

Robustness of Non-Parametric Measurement of Efficiency and Risk Aversion

Daniel M. Settlage¹
Paul V. Preckel²

May 15, 2002

**Selected Paper to be Presented at the Annual Meetings of the
American Agricultural Economics Association
Long Beach, California
July 28 –July 31**

Topic Code: 16 (Research Methods/Econometrics/Statistics)

Short Abstract

This paper examines the performance of a risk-adjusted non-parametric approach to measuring efficiency and risk aversion. Prior work is extended to the case where agent behavior is motivated by expected utility maximization. Results indicate the approach significantly outperforms traditional efficiency measurement methods when applied to risk averse agents.

¹ Department of Agricultural Economics, Purdue University, West Lafayette, Indiana 47905-1145, Voice (765) 494-4257, Fax (765) 494-9176, e-mail settlage@purdue.edu.

² Department of Agricultural Economics, Purdue University, West Lafayette, Indiana 47905-1145, Voice (765) 494-4191, Fax (765) 494-9176, e-mail preckel@purdue.edu.

Robustness of Non-Parametric Measurement of Efficiency and Risk Aversion

The use of frontier functions in the economics literature is pervasive. Frontier functions are employed to measure a wide variety of phenomena, but they are most frequently employed to measure firm or industry efficiency levels (Førsund et al., 1980). Most frontier estimation methods fall into one of two categories: mathematical programming based methods, commonly referred to as data envelopment analysis (DEA), and statistically based techniques such as corrected ordinary least squares or maximum likelihood for estimation of stochastic frontier functions. Although the estimation of frontiers differs based on the method chosen, almost all methods share one common feature, they do not account for the presence of risk averse behavior on the part of the firm (Coelli, 1995).

In a paper by Preckel, Ahmed, and Ehui, (2000) a non-parametric method for the measurement of firm efficiency and risk aversion is developed. This method is a modification of DEA and allows the user to simultaneously determine the firm's efficiency score and level of risk aversion. This study extends the initial method set up by Preckel, Ahmed, and Ehui by extending the Monte Carlo framework to handle a wider variety of agent behavior. The accuracy of both the efficiency scores and risk aversion measurements will be computed and analyzed. The literature on frontier methods, with a few notable exceptions, has been sparse with regard to Monte Carlo comparisons. Most existing studies focus on model performance in the face of differing transformation technologies assuming firms are expected profit maximizers (e.g., Gong and Sickles 1989, 1992 and Settlege 1999). Although producer risk aversion almost certainly plays a role in production and investment decisions, few studies examine the effect that risk aversion has on the ability to correctly measure efficiency.

There also exists a substantial literature on the elicitation of utility functions and risk aversion levels. There are two methods commonly used to elicit utility functions: interview methods and data driven approaches. Hardaker, Huirne, and Anderson (1997) found that most of the interview-based methods encounter problems in the estimation of risk aversion levels. Other methods rely on comparing the actual production results to the production plan that is predicted by modeling the producer. Roughly speaking, the risk aversion parameter is varied until the predicted plan matches the observed plan as closely as possible. These data driven elicitation methods have also been found somewhat lacking. The present method is also a data driven approach. However, it does not require the direct construction of a model of the production possibilities of the producer. Rather, observations of the behavior of other producers are used to define technology.

Efficiency Background and Model Layout

Efficiency is classified in two categories, technical efficiency and allocative efficiency. Technical efficiency deals with the question of whether it is technically feasible for a firm to produce more output given the inputs that the firm used (i.e., did the firm use the inputs it had in the best possible way). Technical efficiency measurement deals only with the physical process of converting inputs into outputs.

Allocative efficiency deals with the question of economical choices of input and output mix based on the prices faced by the firm. Tests of allocative efficiency can focus on measuring cost minimization (input allocation efficiency), revenue maximization (output allocation efficiency), or profit maximization (a simultaneous test of input and output allocation efficiency). In the profit maximization case, allocative efficiency asks whether it is feasible to achieve higher profits given the input and output prices faced by the firm. DEA can be used to measure both the

technical and allocative efficiency of a firm based on his observed actions and prices and the actions and prices of other firms in the sample. For example, the profit maximization test for a firm is:

$$\begin{aligned}
& \text{maximize } \sum_{j=1}^J y_j p_j^0 - \sum_{k=1}^K x_k w_k^0 \\
& \text{subject to :} \\
& \sum_{i=1}^I \lambda_i y_j^i \geq y_j \quad j = 1, \dots, J \\
& \sum_{i=1}^I \lambda_i x_k^i \leq x_k \quad k = 1, \dots, K \\
& \sum_{i=1}^I \lambda_i = 1 \\
& \lambda_i \geq 0 \quad \forall i
\end{aligned} \tag{1}$$

where $(y_1^i, y_2^i, \dots, y_J^i, -x_1^i, -x_2^i, \dots, -x_K^i)$ is the netput vector for the i^{th} observation, p_j^0 and w_k^0 are the output and input prices confronting the particular firm, y_j and x_k are technically feasible levels of outputs and inputs (assuming convexity of the input/output possibilities set, free disposal of inputs and outputs, and not assuming constant returns to scale), and the λ_i 's are weights defining a convex combination of the observed netput vectors. If the firm's maximum feasible profit level (as measured by the optimum value of the objective function) is higher than the firm's observed profit level then the firm is deemed inefficient. If the two profit levels are equal, then the firm is deemed efficient. This is the standard model formulation used to measure profit maximization efficiency in the DEA literature today.

A significant weakness of this method of efficiency measurement is that it assumes that the costs and returns are deterministic. If costs or returns are stochastic and producers are risk averse then this efficiency test will incorrectly attribute risk averse behavior on the part of firms to inefficiency. Preckel, Ahmed, and Ehui have adapted the DEA model to reflect risk aversion

by assuming that producers behave as mean-variance utility maximizers facing certain costs, but uncertain returns. Thus the adapted model formulation in (1) would become:

$$\begin{aligned}
& \underset{\lambda_i}{\text{maximize}} \sum_{j=1}^J y_j p_j^0 - \frac{\rho_0}{2} y_j' \Sigma y_j - \sum_{k=1}^K x_k w_k^0 \\
& \text{subject to :} \\
& \sum_{i=1}^I \lambda_i y_j^i \geq y_j \quad j = 1, \dots, J \\
& \sum_{i=1}^I \lambda_i x_k^i \leq x_k \quad k = 1, \dots, K \\
& \sum_{i=1}^I \lambda_i = 1 \\
& \lambda_i \geq 0 \quad \forall i
\end{aligned} \tag{2}$$

where p_j^0 denotes the mean payoff to the producer for output j , ρ_0 is the individual producer's level of risk aversion and Σ is the covariance matrix of returns faced by the producer. One problem with this formulation is that it treats ρ_0 as observable data. In order to get around this problem, Preckel, Ahmed, and Ehui treat the variance portion of the objective function as a fixed input in a fashion similar to Färe, Grosskopf, and Lee (1990). This dualization of the variance level allows the problem in (2) to be reformulated into the risk-adjusted DEA model (3) below:

$$\begin{aligned}
& \underset{\lambda_i}{\text{maximize}} \sum_{j=1}^J y_j p_j^0 - \sum_{k=1}^K x_k w_k^0 \\
& \text{subject to :} \\
& \sum_{i=1}^I \lambda_i y_j^i \geq y_j \quad j = 1, \dots, J \\
& \sum_{i=1}^I \lambda_i x_k^i \leq x_k \quad k = 1, \dots, K \\
& \sum_{h=1}^J \sum_{j=1}^J \sigma_{hj} y_j y_h \leq V \\
& \sum_{i=1}^I \lambda_i = 1 \\
& \lambda_i \geq 0 \quad \forall i
\end{aligned} \tag{3}$$

where σ_{hj} is the covariance between returns to activities h and j and V denotes the level of portfolio variance observed for the firm. This model formulation asks whether it is feasible to obtain greater profits than were obtained given the constraint that the observed level of variance of the portfolio cannot be exceeded. Now the model is couched in terms of observable variables (assuming the mean and covariances of returns can be estimated) and may be solved using available software. A byproduct of the model's formulation is that the shadow price on the variance constraint allows estimation of the firm's risk aversion. Two times the shadow price on the variance constraint is equal to the estimate of the coefficient of absolute risk aversion for the firm. This enables us to estimate the firm's risk aversion coefficient without resorting to elicitation interviews or directly modeling firms' production possibilities.

A key element of the formulation (3) is the assumption that the true mean and covariance of the returns distribution can be estimated and is equal for all firms. (The assumption of uniformity of the moments of the returns distribution across firms could be relaxed, but is maintained here to simplify the exposition.) The user of this approach should expend substantial effort to correctly characterize the distribution of returns, because inaccuracies in the characterization of these moments will translate to inaccuracies in the efficiency scores. Note that firms operating under incorrect beliefs about the moments of the returns distribution are acting inefficiently. Hence, a key assumption is that the researcher can characterize the distribution of returns at least as accurately as any firm.

Monte Carlo Tests of the Model

A Monte Carlo study is used in this paper to aid in accurately determining the capabilities of the model. There are two separate Monte Carlo sections in this paper. In the first section, all firms will be perfectly efficient risk averse utility maximizers. This will help to determine how

well the model does under ideal data conditions. The second section will examine what happens when firms are both risk averse and inefficient. This section is closer to what is observed in the real world. The Monte Carlo study is similar in spirit to Preckel, Ahmed, and Ehui, but it goes beyond that work by using a more general data generation procedure. The key difference between the data generation procedure used here and the earlier work is that the basis for behavior of firms is expected utility. In the earlier work, firm behavior was based on a mean-variance model. Thus, this work examines the robustness of the measurement method to the nature of behavior.

Data Generation

In order to test for efficiency, the risk-adjusted DEA model requires firm level data on input usage (x), investment decisions (y), and the first two moments of the distribution of returns faced by the firm μ and Σ . In order to generate this data, a calibration procedure is used to generate a discrete sample of returns that exhibit a mean and covariance structure such that a utility maximizing agent will choose a diversified portfolio of investments. Then a firm level expected utility maximizing model is confronted with technology, prices, and the discrete data on returns, and the optimal choices are calculated. The input and investment choices of each agent in the sample are recorded. These input and investment choices, along with the data on the distribution of returns, are used as input data in the Monte Carlo tests of the risk-adjusted DEA model. The first step in this process is the generation of the returns data that agents face.

Generation of Returns Data

The specification of the net returns data follows Preckel, Ahmed, and Ehui closely. This specification relies on the concept of fundamental driving variables. These are underlying variables that drive a particular technology (e.g. rainfall or growing degree days in the case of

crop production). The returns vector is specified as a linear function of the fundamental driving variables ($\tilde{r} = Dv$ where \tilde{r} denotes yield, v denotes the fundamental variables, and D is a matrix of fixed coefficients). The returns vector is transformed into a random variable through the addition of a random disturbance term ($\tilde{r} = Dv + w$ where w is a disturbance vector). In this Monte Carlo there are five investments (\tilde{r} 's) and three fundamental driving variables (v 's). The transformation matrix D is defined as:

$$D = \begin{bmatrix} 8 & 0 & -5 \\ 9 & 4 & 0 \\ -3 & 7 & 0 \\ -4 & 8 & -3 \\ -1 & 2 & 6 \end{bmatrix} \quad (4)$$

A sample of 50 observations of the fundamental driving variables is drawn from a uniform [1,3] distribution. The D matrix is applied to each of the 50 vectors of fundamental variables, and the additive disturbance term is drawn from a uniform [0,1] distribution. The resulting sample is used to compute a mean vector ($\tilde{\mu}$) and correlation matrix ($\tilde{\Lambda}$) for the investments.

Following Preckel, Ahmed, and Ehui and drawing on the model calibration framework of Howitt (1995), the outputs are rescaled via a two-step process. First a calibration problem is solved without regard for variance:

$$\begin{aligned} & \underset{x,y,z \geq 0}{\text{maximize}} \quad \tilde{\mu}^t y - c^t x \\ & \text{subject to:} \\ & y \leq Az \\ & Bz \leq x \\ & e^t y \leq 1 \\ & y_j \leq \frac{1}{J} \end{aligned} \quad (5)$$

where J is the number of investments (5 in this case), A and B are technology matrices, z is a transformation vector, e is a vector with every component equal to 1.0, and c is set equal to 0.5 for each input. Provided that investment returns exceed variable costs, the optimal solution to this problem allocates equal proportions in each investment due to the nature of the last set of constraints. Denote the shadow prices on the last set of constraints as π . A second calibration problem is solved to generate data that exhibit a mean and covariance structure that both ensures a diversified investment strategy and ensures that the vector of returns is non-negative in every state of nature. This calibration problem is:

$$\begin{aligned} & \underset{\sigma, \mu}{\text{minimize}} \quad \sum_{out} (\sigma_i - \tilde{\sigma}_i)^2 + \sum_{out} (\mu_i - \tilde{\mu}_i)^2 \\ & \text{subject to:} \end{aligned}$$

$$\begin{aligned} \mu - \rho^* \sigma \tilde{\Lambda} \sigma \frac{e}{J} &= \tilde{\mu} - \pi & (6) \\ r_i &= (\tilde{r}_i - \tilde{\mu}) \frac{\sigma}{\tilde{\sigma}} + \mu \geq 0 \end{aligned}$$

where σ is a diagonal matrix whose diagonal elements are the standard deviations to be determined, $\tilde{\sigma}$ is the diagonal standard deviation matrix of the original investments, $\tilde{\sigma}_i$ is the standard deviation for the i^{th} output, σ_i is the standard deviation to be determined for the i^{th} output, $\tilde{\mu}_i$ is the mean return for the i^{th} output, μ_i is the mean return to be determined for the i^{th} output, r_i 's are the returns vectors to be solved for, $\tilde{\Lambda}$ is the correlation matrix determined in the previous step, and ρ^* is the coefficient of absolute risk aversion. By choosing a risk aversion coefficient in the center of the range for which we will be generating agent behavior, we increase the likelihood that a diversified investment portfolio will be chosen. A relative risk aversion coefficient (ρ) of 5 is a commonly chosen value to represent a high degree of risk aversion. Setting ρ^* such that it corresponds to a relative risk aversion coefficient of 5 will result in a data

structure such that a risk averse utility maximizer will choose to invest in a proportion of each output. Absolute risk aversion is calculated as ρ/w_0 where w_0 is the initial wealth level. The initial wealth level is set equal to the level of profits the agent in problem 2 received (9.789 in this case). This calibration procedure generates non-negative returns data (r_i 's) that will usually result in a diversified set of input and investment decisions by expected utility maximizing agents.

Behavioral Model of the Firm—Efficient Firms

To increase the generality of behavior over Preckel, Ahmed, and Ehui, the observed behavioral data are generated by modeling risk averse expected utility maximizing agents who face a fixed technology. The agent is modeled using both the expo-power, including its limiting cases (e.g., the power utility). Saha introduced the expo-power utility form in 1993 as a flexible utility form capable of representing a wide variety of risk attitudes. The utility maximization problem is written as follows:

$$\begin{aligned}
 & \underset{y,x,z \geq 0}{\text{maximize}} \sum_{i=1}^N \frac{1}{N} (-\exp(-\beta(r_i^t y - c^t x + w_0)^\alpha)) \\
 & \text{subject to:} \\
 & y \leq Az \tag{7} \\
 & Bz \leq x \\
 & e^t y \leq G
 \end{aligned}$$

where r_i is a vector of returns from the i^{th} state of nature, c is a vector of costs, e is a vector with all components equal to 1, w_0 is the initial wealth of the agent, α and β are parameters of the utility function, G is total available capital, A and B are technology matrices, N is the number of states of nature, and z is a vector of alternative uses of the technologies. The coefficient of

absolute risk aversion for this utility function is $(1 - \alpha + \alpha\beta w_o^\alpha)/w_o$ and the coefficient of relative risk aversion is $(1 - \alpha + \alpha\beta w_o^\alpha)$.

The expo-power utility function is capable of representing the following combinations of risk attitudes with the appropriate parametric restrictions. These risk attitudes are enumerated in Table 1.

Table 1. Risk Attitudes Represented by the Expo-Power Utility Function

	Decreasing	Constant	Increasing
	Relative Risk	Relative Risk	Relative Risk
Decreasing Absolute Risk Aversion	$\alpha < 0, \beta < 0$	Power Utility	$0 < \alpha < 1, \beta > 0$
Constant Absolute Risk Aversion	N/A	N/A	$\alpha = 1, \beta > 0$
Increasing Absolute Risk Aversion	N/A	N/A	$\alpha > 1, \beta > 0$

The cells marked with N/A represent risk attitude combinations that are not feasible for any utility function representing a risk averse agent. Note that the expo-power representation of the DARA/CRRA risk attitude combination is a limiting case that is equivalent to the power utility function. It is also interesting to note that for the CARA/IARA risk attitude combination, the expo-power utility function collapses into the exponential utility function. In addition to the expo-power, the quadratic utility function is capable of representing the IARA/IRRA risk attitude combination.

In order to represent all feasible risk attitude combinations, the DARA/CRRA risk attitude combination will be modeled using the power utility function. The power utility maximization problem can be written as follows:

$$\begin{aligned}
& \underset{y,x,z \geq 0}{\text{maximize}} \sum_{i=1}^N \frac{1}{N} \frac{(r_i^t y - c^t x + w_0)^{(1-\rho)}}{(1-\rho)} \\
& \text{subject to:} \\
& y \leq Az \\
& Bz \leq x \\
& e^t y \leq G
\end{aligned} \tag{8}$$

where ρ is the utility function parameter corresponding to the coefficient of relative risk aversion and all other notation remains the same as in (7).

The risk attitude parameters for both the expo-power and power utility function are calculated in the following manner. For each agent, a coefficient of relative risk aversion (r^*) is drawn from a uniform [1,5] distribution. In the case of the power utility function (DARA/CRRA) the risk aversion parameter ρ is set to equal to r^* . For all other risk attitude combinations the relative risk aversion level $(1 - \alpha + \alpha\beta w_o^\alpha)$ is set equal to r^* . The α coefficient is then drawn from a pre specified distribution, and the β coefficient is solved for as $\beta = (r^* - 1 + \alpha) / \alpha w_o^\alpha$. As before, w_o is the initial wealth level, which is set to 9.789. The distribution α is drawn from varies based upon the specific risk attitude represented by the utility function. For the DARA/DRRA risk attitude α is distributed uniform $[1 - r^*, 0]$. For the DARA/IRRA risk attitude α is distributed uniform $[0, 1]$, and if the resulting β coefficient is less than 0, α is redrawn from a uniform $[1 - r^*, 0]$ distribution. For the CARA/IRRA risk attitude α is set equal to 1, and for the IARA/IRRA risk attitude α is distributed uniform $[1, 1.5]$.

In each case the values of the limits of the α parameter were derived by observing the parametric restriction that α and β must be of the same sign. In addition, α and β must be “small” in absolute value for the problem to be well scaled. For this particular problem, the magnitude of α and β must less than approximately 20 in absolute value for the problem to be

well behaved. Additionally, in some cases the optimal solutions were sensitive to the scaling of the objective function. To improve the scaling of the gradient near the solution, the objective function was rescaled by dividing the objective function by the first derivative of utility evaluated at the initial wealth plus the risk neutral profit level.

Generating the Observed Behavioral Data

The producers behave as expected utility maximizers subject to the technology available and capital constraints. Both the input and investment vectors (x and y) are considered to be observable by the researcher. The technology matrices and utility function parameters are constructed to generate the data but are considered unobservable for the purposes of efficiency testing.

In order to construct a sample of firm level observations, the vector of costs (c), risk aversion level (r^*), and the total available funds (G) are varied 50 times to generate a sample corresponding to 50 firms. The cost vector is drawn from a uniform $[0.45, 0.55]$, the risk aversion level is drawn from a uniform $[1, 5]$, and the total available funds are drawn from a uniform $[0.5, 1.5]$ distribution. The limits on the distribution of the cost vector were chosen by allowing the costs for any given input to vary by 10 percent in either direction from the cost level chosen in the calibration problem (0.5 for all inputs). This simulates firms that face costs that are at a fixed level, but may vary from firm to firm (perhaps based on locale). The coefficient of relative risk aversion is drawn such that it ranges from 1 for a slightly risk averse agent to 5, which is generally considered to be reasonable upper bound on relative risk aversion. The total available funds parameter was drawn such that the available funds would lie within plus or minus 50 percent of the funds available to level chosen for the calibration problem (1.0).

The producer optimization model in (4) and (5) is solved once for each combination of

costs, initial wealth, utility function parameterization, and total available funds. The particular draw of costs, risk aversion, and available funds for any given firm is identical for each of the five utility function parameterizations. Thus, the results from each utility function parameterization are directly comparable. The input and investment choices are recorded for each firm and are treated as data for the subsequent efficiency testing. The sample mean vector and covariance matrix of returns are also treated as observable and used as data in the efficiency test. Each firm in this sample is efficient. Behavioral Model of the Firm—Inefficient Firms

To generate a sample of data in which agents exhibit inefficiency, the optimization problems in (4) and (5) are treated as though they do not have full information. The utility maximization setup is as before, except that agents make their choices based on only a subset of the states of nature. In this case, the number of states of nature that the agent “sees” is a random sample of 40 of the original 50 states of nature. Thus, the agent can be viewed as having only 80 percent of the total available information on the distribution of net returns on which to base decisions. The 40 states of nature available to the firm are drawn as a random sample without replacement from the original 50 states of nature. The firm level optimization models in (4) and (5) are solved and the agent’s input and investment choices are recorded. A new sample of returns is drawn for each firm along with the costs, risk aversion levels, and returns. The researcher is assumed to possess knowledge of all 50 observations of the distribution of net returns on which he bases the sample mean vector and covariance matrix of returns. Thus, the researcher has better information on the true nature of returns than do the firms. This lack of knowledge of all states of nature induces the firm to act in an inefficient manner in the selection of inputs and outputs.

Efficiency Tests

For each set of data, the set where all firms are efficient and the set where firms exhibit inefficiency, two separate efficiency tests are run. First, a traditional DEA efficiency test similar to (1) (with μ being substituted for p_j^0) is run. The formulation of this test is:

$$\begin{aligned} & \text{maximize } \sum_{j=1}^J y_j \mu_j - \sum_{k=1}^K x_k c_k^0 \\ & \text{subject to :} \\ & \sum_{i=1}^I \lambda_i y_j^i \geq y_j \quad j = 1, \dots, J \\ & \sum_{i=1}^I \lambda_i x_k^i \leq x_k \quad k = 1, \dots, K \\ & \sum_{i=1}^I y_i \leq G^0 \\ & \lambda_i \geq 0 \quad \forall i \end{aligned} \tag{9}$$

where a superscript 0 indicates that the values are for the firm for which the testing is being done. The data required to run this efficiency test are the mean returns vector, cost vectors, total available capital, and observed input and output vectors. If the actual expected returns less variable costs are less than the optimal (potential) expected returns less variable costs, then the firm is deemed inefficient, otherwise the firm is efficient. The efficiency measure is calculated as the optimal profit level minus the actual profit level, all divided by the optimal profit level.

The second efficiency test applied to the data is a risk-adjusted test similar to (3). It is formulated as:

$$\begin{aligned}
& \text{maximize } \sum_{j=1}^J y_j \mu_j - \sum_{k=1}^K x_k c_k^0 \\
& \text{subject to :} \\
& \sum_{i=1}^I \lambda_i y_j^i \geq y_j \quad j = 1, \dots, J \\
& \sum_{i=1}^I \lambda_i x_k^i \leq x_k \quad k = 1, \dots, K \\
& \sum_{h=1}^J \sum_{j=1}^J \sigma_{hj} y_j y_h \leq \sum_{h=1}^J \sum_{j=1}^J \sigma_{hj} y_j^0 y_h^0 \\
& \sum_{j=1}^J y_j \leq G^0 \\
& \lambda_i \geq 0 \quad \forall i
\end{aligned} \tag{10}$$

This problem is a more constrained version of (9) and the optimal objective value found for any particular firm in (10) is less than or equal to the optimal objective value found for that firm in problem (9). Thus, we can compare the levels of efficiency predicted by each model to the known level of inefficiency.

Calculating efficiency based only on the profits of the firm (as is done in test (9)) would ignore the fact that firms are risk averse. Thus the efficiency measures (for the true level of efficiency and the level predicted by the risk-adjusted DEA model) are calculated by converting the firm portfolios into certainty equivalents (CE's). Certainty equivalents are a method of placing a certain value on an uncertain gamble (in this case the input and output decisions of the firm) while taking the level of risk aversion into account. The certainty equivalents are calculated as follows. The true CE is the difference between the firm's expected profits for the full sample problem minus one half the true absolute risk aversion coefficient times the level of variance for the full sample problem. Likewise, the observed CE is the difference between the firm's expected profits for the limited sample problem minus one half the true absolute risk aversion coefficient times the level of variance for the limited sample problem. The true level of

efficiency is defined as the true CE minus the observed CE all divided by the true CE. This efficiency level is only observable in a Monte Carlo framework.

The predicted efficiency level is the observable analog of the true efficiency level. Two times the shadow price on the variance constraint in (10) times the initial wealth level provides an estimate of the firm's coefficient of relative risk aversion. Thus, the optimal CE is the difference between the objective value from the efficiency test (10) minus the shadow price on the variance constraint times the level of variance for the firm in the limited sample problem. Likewise, the realized CE is the difference between the firm's expected profits for the limited sample problem minus the shadow price on the variance constraint times the level of variance for the limited sample problem. The estimated level of efficiency is defined as the optimal CE minus the realized CE all divided by the optimal CE. This efficiency level is observable based on the data required to run the risk-adjusted DEA model and is compared to the true efficiency level. In addition, the estimated risk aversion coefficients are compared with the actual risk aversion coefficients to determine the method's accuracy in estimating individual risk aversion levels.

Results and Analysis

Model Performance When Agent Behavior is Efficient

In the first Monte Carlo test all firms exhibit efficient behavior, thus their true levels of efficiency are zero. Table 2 shows the results of the efficiency testing using both standard DEA and the risk-adjusted methods.

Table 2. Efficiency Test Results for Efficient Data (Percent Inefficiency)

Utility Function	Efficiency Test	Standard DEA		Risk-Adjusted DEA	
		Average	Maximum	Average	Maximum
DARA/CRRA (Power)	Standard DEA	12.2%	76.4%	0.0%	0.0%
DARA/DRRA	Standard DEA	12.2%	76.4%	0.0%	0.0%
DARA/IRRA	Standard DEA	12.3%	76.5%	0.0%	0.0%
CARA/IRRA (Exponential)	Standard DEA	12.3%	76.5%	0.0%	0.0%
IARA/IRRA	Standard DEA	12.3%	76.5%	0.0%	0.0%

Without accounting for risk, the standard DEA model indicates that agents are approximately 12 percent inefficient on average. Estimated inefficiencies range from 77 percent to 0 percent. If we correct for risk aversion by using the risk-adjusted DEA model, all estimated inefficiency vanishes. Both the risk-adjusted and standard DEA models seem to represent the various utility forms equally well (or poorly for standard DEA). The results of this portion of the Monte Carlo experiment indicate that the standard DEA model categorizes risk averse behavior as inefficient. The magnitude of inefficiency that the standard model estimates suggests that any attempt to apply standard DEA to risk averse agents will potentially result in gross overestimation of the level of inefficiency present in the sample. In contrast, the risk-adjusted DEA model performed extremely well in this set of experiments, correctly characterizing the efficiency level of every agent.

The risk-adjusted DEA model also performed exceptionally well with regards to the estimation of firm specific risk aversion coefficients (RAC). Table 3a shows the average, minimum, and maximum absolute differences between the true and estimated risk aversion coefficients for the efficient data.

Table 3a. Risk Aversion Estimates for Efficient Data

Utility Function	Absolute Difference Between True and Estimated RAC		
	Average	Minimum	Maximum
DARA/CRRA (Power)	0.1530	0.0005	2.0423
DARA/DRRA	0.1572	0.0005	2.0423
DARA/IRRA	0.1473	0.0005	2.0423
CARA/IRRA (Exponential)	0.1381	0.0010	2.0423
IARA/IRRA	0.1334	0.0018	2.0423

As Table 3a indicates, the average absolute deviation in estimating the firm specific risk aversion coefficient is about 0.146. The minimum deviation is around 0.0005 while the maximum is around 2.04. Again, varying utility forms did not appear to substantially alter the method's ability to predict RAC's. It is important to note that the true risk aversion levels ranged between 1 and 5. Thus these deviations should be viewed with that scaling in mind.

Perhaps a better measure of the risk aversion measurement capabilities of the model is the correlation coefficient. The correlation coefficient measures the strength of the linear relationship between two variables, thus indicating how well the estimated risk aversion levels track the true values. The correlation coefficient is bounded between negative one and one, with a value of one indicating a perfect correlation. Table 3b shows the correlation coefficient between the true and predicted risk aversion levels.

Table 3b. Correlation Coefficients Between True and Estimated Risk Aversion Levels

Utility Function	Correlation Coefficient
DARA/CRRA (Power)	0.9528
DARA/DRRA	0.9518
DARA/IRRA	0.9535
CARA/IRRA (Exponential)	0.9555
IARA/IRRA	0.9569

This table illustrates the strong relationship between the true and estimated risk aversion levels. This indicates that the risk-adjusted DEA model performs exceptionally well with regard to the ability of the model to estimate firm level risk aversion coefficients.

Model Performance When Agent Behavior is Inefficient

In the second Monte Carlo test, firms exhibit inefficient behavior. This is induced by allowing firms to optimize over only a subset of the distribution of net returns rather than the full distribution as in the first Monte Carlo. Table 4a illustrates the efficiency results from both the standard and risk-adjusted DEA models.

Table 4a. Efficiency Test Results for Inefficient Data, Percent Inefficiency

Utility Function	Standard DEA		Risk-Adjusted DEA		True Inefficiency	
	Average	Maximum	Average	Maximum	Average	Maximum
DARA/CRRA (Power)	26.1%	91.5%	21.0%	126.9%	29.0%	109.9%
DARA/DRRA	23.4%	91.5%	16.8%	95.1%	26.1%	100.0%
DARA/IRRA	21.8%	91.6%	17.4%	126.9%	39.1%	109.8%
CARA/IRRA (Exponential)	23.6%	91.6%	16.9%	95.1%	25.9%	100.0%
IARA/IRRA	23.6%	91.6%	16.9%	95.1%	25.8%	100.0%

As Table 4a indicates, the true level of inefficiency for all utility forms hovers between 26 to 39 percent. The standard DEA model appears to be closer to predicting the true level of inefficiency on average. This is somewhat misleading since the error in the pairings between estimated and true inefficiency levels is the error that we are truly concerned with. A better measure of performance would be the absolute difference between the true and estimated efficiency levels, as shown in Table 4b.

Table 4b. Efficiency Test Results for Inefficient Data, Absolute Difference Between True and Predicted Inefficiency Levels

Utility Function	Standard DEA	Risk-Adjusted DEA
	Average	Average
DARA/CRRA (Power)	15.4%	11.6%
DARA/DRRA	15.4%	11.8%
DARA/IRRA	27.1%	24.6%
CARA/IRRA (Exponential)	15.3%	11.5%
IARA/IRRA	15.2%	11.4%

As is shown in Table 4b, the average absolute difference between the true and estimated inefficiencies for the standard DEA models hover at about 15 percent while they are just under 12 percent for the risk-adjusted DEA model. This indicates that the risk-adjusted DEA model comes closer to the true efficiency value on average than the standard DEA model does. For all utility forms, the risk-adjusted DEA method is closer to the true level of inefficiency than the standard DEA model. Both models seem to have a harder time with the DARA/IRRA utility function. As before, the correlation coefficient between the true and estimated levels of efficiency, shown in Table 4c, is a good measure of model performance.

Table 4c. Correlation Coefficient Between True and Estimated Efficiency

Utility Function	Correlation Coefficient	
	Standard DEA	Risk Adjusted DEA
DARA/CRRA (Power)	0.5795	0.7294
DARA/DRRA	0.4830	0.6308
DARA/IRRA	0.1245	0.3154
CARA/IRRA (Exponential)	0.4875	0.6356
IARA/IRRA	0.4881	0.6359

As Table 4c indicates, the risk-adjusted DEA method substantially outperforms the standard DEA method with regard to the correlation of estimated and actual inefficiency. Again, for a reason unknown at the present time, both methods have difficulty with the DARA/IRRA risk attitude.

The risk aversion estimates from by the risk-adjusted DEA method are shown in Table 5.

Table 5. Risk Aversion Estimates for Inefficient Data (Absolute Difference Between True and Estimated RAC)

Utility Function	Average	Minimum	Maximum
DARA/CRRA (Power)	1.5164	0.1586	7.3663
DARA/DRRA	1.4312	0.0240	4.4334
DARA/IRRA	1.5782	0.1619	7.3839
CARA/IRRA (Exponential)	1.4226	0.1563	4.5333
IARA/IRRA	1.4199	0.1515	4.5830

Clearly the risk-adjusted DEA model had a harder time estimating risk aversion levels when the agents in the sample were behaving in a risk averse manner than it did when all agents were efficient. It is important to note that the risk-adjusted DEA model does not provide a reliable estimate of the risk aversion level for a firm that chooses not to produce. For all utility function except DARA/IRRA, one of the fifty firms in the sample chose to produce zero output. For the DARA/IRRA utility form there were eight firms that chose to produce zero output. These firms have been excluded from the calculation of the results in Table 5. Although the results are not as good as the results for the model using efficient data, they are still relatively close to the true values of the risk aversion levels.

Table 6 shows the correlation coefficients between the true and estimated risk aversion coefficients.

Table 6. Correlation Between True and Estimated Risk Aversion Coefficients

Utility Function	Correlation Coefficient
DARA/CRRA (Power)	0.5096
DARA/DRRA	0.4251
DARA/IRRA	0.4295
CARA/IRRA (Exponential)	0.4317
IARA/IRRA	0.4332

The results show that the correlation coefficient lies between 42 and 51 percent. This is fairly good performance in light of the small absolute differences in prediction error. Overall, the results indicate that the risk-adjusted DEA model shows real potential as an alternative method of risk aversion evaluation.

Conclusions

This paper developed a Monte Carlo framework to test a method of risk adjusted efficiency testing and risk aversion measurement pioneered by Preckel, Ahmed, and Ehui. The results of the paper indicate that the model is a viable solution to the problem of efficiency

measurement for risk averse producers. The method also appears to be a promising alternative to the methods of utility elicitation that have previously been used to estimate the risk aversion levels of agents.

There remains much work to be done in the way of model validation. Future work will extend the Monte Carlo to cover more diverse data generation processes. This will include the introduction of skewness into the returns structure that firms face. In addition, several new methods of inducing inefficient behavior on the part of the firm will be examined. The current study restricts the agents in any given sample to behave according to only one utility function. This restriction will be relaxed to allow for varying utility forms and inefficiency generation methods in the same sample of agents. This will test the ability of the risk-adjusted DEA model to perform in an environment where agents behave much more heterogeneously than they currently do.

Potential applications of the model are numerous. It could be used to model the efficiency structure of a set of firms or industry in a more accurate manner. It could also be used to examine the risk preference structure of an industry or group of firms. With panel data, applications could include the study of efficiency and risk aversion change over time or life cycle of an individual or firm. Finally, the model can be used to provide more accurate information to decision makers and to help guide decision makers to make appropriate input and investment choices.

In addition to the Monte Carlo study, a number of empirical applications are planned in order to determine the real world viability of this method. Applications to agricultural banking and mutual fund investment are two such applications that are currently underway. More applications will help shed light on the true capabilities and usefulness of this unique method.

References

- Coelli, T., 1995, 'Recent Developments in Frontier Modeling and Efficiency Measurement,' *Australian Journal of Agricultural Economics*, v. 39, no. 3, pp. 219-245.
- Färe, R., S. Grosskopf, and H. Lee, 1990, 'A Nonparametric Approach to Expenditure-Constrained Profit Maximization,' *American Journal of Agricultural Economics*, v. 72, no. 3, pp. 574-581.
- Førsund, F., K. Lovell, and P. Schmidt, 1980, 'A Survey of Frontier Production Functions and of Their Relationship to Efficiency Measurement,' *Journal of Econometrics*, v. 13, pp. 5-25.
- Gong, B., and R.C. Sickles, 1989, 'Finite Sample Evidence on the Performance of Stochastic Frontier Models Using Panel Data,' *Journal of Productivity Analysis*, v. 1, no. 3, pp. 229-261.
- Gong, B., and R.C. Sickles, 1992, 'Finite Sample Evidence on the Performance of Stochastic Frontiers and Data Envelopment Analysis Using Panel Data,' *Journal of Econometrics*, v. 51, pp. 259-284.
- Hardaker, J.B., R.B.M. Huirne, and J.R. Anderson, 1997, *Coping With Risk in Agriculture*, CAB International, New York, NY.
- Preckel, P., V. M. Ahmed, and S. Ehui, 2000, 'Non-parametric Cross-sectional Approach to Measurement of Risk Aversion,' Selected Paper at the American Association of Agricultural Economics Association, August 2, 2000.
- Saha, A., 1993, 'Expo-Power Utility: A 'Flexible' Form for Absolute and Relative Risk Aversion,' *American Journal of Agricultural Economics*, v. 75, no. 4, pp. 905-913.
- Settlage, D.M., 1999, 'A Comparison of Various Stochastic Frontiers under Differing Data Generation Assumptions,' Unpublished M.S. Thesis, University of Arkansas.