



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

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

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

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

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

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Economic Criteria for Evaluating Commodity Price Forecasts

Jeffrey H. Dorfman and Christopher S. McIntosh

ABSTRACT

Forecasts of economic time series are often evaluated according to their accuracy as measured by either quantitative precision or qualitative reliability. We argue that consumers purchase forecasts for the potential utility gains from utilizing them, not for their accuracy. Using Monte Carlo techniques to incorporate the temporal heteroskedasticity inherent in asset returns, the expected utility of a set of qualitative forecasts is simulated for corn and soybean futures prices. Monetary values for forecasts of various reliability levels are derived. The method goes beyond *statistical* forecast evaluation, allowing individuals to incorporate their own utility function and trading system into valuing a set of asset price forecasts.

Key Words: commodity prices, forecast evaluation, value of information.

In a recent article, Leitch and Tanner provide evidence suggesting that economists would be better off if they evaluated forecasts of economic time series using economic rather than statistical criteria. For example, they suggest that for financial asset prices, profits resulting from trading in the asset based on a set of forecasts would be a better guide to forecast performance than such statistical measures as mean squared error (MSE). Leitch and Tanner go on to show a lack of significant correlation between profits and root mean squared error (RMSE) in forecasts of T-bill rates. In fact, the correlation often has a perverse sign.

Several studies using the approach sug-

gested by Leitch and Tanner have evaluated forecasts of financial asset prices by the profit earned in trading with the forecasts. Brandt and Bessler (1983) compute the net returns to hedging using seven different hog price forecasts. Lukac, Brorsen, and Irwin (1988a, b), in a pair of related papers, compare the performance of 12 technical trading systems by monthly returns in simulated trading. Although none of these studies compute the correlation between profits and MSE, Brandt and Bessler do discuss the differences in rankings of the seven models when ranked by profit versus MSE. Both Figlewski and Urich, and Hein and Spudeck compute profit-related measures for forecasts and find them to be unrelated to point forecast error measures. On a related front, both Park, and McIntosh and Dorfman evaluate livestock price forecasts according to their ability to predict the series' direction of motion; however, they do not connect predictive ability to monetary measures.

Leitch and Tanner found that directional accuracy—the ability of a forecasting model

The authors are associate professors in the Department of Agricultural and Applied Economics at the University of Georgia.

The authors thank David Bessler, Wade Brorsen, Jack Houston, Hal White, seminar participants at UC Davis, and two anonymous reviewers for helpful comments. An earlier version of this paper was presented at the NCR-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management, Chicago.

to predict the upward or downward movement of the series—is significantly correlated with profits. Based on this finding, the authors suggest that this might be a better indicator of forecast quality than MSE.

In this article, we set out to achieve three goals. First, we take Leitch and Tanner's suggestion one step further and propose two metrics for evaluating forecasts based directly on economic criteria: profit and utility of profit. Second, we demonstrate the metrics with an application to agricultural futures markets trading. Third, we investigate the ex post correlation among profit, MSE, directional accuracy, and the two metrics proposed here to examine which measures appear to be most useful to forecast consumers.

In order to properly assess the value of forecasts of commodity price series, the problem of the temporal heteroskedasticity of the price changes must be addressed. That the variance of price changes is not constant for many financial assets is well known and has spawned a large body of literature (a short list includes Bollerslev; Bollerslev, Engle, and Woolridge; Engle, Lilien, and Robins; Fama; and Nelson). This temporal (or conditional) heteroskedasticity implies that *when a forecast system is correct* can have a larger influence on actual trading profits than *how often the forecasts are correct*. It also implies that the statistical expectation of profits from a trading system will not be of a standard form which can be easily handled analytically. To deal with this difficulty, numerical simulations will be employed, in effect, to integrate over the heteroskedasticity, arriving at a relation between profits and directional accuracy which accounts for the time-varying variance of the price changes. This is analogous to the bootstrap simulations used by Brock, Lakonishok, and LeBaron in their evaluation of technical trading rules.

The article proceeds by first providing a description of a simulation method for computing the expected profit from trading in a financial asset market with a given forecasting system and set of trading rules. Next, the simulations are employed to develop the two forecast evaluation metrics. Five years of data on

corn and soybean futures are then used to demonstrate the metrics. Finally, for several simple forecasting models, common measures of forecast accuracy are computed along with the two new metrics, and correlations between actual profit and the various criteria are computed. These results allow comparison with findings of earlier studies (i.e., Brandt and Bessler 1983; Figlewski and Ulrich; Hein and Spudeck; and Leitch and Tanner).

A Method for Profit Simulation

The calculation of the expected returns from forecast trading are complicated by the temporal heteroskedasticity of asset price changes. Because the variance of price changes is not constant over time, analytical methods are not easily applied to the computation of expected profits. Instead, a simulation method was developed which integrates out the heteroskedasticity by randomly varying the distribution of correct and incorrect forecasts across the trading period.

The profit simulations performed in this study are based on the corn and soybean harvest contract futures markets for 1984 through 1988. The harvest contract is simply the contract which matures closest to the crop's harvest date. In order to simulate the profit characteristics of a forecasting system with a given degree of directional accuracy, a set of trading rules must be specified. The rules used here are quite simple and are similar to profit rule A in Leitch and Tanner. Given a forecast that the price will rise (fall), we assume a long (short) position by buying (selling) one contract. If we are already in the desired position, we simply hold the current contract, thus minimizing transactions costs. A forecast of no change would result in remaining in the current position. On the final day of the trading period, the position is closed out. Thus, except on the first and last day of the period, either no transaction takes place (hold pat) or the position is switched from long to short or vice versa.

The costs of trading were calculated assuming a reasonably large trader. Transactions costs of buying or selling were taken to be \$50

per roundturn (one roundturn consists of both buying and selling one contract). The margin requirement was assumed to be 10% of the contract price, and the lost interest income from the margin requirement was calculated using an interest rate of 8%. Profits are accrued daily, since commodity futures contracts are repriced to the market each day. Profits for one year of daily trading (250 days) were calculated as discounted back to the first day of the year using the same 8% interest rate. Each contract represents 5,000 bushels of either corn or soybeans.

Given these trading rules and transactions costs, the simulation of expected profit can proceed. Using the actual daily closing prices, a set of forecasts is generated for a given directional accuracy level using a random number generator to distribute the correct and incorrect forecasts across the 250 trading days. Let f_t be the forecast direction of revision for the futures price on day t , i.e., $f_t = E[\text{sign}(p_t - p_{t-1})]$. Let u_t be a random variable generated from the uniform [0,1] distribution. Then, for a forecasting system with a directional accuracy of d , $0 \leq d \leq 1$, the simulated forecasts are generated by the rule $f_t = \text{sign}[(p_t - p_{t-1}) \times (d - u_t)]$.

This rule will generate forecasts which in the long run correctly forecast the direction of revision (up or down) d percent of the time. While such forecasts are completely artificial (they can only be computed after p_t is known), they serve the necessary purpose of allowing an expected profit to be calculated for forecasts with time-varying payoffs.

The simulations performed in this analysis assume a single year of daily trading. Therefore, a set of 250 f_t s ($t = 1, 2, \dots, 250$) is generated to constitute a single year of simulated trading. This set of f_t s represents the up and down forecasts of a forecast system with a d percent reliability in predicting price change direction. It is used to calculate the discounted profit earned from one year of trading with the forecasts according to the trading rules outlined above.

The profit earned over one year is recorded, and the process is repeated for a given accuracy level d until reliable estimates of the av-

erage profit and the variance of profit are obtained. The expected profit from trading with a forecasting system with an accuracy of d percent is the sample mean of profits obtained in the simulation. That is, $E\pi(d) = (1/n)\sum\pi_i(d)$, where $i = \{1, 2, \dots, n\}$ represents the number of year-long simulations conducted for a given d . In the application presented below, n was set to 5,000. The variance of profit was estimated by the sample variance of the 5,000 simulated profit values; thus, it represents the variability of actual trading profit around the mean due to the temporal distribution of returns, not a day-to-day measure of risk. When the average profit and variance of profit were calculated for a given d , d was changed and the whole process repeated. Forecasts were simulated with directional accuracies ranging from $d = 50\%$ to $d = 80\%$.

Money Metrics from Expected Utility

After completing the simulations, the relation between directional accuracy and expected profit (and variance of profit) has been established for a range of accuracy levels. While it is tempting to think that the profit obtainable using a set of forecasts is the same as the value of the forecasts, this is not the case. Clearly, one would pay forecasters only for the additional profit earnable above what could be made without their help. Further, the risk inherent in the trading process must be considered in the valuation process. Finally, as economists, we are trained to respond that it is the utility of profit which matters, not the profit itself, where the utility of profit accounts for the risks involved and the opportunity costs of partaking in the trading venture.

An Expected Utility Metric

For the reasons noted above, we take an expected utility approach to calculating the information value in a set of forecasts. The utility function used is the negative exponential, $U(\pi) = 1 - \exp(-\phi\pi)$, where ϕ is the Arrow-Pratt absolute risk aversion coefficient. This is a common utility function for evaluating the utility of risky returns (dating back to the 1956

work of Freund).¹ One advantage of this utility function is that the expected utility of a risky return is equal to a function of its mean and variance as long as the return is distributed normally. Maintaining that assumption results in the following relation for expected utility of profit:

$$(1) \quad EU(\pi) = 1 - \exp\left(-\phi\left[E(\pi) - \frac{\phi}{2}\text{var}(\pi)\right]\right).$$

If the average profit and variance of profit calculated from the simulations are inserted into equation (1), the result is the expected utility from using a set of forecasts with a given degree of accuracy. Denote this by $EU(d)$, where d is the percent reliability of the forecasting system being evaluated.

Calculating values of $EU(d)$ for a range of plausible directional accuracy levels provides a utility metric for judging a set of forecasts. Such a criterion is in line with the suggestion of Leitch and Tanner that forecasts be judged on a basis related to their inherent profit potential. In this sense, a utility metric is superior to a statistical metric such as MSE for evaluating economic forecasts. The superiority derives from choosing a metric which measures what the forecast user actually cares about (i.e., gains utility from). However, while an expected utility metric is an appealing way to measure forecast quality, it does not provide direct information on the value of information contained in the forecast set. Therefore, the expected utility metric is best seen as an intermediate step toward a more intuitive metric which makes the link between value and measure more direct.

A Money Metric

A money metric is necessary to price the information contained in a set of forecasts. Yet

such a metric should not depend solely on the profit potential of a forecasting system, since utility may not be a function of profit alone. Therefore, what is needed is a money metric which is a transformation of the expected utility metric derived above. The transformation suggested here is based on a feature of the class of utility function employed. It would be appropriate for any utility function which depends on the moments of the profit distribution.

If one could purchase a set of forecasts (or forecasting system) which had an accuracy level of $d = 100\%$ —i.e., a perfect forecast—there would be no risk. The profit potential of this perfect forecasting system is certain (given the trading rules). The profit distribution is a single point at the maximum profit obtainable for the given trading rules. Because of the lack of risk in using such a forecast, the value of a perfect forecast should be equal to the profit obtainable from the forecasts, $\pi(d)$. A forecast user would be willing to pay up to the full profit potential to obtain such a forecasting system because no risk premium enters into the calculation. This fact fixes one point of the transformation between the expected utility metric and the money metric.²

To obtain the remaining values for the money metric for forecast accuracy from the expected utility metric, one employs the standard conditions for consumer utility maximization when purchasing two or more goods. Denote the money value of a set of forecasts which are correct d percent of the time by $V(d)$. Then, since the forecasts must be consumed in a discrete quantity (either you buy them or you do not), the expected marginal utility from “consuming” a set of forecasts is simply the expected utility from using the forecasts, $MU(d) = EU(d)$. Therefore, using the equilibrium condition that the price ratio of two goods should equal the ratio of their marginal utilities, the value of the information contained in a set of forecasts for any accuracy level d is given implicitly by

¹ Hal White (professor of economics at UC San Diego) has correctly suggested that many professional traders and account managers probably have utility functions which vary considerably from the one used here. The simple function chosen here is still a useful choice for the introduction of this procedure. We agree that more complex functions would be beneficial in applied situations.

² Since all expected utility functions are arbitrary up to a linear transformation, this is equivalent to a choice of the scale of the expected utility function.

$$(2) \quad \frac{V(d)}{MU(d)} = \frac{V(100)}{MU(100)}$$

Solving for $V(d)$, inserting $EU(d)$ for $MU(d)$, and noting that $V(100) = \pi(100)$ gives a simple rule for calculating the value of information for a set of qualitative forecasts which are correct d percent of the time:

$$(3) \quad V(d) = \frac{\pi(100)EU(d)}{EU(100)}$$

Calculating $V(d)$ across a range of values for d provides the money metric. The $V(d)$ s associated with several competing forecasting systems can be computed based on their historic accuracy levels. Instead of the forecasting system with the lowest MSE being declared superior, the $V(d)$ s would serve as the criterion by which the systems are judged. This provides a more logical basis to evaluate economic forecasts. Since the forecasts are intended to be used in an economic arena (investment and speculation), it makes sense to judge forecasts by an economic (and utility) based measure rather than a statistical one.

A Certainty Equivalence Metric

Alternatively, one could transform the expected utility metric to the certainty equivalent of the expected utility metric. The certainty equivalent of a set of forecasts is a direct transformation of the expected utility from the forecasts. The certainty equivalence metric is given by

$$(4) \quad CE(d) = E[\pi(d)] - \frac{\phi}{2} \text{var}[\pi(d)] \\ = -\ln[1 - EU(d)]/\phi.$$

This metric is also a money measure which has the naturally intuitive scale of dollars (or other currency).

The $V(d)$, $CE(d)$, and $EU(d)$ are not invariant to the set of trading rules employed in calculating the expected profits, a feature we believe is an advantage. In particular, the measures calculated in this study depend on the trading rules outlined earlier and the assump-

tion of trading only a single contract. Multiple-contract positions will produce different values of the metrics, as will different methods for deciding when to go long or short. This accounts for the fact that the value of the forecasts to the consumer changes based on how the information is utilized, following the Beckerian view of valuing commodities for the products derived or produced from the initially purchased good or service (Becker). In this sense, the forecasts are purchased in order to "produce" buy and sell decisions which then hopefully produce profit. While the dependence of the metrics on an underlying set of trading rules makes comparison across studies somewhat more difficult than for MSE, a fairly standard set of single-contract trading rules does not seem hard to achieve for academic comparison purposes. User-specific trading rules would still remain important for correctly assessing individual applications.

A Demonstration of the Metrics

To demonstrate how a forecast consumer could use the utility-based money metric in valuing forecasting systems with varying degrees of directional accuracy, an example was developed using 1988 data on the futures markets for corn and soybeans.

The expected utility of forecasts with varying degrees of accuracy d were calculated from the earlier simulations of the average profit and variance of profit and the negative exponential utility function. The risk aversion coefficient ϕ was varied across a range of plausible values to provide some sensitivity analysis to the results. The values of ϕ chosen were 0.0001, 0.00002, 0.00001, 0.000002, and 0.000001. Because the risk aversion coefficient can be interpreted as the inverse of the largest amount one is willing to lose (Pratt), the values of ϕ are all fairly small. They represent willingnesses to lose between \$10,000 and \$1,000,000. It seems unlikely that anyone trading in the commodity futures markets is more risk averse than the more risk averse end of this range.

The values of $EU(d)$ can be used in equation (3) to solve for the value of information

Table 1. Value of Information in 1988 Corn and Soybean Forecasts

<i>d</i>	Corn ($\phi =$)		Soybean ($\phi =$)	
	0.00002	0.000001	0.00002	0.000001
50	-9,563.67	-6,324.04	-21,376.87	-6,682.40
55	-2,706.98	-1,760.06	12,958.03	6,791.70
60	3,195.04	2,498.53	39,234.78	19,756.09
65	8,664.85	6,775.72	60,566.54	33,188.67
70	13,866.62	11,181.54	76,733.34	46,123.33
75	18,450.00	15,393.93	88,969.28	58,567.35
80	22,934.01	19,861.46	98,644.49	71,062.97
100	36,968.60	36,968.60	120,331.08	120,331.08

Note: *d* is the directional accuracy, and ϕ is the risk coefficient.

Table 2. Certainty Equivalent of 1988 Corn and Soybean Forecasts



<i>d</i>	Corn ($\phi =$)		Soybean ($\phi =$)	
	0.00002	0.000001	0.00002	0.000001
50	-6,340.44	-6,180.85	-7,489.91	-6,280.23
55	-1,879.23	-1,698.55	5,157.04	6,420.57
60	2,312.67	2,503.13	17,598.58	18,775.15
65	6,533.92	6,722.55	30,624.46	31,800.31
70	10,907.08	11,060.94	43,398.84	44,474.49
75	15,109.33	15,215.16	55,843.90	56,782.01
80	19,592.90	19,692.63	68,501.36	69,350.07
100	36,970.02	36,975.23	120,341.70	120,361.30

Note: *d* is the directional accuracy, and ϕ is the risk coefficient.

in a set of forecasts, the $V(d)$. The 1988 forecasts for corn and soybean are shown in table 1 for $\phi = 0.00002$ and $\phi = 0.000001$ and for selected values of *d* from 50 to 80%, plus 100% (perfect foresight). The values of $EU(d)$ also can be used to compute the certainty equivalence metric presented in equation (4). These $CE(d)$ values are displayed in table 2 and show the same basic pattern as the $V(d)$.³ In fact, for the corn simulations, the correlation between the $V(d)$ and $CE(d)$ is 0.885. Figure 1 shows the $V(d)$ and $CE(d)$ measures for corn with $\phi = 0.0001$, while figure 2 shows the same graphs for $\phi = 0.000002$. The most important conclusion drawn from these results, and illustrated by comparing the two figures, is that more risk-averse consumers have a lower (more negative) willingness to

pay for forecasts at the low end of the accuracy scale (*d* close to 50%), but place a higher value on accurate forecasts (*d* close to 80%) than do less risk-averse consumers. This makes intuitive sense since the more risk-averse consumer is willing to pay a higher premium to avoid uncertainty. Purchasing a good forecast lowers the consumer's risk, and a risk-averse consumer will trade off the lower resulting net profit against the reduction in uncertainty. A consumer more tolerant of risk (with a small ϕ) has a certainty equivalent closer to the expected value outcome and will not pay as much to move toward the certainty equivalent. As expected, the more risk averse the forecast consumer, the more nonlinear is the relation between directional accuracy and forecast value. As ϕ decreases (implying less risk aversion), the $V(d)$ and $CE(d)$ measures converge and become more linear in *d*. In fact, a graph of $V(d)$ and $CE(d)$ for $\phi = 0.000002$

³ Complete listings for all values of ϕ and *d* in tables 1 and 2 are available from the authors upon request.



Notes: $V(d)$ is value of information, $CE(d)$ is certainty equivalent, d is directional accuracy, and ϕ is the risk aversion coefficient.

Figure 1. $V(d)$ and $CE(d)$ for 1988 corn forecasts with $\phi = 0.0001$

Figure 2. $V(d)$ and $CE(d)$ for 1988 corn forecasts with $\phi = 0.000002$

shows the two lines virtually coincident (figure 2).

Although the less risk-averse consumer places a much lower value on the information in a forecast with an accuracy rate in the 60–80% range, the elasticity of value with respect to accuracy is much higher at $d = 80\%$ for these less risk-averse traders. As the accuracy of the forecasting system approaches 100%, the less risk-averse consumers' valuation is rising much more rapidly than that of the more risk averse. This is because the valuations for all levels of risk aversion must converge to the same point at $d = 100$.

Do the Metrics Work in Practice?

To evaluate the new metrics in a more realistic setting, an experiment was conducted using a set of three competing forecasts for hog prices. The three forecasts were taken from a structural model originally developed by Brandt and Bessler (1981), a state space/time-series model, and an expert forecast issued by the University of Missouri Extension Service (produced by Glenn Grimes and Ron Plain). A set of 68 quarterly forecasts was constructed by each of these three methods, dating from 1976:I through 1992:IV.

Expected profit and variance of profit were

computed using a four-quarter rolling average of trading a single contract (40,000 pounds) as described above based on the direction of the price forecast. Transactions costs were still assumed to be \$50 per roundturn. This reduced the period of evaluation to 1977–92, since 1976 must be used to initialize the expected profit and variance measures. Using the expected profits and variances of profits, the expected utility of profit was computed for each set of forecasts for a variety of values of ϕ and each quarter from 1977:I through 1992:IV (64 periods). These values were then used to construct corresponding series of information value and certainty equivalent metrics. Along with these measures, the MSE of each forecasting series was computed on an identical four-quarter rolling sample basis and the mean profit over the past four quarters was also saved. Thus, in total, four measures of forecast performance were constructed for each of the three forecast series: profit, MSE, V , and CE .

Having constructed these four performance measures, they were then evaluated by a selective trading experiment using these metrics to choose which forecast to believe in each quarter. For example, to evaluate the value of information metric V , trading is conducted (buying or selling one contract) on the advice

Table 3. Trading Evaluation of Four Performance Measures

ϕ	Mean Profit			
	CE	V	MSE	Profit
0.01	802.03	670.16	341.09	668.91
0.005	801.25	668.59	341.09	668.91
0.003	911.41	860.16	341.09	668.91
0.002	879.69	848.13	341.09	668.91
0.001	754.69	742.97	341.09	668.91
0.0005	754.69	754.69	341.09	668.91
0.0001	668.91	668.91	341.09	668.91
0.00001	668.91	668.91	341.09	668.91

Notes: All mean profits are evaluated over 64 quarters, 1977–92, and include transactions costs of \$50 per round-turn. Variances are not displayed, as they were relatively constant with standard deviations in the range of \$2,000–\$2,200 across all methods and risk aversion levels.

of the forecast that has the highest value of V for that period, based on its performance over the past four quarters. In one period, the structural model's forecast might be chosen, the next quarter might find the expert model with the highest V, and so on. Trading is conducted for 64 quarters. Results of such a trading experiment are shown for a variety of ϕ s in table 3.

The mean profits recorded in the trading experiment show that the certainty equivalent metric is the best choice for the application to hog prices. Choosing the forecasting source by its certainty equivalent over the past four quarters results in the highest average trading profit (with no increase in variance). The information value metric is generally the second best measure, with the past year's profit being only the third best indicator of the next quarter's forecast accuracy. Mean squared error is the worst way to evaluate forecasts in terms of linkage to economic performance. This finding is supported by results of earlier research (Leitch and Tanner; Figlewski and Urich; and Hein and Spudeck).

As the risk aversion coefficient ϕ approaches zero, the consumer becomes risk neutral and the CE, V, and profit measures all converge to identical performances as expected. Interestingly, the relative performance of the measures is not monotonic with respect to ϕ , with the mean profit gained by following

the value of information measure V moving up and down several times through the range of ϕ shown in table 3.

Conclusions

Two metrics, based on the expected utility of profit, have been proposed for evaluating forecasts of economic time series. These metrics extend the suggestion of Leitch and Tanner to evaluate economic forecasts by economic rather than statistical criteria.

The metrics were demonstrated using futures markets data for corn and soybeans and in a trading experiment using hog price forecasts. For a set of three hog price forecasting models, the certainty equivalent metric proved the best at choosing a forecast to follow, with the value of information metric being second best. Both of these metrics were superior to using lagged profit from trading to choose forecasts upon which to base future trading. Mean squared error was shown to be far inferior for the purpose of economic evaluation of forecasts.

References

- Becker, G.S. "A Theory of the Allocation of Time." *Econ. J.* 75(1965):493–517.
- Bollerslev, T. "Generalized Autoregressive Conditional Heteroskedasticity." *J. Econometrics* 31(1986):307–27.
- Bollerslev, T., R.F. Engle, and J.M. Woolridge. "A Capital Asset Pricing Model with Time Varying Covariances." *J. Polit. Econ.* 96(1988):116–31.
- Brandt, J.A., and D.A. Bessler. "Composite Forecasting: An Application with U.S. Hog Prices." *Amer. J. Agr. Econ.* 63(1981):135–40.
- . "Price Forecasting and Evaluation: An Application to Agriculture." *J. Forecasting* 2(1983):237–48.
- Brock, W., J. Lakonishok, and B. LeBaron. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." *J. Finance* 47(1992):1731–64.
- Engle, R.F., D.M. Lilien, and R.P. Robins. "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model." *Econometrica* 55(1987):391–407.
- Fama, E.F. "Forward and Spot Exchange Rates." *J. Monetary Econ.* 14(1984):319–38.

- Figlewski, S., and T. Urich. "Optimal Aggregation of Money Supply Forecasts: Accuracy, Profitability, and Market Efficiency." *J. Finance* 38(1983):695-710.
- Freund, R.J. "The Introduction of Risk into a Programming Model." *Econometrica* 24(1956): 253-63.
- Hein, S.E., and R.E. Spudeck. "Forecasting the Daily Federal Funds Rate." *Internat. J. Forecasting* 4(1988):581-91.
- Leitch, G., and J.E. Tanner. "Economic Forecast Evaluation: Profits versus the Conventional Error Measures." *Amer. Econ. Rev.* 81(1991): 580-90.
- Lukac, L.P., B.W. Brorsen, and S.H. Irwin. "Similarity of Computer Guided Technical Systems." *J. Futures Markets* 8(1988a):1-13.
- . "A Test of Futures Market Disequilibrium Using Twelve Different Technical Trading Systems." *Appl. Econ.* 20(1988b):623-39.
- McIntosh, C.S., and J.H. Dorfman. "A Comparison of Two Performance Measures." *Amer. J. Agr. Econ.* 74(1992):209-14.
- Nelson, D.B. "Conditional Heteroskedasticity in Asset Returns: A New Approach." *Econometrica* 59(1991):347-70.
- Park, T. "Forecast Evaluation for Multivariate Time-Series Models: The U.S. Cattle Market." *West. J. Agr. Econ.* 15(1990):133-43.
- Pratt, J.W. "Risk Aversion in the Small and in the Large." *Econometrica* 32(1964):122-36.

