The Empirical Minimum Variance Hedge

Sergio H. Lence and Dermot J. Hayes

Working Paper 93-WP 109
June 1993

Center for Agricultural and Rural Development
Iowa State University
Ames, Iowa 50011-1070

Journal Paper No. J-15025 of the Iowa Agriculture and Home Economics Experiment Station, Ames, Iowa. Project No. 2994.

Sergio H. Lence is a CARD postdoctoral research associate; and Dermot J. Hayes is associate professor of economics and head of the Trade and Agricultural Policy Division of CARD.

THE EMPIRICAL MINIMUM VARIANCE HEDGE

Abstract

Decision making under unknown true parameters (estimation risk) is discussed along with Bayes and parameter certainty equivalent (PCE) criteria. Bayes criterion provides the solution for optimal decision making under estimation risk in a manner consistent with expected utility maximization. The PCE method is not consistent with expected utility maximization, but is the approach commonly used.

Bayes criterion is applied to solve for the minimum variance hedge ratio (MVH) in two scenarios based on the multivariate normal distribution. Simulations show that discrepancies between prior and sample parameters may lead to substantial differences between Bayesian and PCE MVHs. Such discrepancies also highlight the superiority of Bayes criterion over the PCE, in the sense that the PCE method cannot not yield decision rules that contain prior (or nonsample) along with sample information.

THE EMPIRICAL MINIMUM VARIANCE HEDGE

Estimation risk can be defined as a situation in which the joint probability density function (pdf) of the random variables associated with a decision problem is not known with certainty. This is a common occurrence in economics; for example, parameters such as the marginal productivity of fertilizer, the elasticity of demand, and the regression coefficient of futures on cash prices are rarely known by the decision maker. Agents who must make decisions in the presence of random variables are generally confronted with the additional uncertainty of less than perfect knowledge about the pdf governing the distribution of those variables.

Almost all studies involving decision making in the presence of estimation risk implicitly use the "plug-in" or "parameter certainty equivalent" (PCE) approach. This consists of developing the theoretical decision model assuming that the pdf and its parameters are known with certainty. Once the optimal decision rule is derived, the empirical application proceeds by substituting sample estimates for the unknown parameters in the formula for the optimal decision rule. Although it is intuitively appealing and empirically tractable, the parameter certainty equivalent method has no axiomatic foundations and is not consistent with expected utility maximization.

Although the applied section of this paper is concerned with hedging, the shortcomings of the PCE approach can best be demonstrated by means of a speculative example. Assuming that the true parameters are known, theory predicts that a risk-averse individual will speculate if the known futures mean is different from the current futures price. In order to determine the optimal speculative position empirically, the PCE method advocates using the sample futures mean as a substitute for the true but unknown futures mean. Notice that prior information, such as possible strong belief in the efficient market hypothesis, and sample information, such as the standard errors of the estimated parameters, are ignored by the PCE. Taken to its extreme, the PCE would predict a long speculative position whenever the mean of recent futures prices is lower than the current futures price and a short position when the opposite is true. This is clearly a questionable speculative behavior; yet, this procedure is essentially what is followed when using the PCE.

The question addressed in this paper is how the theory itself (and by extension the empirical methodology) is changed if we admit that there is uncertainty about the magnitudes of the true parameters and that information about them can be obtained from both historical data and other sources.

When estimation risk is due to imperfect knowledge about the parameters of the joint pdf (given that the functional form of the pdf is known), the methodology that is consistent with the expected utility paradigm is Bayes decision criterion (DeGroot, Chapters 7 and 8). Bayes criterion takes into account uncertainty regarding the unknown true parameters by assigning a pdf to these parameters and then integrating over the parameter space.

Bayes decision criterion has been thoroughly studied in statistics (Raiffa and Schlaifer, DeGroot, Berger, Klein et al.). It has also been applied to solve important problems in finance, such as security market equilibrium (Bawa; Barry and Brown; Coles and Loewenstein), portfolio choice (Klein and Bawa 1976 and 1977; Bawa; Stephen Brown; Chen and Brown; Alexander and Resnick; Jorion 1985 and 1986; Frost and Savarino; Cheung and Kwan), and option pricing (Boyle and Ananthanarayanan). However, Bayes criterion has been largely ignored in agricultural economics. A possible exception is a study by Dixon and Barry, who modeled the allocation of funds by an agricultural bank among three assets and concluded that estimation risk influenced the portfolio's composition.

Only in recent years have agricultural economists shown some interest in estimation risk (e.g., Dixon and Barry; Pope and Ziemer; Collender and Zilberman; Collender; Chalfant, Collender, and Subramanian). Pope and Ziemer examined the performance of alternative estimation methods for second-degree stochastic efficiency analysis. Using Monte Carlo experiments, they found that the plug-in approach generally performed no better than the empirical distribution function, and that the empirical distribution generally led to more correct rankings under small sample sizes. Collender and Zilberman analyzed the optimal land allocation problem under alternative joint pdfs for crop returns. They concluded that farmers with different opinions regarding the joint pdf of crop returns will both allocate and value land differently, even if they have the same degree of absolute risk aversion and identical opinions about the mean and the variance of crop returns. Collender addressed the decision maker's ability to distinguish among different farm plans based on their sample means and variances. In his application, Collender found that it may be statistically impossible to distinguish among most estimated mean-variance combinations lying at the efficient frontier at reasonable levels of significance, even for large sample sizes. The work by Chalfant, Collender, and Subramanian studies the sampling properties of the portfolio allocations based on the PCE approach. They show that allocation decisions from the PCE method are biased and inefficient. They also propose an alternative approach which leads to decisions that are unbiased and that in addition have lower variance.

The goals of this study are twofold. First, the Bayesian approach is discussed in general terms and is compared to the PCE approach. Second, Bayes criterion is applied to the estimation of the minimum variance hedge ratio. The issue addressed here is related but different from that studied by Collender and Zilberman. These authors analyzed the problems associated with using an incorrect functional form for the joint pdf, assuming perfect knowledge about the parameters. Here, we are concerned with the problems associated with less than perfect knowledge about the parameters, assuming perfect knowledge about the functional form of the joint pdf.

Decision Making under Uncertainty

The standard optimization problem under uncertainty can be represented by the following expression

$$(2.1) \quad \max_{\underline{d} \in D} \mathrm{E}_{\underline{y}|\underline{\theta}}(U) = \max_{\underline{d} \in D} \int_Y U[R(\underline{d}, \underline{y})]\, p(\underline{y}|\underline{\theta})\, d\underline{y}$$

where $\mathrm{E}(\cdot)$ is the expectation operator, $U[R(\underline{d}, \underline{y})]$ is a von Neumann-Morgenstern utility function, $R(\underline{d}, \underline{y})$ is a function of a vector of decision variables $\underline{d}$ and a $(k \times 1)$ vector of future random variables $\underline{y} \equiv \underline{x}_{T+1}$ related to the decision problem, $p(\underline{y}|\underline{\theta})$ is the joint pdf of $\underline{y}$ given the vector of parameters $\underline{\theta}$, Y is the domain of $\underline{y}$, and D is the feasible decision set. The joint pdf will generally depend on $\underline{d}$ (Klein et al.), but for notational convenience we will denote it by $p(\underline{y}|\underline{\theta})$ rather than by $p(\underline{y}|\underline{\theta}, \underline{d})$ in the exposition.

The decision problem represented by expression (2.1) is the basic paradigm of expected utility theory, and it provides the framework used to develop the theories of the firm and the consumer under uncertainty (Hey). An important underlying assumption of (2.1) is that $p(\underline{y}|\underline{\theta})$ is perfectly known. However, there are many real-world situations in which this assumption is not valid, in which case there exists *estimation risk* (Bawa, Brown, and Klein). Estimation risk may arise because of less than perfect knowledge about either (i) the functional form of $p(\underline{y}|\underline{\theta})$, or (ii) the parameters contained in the vector $\underline{\theta}$ (given that the function $p(\underline{y}|\underline{\theta})$ is known with certainty). Although case (i) is relevant in certain situations (Bawa; Collender and Zilberman), in this paper we are concerned only with case (ii). In other words, we will define estimation risk as the situation where the decision maker knows the functional form of the joint pdf $p(\underline{y}|\underline{\theta})$ with certainty, but has less than perfect knowledge about the parameters in $\underline{\theta}$. Consequently, we will refer to the absence of estimation risk as a case of perfect parameter information (PPI).[1]

If $\underline{\theta}$ in (2.1) is not known with certainty, then $E_{\underline{y}|\underline{\theta}}(U)$ is not known either because the expectation is a function of $\underline{\theta}$; therefore, $E_{\underline{y}|\underline{\theta}}(U)$ cannot be maximized. Bayes decision criterion provides a remedy to this situation in a manner that is consistent with the axioms of expected utility theory (DeGroot, Chapters 7 and 8). The solution consists of taking into account the uncertainty about the parameters by postulating a joint pdf of $\underline{\theta}$ and integrating over the parameter space, i.e., the decision problem is

$$(2.2) \quad \max_{\underline{d} \in D} E_{\underline{\theta}}[E_{\underline{y}|\underline{\theta}}(U)] = \max_{\underline{d} \in D} \int_{\Theta} \{ \int_{Y} U[R(\underline{d}, \underline{y})]\, p(\underline{y}|\underline{\theta})\, d\underline{y} \}\, p(\underline{\theta}|\underline{X}, I_T)\, d\underline{\theta}$$

[1]Note that decisions based on the PPI need not be similar to those based on the PCE. The former assumes perfect prior knowledge about the parameters, whereas the latter assumes perfect confidence in the quality of the sample information. Because there is no need for the sample information in one scenario to be identical to the prior information used in the other, the resulting decisions may be different.

where $p(\underline{\theta}|\underline{X}, I_T)$ is the *posterior* pdf of $\underline{\theta}$ given the sample data matrix $\underline{X}$ and the prior (nonsample) information I_T, and Θ is the domain of $\underline{\theta}$. The sample data matrix $\underline{X} \equiv (\underline{x}_1, \ldots, \underline{x}_T)'$ is a $(T \times k)$ matrix of T past realizations of $\underline{x}$.

The posterior pdf $p(\underline{\theta}|\underline{X}, I_T)$ contains all the information available regarding the parameter vector $\underline{\theta}$ at the decision time T. This pdf conveys all the sample and nonsample information about $\underline{\theta}$ because it is obtained by application of Bayes theorem as follows:[1]

$$(2.3) \quad p(\underline{\theta}|\underline{X}, I_T) \propto p(\underline{\theta}|I_T)\, p(\underline{X}|\underline{\theta})$$

where $\propto$ denotes proportionality, $p(\underline{\theta}|I_T)$ is the *prior* pdf of $\underline{\theta}$, and $p(\underline{X}|\underline{\theta})$ is the *likelihood function*. The prior pdf represents the decision maker's prior (nonsample) information about $\underline{\theta}$; this pdf reflects the probabilities of different values of $\underline{\theta}$ assigned by the agent based on his practical experience, knowledge and beliefs. And according to the *Likelihood Principle*, all relevant experimental information about $\underline{\theta}$ after $\underline{X}$ is observed is contained in the likelihood function for the observed $\underline{X}$ (Berger, p. 28). By combining both sample and nonsample information, the posterior pdf provides a better assessment about the parameter vector than either the prior pdf or the likelihood function alone.

Expression (2.2) can be alternatively stated as

$$(2.4) \quad \max_{\underline{d} \in D} E_{\underline{\theta}}[E_{\underline{y}|\underline{\theta}}(U)] = \max_{\underline{d} \in D} \int_Y U[R(\underline{d}, \underline{y})]\, p(\underline{y}|\underline{X}, I_T)\, d\underline{y}$$

where $p(\underline{y}|\underline{X}, I_T)$ is the *predictive* pdf of $\underline{y}$.[2] Expression (2.4) facilitates the comparison of Bayes criterion with the PPI case (2.1). It can be observed that the only difference between the right-

---

[1]Recall that $p(a, e) = p(e)\, p(a|e) = p(a)\, p(e|a)$, and therefore $p(a|e) = p(a)\, p(e|a)/p(e) \propto p(a)\, p(e|a)$, where $p(a, e)$ is the *joint* pdf of any pair of random variables a and e, $p(a|e)$ and $p(e|a)$ are the *conditional* densities, and $p(a)$ and $p(e)$ are the *marginal* densities.

[2]Expression (2.4) is obtained by reversing the order of integration of (2.2), noting that $U[R(\underline{d}, \underline{y})]$ is independent from $\underline{\theta}$, and using the fact that $p(\underline{y}|\underline{X}, I_T) = \int_\Theta p(\underline{y}|\underline{\theta})\, p(\underline{\theta}|\underline{X}, I_T)\, d\underline{\theta}$.

hand sides of expressions (2.1) and (2.4) is that the joint pdf of $\underline{y}$ in the former is $p(\underline{y}|\underline{\theta})$, whereas in the latter it is $p(\underline{y}|\underline{X}, I_T)$. When the parameter vector $\underline{\theta}$ is known with certainty, the sample $\underline{X}$ adds no information about the parameters; therefore, the decision maker can ignore $\underline{X}$ and proceed to make decisions based on the joint pdf $p(\underline{y}|\underline{\theta})$ as indicated by (2.1). In the more common situation characterized by imperfect knowledge regarding $\underline{\theta}$, however, it is unreasonable to ignore either the prior or the sample information. In this case, by employing the predictive pdf the decision maker uses all the available information.

As mentioned in the introduction, the standard approach employed in studies involving estimation risk is the PCE. Letting $\underline{\hat{\theta}}(\underline{X})$ denote the sample point estimate of the unknown parameter vector $\underline{\theta}$, this method can be stated as

$$(2.5)\quad \max_{\underline{d} \in D} E_{\underline{y}|\underline{\theta} = \underline{\hat{\theta}}}(U) = \max_{\underline{d} \in D} \int_Y U[R(\underline{d}, \underline{y})]\, p[\underline{y}|\underline{\hat{\theta}}(\underline{X})]\, d\underline{y}$$

Simply put, in the PCE the sample point estimate $\underline{\hat{\theta}}(\underline{X})$ replaces the unknown vector $\underline{\theta}$ in (2.1), i.e., the parameter estimates are taken as if they were known with certainty. Solving the decision problem by means of the PCE is generally much easier than doing so using Bayes criterion, but the PCE has no axiomatic foundations. Klein et al. analyzed the necessary and sufficient conditions for the PCE approach to yield the optimal solution (i.e., the Bayesian solution). They show that these conditions are very restrictive and seldom fulfilled by the pdfs commonly used in economic studies. Moreover, they also show that the loss in utility from using the PCE rather than Bayes criterion may be very large.

An alternative decision rule in the presence of estimation risk has been recently suggested by Chalfant, Collender, and Subramanian. These authors show that in a mean-variance framework the PCE portfolio will have, on average, a greater return and a greater variance than the optimal PPI portfolio. As an alternative to the PCE, they propose a portfolio based on sample mean and variance estimates but in such a way that it is the same, on average, as the optimal PPI portfolio in the absence of estimation risk. They prove that such a portfolio is superior to that obtained using the PCE because it yields greater expected utility. One limitation of this decision rule, however, is that it does not necessarily maximize expected utility. In this regard, Bawa, Brown, and Klein have shown that Bayes criterion yields, on average, the maximum expected utility.

The Minimum Variance Hedge Ratio

A typical problem involving estimation risk is that of calculating the minimum variance hedge ratio (MVH). The MVH is the ratio between the futures and the cash positions that minimizes the variance of income, given the agent's cash position. The MVH is an important paradigm in the theory of hedging, and dominates the applied hedging literature.

Reduced to its essentials, the derivation of the MVH is as follows. Consider a decision maker at decision date T whose random terminal income π_{T+1} equals the returns from his cash and futures positions, i.e.,

$$(3.1)\quad \pi_{T+1} = p_{T+1}\, Q - (f_{T+1} - f_T)\, F$$

where p_{T+1} is the random cash price at date T+1, Q is the amount of product sold at date T+1, f_{T+1} is the random futures price prevailing at date T+1 for delivery at some date $T+t \geq T+1$, f_T is the current futures price for delivery at date $T+t$, and F is the amount sold in the futures market at date T and purchased at date T+1. The decision problem consists of selecting the hedge F that minimizes the variance of terminal income, given the cash position Q:

$$(3.2)\quad \min_F \mathrm{Var}_T(\pi_{T+1}) = \min_F[Q^2\, \mathrm{Var}_T(p_{T+1}) - 2\, Q\, F\, \mathrm{Cov}_T(p_{T+1}, f_{T+1}) + F^2\, \mathrm{Var}_T(f_{T+1})]$$

The subscripts in the variance and covariance operators denote that they are conditional on the information at date T. The first order condition (FOC) corresponding to (3.2) is

$$(3.3) \quad \frac{\partial \mathrm{Var}_T(\pi_{T+1})}{\partial F} = -2\,Q\,\mathrm{Cov}_T(p_{T+1}, f_{T+1}) + 2\,F\,\mathrm{Var}_T(f_{T+1}) = 0$$

which can be solved for the variance-minimizing hedging position[1]

$$(3.4) \quad F_{PPI} = \frac{\mathrm{Cov}_T(p_{T+1}, f_{T+1})}{\mathrm{Var}_T(f_{T+1})}\,Q$$

The ratio $\mathrm{Cov}_T(p_{T+1}, f_{T+1})/\mathrm{Var}_T(f_{T+1})$ is the MVH. It has been shown that the MVH is the optimal hedge ratio if the current futures price f_T is an unbiased predictor of the posterior futures price f_{T+1}, regardless of the decision maker's degree of absolute risk aversion (Benninga, Eldor, and Zilcha). In addition, the MVH is the optimal hedge ratio for extremely risk-averse decision makers (Kahl). Because of these attributes, and also because of the *apparent* easiness of the empirical calculation, the estimation of the MVH has been the focus of numerous applied studies. Implicitly or explicitly, all such studies use the PCE approach, that is, they estimate F_{PPI} by means of F_{PCE}:

$$(3.5) \quad F_{PCE} = \frac{\hat{\sigma}_{pf}}{\hat{\sigma}_{ff}}\,Q$$

where $\hat{\sigma}_{pf}$ and $\hat{\sigma}_{ff}$ stand for the sample estimates of $\mathrm{Cov}_T(p_{T+1}, f_{T+1})$ and $\mathrm{Var}_T(f_{T+1})$, respectively.

Different methods have been applied to obtain the MVH estimate $\hat{\sigma}_{pf}/\hat{\sigma}_{ff}$. A popular technique consists of regressing cash on futures prices employing historical data, and using the futures price regression coefficient as the estimated MVH. Examples of studies applying this approach (or some variation of it) are Ederington; Hayenga and DiPietre; Stewart Brown; Witt, Schroeder, and Hayenga; Wilson; and Myers and Thompson. Other authors advocate the use of GARCH models (Baillie and Meyers) and conditional forecasts (Peck) to estimate the MVH. For illustrative purposes, in Table 1 we reproduce MVH estimates for a set of agricultural

---

[1]The second order condition for a minimum is always satisfied because $\mathrm{Var}_T(f_{T+1}) > 0$.

Table 1. Estimates of the minimum variance hedge ratio

Study	Corn	Soybeans	Wheat	Beef	Cows	Hogs
Myers and Thompson	0.85-1.04	0.87-1.12	0.61-1.10			
Ederington	0.76-1.02		0.78-0.92			
Baillie and Myers[a]	0.09-1.53	0.22-1.17		(-0.06)-0.40		
Mathews and Fackler					0.55-1.04	0.86-0.95
Mathews and Holthausen					0.84	0.93

[a]The estimates by Baillie and Meyers correspond to the ratio of the *values* of the futures and cash positions.

commodities. The different MVH estimates correspond to alternative methods and/or data sets.

Even though the collection in Table 1 is by no means exhaustive, the large range of the estimated MVHs for each particular commodity is striking. The values reported in Table 1 suggest that estimation risk is important in the case of the MVH. But if estimation risk exists, then the sample estimate of the MVH need not lead to the optimal decision. Moreover, the stated properties of the MVH (i.e., optimality under unbiased futures prices or extreme risk aversion) hold under PPI conditions, but they do not necessarily hold in the presence of estimation risk.

Using Bayes criterion to derive the MVH in the presence of estimation risk

In the PPI case, utility is given by $U[R(\underline{d}, \underline{y})] = - Var_T(\pi_{T+1})$, with $R(\underline{d}, \underline{y}) = \pi_{T+1}$, $\underline{d} = F$, and $\underline{y} = (p_{T+1}, f_{T+1})'$. To be consistent with the variance-minimizing principle that underlies the MVH, we will assume that the utility under estimation risk is $U[R(\underline{d}, \underline{y})] = - \sigma_{\pi\pi P}$, with $\sigma_{\pi\pi P}$ being the *predictive* variance of income conditional on the information at date T. Then, applying a derivation analogous to that of expressions (3.1) through (3.4), the futures position that minimizes the predictive variance of income is

$$(3.6) \quad F_{BAY} = \frac{\sigma_{pfP}}{\sigma_{ffP}} Q$$

where σ_{pfP} is the predictive covariance between futures and cash prices, and σ_{ffT} is the predictive variance of futures prices. We use the subscript BAY to identify the hedging solution obtained by application of Bayes criterion (i.e., by using the predictive pdf). The hedge F_{BAY} is the counterpart of F_{PPI} in the presence of estimation risk.

The exact form of the ratio $\sigma_{pfP}/\sigma_{ffP}$ depends on the functional form of the predictive pdf. In this paper, we will analyze two particular scenarios based on the multivariate normal distribution. For our purposes, the most important result regarding the multivariate normal distribution is the relationship between prior, posterior, and predictive pdfs. These are available in Aitchison and Dunsmore, and summarized below.

Assume that the as yet unobserved $(k \times 1)$ random vector $\underline{y}$ is k-variate normal and independently distributed with unknown mean vector $\underline{\mu}$ and unknown covariance matrix $\underline{\Sigma}$. Let the sample data consist of the $(T \times k)$ matrix $\underline{X}$ of T past observations, with sample estimates of the mean vector and the covariance matrix given by

$$\hat{\underline{\mu}} = \underline{1}' \, \underline{X} / T$$

$$\hat{\underline{\Sigma}} = (\underline{X} - \underline{1}\,\hat{\underline{\mu}}')' \, (\underline{X} - \underline{1}\,\hat{\underline{\mu}}')/(T - 1)$$

where $\underline{1}$ is a $(T \times 1)$ vector of ones. Let the prior information regarding the unknown mean vector $\underline{\mu}$ and the unknown covariance matrix $\underline{\Sigma}$ be represented by the k-variate normal pdf (3.7) and the k-variate Wishart pdf (3.8), respectively.

$$(3.7) \quad p(\underline{\mu} | \mathrm{I}_T) = f_N^{(k)}(\underline{\mu} | \underline{\mu}_0, \underline{\Sigma}/\tau)$$

$$(3.8) \quad p(\underline{\Sigma}^{-1} | \mathrm{I}_T) = f_W^{(k)}(\underline{\Sigma}^{-1} | \underline{\Sigma}_0^{-1}/\nu, \nu)$$

Then, the predictive pdf of $\underline{y}$ is k-variate Student-t with $(\nu + T)$ degrees of freedom as follows

$$(3.9) \quad p(\underline{y} | \underline{X}, \mathrm{I}_T) = f_S^{(k)}(\underline{y} | \underline{\mu}_P, \underline{\Sigma}_P, \nu + T)$$

where: $\underline{\mu}_P = \omega_\tau \, \underline{\mu}_0 + (1 - \omega_\tau)\, \hat{\underline{\mu}}$

$$\underline{\Sigma}_P = \left(\frac{\nu + T}{\nu + T - 2}\right)\left(1 + \frac{1}{\tau + T}\right) [\omega_\nu \, \underline{\Sigma}_0 + (1 - \omega_\nu)\, \hat{\underline{\Sigma}} + \omega_\tau (1 - \omega_\nu)\, (\hat{\underline{\mu}} - \underline{\mu}_0)\, (\hat{\underline{\mu}} - \underline{\mu}_0)']$$

$$\omega_\tau = \tau/(\tau + T)$$

$$\omega_\nu = \nu/(\nu + T)$$

The prior mean vector μ_0 represents the decision maker's beliefs regarding the unknown mean vector μ. The degree of confidence in this prior mean vector is measured by the positive scalar τ. The situation of complete knowledge about μ is represented by the limit of τ as it approaches infinity, and zero prior knowledge is represented by the limit of τ approaching zero. Complete knowledge regarding μ means that the agent is completely certain that $\mu = \mu_0$, in which case $\mu_P = \mu_0$. It can be seen that the predictive mean μ_P is a weighted average of the prior and the sample means. As expected, the weight given to the prior (ω_τ) decreases as the amount of sample information (T) increases relative to the amount of prior information (τ), thus yielding a predictive mean μ_P closer to the sample mean $\hat{\mu}$.

The Wishart distribution (3.8) used for the prior of the covariance matrix is, loosely speaking, a multivariate generalization of the Chi-square distribution; the Wishart distribution applies to a covariance matrix rather than to a scalar variance. The prior covariance matrix Σ_0 denotes the decision maker's prior beliefs about the unknown covariance matrix Σ. The positive scalar ν measures the decision maker's degree of confidence in Σ_0; the degree of confidence in Σ_0 increases with ν. Complete certainty (lack of knowledge) about Σ is modeled by letting ν approach infinity (zero). The predictive covariance matrix Σ_P is a weighted average of the prior and the sample covariance matrices plus a term involving the difference between the sample and the prior mean vectors. The weight assigned to the prior covariance (ω_ν) approaches unity as the decision maker becomes more confident about Σ_0 (i.e., as ν increases); therefore, the predictive covariance matrix Σ_P approaches the prior covariance matrix when the agent is very confident about Σ_0. In the case of perfect knowledge about Σ, ν tends to infinity and ω_ν equals unity; hence, $\Sigma_P = \Sigma_0 = \Sigma$. In this limiting instance, the predictive pdf of y becomes k-variate normal because the k-variate Student-t pdf approaches the former pdf as the degrees of freedom tend to infinity.

Note that Σ_P is also affected by the difference between the prior and the sample mean vectors. The reason for this secondary effect is that the discrepancy between the prior and the sample mean vectors adds another source of variability. The importance of this additional variability is greater when the relative confidence in the prior mean vector and the sample covariance matrix are both large [i.e., ω_τ and $(1 - \omega_v)$ close to one, respectively]. The greater the confidence in the prior mean vector, the greater its relevance; similarly, the greater the relative quality of the sample covariance matrix, the greater the accuracy of the sample mean vector.

We will now apply the stated result concerning the predictive pdf $p(\underline{y}|\underline{X}, I_T)$ of a k-variate normal joint pdf $p(\underline{y}|\underline{\theta})$ to obtain the solutions for F_{BAY} under two alternative scenarios, namely,

Case (i). Cash and futures prices are independently bivariate normally distributed with unknown mean vector $\underline{\mu}_p$ and unknown covariance matrix $\underline{\Sigma}_p$:

$$(3.10)\quad p(\underline{y}|\underline{\theta}) = f_N^{(2)}(\underline{p}|\underline{\mu}_p, \underline{\Sigma}_p)$$

where: $\underline{p} = (p_{T+1}, f_{T+1})'$

$$\underline{\mu}_p = (\mu_p, \mu_f)'$$

$$\underline{\Sigma}_p = \begin{pmatrix} \sigma_{pp} & \sigma_{pf} \\ \sigma_{pf} & \sigma_{ff} \end{pmatrix}$$

Case (ii). The natural logarithms of cash and futures prices are independently bivariate normally distributed with unknown mean vector $\underline{\mu}_l$ and unknown covariance matrix $\underline{\Sigma}_l$:

$$(3.11)\quad p(\underline{y}|\underline{\theta}) = f_N^{(2)}(\underline{l}|\underline{\mu}_l, \underline{\Sigma}_l)$$

where: $\underline{l} = (l_p, l_f)'$

$$l_p = \ln(p_{T+1})$$

$$l_f = \ln(f_{T+1})$$

$$\underline{\mu}_l = (\mu_{l_p}, \mu_{l_f})'$$

$$\underline{\Sigma}_l = \begin{pmatrix} \sigma_{l_p l_p} & \sigma_{l_p l_f} \\ \sigma_{l_p l_f} & \sigma_{l_f l_f} \end{pmatrix}$$

Cases (i) and (ii) were chosen because they have been frequently used in modeling the MVH, and also because they yield closed-form solutions for the ratio $\sigma_{pfP}/\sigma_{ffP}$. It can be argued that it is more realistic to hypothesize that the net returns -not the price levels- are normally distributed, but it is straightforward to show that the solution under such assumption is essentially the same as that for Case (i). Similarly, if it is assumed that the logarithms of returns rather than the logarithms of prices are normally distributed, the solution is basically the same as that for Case (ii). Hence, to save space we will concentrate our attention on Cases (i) and (ii) only.

Case (i). When prices are bivariate normally distributed as stated in (3.10), the ratio of the predictive covariance to the predictive variance is

$$(3.12)\quad \frac{\sigma_{pfP}}{\sigma_{ffP}} = \frac{\omega_\nu\, \sigma_{pf0} + (1-\omega_\nu)\, \hat{\sigma}_{pf} + \omega_\tau\, (1-\omega_\nu)\, (\hat{\mu}_p - \mu_{p0})\, (\hat{\mu}_f - \mu_{f0})}{\omega_\nu\, \sigma_{ff0} + (1-\omega_\nu)\, \hat{\sigma}_{ff} + \omega_\tau\, (1-\omega_\nu)\, (\hat{\mu}_f - \mu_{f0})^2}$$

where the subscript 0 identifies the priors and the hat denotes the sample estimates.[1]

As it was discussed before, both ω_τ and ω_ν approach one when the decision maker is much more confident about his prior beliefs than about the sample estimates. In such circumstances, the ratio $\sigma_{pfP}/\sigma_{ffP}$ in (3.12) simplifies to $\sigma_{pf0}/\sigma_{ff0}$, which is the same as the PPI minimum variance hedge ratio (3.4). The opposite situation arises when the sample size is so large compared to the strength of the prior beliefs that the relative weights ω_τ and ω_ν both approach zero. In this instance, the ratio $\sigma_{pfP}/\sigma_{ffP}$ in (3.12) collapses to $\hat{\sigma}_{pf}/\hat{\sigma}_{ff}$ and is equal to the ratio found using the PCE approach (3.5). Case (i) may convey the incorrect idea that it is always true that $F_{BAY} = F_{PCE}$

---

[1]Expression (3.12) was obtained by using the following result concerning the multivariate Student-t pdf (Zellner, p. 388): If $\underline{w} = \underline{L}\,\underline{y}$, where $\underline{L}$ is a $(k_1 \times k)$ matrix of rank $k_1 \leq k$ and $\underline{y}$ is a $(k \times 1)$ random vector distributed as the k-variate Student-t shown in (3.9), then $\underline{w}$ is distributed as k_1-variate Student-t with pdf $f_S^{(k_1)}(\underline{w} \mid \underline{L}\, \mu_P, \underline{L}\, \Sigma_P\, \underline{L}', \nu + T)$.

when prior information is negligible (i.e., when both ω_τ and ω_v tend to zero). Case (ii) below shows that the equivalence between the Bayesian and the PCE solutions in the absence of prior information does not hold in general.

It is interesting to note that even if the decision maker's prior beliefs about the variance and covariance are such that ω_v tends to zero, the ratios from the Bayesian criterion (3.12) and the PCE (3.5) may still be different. In general, this will happen if the individual attaches some weight to his prior beliefs about the mean vector (i.e., $\omega_\tau > 0$) and the sample futures mean is not the same as the prior futures mean (i.e., $\hat{\mu}_f \neq \mu_{f0}$). In such event, the ratio $\sigma_{pfP}/\sigma_{ffP}$ in (3.12) will be

$$(3.13)\quad \frac{\sigma_{pfP}}{\sigma_{ffP}} = \frac{\hat{\sigma}_{pf} + \omega_\tau (\hat{\mu}_p - \mu_{p0})(\hat{\mu}_f - \mu_{f0})}{\hat{\sigma}_{ff} + \omega_\tau (\hat{\mu}_f - \mu_{f0})^2}$$

which is different from $\hat{\sigma}_{pf}/\hat{\sigma}_{ff}$. The reason for this result is that the discrepancy between prior and sample means is an additional source of variability.

Case (ii): Under the assumptions stated in (3.11), the ratio of the predictive covariance to the predictive variance can be shown to equal[1]

$$(3.14)\quad \frac{\sigma_{pfP}}{\sigma_{ffP}} = \frac{\exp(\phi_p)\exp(\phi_{pp}/2)\,[\exp(\phi_{pf}) - 1]}{\exp(\phi_f)\exp(\phi_{ff}/2)\,[\exp(\phi_{ff}) - 1]}$$

[1]To derive expression (3.14), three statistical results were employed. First, the k-variate Student-t pdf in (3.9) was approximated by the k-variate Normal pdf $f_N^{(k)}(\underline{y}|\underline{\mu}_P, (v+T)/(v+T-2)\,\Sigma_P)$ (Johnson and Kotz, p. 101). This approximation was done because, strictly speaking, the predictive covariance and variance of prices does not exist if the logarithms of prices follow a bivariate Student-t distribution.

Second, that if $\underline{v} = \underline{L}\,\underline{z}$, where $\underline{L}$ is a $(k_1 \times k)$ matrix of rank $k_1 \leq k$ and $\underline{z}$ is a $(k \times 1)$ random vector distributed as k-variate Normal with mean vector $\underline{\mu}_z$ and covariance matrix Σ_z, then $\underline{v}$ is distributed as k_1-variate Normal with mean vector $\underline{L}\,\underline{\mu}_z$ and covariance matrix $\underline{L}\,\Sigma_z\,\underline{L}'$ (Zellner, p. 382).

Third, that if $\underline{u} \equiv (u_1, \ldots, u_k)'$ is a $(k \times 1)$ positive random vector such that $\underline{e} = \ln(\underline{u})$ is k-variate Normally distributed with mean vector $\underline{\mu}_e \equiv (\mu_{e_1}, \ldots, \mu_{e_k})'$ and covariance matrix $\Sigma_e \equiv (\sigma_{e_ie_j})$, then (Press, p. 149)

$$\mathrm{cov}(u_i, u_j) = \exp[\mu_{e_i} + \mu_{e_j} + (\sigma_{e_ie_i} + \sigma_{e_je_j})/2 + \sigma_{e_ie_j}] - \exp[\mu_{e_i} + \mu_{e_j} + (\sigma_{e_ie_i} + \sigma_{e_je_j})/2];\ i, j = 1, \ldots, k$$

where: $\phi_i = \omega_\tau \mu_{l_{i0}} + (1 - \omega_\tau) \hat{\mu}_{l_i}$; i, j = p, f

$$\phi_{ij} = \left(\frac{\nu + T}{\nu + T - 2}\right)^2 \left(1 + \frac{1}{\tau + T}\right)$$

$$[\omega_\nu \sigma_{l_i l_{j0}} + (1 - \omega_\nu) \hat{\sigma}_{l_i l_j} + \omega_\tau (1 - \omega_\nu) (\hat{\mu}_{l_i} - \mu_{l_{i0}}) (\hat{\mu}_{l_j} - \mu_{l_{j0}})];\ i, j = p, f$$

and exp(·) denotes the base of the natural logarithms raised to the power (·).

If prices are distributed as shown in (3.11), the PCE estimator of the MVH can take several forms depending on the technique used for estimation. For example, employing the nonparametric unbiased estimators of the covariance and variance yields

$$(3.15)\quad \frac{\hat{\sigma}_{pf}}{\hat{\sigma}_{ff}} = \frac{\sum_{t=1}^{T} (p_t - \hat{\mu}_p)(f_t - \hat{\mu}_f)}{\sum_{t=1}^{T} (f_t - \hat{\mu}_f)^2}$$

whereas the maximum likelihood estimator gives

$$(3.16)\quad \frac{\hat{\sigma}_{pf}}{\hat{\sigma}_{ff}} = \frac{\exp(\hat{\mu}_{l_p}) \exp[(1 - 1/T)\, \hat{\sigma}_{l_p l_p}/2]\, \{\exp[(1 - 1/T)\, \hat{\sigma}_{l_p l_f}] - 1\}}{\exp(\hat{\mu}_{l_f}) \exp[(1 - 1/T)\, \hat{\sigma}_{l_f l_f}/2]\, \{\exp[(1 - 1/T)\, \hat{\sigma}_{l_f l_f}] - 1\}}$$

Still other PCE estimator is given by the uniformly minimum variance unbiased estimator (Shimizu), which is a complex function of the sample estimates $\hat{\mu}_{l_p}$, $\hat{\mu}_{l_f}$, $\hat{\sigma}_{l_p l_p}$, $\hat{\sigma}_{l_p l_f}$, and $\hat{\sigma}_{l_f l_f}$. In general, none of these estimators is identical to the Bayesian ratio (3.14), even in the limiting case in which the decision maker has no prior information (i.e., $\omega_\nu = \omega_\tau = 0$). This result shows that it is not always true that $F_{BAY} = F_{PCE}$ when there is no prior information about the parameters. Therefore, the equivalence between the Bayesian and the PCE solutions in the absence of any prior information found in Case (i) depicts a very special situation and cannot be generalized.

It is worth noting, however, that the PPI MVH in (3.4) is nested in the Bayesian solution in both Cases (i) and (ii) [i.e., in expressions (3.12) and (3.14)]. This observation can be easily proven by increasing both ν and τ to infinity, which represents the case of complete confidence about the priors. This result always holds because the posterior pdf under PPI gives probability one to the parameters being equal to the priors; therefore, integrating over the parameter space does not make any difference and expression (2.2) collapses to (2.1).

In Tables 2, 3 and 4, we report the results of simulations regarding the PPI ratio in (3.4), the PCE ratio in (3.5), and the Bayesian ratio for Case (i) (3.12).[1,2] The parameter values in the simulations are arbitrary but realistic; the means and variances employed reflect a coefficient of variation of approximately 16 percent, and the differences between the prior and the sample price means are either zero or one standard deviation.

Table 2 contains the results for the case in which the prior and the sample futures means are identical. The values in this table are unchanged regardless of the difference between the prior and sample cash means. The reason for this result is that the additional variability caused by the difference between the prior and sample means vanishes if the prior and sample futures means are identical.[3] The top and bottom four rows of Table 2 show that when the prior and sample means, variances, and covariances are all equal then all three methods provide similar results. The fifth row shows how the results change if the sample covariance is one half as large as the prior, and the individual has three times as much confidence in the sample as in the prior. In the PPI and PCE scenarios we ignore these confidence weights. The PPI individual ignores the sample information and places a full hedge. The PCE individual places full weight on the sample and hedges only half as much as under PPI. If the weights on the prior and sample information change then the Bayesian position changes but the other two do not.

Table 3 summarizes the results for the case in which the prior exceeds the sample futures mean by one standard deviation and the sample cash mean is identical to the prior cash mean. The top line of this table presents an example in support of the most surprising results of these

---

[1]The simulations consisted of solving for the MVHs in expressions (3.4), (3.5), and (3.12) under particular values of the independent variables.

[2]The results for Case (ii) are similar in terms of the differences among the solutions for the three alternative methods.

[3]The difference between the prior and sample cash means affects the predictive cash variance, but this is not involved in the MVH.

Table 2. Minimum variance hedge ratios for prior and sample futures means (μ_{f0} and $\hat{\mu}_f$) equal to 10, and prior and sample futures variances (σ_{ff0} and $\hat{\sigma}_{ff}$) equal to 2.56

Covariance		Relative Strength of Prior		Min. Variance Hedge Ratio Corresponding to		
Prior	Sample	Mean	Variance	Bayes Crit.	PCE	PPI
2.56	2.56	0.25	0.25	1.00	1.0	1.0
2.56	2.56	0.25	0.75	1.00	1.0	1.0
2.56	2.56	0.75	0.25	1.00	1.0	1.0
2.56	2.56	0.75	0.75	1.00	1.0	1.0
2.56	1.28	0.25	0.25	0.62	0.5	1.0
2.56	1.28	0.25	0.75	0.88	0.5	1.0
2.56	1.28	0.75	0.25	0.62	0.5	1.0
2.56	1.28	0.75	0.75	0.88	0.5	1.0
1.28	2.56	0.25	0.25	0.88	1.0	0.5
1.28	2.56	0.25	0.75	0.62	1.0	0.5
1.28	2.56	0.75	0.25	0.88	1.0	0.5
1.28	2.56	0.75	0.75	0.62	1.0	0.5
1.28	1.28	0.25	0.25	0.50	0.5	0.5
1.28	1.28	0.25	0.75	0.50	0.5	0.5
1.28	1.28	0.75	0.25	0.50	0.5	0.5
1.28	1.28	0.75	0.75	0.50	0.5	0.5

Table 3. Minimum variance hedge ratios for prior futures mean (μ_{f0}) equal to 11.5, sample futures mean ($\hat{\mu}_f$) equal to 10, prior and sample cash means (μ_{p0} and $\hat{\mu}_p$) equal to 10, and prior and sample futures variances (σ_{ff0} and $\hat{\sigma}_{ff}$) equal to 2.56

Covariance		Relative Strength of Prior		Min. Variance Hedge Ratio Corresponding to		
Prior	Sample	Mean	Variance	Bayes Crit.	PCE	PPI
2.56	2.56	0.25	0.25	0.86	1.0	1.0
2.56	2.56	0.25	0.75	0.95	1.0	1.0
2.56	2.56	0.75	0.25	0.67	1.0	1.0
2.56	2.56	0.75	0.75	0.86	1.0	1.0
2.56	1.28	0.25	0.25	0.54	0.5	1.0
2.56	1.28	0.25	0.75	0.83	0.5	1.0
2.56	1.28	0.75	0.25	0.42	0.5	1.0
2.56	1.28	0.75	0.75	0.75	0.5	1.0
1.28	2.56	0.25	0.25	0.75	1.0	0.5
1.28	2.56	0.25	0.75	0.60	1.0	0.5
1.28	2.56	0.75	0.25	0.59	1.0	0.5
1.28	2.56	0.75	0.75	0.54	1.0	0.5
1.28	1.28	0.25	0.25	0.43	0.5	0.5
1.28	1.28	0.25	0.75	0.47	0.5	0.5
1.28	1.28	0.75	0.25	0.34	0.5	0.5
1.28	1.28	0.75	0.75	0.43	0.5	0.5

Table 4. Minimum variance hedge ratios for prior futures mean (μ_{f0}) equal to 11.5, sample futures mean ($\hat{\mu}_f$) equal to 10, prior cash mean (μ_{p0}) equal to 11.5, sample cash mean ($\hat{\mu}_p$) equal to 10, and prior and sample futures variances (σ_{ff0} and $\hat{\sigma}_{ff}$) equal to 2.56

Covariance		Relative Strength of Prior		Min. Variance Hedge Ratio Corresponding to		
Prior	Sample	Mean	Variance	Bayes Crit.	PCE	PPI
2.56	2.56	0.25	0.25	1.00	1.0	1.0
2.56	2.56	0.25	0.75	1.00	1.0	1.0
2.56	2.56	0.75	0.25	1.00	1.0	1.0
2.56	2.56	0.75	0.75	1.00	1.0	1.0
2.56	1.28	0.25	0.25	0.68	0.5	1.0
2.56	1.28	0.25	0.75	0.88	0.5	1.0
2.56	1.28	0.75	0.25	0.75	0.5	1.0
2.56	1.28	0.75	0.75	0.89	0.5	1.0
1.28	2.56	0.25	0.25	0.89	1.0	0.5
1.28	2.56	0.25	0.75	0.64	1.0	0.5
1.28	2.56	0.75	0.25	0.92	1.0	0.5
1.28	2.56	0.75	0.75	0.68	1.0	0.5
1.28	1.28	0.25	0.25	0.57	0.5	0.5
1.28	1.28	0.25	0.75	0.53	0.5	0.5
1.28	1.28	0.75	0.25	0.66	0.5	0.5
1.28	1.28	0.75	0.75	0.57	0.5	0.5

simulations. This row compares the three optimal hedges when the prior and sample variances and covariances are similar, but where the prior and sample futures means are different. Intuitively, one would not have expected the hedge ratio to change under these circumstances, and yet the Bayes criterion ratio is different from the PCE or PPI. The intuition here is that the difference between the sample and prior futures means is another source of variability which increases the predictive futures variance. This result is important because it represents a realistic scenario where the decision maker is naive about the variance and covariance terms but has prior about the futures mean (e.g., the decision maker might have insider information about the futures mean).

Table 4 reports the results from the simulations assuming that each prior mean is one standard deviation larger than the respective sample mean. The Bayes MVHs in this table are larger than those in Table 3 because the numerator in expression (3.12) increases with the difference between the cash means. The Bayes MVHs in the top four rows are the same as the PCE and PPI MVHs because of a very special scenario: the differences between the prior and sample means are identical, and the PPI and PCE MVHs are identical and equal to one.

Summary and Conclusions

Decision making models generally assume perfect parameter information (PPI), i.e., that the true parameters characterizing the joint probability density function (pdf) of the relevant random variables are known. In most applications, however, the true parameters are not known, that is, there is estimation risk.

Bayes decision criterion provides a way of dealing with estimation risk in a manner consistent with expected utility maximization. This approach assigns a pdf for the unknown true parameters based on sample and prior information, and uses this pdf to integrate the original objective function over the parameter space; optimization is then performed over the resulting integral. Bayes criterion has been used in statistics and finance, but has been neglected in agricultural economics. The standard technique employed in agricultural economics is the parameter certainty equivalent (PCE). The PCE consists of substituting the sample estimates for the unknown true parameters in the PPI decision rule. The PCE approach is easier to implement than Bayes criterion, but not consistent with expected utility maximization. Moreover, the PCE decision rules are generally different from the Bayesian decision rules.

In this paper, estimation risk is discussed in general terms and Bayes criterion is presented and its properties are compared with the PCE. Bayes criterion is then applied to obtain the solution of the minimum variance hedge ratio (MVH).

The MVH is the ratio of futures to cash positions that minimizes the variance of income, given a particular cash position. Empirical estimation of the MVH has been the subject of many studies employing the PCE approach, and is a clear example of a problem involving estimation risk. Bayesian MVHs are derived for two scenarios, one assuming that cash and futures prices are bivariate normally distributed, and the other assuming that the logarithms of cash and futures prices are bivariate normally distributed. In both scenarios, the Bayesian solutions to the MVH are functions of weighted averages of sample and prior parameters, with weighing factors that depend on the relative qualities of the sample and prior information.

Simulation results reveal that discrepancies between prior and sample parameters may lead to substantial differences among the PPI, PCE, and Bayesian solutions for the variance-minimizing hedge ratio. Such discrepancies also highlight the superiority of Bayes criterion over the PCE or the PPI approaches in the sense that neither of the latter two methods yield decision rules which combine sample and prior (or nonsample) information.

The PCE decision rule only takes into consideration sample information and neglects any prior information, and the opposite is true of the PPI decision rule. Only the Bayesian decision rule blends both types of information in a manner consistent with expected utility theory. This is a highly desirable property of the Bayesian approach because it is common for the decision maker to have available prior information but not be completely certain about it. In the MVH example, prior information could well be represented by the agent's knowledge about the market, by insider information, or by opinions from market experts. Neither of these sources of information matters when using the PCE technique, e.g., the PCE MVH remains unmodified even if the decision maker receives insider information from a reliable source. This characteristic is clearly unacceptable, yet it is intrinsic to any PCE decision rule.

References

Aitchison, J., and I. R. Dunsmore. *Statistical Prediction Analysis*. New York: Cambridge University Press, 1975.

Alexander, Gordon J., and Bruce G. Resnick. "More on Estimation Risk and Simple Rules for Optimal Portfolio Selection." *J. Finan.* 40(1985):125-33.

Baillie, Richard T., and Robert J. Myers. "Bivariate GARCH Estimation of the Optimal Commodity Futures Hedge." *J. Applied Econometrics* 6(1991):109-24.

Barry, Christopher B., and Stephen J. Brown. "Differential Information and Security Market Equilibrium." *J. Finan. and Quant. Anal.* 20(1985):407-22.

Bawa, Vijay S. "Portfolio Choice and Capital Market Equilibrium with Unknown Distributions." *Estimation Risk and Optimal Portfolio Choice*, eds. Bawa, V., S. Brown, and R. Klein, Chapter 7. Amsterdam: North-Holland Publishing Co., 1979.

Bawa, V., S. Brown, and R. Klein. "Estimation Risk: An Introduction." *Estimation Risk and Optimal Portfolio Choice*, eds. Bawa, V., S. Brown, and R. Klein, Chapter 1. Amsterdam: North-Holland Publishing Co., 1979.

Benninga, S., R. Eldor, and I. Zilcha. "The Optimal Hedge Ratio in Unbiased Futures Markets." *J. Futures Mkts.* 4(1984):155-59.

Berger, James O. *Statistical Decision Theory and Bayesian Analysis, 2nd edition*. New York: Springer-Verlag, 1985.

Boyle, Phelim P., and A. L. Ananthanarayanan. "The Impact of Variance Estimation in Option Valuation Models." *J. Finan. Econ.* 5(1977):375-87.

Brown, Stephen J. "Estimation Risk and Optimal Portfolio Choice: The Sharpe Index Model." *Estimation Risk and Optimal Portfolio Choice*, eds. Bawa, V., S. Brown, and R. Klein, Chapter 9. Amsterdam: North-Holland Publishing Co., 1979.

Brown, Stewart L. "A Reformulation of the Portfolio Model of Hedging." *Amer. J. Agr. Econ.* 67(1985):508-12.

Chalfant, J. A., R. N. Collender, and S. Subramanian. "The Mean and Variance of the Mean-Variance Decision Rule." *Amer. J. Agr. Econ.* 72(1990):966-74.

Chen, Son-Nan, and Stephen J. Brown. "Estimation Risk and Simple Rules for Optimal Portfolio Selection." *J. Finan.* 38(1983):1087-93.

Cheung, C. Sherman and Clarence C. Y. Kwan. "A Note on Simple Criteria for Optimal Portfolio Selection." *J. Finan.* 43(1988):241-45.

Coles, Jeffrey L., and Uri Loewenstein. "Equilibrium Pricing and Portfolio Composition in the Presence of Uncertain Parameters." *J. Finan. Econ.* 22(1988):279-303.

Collender, Robert N. "Estimation Risk in Farm Planning under Uncertainty." *Amer. J. Agr. Econ.* 71(1989):996-1002.

Collender, Robert N., and David Zilberman. "Land Allocation under Uncertainty for Alternative Specifications of Return Distributions." *Amer. J. Agr. Econ.* 67(1985):779-86.

DeGroot, Morris H. *Optimal Statistical Decisions.* New York: McGraw-Hill, 1970.

Dixon, Bruce L., and Peter J. Barry. "Portfolio Analysis Considering Estimation Risk and Imperfect Markets." *West. J. Agr. Econ.* 8(1983): 103-11.

Ederington, Louis H. "The Hedging Performance of the New Futures Markets." *J. Finan.* 34(1979):157-70.

Frost, Peter A., and James E. Savarino. "An Empirical Bayes Approach to Efficient Portfolio Selection." *J. Finan. and Quant. Anal.* 21(1986):293-305.

Hayenga, Marvin L., and Dennis D. DiPietre. "Cross-Hedging Wholesale Pork Products Using Live Hog Futures." *Amer. J. Agr. Econ.* 64(1982):747-51.

Johnson, Norman L., and Samuel Kotz. *Continuous Univariate Distributions, Volume 2.* New York: John Wiley and Sons, 1970.

Jorion, Philippe. "International Portfolio Diversification with Estimation Risk." *J. Bus.* 58(1985):259-78.

Jorion, Philippe. "Bayes-Stein Estimation for Portfolio Analysis." *J. Finan. and Quant. Anal.* 21(1986):279-92.

Kahl, Kandice H. "Determination of the Recommended Hedging Ratio." *Amer. J. Agr. Econ.* 65(1983):603-05.

Klein, R. W., and V. S. Bawa. "The Effect of Estimation Risk on Optimal Portfolio Choice." *J. Finan. Econ.* 3(1976):215-31.

Klein, R. W., and V. S. Bawa. "The Effect of Limited Information and Estimation Risk on Optimal Portfolio Diversification." *J. Finan. Econ.* 5(1977):89-111.

Klein, R. W., L. C. Rafsky, D. S. Sibley, and R. D. Willig. "Decisions with Estimation Uncertainty." *Econometrica* 46(1978):1363-87.

Mathews, Kenneth H., and Paul L. Fackler. "Hedge Ratios in the Presence of Heteroskedasticity." Presented in the Annual Meeting of the American Agricultural Economics Association, Manhattan, Kansas, August 4-7, 1991.

Mathews, Kenneth H., and Duncan M. Holthausen. "A Simple Multiperiod Minimum Risk Hedge Model." *Amer. J. Agr. Econ.* 73(1991):1020-26.

Myers, Robert J., and Stanley R. Thompson. "Generalized Optimal Hedge Ratio Estimation." *Amer. J. Agr. Econ.* 71(1989):858-68.

Peck, Anne E. "Hedging and Income Stability: Concepts, Implications, and An Example." *Amer. J. Agr. Econ.* 57(1975):410-19.

Pope, Rulon D., and Rod F. Ziemer. "Stochastic Efficiency, Normality, and Sampling Errors in Agricultural Risk Analysis." *Amer. J. Agr. Econ.* 66(1984):31-40.

Press, James S. *Applied Multivariate Analysis: Using Bayesian and Frequentist Methods of Inference, 2nd. edition.* Malabar, Florida: Robert E. Krieger Publishing Company, 1982.

Raiffa, Howard, and Robert Schlaifer. *Applied Statistical Decision Theory.* Boston: Harvard University Press, 1961.

Shimizu, Kunio. "Point Estimation." *Lognormal Distributions--Theory and Applications*, eds. Crow, E., and K. Shimizu, Chapter 2. New York: Marcel Dekker, 1988.

Wilson, William W. "Price Discovery and Hedging in the Sunflower Market." *J. Futures Mkts.* 9(1989):377-91.

Witt, H., T. Schroeder, and M. Hayenga. "Comparison of Analytical Approaches for Estimating Hedge Ratios for Agricultural Commodities." *J. Futures Mkts.* 7(1987):135-46.

Zellner, Arnold. *An Introduction to Bayesian Inference in Econometrics.* New York: John Wiley and Sons, Inc., 1971.