expected pricing performance for each advisory programs under each estimation model. The dots above the lower gray line are for programs expected to outperform the market benchmark by more than 1%. The dots above the higher gray line are for programs expected to outperform the market benchmark by more than 5%.

Figure 6. Expected Pricing Performance for Corn and Soybeans Advisory Programs, Traditional and Bayesian Point Estimates

Panel A. AgResource



Panel B. AgriVisor (aggressive cash)

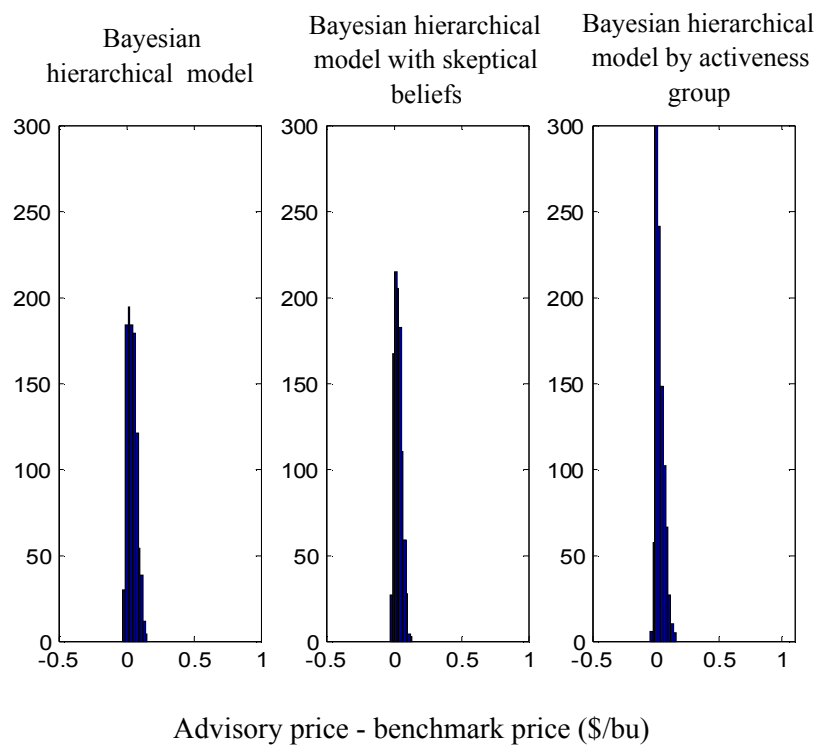
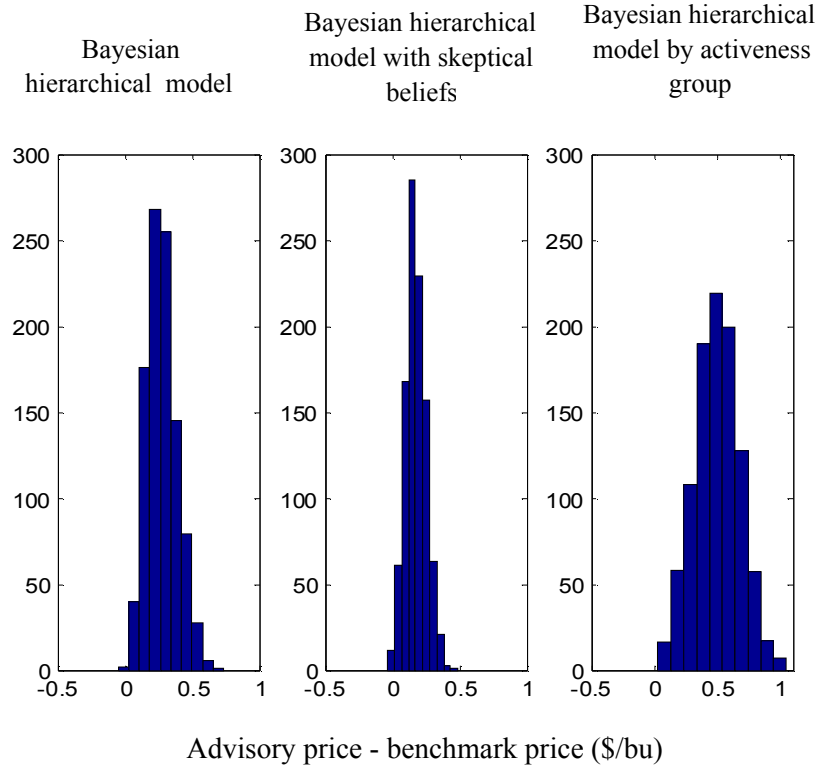


Figure 7. Simulated Values of Expected Performance for the Top Performing Programs in Corn

Panel A. AgResource



Panel B. Northstar Commodity

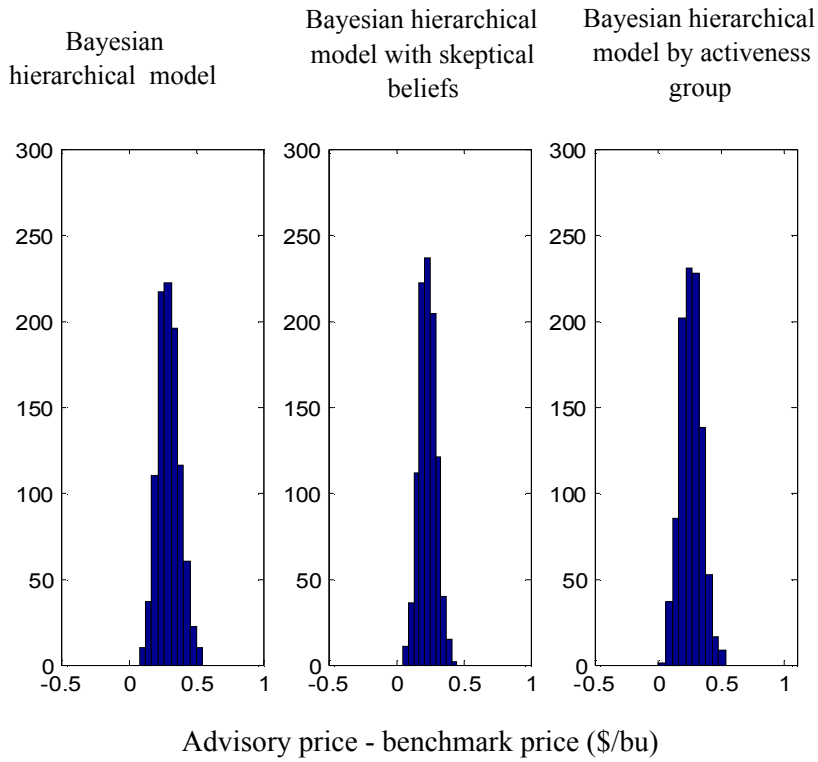


Figure 8. Simulated Values of Expected Performance for the Top Performing Programs in Soybeans