

GENERALIZED HEDGE RATIO ESTIMATION WITH AN UNKNOWN MODEL

Jeffrey H. Dorfman
Department of Agricultural & Applied Economics
University of Georgia
Athens GA 30602
email: jdorfman@agecon.uga.edu

and

Dwight R. Sanders
Department of Agribusiness Economics
Southern Illinois University
Carbondale IL 62901
email: DwightS@siu.edu

May 31, 2005

* We wish to thank Gary Koop and participants at the 2004 NCR-134 Conference for helpful discussions and comments.

GENERALIZED HEDGE RATIO ESTIMATION WITH AN UNKNOWN MODEL

Abstract

Myers and Thompson (1989) pioneered the concept of a generalized approach to estimating hedge ratios, pointing out that the model specification could have a large impact on the hedge ratio estimated. While a huge empirical literature exists on estimating hedge ratios, the literature is lacking a formal treatment of model specification uncertainty. This research accomplishes that task by taking a Bayesian approach to hedge ratio estimation, where specification uncertainty is explicitly modeled. Specifically, we present a Bayesian approach to hedge ratio estimation that integrates over model specification uncertainty, yielding a hedge ratio estimator that is robust to possible model specification because it is an average across a set of hedge ratios conditional on different models. Model specifications vary by exogenous variables (such as exports, stocks, and interest rates) and lag lengths included. The methodology is applied to data on hedging of corn and soybeans and on cross-hedging of corn oil using soybean oil futures. Results show the potential benefits and insights gained from such an approach.

JEL Categories: C11, C51, Q14

Key Words: Bayesian Econometrics, Corn, Futures Markets, Hedge Ratios, Model Specification, Soybeans.

1. Introduction

Myers and Thompson (1989) pioneered the concept of a generalized approach to estimating hedge ratios. They pointed out that the form of the equation to use in estimation is dependent upon assumptions concerning the stochastic prices whose risks are being managed. A huge literature exists on estimating hedge ratios under different model assumptions, adding generalizations to ARCH or GARCH errors (Baillie and Myers, 1991), parameter estimation uncertainty (Lence and Hayes, 1994a), and many other features of model specification (Witt, Schroeder, and Hayenga, 1987; Vukina, 1992).

However, the literature is lacking a formal treatment of model specification uncertainty as the central issue in hedge ratio estimation. In Myers and Thompson's original work, they suggest that the optimal hedge ratio, β , should be estimated with an OLS regression: $p_t = \beta f_t + \alpha X_{t-1}$, where p_t is the cash price level, f_t is the futures price level, and X_{t-1} is a vector of variables known at time $t - 1$ that help predict p_t and f_t . While Myers and Thompson suggest that X_{t-1} include lagged values of p_t and f_t , production, storage, exports, and consumer income, they likewise admit that "model specification is somewhat ad hoc with economic theory, hypothesis testing, and common sense used as guidelines" (p. 864). Furthermore, the authors readily acknowledge that model specification is perhaps the most difficult aspect of estimating generalized hedge ratios. As a practical matter, Financial Accounting Standard 133, has increased the need for risk managers to more carefully document hedging effectiveness and methods for determining hedged quantities and ratios (International Treasurer). A general method for determining hedge ratios may help risk managers more easily verify their methods and meet regulatory requirements for hedge accounting.

This research seeks to formally address the specification problem by taking a Bayesian approach to hedge ratio estimation, where model uncertainty is a given. Specifically, we present a Bayesian approach to hedge ratio estimation that integrates over model specification uncertainty. This yields a hedge ratio estimator that is robust to possible model specification because it is an average across a set of hedge ratios conditional on different models.

Formally, we consider a set of 64 possible model specifications and estimate the posterior distribution of the hedge ratio and the posterior odds in favor of the model for each model in that set. The distributions and model odds are then used to construct the marginal distribution of the robust hedge ratio, integrating out the model uncertainty. The integration with respect to the model uncertainty, which yields the marginal posterior distribution, is accomplished by computing a weighted average of the 64 conditional (on model specification) distributions where the weights are equal to the model odds. A single robust hedge ratio can then be chosen using any desired loss function; for example, a quadratic loss function will produce a posterior point estimator for the hedge ratio equal to the mean of the marginal posterior distribution.

The methodology is applied to hedging for corn and soybeans and to cross-hedging for corn oil using soybean oil futures. Model specifications vary by exogenous variables (exports, stocks, and interest rates) and lag lengths included. Importantly, the research presents a different approach to estimating hedge ratios, which may protect practitioners against model specification errors. In simulations using our application, risk management performance of the robust hedge ratio appears to be as good as alternatives, although significant improvement is not found in the regular hedging examples. However, in the cross-hedging example significant improvement in risk management is demonstrated. This

highlights an important aspect of model specification uncertainty: one is never sure when the model being used is wrong.

2. Literature Review and Problem Overview

Myers and Thompson (1989) generalized the estimation of hedge ratios to account for conditioning information that is available at the time a hedging decision is made. The authors demonstrate that the traditional approach of using a simple regression of cash price levels on futures price levels or cash price changes on futures price changes are correct only under a very restrictive set of assumptions. A regression approach is suggested where the cash price level is regressed against the futures price level plus a set of conditioning variables. Myers and Thompson suggest the conditioning variables include lags of futures and cash prices, plus any variables thought to influence prices such as stocks, exports, and storage costs. In an example using corn and soybeans, the authors show that the generalized robust hedge ratio can vary substantially from the unconditional ratio estimated with price levels; but, they argue that the unconditional ratio estimated with price changes may provide a reasonable estimate of the generalized hedge ratio. The authors urge researchers to extend the methodology to allow for conditional heteroscedastic shocks, and to use out-of-sample data to compare performance among the different approaches to estimating hedge ratios.

Baillie and Myers (1991) apply bivariate GARCH models to estimated time-varying hedge ratios. That is, the hedge ratio is defined as the conditional covariance between cash and futures prices divided by the conditional variance of futures prices, where the time variation in the conditional covariance matrix is modeled using a GARCH specification. The authors find that hedge ratios are time-varying and nonstationary. Furthermore, the GARCH hedge ratios outperform constant (unconditional) hedge ratios in out-of-sample

tests. Despite this advance in estimation techniques, the authors do not generalize the hedging regression in the sense of Myers and Thompson to include conditioning variables.

Researchers have extended the procedure of Baillie and Myers to areas such as simultaneously determined hedge ratios (Garcia, Roh, and Leuthold, 1995). Others have delved into whether hedge ratios should be estimated with price levels, price changes, or returns (Witt, Schroeder, and Hayenga, 1987) or whether or not the use of hedge ratios out-perform naive unit-for-unit hedging (Jong, De Roon, Veld, 1997; Collins, 2000). Still, the use of the simple (unconditional) hedge ratio (usually estimated in price changes) is pervasive in the literature (e.g., Ferguson and Leistikow, 1998). This may stem from the inherent problems in specifying the generalized model of Myers and Thompson, and the potential sensitivity of hedge ratios to model specification. Here, we pose one potential solution to this dilemma.

By taking a Bayesian approach to hedge ratio estimation, model uncertainty is treated similarly to a parameter to be estimated and one can integrate over model specification uncertainty. This yields a robust hedge ratio estimator that is robust to possible model specification because it is an average across a set of hedge ratios conditional on different models. Such an approach was first undertaken empirically in economics by Poirier (1991), who considered 147 different macroeconomic models. Poirier tested important macroeconomic hypotheses such as money neutrality while removing the potential influence of model specification by deriving results that were averaged across a large set of possible models differing in both included variables and identifying restrictions. Considering that Alston and Chalfant (1993) showed how important model specification can be to the results of applied econometrics in agricultural economics, it is surprising that more work on model specification uncertainty has not appeared in the agricultural economics literature. A rare

example is found in Dorfman and Lastrapes (1996) who use an approach similar to Poirier’s to estimate agricultural price responses to monetary policy.

3. Modeling and Estimation Issues

In this section we will show the model specifications used, the methodology for handling model specification uncertainty, and the process used to accomplish the Bayesian estimation of the robust hedge ratio. The important parts of the robust estimation approach are the set of models considered and the assumptions made for the likelihood functions and prior distributions of the unknown parameters. Given those details and the data, Bayes’ Theorem leads us through a straightforward process which optimally combines this (researcher-specified) information with the information in the data to yield the posterior distributions of model odds, regression parameters, and any other features of interest in our models. Further details for handling model choice and comparison in a Bayesian framework using the approach here can be found in Koop (2003, pp38-43) which contains an easy to follow exposition of the process.

3.1. Assumptions and Statistical Mechanics

First, we need to describe the process by which Bayesian statistics handles model specification uncertainty. To begin the estimation process, define the set of models to be considered, $\mathcal{M} = \{M_j, \quad j = 1, \dots, M\}$, here all assumed to be linear regression models:

$$y = X_j\beta_j + \epsilon_j, \quad j = 1, \dots, M, \quad (3.1)$$

where y is the vector of observations on the dependent variable assumed for simplicity here not to vary across models, X_j is the matrix of regressors for the j th model considered, ϵ_j is the random error term vector for the j th model, and j indexes the models in the set of M models considered. Given that the dependent variable is here assumed identical

in all models, the differences in models are all confined to the regressor matrix X which is allowed to vary both in the number of regressors, k_j , and in the particular regressors included (which could include variation in variables included and/or transformations of variables such as logs versus levels).

The prior distributions on the regression parameters β_j are specified as

$$p(\beta_j) \sim N(b_{0j}, \sigma_j^2 V_{0j}), \quad j = 1, \dots, M, \quad (3.2)$$

where N stands for the (multivariate) normal distribution, b_{0j} is the prior mean of the j th model's regression parameters and $\sigma_j^2 V_{0j}$ is the prior covariance matrix. The variable σ_j^2 also needs a prior distribution which is specified as an inverse-gamma, or equivalently for its inverse as

$$p(\sigma_j^{-2}) \sim G(s_{0j}^{-2}, d_{0j}), \quad j = 1, \dots, M, \quad (3.3)$$

where G stands for the gamma distribution, s_{0j}^{-2} is the prior mean for the inverse error variance, and d_{0j} is the prior degrees of freedom parameter which controls the tightness (or informativeness) of the prior distribution—higher values of d_{0j} imply a more informative prior (Koop, 2003).

The likelihood function for each model is assumed to follow a standard form based on identically and normally distributed random error terms ϵ_j . While there is some evidence of commodity prices following non-normal distributions and having nonconstant variances (cf. Baillie and Myers, 1991), this assumption allows analytical derivation of the form of each model's posterior distribution and of the model's marginal posterior odds. The likelihood function is therefore specified in the form

$$L_j(y|\beta_j, \sigma_j^2, X_j) = (2\pi\sigma_j^2)^{-n/2} \exp\{-0.5(y - X_j\beta_j)' \sigma_j^{-2} (y - X_j\beta_j)\}, \quad j = 1, \dots, M. \quad (3.4)$$

Given the prior distributions and likelihood functions above, the joint posterior distribution for β_j and σ_j^2 can be derived by Bayes' Theorem which states the posterior distribution is proportional to the prior distribution times the likelihood function. Doing the math, one finds the joint posterior is given by

$$p(\beta_j, \sigma_j^2 | y, X_j) \sim NG(b_{pj}, V_{pj}, s_{pj}^2, d_{pj}), \quad j = 1, \dots, M, \quad (3.5)$$

where

$$V_{pj} = (V_{0j}^{-1} + X_j' X_j)^{-1}, \quad (3.6)$$

$$b_{pj} = V_{pj}(V_{0j}^{-1} b_{0j} + (X_j' X_j) \hat{\beta}_j), \quad (3.7)$$

$$d_{pj} = d_{0j} + n_j, \quad (3.8)$$

and

$$s_{pj}^2 = d_{pj}^{-1} [d_{0j} s_{0j}^2 + (n_j - k_j) s_j^2 + (\hat{\beta}_j - b_{0j})' (V_{0j} + (X_j' X_j)^{-1})^{-1} (\hat{\beta}_j - b_{0j})], \quad (3.9)$$

where NG stands for the joint normal-gamma distribution, $\hat{\beta}_j$ and s_j^2 are the standard OLS quantities and n_j and k_j are the rows and columns of X_j , respectively. Essentially what equations (3.5) to (3.9) accomplish is an optimal combination of the information in the prior distribution with the information contained in the data as measured through the likelihood function. The result is a posterior distribution with parameters that are weighted averages of the parameters of the prior distribution and the parameters of the distribution of a maximum likelihood estimator based on the data.

Most research interest focuses on the posterior estimate of β_j or a subset of those regression parameters (like the hedge ratio). Because of this focus, it makes sense to

derive the marginal posterior distribution of β_j by integrating out the variance parameter σ_j^2 to yield

$$p(\beta_j|y, X_j) \sim t(b_{pj}, s_{pj}^2 V_{pj}, d_{pj}), \quad j = 1, \dots, M, \quad (3.10)$$

where t stands for the multivariate Student's t -distribution. The marginal posterior distribution of a particular element of β_j also follows a t -distribution with posterior mean and variance as in the multivariate distribution above.

Now, introduce the apparatus for handling model specification uncertainty. Begin with a discrete prior weight on each model,

$$p(M_j) = \mu_j, \quad \sum_{j=1}^M \mu_j = 1. \quad (3.11)$$

These weights may be uninformative in the sense of treating all models equally or may be weighted to display a preference for certain models. In the uninformative case, $\mu_j = 1/M, \forall j$. Next, using the above results on the posterior distributions shown in (3.5), derive the marginal likelihood functions by integrating out the parameter uncertainty to leave a conditional likelihood for each model,

$$p(y_j|M_j) = c_j[|V_{pj}|/|V_{0j}|]^{1/2}(d_{pj}s_{pj}^2)^{-d_{pj}/2}, \quad (3.12)$$

where

$$c_j = \frac{\Gamma(d_{pj}/2)(d_{0j}s_{0j}^2)^{d_{0j}/2}}{\Gamma(d_{0j}/2)\pi^{n/2}}, \quad (3.13)$$

and $\Gamma(\cdot)$ is the Gamma function. Combining these two equations, (3.11) and (3.12), one can derive the posterior probability of each model

$$p(M_j|y_j) \propto \mu_j[|V_{pj}|/|V_{0j}|]^{1/2}(d_{pj}s_{pj}^2)^{-d_{pj}/2} = \mu_j p(y_j|M_j), \quad j = 1, \dots, M. \quad (3.14)$$

Normalizing the values in (3.14) by dividing each value by the sum across all M models will ensure that the posterior model probabilities will sum to unity. Denote these normalized posterior probabilities by

$$\omega_j = \frac{\mu_j p(y_j | M_j)}{\sum_{j=1}^M \mu_j p(y_j | M_j)}, \quad j = 1, \dots, M. \quad (3.15)$$

These posterior model probabilities are the key to the handling of model uncertainty.

3.2. Robust Bayesian Parameter Estimation: Accounting for Model Uncertainty

Given the normalized posterior model probabilities, the next step is to derive the marginal posterior distribution, removing the conditioning on the model specification. This is done by integrating over the models in the set \mathcal{M} , essentially creating a single posterior distribution for β that is a weighted average of the posteriors for each model specification. Thus, the full marginal posterior distribution of the regression parameters, β , accounting for all the possible models, is a mixture distribution, in this case, a mixture of t -distributions:

$$p(\beta | y, X) \sim \sum_{j=1}^M \omega_j t(b_{pj}, s_{pj}^2 V_{pj}, d_{pj}). \quad (3.16)$$

Note that the subscript has been dropped from the parameter vector β since we are no longer conditioning on the model specification.

If a point estimate of β is desired as opposed to the entire posterior distribution, a Bayesian uses a loss function to derive the optimal point estimator given the distribution of the parameters of interest (cf. Zellner, 1971). If one uses a quadratic loss function,

$$\mathcal{L}(\bar{\beta}) = (\bar{\beta} - \beta)'(\bar{\beta} - \beta), \quad (3.17)$$

where $\bar{\beta}$ is the chosen point estimator and β is the unknown vector being estimated, then the optimal point estimator is the vector that minimizes the expected value of the loss

function in (3.17) where the expectation is taken with respect to the posterior distribution shown in (3.16). Thus, the optimal estimator $\bar{\beta}$ is the solution to

$$\operatorname{argmin}_{\bar{\beta}} E[\mathcal{L}(\bar{\beta})] = \int (\bar{\beta} - \beta)'(\bar{\beta} - \beta)p(\beta|y, X)d\beta. \quad (3.18)$$

The optimal estimator with respect to the quadratic expected loss shown above is the mean of the posterior distribution given in (3.16). Given the symmetry of the t -distribution, the mean of this mixture distribution is the weighted average of the individual means where the weights are the $\{\omega_j\}$ that represent the posterior model probabilities. Thus, the robust estimator accounting for the model specification uncertainty under the expected loss described in (3.18) is given by

$$\bar{\beta} = \sum_{j=1}^M \omega_j b_{pj}, \quad (3.19)$$

recalling that b_{pj} is the mean of each model's posterior distribution as given in (3.10) which is the optimal estimator $\bar{\beta}_j$ for the quadratic loss function. This is the estimator used in this paper; researchers can easily employ different loss functions better suited to particular applications to derive alternative estimators which are optimal for the loss functions so employed. For example, an absolute loss function results in the posterior median being the appropriate point estimator. In applications such as hedging, the loss function could also be designed to provide an estimator with optimal characteristics relative to the potential costs from hedging with an incorrect hedge ratio (cf. Lence and Hayes, 1994b).

An interesting point about robust estimators such as the one defined in equation (3.19) is that they are averages of parameters estimated from different models, not all of which can be "true." If one takes a classical point of view, some of these models have likely excluded variables that belong in the model, implying a possibility of biased parameters being included in the subsequent average, robust estimator. Because the concept of unbiasedness

of estimators is not a particular concern to Bayesians (in fact, the concept does not really translate to the Bayesian paradigm), this is not seen as a large problem by Bayesians. Recall that almost all Bayesian point estimators are averages across values of a posterior distribution (see, equation (3.18), for example) and one can see how it is possible to motivate averaging across a range of values that cannot all possibly be correct.

4. The Data

We use data on three commodities in our examples: corn, soybeans, and corn oil. Corn oil does not have a futures contract, so we estimate a cross-hedge using the soybean oil contract. In all three applications we assume that hedges are held for one month in the nearby contract. For corn and soybean the cash prices are those reported for Central Illinois by the Illinois Agricultural Statistics Service. For corn oil, the cash price is a wet mill crude corn oil price. Cash and futures prices are collected on the last business day of each month.

For corn, possible variables to include in the model are corn exports, interest rates, and lags of the cash and futures prices. For soybeans, the possible variables are soybean stocks (at mills), soybean crushings, and lags of the cash and futures prices. For corn oil, the possible variables are soybean oil stocks, interest rates, and lags of the cash and futures prices. Note the lagged futures prices are carefully constructed for each expiring contract such that for any given observation at time t , the futures price at time $t - n$ represents the n period lagged price for the same expiring contract. The data in all cases spans from January 1975 through April 2003, resulting in 340 observations.

5. Empirical Results

5.1. The Models and Priors

Given the data described above, we considered 64 distinct models for each commodity. All models have the cash price as the dependent variable. The contemporaneous futures price and twelve monthly dummy variables to model seasonality are included in all models as regressors. Two exogenous variables were considered for inclusion in each model (listed above in the data section). The inclusion of none, one, or both exogenous variables gives four possible specifications with respect to exogenous variables.

To account for possible dynamic effects in the stochastics of the cash and future prices, including possible nonstationarity, up to three lagged values of both prices were considered for inclusion. The lags were only included in complete ordered sets; for example, for the cash price the options were: no lags, $[p_{t-1}]$, $[[p_{t-1}, p_{t-2}]$, and $[p_{t-1}, p_{t-2}, p_{t-3}]$. That is, no “holes” were allowed in the lag structure. This uncertainty over lagged prices in the model adds four possible lag specifications for the cash price and four possible lag specifications for the futures price.

Allowing all possible combinations of these three dimensions of model specification yields the $(4 \times 4 \times 4 =)$ 64 total model specifications for each commodity studied here. Since all models contain twelve monthly dummies and the current futures price, the smallest model has 13 regressors and the largest has 21 (the 13 always included plus two exogenous variables, 3 lagged cash prices, and 3 lagged futures prices). Some of the models are nested within others, some are not. Thus, classical statistics does not have an exact or Fisher-type test for deciding among or ranking these models, making this set of models a good application for the Bayesian approach.

Given these 64 models for each commodity, and the data described in section 4 above, only the prior distributions still need specification to allow completion of the estimation process. The models each receive equal prior weights; that is, $\mu_j = 1/64 \quad \forall j$. The priors on the regression parameters follow distributions as described in equations (3.2) and (3.3). The dimension of the priors depends on the exact model specification, but priors on parameters associated with specific regressors do not change with model specification (i.e., if the variable is in the model, its prior is the same every time). For all three models, the prior means that do not equal zero for regression parameters were chosen to match the authors' consensus prior belief about likely parameter values by thinking in terms of elasticities, the speed of adjustment of prices, market efficiency, and common hedge ratios employed in practice. Priors, by definition, are subjective; however, we later include sensitivity analysis of our priors which show the results are not overly sensitive to the prior (although not completely insensitive either). The largest model, with all possible regressors included, is used to detail the prior. The order of regressors for the purposes of displaying these priors is $[f_t, X_{1,t-1}, X_{2,t-1}, p_{t-1}, p_{t-2}, p_{t-3}, f_{t-1}, f_{t-2}, f_{t-3}, D]$, where D is the matrix holding the 12 monthly dummies.

For the corn model, the prior means are set to

$$b_0 = [0.95, 0.3, -0.025, 0.9, 0, 0, -1.0, 0 \dots, 0]', \quad (5.1)$$

where nonzero prior means are employed only for (in order) the hedge ratio, the two exogenous variables (corn exports, then interest rate), and the first lags of both cash and future prices. Thus, the prior hedge ratio is 0.95, corn exports are assumed, *a priori*, to increase the cash price while higher interest rates lower it, and the price dynamics of the prior are for high, positive autocorrelation in cash prices and a unit root in the futures prices. The prior variance matrix V_{0j} is a k -dimensional diagonal matrix with ones on the

diagonal except for the five elements with nonzero prior means. The diagonal elements for those five parameters are set to 0.01, 0.25, 4.0, 0.25, and 0.25, respectively.

For the soybean model, the prior means are set to

$$b_0 = [0.95, -0.2, -0.2, 0.9, 0, 0, -1.0, 0 \dots, 0]', \quad (5.2)$$

where nonzero prior means are employed only for the same five regressors. The prior hedge ratio is again set equal to 0.95 and identical price dynamics are assumed. The priors on the exogenous variables assume higher stocks and crush both lower the cash price. The prior variance matrix V_{0j} is again a k -dimensional diagonal matrix with ones on the diagonal except for the five elements with nonzero prior means. The diagonal elements for those five parameters are set to 0.01, 0.25, 0.25, 0.25, and 0.25, respectively.

For the corn oil model, the prior means are set to

$$b_0 = [0.95, -0.5, -0.5, 0.9, 0, 0, -1.0, 0 \dots, 0]', \quad (5.3)$$

where nonzero prior means are employed only for the same five regressors. The prior hedge ratio is again set equal to 0.95 and identical price dynamics are assumed. The priors on the exogenous variables assume higher soybean oil stocks and interest rates both lower the cash price. The prior variance matrix V_{0j} is again a k -dimensional diagonal matrix with ones on the diagonal except for the five elements with nonzero prior means. The diagonal elements for those five parameters are set to 0.01, 0.25, 4.0, 0.25, and 0.25, respectively.

For all models, the remaining prior parameters are set to

$$s_{0j}^2 = 1, \quad d_{0j} = 15, \quad j = 1, \dots, M. \quad (5.4)$$

This completes the specification of all features of the estimation process. To derive the results, the marginal posterior distribution of the regression parameters β_j is computed

using the above prior values for each model according to equation (3.10), yielding a posterior mean conditional on each model specification. The posterior model weights for all the models are then calculated using equations (3.12) and (3.15). These two sets of results are combined according to (3.19) to arrive at the marginal posterior point estimator $\bar{\beta}$, which is the robust hedge ratio. This robust hedge ratio, $\bar{\beta}$, optimally accounts for the model specification uncertainty, but is simple to calculate as the weighted average of the 64 individual coefficients for f_t from the regression models.

Because of the prior distributions chosen and the specification of the likelihood function, all the results in this application can be derived analytically. Thus, numerical methods were not necessary to approximate the posterior distribution of the regression parameters or to compute the model odds. However, if other likelihood functions (with assumptions of non-normal residuals) or priors were used, numerical methods would allow approximation of the analogs to all the expressions here and the same general process to be followed.¹

5.2. Posterior Model Probabilities for Corn and Soybeans

The first results worth investigating are the posterior model probabilities. If the ω_j are concentrated tightly over one (or similar) models, then model specification uncertainty is not a significant problem in estimating hedge ratios. However, if the model probabilities are spread over many models, model specification for hedge ratio estimation needs more attention. Obviously, after examining the model odds, we also need to determine if hedge

¹ An important type of different distribution for the residuals would be if the models suffer from autocorrelation. Given that this is time series data and lagged dependent variables (plus several lags of future prices) are included in the model, this is a possible problem. While autocorrelated errors could be modeled and the likelihood function adjusted, we stayed with the simple, white-noise error assumption for all models. Checks on the models showed no autocorrelation in the soybean and corn oil data, using both the standard Durbin-Watson test and the m test which accounts for the lagged dependent variables present in the regressor matrix. The corn models showed some evidence of autocorrelation using the m test, but not the Durbin-Watson test. We choose, therefore, to present results for all models based on the assumption of no autocorrelation.

ratios vary across model specifications because if estimated hedge ratios are (relatively) constant across specifications, then model specification does not matter.

For corn, 19 models receive at least 1% of the posterior model probabilities, indicating significant model specification uncertainty. In fact, six models have posterior probabilities of over 5% and the most likely model has only a 16.6% posterior probability in its favor. This most likely model has no exogenous variables, one lag of cash price and one lag of the future price.

For soybeans, only five models have at least 1% of the posterior model probabilities and two models combine for 80% of the probability, suggesting much less model specification uncertainty than for corn. The most likely model has 49.4% of the posterior probability and contains soybean stocks and no lagged prices at all.

The 64 models are too many to display the individual model weights in a meaningful table, so instead we present in table 1 the marginal probabilities of each model specification feature (such as probability of one lag of cash price, etc.). Each of these probabilities is the sum of posterior probabilities for all the individual models that share the named model specification feature (such as all models with both exogenous variables included). These probabilities measure the posterior support for that feature of the model specification while controlling for the effect of all other model specification choices. Given our set of models, each of these marginal probabilities contains 16 separate models. The results in table 1, columns 2 and 4 contain these marginal probabilities for each model feature. The values are to be interpreted as the posterior support in favor of the models containing that feature. For example using the first row of column 2, we would say that 56.6% of the posterior support is placed on corn models with no exogenous variables implying that such models are slightly favored relative to all the possible models with one or two exogenous variables.

The results for the corn model (table 1, column 2) clearly show posterior support in favor of either no exogenous variables or the inclusion of the interest rate. Models with one lag of cash price are most favored, with considerable support for two lags as an alternative. The same is true for lags of the futures price. The results for the soybean model (table 1, column 4) show strong posterior support for the inclusion of soybean stocks, with no exogenous variables as a strong second option. The soybean results show overwhelming posterior support for the exclusion of all cash and futures price lags with approximately 92% support in favor of leaving out all lagged prices. Overall, the model probability results show more uncertainty over the corn model features than the soybean model.

5.3. Robust Hedge Ratios

Moving on to the estimated hedge ratios, the results in table 1 make clear that the hedge ratio does vary with model specification. Table 2 displays the marginal posterior results for both commodities. For corn, we see the robust hedge ratio after accounting for model specification uncertainty is 0.941 with a range over the 64 models from 0.900 to 0.990. Checking table 1 shows that the most important aspect of model specification with respect to the corn hedge ratio is due to the presence or absence of lagged prices. Columns 3 and 5 in table 1 display the hedge ratios conditional on a particular model specification feature. If no price lags are present, the conditional hedge ratio goes up over 0.98. As long as at least one lagged price (cash or futures) is present, the conditional hedge ratio falls back into the neighborhood of 0.94, right where the robust hedge ratio lies. With a range from 0.900 to 0.990, clearly these conditional hedge ratios hide much of the variation across individual models, but they are useful for identifying the features of model specification that have the most influence on the estimated hedge ratios.

For the soybean model, the results in table 2 show the robust hedge ratio equals 0.985 with a range across models from 0.918 to 0.991. The results in table 1 show the conditional soybean hedge ratio moves much more with conditioning on different model specification features than in the case of corn. The presence or absence of exogenous variables matters as does the presence/absence of both lagged cash and futures prices. Once lagged prices are included, the number of lags does not impact the conditional hedge ratios. Again, the conditional hedge ratios hide the full variation in hedge ratios estimated across all 64 models. Here, with soybeans, we find that the estimated hedge ratio is sensitive to a wider range of model specification issues. Interestingly, while the soybean hedge ratio is more sensitive to model specification, the results reveal less uncertainty about the correct model for soybeans with two models dominating the posterior model probabilities.

Examining the posterior standard deviation of the robust hedge ratios in table 2 reveals them to be very small relative to the hedge ratios (0.006 and 0.005, respectively, for corn and soybeans). This implies high statistical precision. For example, for the corn model a Bayesian analog to the familiar confidence interval, the 95% posterior density region (PDR), would be ± 0.012 around the Bayesian robust point estimator, suggesting we have successfully identified the central tendency of the hedge ratio relative to the variation in both the model specification and the data. Clearly, this PDR is so small that most of the individual model estimates would be outside this range.

5.4. Risk Reduction Performance for Corn and Soybeans

Following the suggestion of Myers and Thompson and the methodology of Baillie and Myers, the Bayesian hedge ratios are compared with traditional hedge ratios from various models in an out-of-sample simulation. The hedge ratios are estimated first with monthly data from 1975 through December of 1999, then the effectiveness of monthly

hedges are simulated using data for 2000. The models are then re-estimated adding twelve more months of data (through December of 2000), and the resulting hedge ratios used for simulated hedging in 2001, and so forth. The result is 40 simulated monthly hedges from January 2000 through April 2003.

A total of seven hedge ratios are compared in the simulations. Standards for comparison are provided by the traditional constant hedge ratio (estimated with price levels) and a naïve one-to-one hedge. Along with the robust Bayesian hedge ratio presented in this paper, four other hedge ratios are considered from the 64 models estimated. These four are the largest and smallest hedge ratios estimated from an individual model among the set considered and the hedge ratios from the models that receive the largest and smallest posterior probability weight in the model specification part of our process.

In the spirit of a risk minimizing hedge, the simulation procedure calculates the variability in the portfolio consisting of a cash position and the chosen futures hedge. For corn and soybeans, variability is measured in cents per bushel as the change in cash price minus the change in the hedge value. This approach closely reflects the change in economic value of the hedgers overall position. The standard deviation and risk reduction relative to an unhedged position are presented in table 3.

For corn the monthly standard deviation falls from 13.00 cents per bushel to 4.00 cents per bushel for the Bayesian robust hedge, a 90.5% reduction in risk (variance) from the unhedged position (results are shown in table 3). All of the corn hedge ratios reduce risk by similar amounts, ranging from 89.7% for the Bayesian least likely model to a high of 91.1% for the unitary hedge ratio. In fact, there is no statistical difference in the risk reduction performance across the seven hedge ratios (tested using F-tests).

Soybean hedges reduced risk from 26.08 cents per bushel for an unhedged position to 8.87 cents for the Bayesian robust hedge ratio. Again, given the similarity of the hedge ratios, it is not surprising that the performance across hedge ratios is very similar. The monthly standard deviation of the hedged positions is very close to 9 cents per bushel for all seven hedge ratios (table 3). The Bayesian minimum hedge ratio actually provides the greatest risk reduction at 89.2% while the unitary hedge ratio is the least effective with an 88.1% reduction in risk from an unhedged position. The robust hedge ratio has a risk reduction performance in the middle of the seven hedge ratios tested. Again, the different hedge ratios do not produce statistically different risk reduction levels using F-tests for equality of variance.

The presented results are consistent with those of Baillie and Myers in that for some commodities, such as corn and soybeans, more advanced hedge ratio estimation techniques may not significantly increase hedge effectiveness. This may or may not be the case with less standardized commodities (such as slaughter cattle) or when estimating cross-hedge ratios (such as hedging cottonseed meal with soybean meal futures). We explore this below with an application to a corn oil cross hedge using soybean oil futures.

In the presented results, it is noteworthy that the one-for-one textbook hedge is the most effective out-of-sample for corn and does not produce statistically different results for soybeans. The results are particularly interesting in corn, where the Bayesian minimum hedge ratio averaged 0.900 over the simulation period. In this case, over a three year period, using a hedge ratio of 0.90 and 1.00 did not produce statistically different results. This is consistent with Jong, Roon, and Veld, who find that naïve one-for-one hedging may perform equally well to estimated ratios in practice.

5.5. The Impact of the Priors

Many researchers criticize Bayesian approaches due to the influence of subjective prior information on the posterior distribution (and through that, the “estimators”). The prior distribution effects the posterior distribution in two ways: through the prior mean and the prior variances. Obviously, changing the prior mean will change the posterior mean since the posterior mean is a weighted average of the prior mean and the standard, likelihood based estimator as shown in equation (3.7). Sensitivity analysis not reported in detail here showed that changes in the prior means did indeed result in changes in the point estimator for the robust hedge ratio, with the estimated soybean hedge ratio varying from about 0.95 to 0.99 as the prior means for all parameters were varied over a fairly wide range of values (such as from 0.50 to 1.00).

The impact of the prior means is reasonable and easy to evaluate since their effects are fairly transparent and the prior means should always be clearly stated by Bayesian researchers. Evaluation of the role of prior variances is somewhat harder to determine from simple inspection. In particular, the prior mean matters less if the prior variances are large enough to allow the data to contribute the majority of the information in the posterior distribution. If changes in the prior variances do not result in disproportionate changes in the posterior distribution, then one might reasonably conclude that the prior’s influence on the posterior distribution is reasonable. To evaluate our prior, we repeated the analysis for corn and soybeans with four different prior variance matrices, all scalar multiples of the base prior distribution. The results of this sensitivity analysis for soybeans are presented in table 4 with similar results obtained for corn. From the small changes in the results that occurs with fairly large changes in the prior variances, we conclude that the prior distribution used here is performing satisfactorily.

5.6. Cross-Hedging Results

The results for corn and soybeans provide evidence of the ability of the Bayesian robust hedge ratio estimation to identify model specification features and to estimate reasonable hedge ratios, but did not provide gains in risk reduction. To examine if such evidence could be provided in a case with greater model specification uncertainty, we also modeled a cross-hedge for corn oil using soybean oil futures.

The corn oil cross-hedge model identified four models with greater than 1% posterior support with 89% of the weight on the simplest model—no lags and no exogenous variables. Table 5 displays the results of the posterior probabilities by model feature. These probabilities, and the cross-hedge ratios associated with the different model specifications show that while the Bayesian procedure easily identifies the model with support from the data, an incorrect model can have severe consequences. While the specification with respect to the soybean oil stocks and interest rates does not have a significant impact on the cross-hedge ratio, the presence of lagged cash or soybean oil futures prices has an enormous impact. With lagged futures prices included the estimated hedge ratio drops from 0.95 to 0.52; with lagged cash prices included the estimated hedge ratio drops from 0.95 to 0.26. These are obviously economically significant changes in risk management programs. Table 6 displays the robust Bayesian cross-hedge ratio and the range of estimates found over the 64 models. Note that the posterior standard deviation of the Bayesian estimator converts to a 95% posterior density region that is ± 0.139 around the point estimator, or (0.789, 1.067). That range contains only four of the 64 models. Thus, while the hedge ratio estimates vary widely across models, the Bayesian procedure computes a robust hedge ratio that is statistically precise enough to exclude many of the individual model hedge ratios estimated. Also, it is worth noting that the models with estimated hedge ratios

that fall outside that range also have small posterior standard deviations, meaning that individual models can estimate hedge ratios that vary widely yet are all statistically precise *conditional on the model specification*.

When the robust cross-hedge is compared to other possible choices in terms of risk reduction performance, we find better results than were found for corn and soybean hedging. Table 7 reports these results, showing the robust Bayesian hedge to have the best risk reduction results, with a 20.7% reduction. This compares to a 16.7% reduction for the unitary hedge, and a 14.4% reduction if we had chosen the minimum cross-hedge ratio from among the 64 models. Both of these differences are significant at the 5% level. The minimum estimated hedge ratio is not an odd model either; any model with neither interest rates or soybean oil stocks, no lagged futures prices, and any set of lagged cash prices produces a small hedge ratio of around 0.25, so a researcher might choose such a specification *a priori*.

While the Bayesian approach would guide a researcher away from these models, without it a model with lagged prices is quite likely to be specified. This will result in either a cross-hedge ratio of 0.25 if only lagged cash prices are included or a cross-hedge ratio around 0.44 if both cash and futures lags are included. Such departures from the robust cross-hedge ratio of 0.928 will cause the risk management to differ by economically significant amounts.

6. Conclusions

Since Myers and Thompson raised the important question of model specification's influence on estimated hedge ratios, much work has been done on estimating hedge ratios while little work has been done on the issue of model specification uncertainty in hedging

models. We have returned to this important topic and introduced a systematic approach to model specification uncertainty using Bayesian inference to treat the uncertainty like other uncertain parameters. This allows the model specification uncertainty to be integrated out of the estimation and inference problems and marginal statistical inferences to be made that optimally account for the relative probabilities of the different models considered. The approach also allows for inference concerning the uncertain model specification itself, providing probability measures of support for various models, variables, and dynamic specifications. These empirical results can guide future researchers in the direction of the models which enjoyed the most support from previous research.

The Bayesian robust estimation approach was applied to data for corn, soybeans, and corn oil. Robust hedge ratios were computed, along with the posterior probabilities of individual models and model specification features. The individual models produced hedge ratios which varied relatively widely, while the robust hedge ratios integrating over model uncertainty were quite statistically precise. The model specification results also identified which features of the model specification had significant impacts on the estimated hedge ratio. We found that for corn, the specification of exogenous variables was crucial to the posterior support of the model, but it was the presence or absence of lagged prices that had the biggest impact on the level of the estimated hedge ratio. For soybeans, we found that all aspects of model specification had the potential to significantly impact the estimated hedge ratio. For the corn oil cross-hedge, lagged prices caused large drops in the estimated cross-hedge ratio. The process identified few credible models for soybeans, with two out of the 64 models gathering 80% of the posterior probability. Similarly, for corn oil only four models contained over 98% of the posterior probability. For corn, many models enjoyed relatively comparable posterior support.

The wide range of the estimated hedge ratios makes clear that model specification is an important issue in hedging models and that Myers and Thompson were right to raise the issue. The approach taken here allows a researcher to avoid choosing a single, potentially incorrect, model specification (or to scientifically identify supported model specifications without incurring pre-test bias). This is an important ability given that the empirical results show the estimated hedge ratio can change by a magnitude of three in the corn oil cross-hedge model depending on which model is selected. By incorporating 64 possible models and integrating across that model specification uncertainty, the resulting robust hedge ratios are not only robust but quite stable.

We believe that the approach demonstrated here has great potential to provide better (more robust) estimators of hedge ratios and other important economic parameters. Given that hedge ratios are designed to reduce risk, the ability to reduce the risk of estimation biases due to model specification seems attractive. While the risk reduction performance of the robust hedge ratios for corn and soybeans was not significantly better than that of other hedge ratios, it was not worse either. For the corn oil cross-hedging model, the Bayesian robust hedge ratio did have the best risk reduction performance, showing that in a commodity where the estimated hedge ratios vary more across models, using the Bayesian approach can make an important difference in risk reduction. Since one can rarely be sure, *a priori*, which case one has, the robust estimation procedure presented here seems to have a useful advantage. We also think the information contained in the posterior probabilities of the individual models and the model specification features can help guide researchers for future investigations concerning hedging models and what factors influence price levels in commodity markets.

References

- Alston, J. M., and J. A. Chalfant. "The Silence of the Lambdas: A Test of the Almost Ideal and Rotterdam Models Specification Paper." *American Journal of Agricultural Economics* 75 (May 1993):304-313.
- Anderson, R.W. and J-P. Danthine. "Cross Hedging." *Journal of Political Economy* 89 (1981):1182-1196.
- Baillie, R. T., and R. J. Myers. "Bivariate GARCH Estimation of the Optimal Commodity Futures Hedge," *J. Appl. Econometrics* 6 (April-June 1991):109-124.
- Brorsen, B.W., D.W. Buck, and S.R. Koontz. "Hedging Hard Red Winter Wheat: Kansas City versus Chicago" *Journal of Futures Markets* 18 (1998):449-466.
- Collins, R.A. "The Risk Management Effectiveness of Multivariate Hedging Models in the U.S. Soy Complex," *Journal of Futures Markets* 20 (2000):189-204.
- Dorfman, J. H., and W. D. Lastrapes. "The Dynamic Responses of Crop and Livestock Prices to Money-Supply Shocks: A Bayesian Analysis Using Long-Run Identifying Restrictions," *American Journal of Agricultural Economics* 78 (August 1996):530-541.
- Ferguson, R. and D. Leistikow. "Are Regression Approach Futures Hedge Ratios Stationary," *Journal of Futures Markets* 18 (1998):851-866.
- Garcia, P., J.S. Roh, and R.M. Leuthold. "Simultaneously Determined, Time-Varying Hedge Ratios in the Soybean Complex," *Applied Economics* 27 (1995):1127-1134.
- Hull, J.C. *Fundamentals of Futures and Options Markets*. (4th ed.) New Jersey: Prentice Hall, 2002.
- International Treasurer. "FAS 133: The Impact on Risk Management Strategy & Technique," *The Journal of Global Treasury and Financial Risk Management*. (1998).
- Jong, A.D., F. De Roon, and C. Veld. "Out-of-Sample Hedging Effectiveness of Currency Futures for Alternative Models and Hedging Strategies," *Journal of Futures Markets* 17 (1997):817-837
- Koop, G. *Bayesian Econometrics*. Chichester, England: Wiley, 2003.
- Lence, S. H., and D. J. Hayes. "The Empirical Minimum-Variance Hedge," *American Journal of Agricultural Economics* 76 (Feb. 1994a):94-104.
- Lence, S. H., and D. J. Hayes. "Parameter-based Decision Making Under Estimation Risk: an Application to Futures Trading," *Journal of Finance* 49 (1994b):345-357.
- Myers, R. J., and S. R. Thompson. "Generalized Optimal Hedge Ratio Estimation," *American Journal of Agricultural Economics* 71 (Nov. 1989): 858-868.
- Poirier, D. J. "A Bayesian View of Nominal Money and Real Output through a New Classical Macroeconomic Window (with discussion)," *Journal of Business & Economic Statistics* 9 (1991):125-148.
- Shafer, C.E. "Hedge Ratios and Basis Behavior: An Intuitive Insight?" *Journal of Futures Markets* 13 (1993):837-847.

Vukina, T. "Hedging With Forecasting: A State-Space Approach To Modeling Vector-Valued Time Series," *Journal of Futures Markets* 12 (1992):307-328.

Witt, H.J., T.C. Schroeder, and M.L. Hayenga. "Comparison of Analytical Approaches for Estimating Hedge Ratios in Agricultural Commodities," *Journal of Futures Markets* 7 (1987):135-146.

Zellner, A. *An Introduction to Bayesian Inference in Econometrics*. New York: John Wiley & Sons, 1971.

Table 1: Model Feature Posterior Probabilities

Feature	Corn Model Post. Prob.	Corn Model Hedge Ratio	Soybean Model Post. Prob.	Soybean Model Hedge Ratio
no x	0.566	0.941	0.369	0.977
X_1	0.005	0.940	0.504	0.990
X_2	0.424	0.941	0.079	0.981
both X	0.004	0.940	0.048	0.990
no p lags	0.027	0.984	0.918	0.987
p_{t-1}	0.574	0.940	0.009	0.957
p_{t-2}	0.288	0.940	0.062	0.956
p_{t-3}	0.111	0.940	0.011	0.956
no f lags	0.027	0.984	0.918	0.987
f_{t-1}	0.502	0.939	0.011	0.957
f_{t-2}	0.357	0.941	0.061	0.956
f_{t-3}	0.114	0.941	0.011	0.956

Note: For corn X_1 = corn exports, X_2 = interest rate. For soybeans, X_1 = soybean stocks, X_2 = soybean crush.

Table 2: Robust Hedge Ratios

	Corn	Soybean
Robust Hedge Ratio	0.941	0.985
Standard deviation	0.006	0.005
Minimum Hedge Ratio	0.900	0.918
Maximum Hedge Ratio	0.990	0.991

Table 3: Risk Reduction Performance of Alternative Hedge Ratios

Hedge Ratio	Corn Model Standard Deviation	Corn Model Percent Reduction	Soybean Model Standard Deviation	Soybean Model Percent Reduction
No Hedge	13.00	0.0	26.08	0.0
Robust	4.00	90.5	8.87	88.4
Maximum	3.90	91.0	8.94	88.3
Minimum	4.18	89.7	8.58	89.2
Most Likely	3.97	90.7	8.87	88.4
Constant	3.91	90.9	8.85	88.5
Unitary	3.88	91.1	9.01	88.1

Note: Standard deviation is measured with price changes over monthly horizons in cents per bushel.

Table 4: Sensitivity Analysis of the Prior Variances for the Soybean Model

Multiple of Base Variance	Robust Hedge Ratio	Posterior Weight on Most Likely Model	Posterior Weight on Models with No Exogenous Variables
0.5	0.979	0.595	0.626
1.0	0.983	0.429	0.464
2.0	0.984	0.334	0.420
5.0	0.979	0.482	0.540

Table 5: Cross-Hedge Model Feature Posterior Probabilities

Feature	Post. Prob.	Hedge Ratio
no x	0.974	0.928
X_1	0.016	0.935
X_2	0.010	0.898
both X	0.000	0.912
no p lags	0.975	0.945
p_{t-1}	0.024	0.260
p_{t-2}	0.001	0.260
p_{t-3}	0.000	0.265
no f lags	0.941	0.954
f_{t-1}	0.057	0.526
f_{t-2}	0.002	0.522
f_{t-3}	0.000	0.522

Note: Corn oil cross-hedge using soybean oil futures. X_1 = soybean oil stocks, X_2 = interest rate.

Table 6: Robust Corn Oil Cross-Hedge Ratios

Corn Oil vs. Soybean Oil	
Robust Hedge Ratio	0.928
Standard deviation	0.033
Minimum Hedge Ratio	0.245
Maximum Hedge Ratio	0.979

Table 7: Risk Reduction Performance of Alternative Cross-Hedge Ratios

Hedge Ratio	Standard Deviation	Percent
No Hedge	1.648	0.0
Robust	1.467	20.7
Maximum	1.484	18.9
Minimum	1.525	14.4
Most Likely	1.478	19.5
Constant	1.478	19.5
Unitary	1.504	16.7

Note: Corn oil cross-hedge using soybean oil futures. Standard deviation is measured with price changes over monthly horizons in cents per pound.