Optimal Aggregation of Money Supply Forecasts: Accuracy, Profitability and Market Efficiency

Author(s): Stephen Figlewski and Thomas Urich

Source: *The Journal of Finance*, Vol. 38, No. 3 (Jun., 1983), pp. 695-710
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/2328077
Accessed: 27/11/2009 08:02

THE JOURNAL OF FINANCE • VOL. XXVIII, NO. 3 • JUNE 1983

Optimal Aggregation of Money Supply Forecasts: Accuracy, Profitability and Market Efficiency

STEPHEN FIGLEWSKI and THOMAS URICH*

ABSTRACT

We present a general procedure for aggregating expert forecasts which exploits regularities in the structure of information within the forecaster population. Specific information structures lead to aggregation methods which adjust for additive bias, differences in individual accuracy, and correlation among forecasts. As an application, we construct composite predictions of the weekly change in the money supply from forecasts made by twenty major securities dealers, for which high positive correlation is found to be a significant characteristic. Due to instability in the information structure, our methods cannot improve on the accuracy of a simple average in this case. However, they do capture information about the correlation among money supply forecasts which is not fully impounded in short-term interest rates. Forecasts from our models accurately predict the direction of price changes for Treasury bills and Treasury bill futures after a money supply announcement.

THE ADJUSTMENT OF PRICES to information in financial markets has been studied at length, mostly from the perspective of the Efficient Markets Hypothesis. Market efficiency tests typically look at the behavior of prices, the end product of a market's information processing, without considering the actual mechanism by which information enters prices. Examining this complex process is difficult because expectations formation is inherently subjective. This is compounded by the fact that much of the information used by market participants is itself subjective, consisting of predictions and opinions of professional analysts and other investors. Many participants rely almost exclusively on this type of "second hand" information.

For available information to be accurately incorporated into prices three conditions must be satisfied. First, those with access to objective information must use it correctly in making forecasts. Second, traders who rely on forecasts made by others must be able to evaluate the information they contain accurately. In particular, they must be able to form an accurate composite forecast from a set of diverse individual predictions. Finally, the market must aggregate investors' expectations, as manifested in their demand functions, in such a way that each bit of information receives the appropriate weight in the market clearing price.

This paper deals with the last two aspects of a market's information processing. We begin by examining methods for aggregating individual forecasts into a

* Associate Professor of Finance, New York University Graduate School of Business Administration, and Assistant Professor of Finance, Baruch College, City University of New York. This research was done during 1981–2 while Stephen Figlewski was a Batterymarch Fellow. We would like to thank Roy Radner and Roger Klein for many helpful discussions and Michael Brennan for excellent editorial suggestions.

composite that captures as much information as possible. Then market prices are tested to see if they reflect all of the information that can be obtained in this way. Specifically, we are concerned with whether forecasts of money supply changes made by major securities dealer firms contain information that is not fully incorporated in yields on Treasury bills and Treasury bill futures.

The next section describes the data. In Section II, we present a general procedure which uses the estimated stochastic structure of the forecast errors to aggregate individual forecasts into an unbiased minimum mean squared forecast error composite. Since this error structure must be estimated from past data, it may be advantageous to impose constraints on its form. Section III describes several constrained models and tests their restrictions in our sample. Under the most extensive set of constraints, forecast errors are assumed to have zero mean and to be independent and identically distributed. In this case the optimal forecast is the simple average. However, the hypothesis of independence among forecast errors is overwhelmingly rejected for our data–errors are highly positively correlated with one another.

The forecasting performance of the models is analyzed in Section IV. Because the estimated error structure was not completely stable over time, the models which adjusted for correlation did not achieve lower mean squared forecast error than the simple average in out-of-sample tests. Even so, we find in Section V that forecasts from these models, while less accurate than the simple mean, do contain information which is not fully reflected in prices in the money market, and is therefore economically valuable. The final section contains concluding comments.

I. The Data

Since mid-1977, Money Market Services Inc. has conducted weekly surveys of U.S. Government security dealer firms for their forecasts of the money supply. During the period covered by our sample, the money supply was announced on Thursday afternoon. The 50 to 60 individual respondents were asked for their forecasts of the change in M1, first on Tuesday morning and then again on Thursday morning. The Thursday data are less complete than the Tuesday data and preliminary tests showed similar results for the two data sets. We therefore focus our attention solely on the Tuesday forecasts. Our sample consists of the responses for the twenty forecasters we are able to identify individually and follow over the sample period. Seventeen of these represented primary U.S. Government security dealer firms. There are 118 weeks in the sample, spanning the period from October 1977 to January 1980.

This survey provides information about money supply expectations of some of the largest and most sophisticated participants in the money market. Earlier analysis by Urich and Wachtel [9, 10] found that the mean forecast from the survey was clearly superior to the predictions from an ARIMA model, in contrast to what has been found in other studies of forecast data e.g., Pearce [7]. Also, the survey mean appears to be a good measure of the market's expectation of the change in the money supply which is incorporated in short-term interest rates. The monetary surprise, as measured by the difference between the actual change and the mean expectation from the survey, was a significant determinant of the change in Treasury bill yields immediately after the money supply announcement, while the mean expectation itself had no explanatory power. The median forecast (though not the mean) is publicly available on various computerized data systems and is widely followed by the financial community.

II. Aggregating Diverse Forecasts with Limited Information

This section presents a general procedure for aggregating a collection of diverse forecasts into a composite which captures as much information as possible, but does not require independent knowledge beyond a past history of forecasts and realizations.[1]

Suppose an investor regularly has access to point forecasts F_{it} of M_t, the weekly change in the money supply at time t, where $i = 1, \cdots I$ designates the producer of the forecast. The investor's objective in the current period (T) is to use the vector of forecasts $\underline{F}_T$ to form an optimal composite prediction M_T^* of M_T. We shall consider composite forecasts that minimize mean squared forecast error, and are linear, i.e., of the form $M_T^* = W_0 + \underline{W}'\underline{F}_T$ and unbiased over all realizations of $(M_T, \underline{F}_T)$.

We assume that the only relevant data available to the investor are the vector of current period forecasts $\underline{F}_T$ and a set of T observed forecast vectors and actual money supply changes from earlier periods, $\underline{F}_0, \cdots \underline{F}_{T-1}, M_0, \cdots M_{T-1}$. Even if investors do have knowledge of fundamental factors beyond present and past expert forecasts, it is reasonable for them to assume that any information they may have is also known to the experts and has been appropriately included in their forecasts. In that case, the conditional probability distribution of M_T given $\underline{F}_T$ is the same as the conditional distribution given $\underline{F}_T$ plus the investor's additional information, so that only the forecasts need be analyzed. We do assume, however, that the experts report personal forecasts based on their own information, rather than composite forecasts. Thus the investor may derive valuable information from the *set* of expert forecasts which has not been accurately impounded in the individual predictions.

Economic time series are in general nonstationary processes. Because seasonal factors, money growth targets, operating procedures, and numerous other factors change over time, the joint probability distribution of M and $\underline{F}$ for one week cannot be assumed to be exactly the same as that for the previous week. To deal with this problem, we focus instead on the *information generating process*, which will be defined in terms of the probability distribution of forecast errors. It is this process that will be treated as being time-invariant. For example, we assume a forecaster whose good information has led to small forecast errors in the past will continue to be accurate in the future, and when two experts make similar predictions, their positive correlation will persist over time.[2]

[1] The procedure presented here has been applied successfully to aggregating expectations about the Consumer Price Index drawn from the well-known Livingston survey. See Figlewski [3].

[2] Clearly, the distribution of forecast errors may also be nonstationary. In that case, further procedures to minimize the effects of changes over time may be appropriate.

We first make an assumption about the properties of individual forecasts which amounts to assuming that forecasters are rational except for a possible additive bias.

$$\textit{Assumption 1: } M_t = F_{it} + \varepsilon_{it}, \qquad E[\varepsilon_{it}] = b_i \quad \text{for all } i.$$

Although full rationality would require that all b_i be zero, since a consistent bias ought to be corrected over time, we prefer to treat unbiasedness as a testable proposition at this stage rather than imposing it *a priori.*

Our second assumption concerns the joint distribution of the errors the forecasters make in predicting the money supply.

> *Assumption 2:* The vector of forecast errors $\underline{\varepsilon}_t \equiv M_t\underline{1} - \underline{F}_t$, is distributed as a multivariate normal with mean vector $\underline{b}$ and variance-covariance matrix Ω, which has full rank.

The term $\underline{1}$ denotes a vector of 1's.

From Assumption 1 we have

$$E[\underline{F}_T] = E[M_T]\,\underline{1} - \underline{b}$$

Imposing the condition of unbiasedness on the composite forecast M_T^* yields

$$E[M_T] = E[M_T^*] = E(W_0 + \underline{W}'\underline{F}_T],$$
$$\text{or } (W_0 - \underline{W}'\underline{b}) + (\underline{W}'\underline{1} - 1)E[M_T] = 0 \quad (1)$$

Since W_0 and $\underline{W}$ must not depend on the unknown value of $E[M_T]$, this equation yields two constraints:

$$\underline{W}'\,\underline{1} = 1$$
$$W_0 = \underline{W}'\underline{b} \quad (2)$$

The first constraint implies that the optimal predictor will be a weighted average of the forecasts, with weights which sum to 1. The second implies that one can correct the forecasts for bias first and then average, i.e., $M_T^* = \underline{W}'(\underline{F}_T + \underline{b})$.

Since M_T^* is unbiased, mean squared forecast error is just the forecast variance of M_T^*, given by

$$\begin{aligned} E[(M_T - M_T^*)^2] &= E[(M_T - \underline{W}'(\underline{F}_T + \underline{b}))(M_T - \underline{W}'(\underline{F}_T + \underline{b}))'] \\ &= \underline{W}'E[\underline{\varepsilon}_T - \underline{b})(\underline{\varepsilon}_T - \underline{b})']\underline{W} \\ &= \underline{W}'\Omega\underline{W} \end{aligned}$$

where the second equality uses the constraint that $\underline{W}'\underline{1} = 1$.

Thus, the optimal weights $\underline{W}^*$ solve the minimization problem

$$\underset{\underline{W}}{\text{Minimize}}\ \underline{W}'\Omega\underline{W} \text{ subject to } \underline{W}'\,\underline{1} = 1 \quad (3)$$

which has the solution[3]

$$\underline{W}^* = (\underline{1}'\Omega^{-1}\underline{1})\Omega^{-1}\underline{1} \tag{4}$$

Thus,

$$M_T^* = (\underline{1}'\Omega^{-1}\underline{1})^{-1}\underline{1}'\Omega^{-1}(\underline{F}_T + \underline{b}) \tag{5}$$

Because of the stationarity assumption, $\underline{b}$ and Ω can be estimated using historical data on M and $\underline{F}$, so that M_T^* can be computed by the investor from the information he has available.

Equation (5) gives the minimum mean squared error linear predictor for the case of a general error bias vector $\underline{b}$ and covariance matrix Ω. This expression reduces to one of the more familiar aggregation procedures when $\underline{b}$ and Ω have certain special forms. Clearly, when all forecasters are assumed to be unbiased, $\underline{b}$ drops out. If, in addition, all forecasts are independent and identically distributed, Ω is proportional to the identity matrix and (5) reduces to taking a simple average. The more general case of independent errors with different variances leads to a weighted average in which each forecaster's weight is proportional to his precision (the inverse of his error variance). Thus, these aggregation methods can be thought of as special cases of (5) when the error structure is constrained to be a specific form. In the following sections, we shall test whether such constraints are valid for our sample of money supply forecasts and examine the effects on forecast performance of relaxing them selectively.

III. The Structure and Stability of Money Supply Forecasts

Equation (5) shows how to construct the minimum variance linear unbiased predictor given a set of forecasts for the current period and the parameters of the forecast error distribution, $\underline{b}$ and Ω. In turning (5) into an operational system, it may be appropriate to impose constraints on these parameters for various reasons. This section will examine several types of constraints and present evidence on their validity for our data sample.

One reason to constrain the error distribution is to impose *a priori* beliefs about forecasting behavior. In particular, one may want to eliminate constant bias terms as being too contrary to rationality. In this case $\underline{b}$ will be constrained to be zero.

[3] From the Lagrangian equation, the first order conditions are

$$\underline{0} = 2\Omega\underline{W} - \lambda\underline{1}$$

$$0 = \underline{W}'\underline{1} - 1$$

Premultiplying the first condition by $\frac{1}{2}\underline{1}'\Omega^{-1}$ and rearranging gives

$$\underline{W}'\underline{1} = \frac{\lambda}{2}\,\underline{1}'\Omega^{-1}\underline{1} = 1$$

$$\lambda = 2(\underline{1}'\Omega^{-1}\underline{1})^{-1}$$

Substituting in the first order condition and premultiplying by Ω^{-1} yields (4).

It also may be desirable, and in some cases necessary, to impose constraints in order to reduce the number of parameters to be estimated. Since there are $(I^2 + I)/2$ distinct elements in a general Ω matrix, a large amount of data from earlier periods is needed when I is large. An extreme case is when there are fewer periods in the sample than there are forecasters, and an Ω matrix estimated without constraints will be singular. Even when enough past data are available to fit an unconstrained matrix, the potential gains in forecasting accuracy from allowing a general error structure must be weighed against the increased effect of sampling error when a larger number of parameters are estimated from a limited data sample.

Instability of the error structure may also be mitigated by constraining its form to require fewer parameters. We have assumed that the error distribution was stationary, but at best this will only be approximately true, and the assumption becomes less tenable the longer the span of time that is being considered. Forecasters can change their behavior over time as they refine their techniques. Inaccurate forecasters in particular will try to improve, and failing that, may be replaced. Also, when forecasts are identified by firm and not individually, normal turnover in personnel can induce instability.

The most common and the easiest procedure to aggregate diverse forecasts is by taking a simple average. The average will also be the minimum mean squared error composite if all forecasts are unbiased, identically distributed, and independent, i.e., under an extensive set of restrictions on the structure of forecast errors.[4] We now examine a series of models in which these restrictions are progressively relaxed.

Unbiasedness is a restriction on the mean of the error distribution and can be imposed independently of the covariance structure. Therefore, each covariance model was fitted in both "uncorrected" and "bias corrected" form.

The model in which Ω is assumed to be proportional to the identity matrix will be called the mean model. In moving towards more general cases, the first restriction to be relaxed is that of equal variances. When forecast errors are independent but some forecasts are more accurate than others, Ω becomes a diagonal matrix. We refer to this as the diagonal model.

Of the restrictions implied by the mean model, probably the most likely to be incorrect in practical forecasting applications is that of independence. To a large extent, forecasters share the same information and methods of analysis, and a good portion of every forecaster's error is due to events which were not fully anticipated by any of them. This is especially true in our sample. The 20 money supply forecasters were a very homogeneous group. The smallest (most negative) estimated bias was -0.45 ($-\$450$ million) and the largest was .05. The small size of these numbers can be seen in the fact that a bias of -0.45 added only .05 to root mean squared forecast error (RMSE). Differences in individual accuracy were also rather unimportant. The lowest RMSE was 1.64 and the largest was

[4] In fact, the simple average is the optimal composite under more general conditions, namely, if $\underline{b}'\underline{1} = 0$ and $\underline{1}$ is an eigenvector of Ω. The first condition implies that the average is unbiased and the second can be used in (4) to yield $\underline{W}^* = (1/I)\underline{1}$. The intuitive interpretation of these more general conditions is not clear. However, one important case is when $\Omega = \theta^2\underline{1}\,\underline{1}' + \sigma^2 I$, i.e. all forecasters have both equal forecast variances and equal covariances with every other forecaster.

1.93, which is a considerably smaller range than has been found in other forecast data sets. However, all of the forecast errors showed a large degree of positive correlation. Of 190 pairwise correlation coefficients, the smallest was .65 while the largest was .88.

This suggests that allowing the off-diagonal elements of the Ω matrix to be nonzero may make a substantial difference in the composite forecasts. Estimating the Ω matrix with no constraints (the unconstrained model) is clearly one solution. In this case that involves fitting 210 distinct elements of Ω plus 20 bias terms, and the associated problems of sampling error and instability. We therefore include another constrained model, the single-index model, which represents an intermediate step between the diagonal and the fully unconstrained models.

In the single-index model, each forecaster's error is assumed to be made up of one part which is independent of all other forecasters' errors plus a part which is proportional to a common error. That is, $\epsilon_{it} = \eta_{it} + \beta_i\delta_t$, where η_{it} is the independent error and δ_t the common error. If $\mathrm{Var}[\delta_t] = \theta^2$, the covariance between ϵ_i and ϵ_j is then $\beta_i\beta_j\theta^2$. Letting D be a diagonal matrix with $\mathrm{Var}[\eta_i]$ on the diagonal, we can write $\Omega = D + \theta^2\underline{\beta}\underline{\beta}'$. The single-index model allows nonzero covariances among errors but constrains them to arise only from mutual correlation with the single index δ. This reduces the number of distinct parameters in Ω from $(I^2 + I)/2$ to $2I + 1$: I independent variances, I β_i's and θ^2, 41 parameters in the present case. Actually, because of a linear dependence among the β_i and θ^2, only 40 parameters are independent.

The reader may recognize that the single-index model is identical to the "market" model in financial theory, where θ^2 represents market risk, the diagonal elements of D represent nonmarket risk, and β_i measures the sensitivity of the individual security (forecaster) to the common part of the variance. To impose the single-index constraint, we have taken the forecast error of the simple average to represent the common portion of individual errors. Thus, each forecaster's error is split up into one part which is perfectly correlated with the average error, and one part which is orthogonal. The parameter β_i is given by $\mathrm{Cov}[\epsilon_i, \epsilon_A]/\mathrm{Var}[\epsilon_A]$ where ϵ_A refers to the average error.

This discussion brings out the essential similarity between our procedures for constructing a minimum mean squared error composite forecast, given the probability distribution of forecast errors, and the portfolio problem of constructing a minimum variance asset portfolio, given a distribution of security returns. Although the two problems are not identical, since we make use of characteristics of rationally formed forecasts which need not apply to securities, much of the intuition is the same in both cases.

The validity of the various constraints we have discussed can be tested using a likelihood ratio test. Under the assumption of multivariate normality the log-likelihood function is given by

$$\log(L) = -\frac{IT}{2}\log(2\pi) - \frac{T}{2}\log|\Omega| - \frac{1}{2}\sum_{t=1}^{T}(\underline{\epsilon}_t - \underline{b})'\Omega^{-1}(\underline{\epsilon}_t - \underline{b}) \qquad (6)$$

Let L_c be the maximized value of the likelihood function when constraints are imposed on $\underline{b}$ or Ω, and L_u be the maximized value without the constraints. Under the hypothesis that the constrained model is the true model, the statistic

$-2(\log(L_c) - \log(L_u))$ is distributed asymptotically as chi-squared with degrees of freedom equal to the number of constraints, which in this case is the difference in the number of independent parameters estimated in the two models.

Table I shows the values of the log-likelihoods for the various models. Consider first the unconstrained model in its uncorrected and bias corrected form. Minus twice the difference in log-likelihoods is 42.0 which is significant at the .005 level for a chi-squared distribution with 20 degrees of freedom. Thus the hypothesis that all b_i are zero is strongly rejected for the unconstrained model, and for the others as well.

Testing the constraints on the form of the Ω matrix yields the following results. Estimating 169 additional parameters increases the log-likelihood of an unconstrained model over that of a single-index model by an amount which is significant at the .001 level. But the difference is small compared to the increase in log-likelihood of the single-index over the diagonal model. The latter is not found to be significantly different from a simple average in either uncorrected or bias corrected form. We find that in most cases relaxing a constraint yields a statistically significant increase in log-likelihood, but the largest difference occurs when nonindependence among errors is allowed as one goes from the diagonal to the single-index model.

The same test can be easily applied to examine the error structure's stability over time. Each model is fitted over the whole sample period to produce the log-likelihood when the parameters are constrained to be the same throughout. Then the sample is split in half and the parameters estimated separately for the two parts. The unconstrained log-likelihood is the sum of the log-likelihoods from the subsamples. Under the joint null hypothesis that the specified form of the model represents the true error structure and that its parameters are the same in both periods, the likelihood ratio test becomes a test of stability.

Both the simple average and the diagonal models appear to be stable over time, whether or not they are corrected for bias. The models which fit covariances, however, are not completely stable, with the hypothesis of stability being rejected at the .05 level for the single index and the .001 level for the unconstrained model.

Table I

Log-Likelihood of Constrained Error Structures

	Uncorrected		Bias Corrected	
	Log (L)	Number of parameters	Log (L)	Number of parameters
Unconstrained	−2952.4	210	−2931.4	230
Single-Index	−3112.6	41	−3086.7	61
Diagonal	−4722.0	20	−4690.9	40
Mean	−4726.2	1	−4695.4	21

Note: Estimation of the single-index models uses one degree of freedom less than the number of parameters.

Table II

Post-Sample Forecasting Performance of Composite Predictions

	Mean error	Standard deviation	Root mean squared error	Full sample weights Maximum	Minimum
Uncorrected					
Mean	−.111	1.605	1.595	.050	.050
Diagonal	−.111	1.603	1.594	.060	.043
Single-Index	−.078	1.696	1.683	.345	−.237
Unconstrained	−.108	1.804	1.792	.560	−.456
Bias corrected					
Mean	.326	1.598	1.617	.050	.050
Diagonal	.326	1.596	1.616	.060	.044
Single-Index	.351	1.695	1.717	.325	−.259
Unconstrained	.305	1.784	1.795	.490	−.413

Notes: Figures represent performance over 59 post-sample periods. Model parameters were re-estimated every period using all previous data. Amounts are in $ billions. Full sample weights were calculated from model parameters estimated over all 118 periods.

IV. Forecasting Accuracy of Composite Forecasts

The fully unconstrained model with bias adjustment explains the data significantly better than any constrained model in tests based on analysis of the whole sample. However, formal stability tests showed that the single-index and unconstrained parameter estimates changed significantly over time. Since the error structure must be estimated from past data, the fitted parameters will generally be somewhat different from the true parameters for the current period. If the difference is substantial, forecasts from the unconstrained model may be less accurate than those from one with more restrictions and greater stability, despite their theoretical superiority.

The true test of a forecasting strategy is how well it performs out of sample. To look at the operational forecasting performance of the various models, we constructed series of post-sample predictions and compared them on the basis of mean forecast error, standard deviation, and root mean squared error. Beginning in Period 59, the model parameters were estimated on data from all earlier periods, the weights were constructed, and composite forecasts were calculated for the next period. The parameters were then refitted adding Period 60 to the sample. The process of fitting the models on all previous data and forecasting one period ahead was continued until post-sample forecasts for the whole second half of the data sample had been created.

Table II shows the summary statistics for the various models. As might be anticipated from previous results, the mean and diagonal models are virtually identical, with a root mean squared error of about $1.7 billion in both uncorrected and bias corrected form. The single-index model did distinctly less well, and the unconstrained model was the least accurate in the post-sample. Apparently parameter instability was sufficiently great to offset the potential increase in accuracy afforded by allowing a more detailed error structure.

This result contrasts with what was found in tests on the Livingston inflation expectations survey. While unconstrained models could not be estimated on that data set, both the diagonal and single-index models were substantially more accurate than the simple mean.[5] An important reason for the difference is that the money supply forecasters are an exceptionally homogeneous group. This limits the ability of these techniques to exploit differences in accuracy and correlation within the forecaster population.

One possibility for reducing the problem of parameter instability is to eliminate old data in the estimation. We tried this by dropping a sample point each time the models were refitted and using only 59 periods for estimation. The differences in forecast performance were small except for the unconstrained model, for which RMSE increased to 1.999 and 2.065 in uncorrected and bias corrected form, respectively.

Comparing the bias corrected and uncorrected versions of the models shows them to be very similar. It is interesting that in all cases the mean error of the uncorrected model is smaller in magnitude than that of the corrected model, though this has not much effect on root mean squared error.

One way to see the differences in the models and to understand why the single-index and unconstrained models are less accurate is to compare the weights they produce. The last two columns in Table II show the largest and the smallest weights for each model. Since the weights change when a model is refitted, we have chosen to report only those derived from parameters estimated on all sample data. These tend to be less extreme than weights from shorter estimation periods. The simple average naturally gives equal weights of 1/20 to each forecaster, and the diagonal model hardly deviates from this. The two models which fit covariances, however, allow negative weights and show a much wider spread. These models are trying to make use of the small differences in correlation among individuals to reduce the effect of the common portion of their forecast errors. Forecasts from an expert who is highly correlated with others' errors will be assigned a negative weight.[6] But to the extent that differences in correlations are small and not completely stable over time, the relatively extreme weights produced by the more detailed models can lead to less accurate post-sample forecasts.

V. Composite Forecasts and Interest Rate Movements

In this section we examine whether yields in the Treasury bill and Treasury bill futures markets fully incorporate all of the information about money supply changes which can be obtained by our composite forecasting methods. The forecast errors made by individuals in our sample are highly correlated with one another. Even though the error structure is not stable enough over time for the single-index and unconstrained models which attempt to correct for correlation to produce more accurate forecasts than the simple mean, they may still contain economically valuable information that has not been impounded in prices. This

[5] See Figlewski [3].

[6] This is like holding short positions in high beta stocks and long positions in low beta stocks to minimize returns variance in an equity portfolio.

is especially likely if investors concentrate primarily on obtaining the most accurate forecast of money growth and only pay attention to the average forecast.

If our aggregation procedures are able to capture information that the market has not fully discounted, the difference between a forecast from a composite model and the market's expectation of the money supply change should explain changes in yields immediately after the money supply is announced. Such a test is clearly impossible, however, since the market's expectation is unobservable. Instead, we base our tests on the differences between the average forecast and the single-index and unconstrained composites. Recent articles by Urich [8] and Urich and Wachtel [9] have shown that the interest rate movement after the money supply announcement is positively related to the unanticipated change in the money supply, as measured by the difference between the announced figure and the average forecast. The rationale for the direction of response is that market participants take current money growth as an indicator of future Federal Reserve policy. For example, an unanticipated increase in the money supply tends to push interest rates up because investors expect the Federal Reserve will have to tighten credit conditions in the future to meet its money growth targets.

These results suggest that the mean forecast from the survey may be a reasonable proxy for the market's expectation. However, the validity of our test does not depend on that being true. If market yields before the announcement do accurately discount all of the information contained in the set of money supply forecasts, no combination of them will have any power to explain the subsequent movement. But if the market expectation does not fully adjust for the effect of correlation in forecast errors, by looking at the difference between the mean forecast, which takes no account of correlation, and a composite forecast which does, we can obtain a proxy for the information the market price does not contain, regardless of whether the market's expectation is close to the mean forecast.

Information enters prices when investors trade on it in hopes of earning a profit, and information which is not accurately impounded will give rise to profitable trading opportunities. Thus our test of the ability of the composite models to predict interest rate changes takes the form of a trading rule. If the composite forecast predicts smaller money supply growth than the average forecast, Treasury bills or Treasury bill futures are purchased on Thursday afternoon just prior to the announcement and resold at the beginning of trading on Friday. When the composite forecast is above the average, a short position is taken before the announcement and covered on Friday morning. Profits earned by the rule indicate to what extent the composite forecast contains information that is not fully discounted in prices.

We examined the profitability of this strategy over Periods 60 through 118 using the post-sample composite forecasts which were analyzed in the previous section. Three maturities of Treasury bills and four futures contracts on 3-month Treasury bills were examined. The Treasury bills were the most recently issued three month, six month, and year bills. The futures contracts used were the nearest four International Monetary Market contracts. Thus, the first contract, labelled F1, matured within three months, the second, F2, in three to six months, and so forth. Profits were measured by the change in yield in basis points (bid,

Table III

Average Profit per Trade of Trading Rules based on Differences between Composite and Average Forecasts

Required difference	Method	Number of trades	Treasury bills			Treasury bill futures			
			3 Month	6 Month	1 Year	F1	F2	F3	F4
0	Perfect	54	3.8*	4.4*	4.6*	6.0*	6.7*	6.6*	6.2*
	Forecast		(2.69)	(3.29)	(3.11)	(3.97)	(3.63)	(3.81)	(4.26)
	Single-index	54	1.5	1.4	1.5	3.9*	4.2*	3.3*	3.1*
	Uncorrected		(.98)	(.98)	(.92)	(2.42)	(2.11)	(1.74)	(1.93)
	Unconstrained	54	1.0	.7	.9	3.1*	2.8	2.6	2.2
	Uncorrected		(.68)	(.48)	(.57)	(1.88)	(1.40)	(1.36)	(1.31)
	Single-Index	54	−.8	−.6	−.1	−.5	−.1	.4	−.2
	Bias corrected		(.56)	(.40)	(.08)	(.31)	(.06)	(.22)	(.13)
	Unconstrained	54	.2	.7	1.2	1.3	2.5	2.4	1.9
	Bias corrected		(.11)	(.50)	(.78)	(.77)	(1.21)	(1.22)	(1.14)
	Median	54	−1.8	−2.6*	−1.8	−3.1*	−3.6*	−3.0	−2.6
			(1.24)	(1.80)	(1.16)	(1.84)	(1.79)	(1.56)	(1.56)
300	Perfect	44	4.6*	5.3*	5.5*	7.2*	8.3*	8.0*	7.3*
	Forecast		(2.77)	(3.31)	(3.14)	(4.10)	(3.85)	(3.95)	(4.38)
	Single-index	29	1.5	1.3	2.0	4.0	3.7	3.5	2.1
	Uncorrected		(.63)	(.55)	(.77)	(1.59)	(1.21)	(1.17)	(.92)
	Unconstrained	35	2.4	2.1	2.7	4.9*	5.1*	4.7*	3.9*
	Uncorrected		(1.16)	(1.04)	(1.18)	(2.24)	(2.02)	(1.88)	(2.06)
	Single-Index	38	.3	.6	1.2	1.5	2.5	2.6	2.1
	Bias Corrected		(.17)	(.32)	(.53)	(.67)	(.93)	(1.00)	(.97)
	Unconstrained	41	.1	.6	1.2	2.1	3.2	2.8	2.4
	Bias Corrected		(.05)	(.31)	(.60)	(.96)	(1.23)	(1.11)	(1.18)
600	Perfect	37	5.1*	6.0*	6.2*	7.8*	9.0*	8.7*	8.1*
	Forecast		(2.65)	(3.25)	(3.06)	(3.84)	(3.58)	(3.65)	(4.14)

Single-Index	20	3.8	3.3	4.0	5.9*	5.9	5.8	4.1
Uncorrected		(1.13)	(.99)	(1.07)	(1.68)	(1.42)	(1.41)	(1.33)
Unconstrained	26	3.2	2.7	3.6	6.4*	7.5*	6.7*	5.0*
Uncorrected		(1.25)	(1.06)	(1.24)	(2.41)	(2.66)	(2.41)	(2.20)
Single-Index	25	−.3	.1	.1	2.0	2.2	2.1	1.2
Bias Corrected		(.10)	(.03)	(.17)	(.64)	(.64)	(.59)	(.47)
Unconstrained	27	.5	.2	.9	2.2	2.3	2.5	1.6
Bias Corrected		(.06)	(.06)	(.31)	(.77)	(.72)	(.78)	(.64)

Notes: Profit figures are changes in yield in basis points. *t*-statistics in parentheses.
* statistically significant at 5% level in one-tailed test.

on discount basis) from 3:30 P.M. Thursday to 10:30 A.M. Friday for the cash bills, and from the Thursday settlement price to Friday's opening price for the futures. Weeks in which Thursday or Friday was a holiday were dropped from the sample.

Table III shows the results of these tests for the single-index and unconstrained models in uncorrected and bias corrected form. In the first panel, a position is assumed to be taken in every period, while in the second and third panels we try to reduce the effects of random noise by requiring that the difference between the composite forecast and the average exceed a threshold of $300 million and $600 million, respectively, before initiating a trade. As a base to gauge how well these rules capture the interest rate change after a money supply announcement, the first line in each panel shows the profit rate of a "perfect forecast," i.e., a rule based on the difference between the actual money supply change and the mean forecast. For example, knowing the figure that would be announced and trading in the nearest futures contract according to the rule we have specified would have yielded an average of 6.0 basis points per week.

In nearly all cases, the composite models successfully predicted interest rate movements, with the largest profits being shown by the uncorrected models in the futures market. For instance, the single-index uncorrected model produced an average return of 3.9 basis points per week in the F1 futures contract (out of a possible 6.0 earned by the perfect forecast). This figure was statistically significant at the 1% level in a one-tailed test.

Contrasting the various models, markets, and strategies leads to several general conclusions. In virtually all cases, the uncorrected models were more successful than their corrected counterparts. This provides further evidence that attempting to correct the forecasts for bias is not appropriate. For the uncorrected models, requiring a threshold difference between the composite and the average reduced the number of trades, as one expects, and increased their profit rate substantially. It also improved the results for the uncorrected unconstrained model relative to the single-index, probably because the unconstrained forecasts are noisier.

A useful benchmark for evaluating the attractiveness of these trading rules is to compare them with a buy and hold strategy. With no threshold requirement, our strategies earned higher mean returns in every case. In 21 of 28 cases, the standard deviation of returns was also lower than for buy and hold. In the other 7 cases, the largest increase was only 0.1 basis point.

Profits for trades in the futures market were substantially larger than in the cash bills. Looking at the performance of the perfect forecast, we see that money supply changes appear to affect the futures market more than the cash market, but that does not account for all of the difference. Reasons for this result could include differences in the way the two markets incorporate money supply expectations from the government bond dealer community, differences in their speed of reaction to new information, or the slight difference in timing between cash and futures quotes in our data.

The trading performance of one other "composite" model, the median, is shown in Table III since the median, not the mean, is the statistic which is reported to the financial community by Money Market Services. The sizeable losses generated by the trading rule suggest that the median may not be the best summary statistic to publish.

Overall, the results in Table III indicate that the information that can be obtained from the set of money supply forecasts by comparing the single-index or unconstrained model forecasts with the mean is not fully incorporated in short term interest rates.[7] An especially significant result which is not visible in the table is that the composite forecasts exhibited a higher percentage of correct trades in periods when there was a large interest rate movement after the announcement. For example, both the single-index and unconstrained uncorrected methods produced correct decisions in 18 of 27 periods when the F2 futures contract moved 5 basis points or more; they were profitable 12 and 11 times, respectively, in 13 periods when the rate changed 10 or more basis points. Large price adjustments may well be the result of widely shared expectations errors, which these procedures are designed to correct for.

A final issue which should be raised concerns transactions costs. It is not inconsistent with a fully efficient market for prices to be slightly out of line with available information if the discrepancy is small relative to the transactions cost that would be incurred in trying to exploit it. For the small investor, both the bid-ask spread in cash bills, and the commission on a futures transaction are on the order of two to three basis points. Costs of this size would substantially reduce the attractiveness of this kind of trade. However, the costs to a securities dealer or a member of the futures exchange for the same trade are negligible. In particular, securities dealers adjust their positions in advance of the money supply announcement anyway. The marginal cost of taking a different position may well be zero. In any case, our objective in this research has not been to uncover evidence of major market inefficiency, but rather to explore how information is aggregated in a financial market. The results suggest that market prices may not accurately discount all available information when investors are subject to widely shared expectations errors.

VI. Concluding Comments

The process of price formation in a financial market involves information aggregation at several levels. In forming expectations about future prices, individual investors must combine diverse information in the form of forecasts from well-informed sources. These expectations enter investors' demand functions and are aggregated by the market into a single price which, to some extent, reflects all of the information held by the market population. This paper has explored statistical procedures for aggregating diverse subjective forecasts optimally, based on the variances, covariances, and biases in the set of forecast errors. Neither individual investors nor the market are likely to combine forecasts in the ways

[7] An alternative test of the ability of the composite forecasts to explain interest rate movements is simply to regress the change in the rate on the difference between the composite and the average forecast. When we do this, of 28 regression coefficients 15 are statistically significant at the 5% level in a one-tailed test; 24 are significant at the 10% level. These regression results support the results of Table III, from a slightly different perspective.

we suggested. Individuals are, if anything, most likely to use a simple average, which produces accurate forecasts but does not adjust for correlation among forecast errors. Financial markets aggregate investors' expectations by weighting them according to the size of their transactions. The trading rule results of Section V indicate that our composite forecasting methods could capture information about the structure of expectations within the market population which was not fully incorporated in short-term interest rates.

REFERENCES

1. Eugene Fama. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25 (May 1970), 333–417.
2. Stephen Figlewski. "Subjective Information and Market Efficiency in a Betting Market." *Journal of Political Economy* 87 (February 1979), 75–88.
3. ———. "Optimal Price Forecasting Using Survey Data." *Review of Economics and Statistics*, forthcoming (1983).
4. C. W. J. Granger and Paul Newbold. *Forecasting Economic Time Series*. New York: Academic Press, 1977.
5. J. Jaffe and R. Winkler. "Optimal Speculation Against an Efficient Market." *Journal of Finance* 31 (March 1976), 49–61.
6. Charles Nelson. "The Prediction Performance of the FRB-MIT-Penn Model of the U.S. Economy." *American Economic Review* 62 (September 1972), 902–15.
7. Douglas Pearce. "Comparing Survey and Rational Measures of Expected Inflation: Forecast Performance and Interest Rate Effects." *Journal of Money, Credit and Banking* 11 (November 1979), 447–56.
8. Thomas Urich. "The Information Content of Weekly Money Supply Announcements." *Journal of Monetary Economics* 10 (July 1982).
9. ——— and Paul Wachtel. "Market Response to the Weekly Money Supply Announcement in the 1970's." *Journal of Finance* 36 (December 1981).
10. ——— and ———. "The Structure of Expectations of the Weekly Money Supply Announcement." Manuscript (1981).