To be published in:
'Oxford Handbook' on Economic Forecasting,
Eds., M.Clements and D. Hendry, Oxford

ANALYZING THREE-DIMENSIONAL PANEL DATA OF FORECASTS

Antony Davies
Duquesne University

Kajal Lahiri*
University at Albany, SUNY

Xuguang Sheng
American University

* Corresponding author. Mail address: Department of Economics, University at Albany, SUNY, Albany, NY 12222. USA. Email: klahiri@albany.edu.

Introduction

Multi-dimensional panel data of survey forecasts predate econometric methodologies for extracting diverse macroeconomic information from these rich sources of data. The Livingston Survey (LS), instituted in 1946, the Survey of Professional Forecasters (SPF), instituted in late 1968, and the Blue Chip Economic Indicators (BCEI), instituted in August 1976, are three-dimensional panel data sets in which multiple forecasters forecast macroeconomic variables for multiple target dates and at multiple forecast horizons. The SPF and the LS have longer histories than BCEI, though their forecast panels are anonymous and the forecasts are reported relatively infrequently (quarterly for SPF and semi-annually for LS). In contrast, because BCEI forecasters are not anonymous, researchers have suggested that these forecasters have greater incentive to produce accurate forecasts. Maddala (1990), Zarnowitz and Braun (1993) and Fildes and Stekler (2002) contain reviews of studies using these data sets. Earlier attempts to analyze these data sets involved testing the rational expectations hypothesis by pooling the data or collapsing one of the three dimensions either through elimination or aggregation. Different approaches include modeling only one forecaster at a time and thus reducing the data set to the two dimensions of targets and horizons (Batchelor and Dua 1991), modeling a single horizon thereby reducing the data set to the two dimensions of forecasters and targets (Swinder and Ketcher 1990; Keane and Runkle 1990), or by averaging individual forecasts into a single consensus forecast thereby reducing the dimensions to targets and horizons (De Bont and Bange 1992). Figlewski and Wachtel (1983) point out that collapsing the individuals dimension by aggregating forecasters into a consensus can mask private information and so may result in inconsistent parameter estimates.

Davies and Lahiri (1995) developed an econometric framework for analyzing multi-dimensional panel data of forecasts. By creating a general model that described the process by which forecasts were generated and actuals were realized, they were able to show that forecast errors have two distinct components: shocks (i.e., errors that are generated external to the forecasters and that are, by definition, unpredictable) and idiosyncratic errors (i.e., errors that are generated by and specific to the individual forecasters at individual points in time). With the assumption of homoskedasticity, they constructed a covariance matrix of forecast errors involving only N+1 variance terms that would otherwise require estimation of $(NTH)(NTH+1)/2$ terms (number of forecasters x number of target dates x number of forecast horizons). The general model suggested a complex pattern to the forecast error covariance that was a function of the variance of news and the forecasters' idiosyncratic variances. Davies and Lahiri (1999) further generalized their framework by allowing for the variance of shocks to change over time (i.e., conditionally heteroskedastic).

The purpose of this chapter is to illustrate how frameworks built around multi-dimensional panel data of forecasts can be used not only to test the rational expectations hypothesis correctly, but also to study alternative expectations formation mechanisms, to distinguish anticipated from unanticipated shocks, and to distinguish forecast uncertainty from disagreement.

The General Case of Rational, Implicit, and Adaptive Expectations

Muth's (1961) traditional rational expectations framework treats the forecast for target period t, F_t, as predetermined in repeated samples thereby attributing all stochastic components to the process that generates the target variable at time t, A_t, such that (for a normally distributed error, η_t)

$$A_t = \alpha + \beta F_t + \eta_t \,. \tag{0.1}$$

Muth's test for rationality is actually a test for unbiasedness where the forecaster is found to be unbiased when $\alpha = 0$ and $\beta = 1$. Nordhaus (1987) builds on Muth by defining strong efficiency as the state in which and all information available to the forecaster at the time the forecast was made is incorporated into the forecast. Combining Muth's unbiasedness condition with Nordhaus' efficiency condition gives us the modern rational expectations model

$$A_t = \alpha + \beta F_t + \gamma X_t + \eta_t \tag{0.2}$$

where rationality, the combination of unbiasedness and efficiency, requires $\alpha = 0$, $\beta = 1$, and $\gamma = 0$. The variable X_t represents information available to the forecaster at the time the forecast was made. Whereas Muth's unbiasedness condition can be easily tested, Nordhaus' efficiency condition can only be rejected since, strictly speaking, it requires testing all information that was available to the forecaster. Finding $\gamma \neq 0$ for specific information is sufficient to reject efficiency, but finding $\gamma = 0$ for specific information is necessary but not sufficient to fail to reject efficiency. Nordhaus offers a test of weak efficiency in which the information available to the forecaster is replaced with the forecaster's past forecasts.

What distinguishes the rational expectations model from other expectations models is that the forecasts errors in the former are analyzed conditional on a given set of forecasts, implying that the variance of the target variable exceeds the variance of the forecasts. Mill's (1957) implicit expectations framework, which found many empirical applications prior to the rational expectations era, treats the target variable as fixed in repeated samples such that

$$F_t = \alpha + \beta A_t + \gamma X_t + \eta_t \tag{0.1}$$

where the stochastic component is attributed to the forecasts, see Lovell (1986). The implication here is that the variance of the forecasts exceeds the variance of the target variable.[1] Mincer's (1969) adaptive expectations framework, a special case of extrapolative expectations, models the forecast revision at horizon h as a function of the last realized forecast error such that[2]

$$F_{t+h,h} - F_{t+h,h+1} = \beta_h \left(A_t - F_{t,1} \right) + \eta_t \tag{0.2}$$

where $F_{t,h}$ is the forecast for target period t made h periods prior to the realization of the target, and $\beta = 1$ implies a forecast that fully incorporates information from the most recently realized forecast error.

The Davies-Lahiri framework assigns stochastic components to both the target variable and the forecasts. They define all shocks as unforecastable in that, by definition, shocks cannot be anticipated by rational forecasters. These shocks can occur at any point from a horizon h periods prior to the realization of the target variable at period t until the end of period t. A rational forecaster standing h periods prior to the end of period t would have available to him two types of information: the value of the target at the time the forecast is made, A_{th}^*, and the (correctly perceived) impact of information available h periods prior to the end of period t on the target variable, γ_{th}. The latter can be described as a "rationally anticipated change." Combining these two pieces of information yields the rational forecaster's forecast of the actual at the end of period t, $A_{th}^* + \gamma_{th}$. When the rational forecaster is wrong, he is so because of (unforecastable) shocks, λ_{th}, that occurred between the time at which the forecast was made and the time at which the actual was realized. The actual at the end of period t can be modeled as the actual as it

[1] It is interesting to note that the two dominant approaches to evaluate probability forecasts, viz., due to Murphy (1972) and Yates (1982), differ in a parallel fashion, See Murphy and Winkler (1992).
[2] See Pesaran and Weale (2006) for detailed descriptions of this class of models.

existed h periods prior to t plus changes rationally anticipated to occur and shocks that did occur over the period:

$$A_t = A^*_{th} + \gamma_{th} + \lambda_{th} \quad (0.3)$$

Note that the shocks can take the form of changes in the actual that were not rationally anticipated or changes in the actual that were rationally anticipated and yet did not occur.

The i^{th} rational forecaster standing at a horizon h periods prior to the end of period t, would generate a forecast, F_{ith}, for the target variable at the end of period t. By definition, all rational forecasters would generate the same forecast, $\tilde{F}_{th}$, where

$$\tilde{F}_{th} = A^*_{th} + \gamma_{th} \quad (0.4)$$

Forecasters who are not rational will generate forecasts that deviate from $\tilde{F}_{th}$ due to bias, idiosyncratic errors, heterogeneous interpretation of public information, or private information. Valchev and Davies (2009) use this framework as the basis of a behavioral model describing the interactions of bureaucrats who can control private information and politicians who attempt to forecast bureaucrats' behaviors. Lahiri and Sheng (2008, 2010b) have shown that a large part of disagreement among forecasters is due to the fact that they interpret public information differentially. Let the bias, ϕ_{ih}, vary across individuals and horizons, and the idiosyncratic errors and private information, ε_{ith}, vary across individuals, horizons, and target periods. The i^{th} forecaster standing at horizon h will generate a forecast F_{ith} for the target variable at the end of period t where

$$F_{ith} = A^*_{th} + \gamma_{th} + \phi_{ih} + \varepsilon_{ith} \quad (0.5)$$

The framework identifies two mutually orthogonal stochastic components: λ_{th} and ε_{ith}. Because one component is part of the process that generates the target variable while the other is part of the process that generates the forecast, the correct way to construct the expectations model is

$$A_t - F_{ith} = -\phi_{ih} + \lambda_{th} - \varepsilon_{ith} \quad (0.6)$$

Within this framework, the traditional rational expectations model becomes a special case wherein $\varepsilon_{ith} = 0 \ \forall \ i, t, h$. Implicit expectations becomes the special case of $\lambda_{th} = 0 \ \forall \ t, h$. Holding the target period constant, the adaptive expectations model becomes (where the forecast $F_{i,t,h+1}$ is made one period prior to the forecast F_{ith})

$$F_{i,t+h,h} - F_{i,t+h,h+1} = \beta_h \left(A_t - F_{i,t,1} \right) + \eta_{ith} \quad (0.7)$$

Solving (0.3) for A^*_{th}, plugging the result into (0.5), then substituting the resulting right hand side of (0.5) for F_{ith} in (0.7), we have:

$$\begin{aligned} &A_{t+h} - \lambda_{t+h,h} + \phi_{ih} + \varepsilon_{i,t+h,h} - \left(A_{t+h} - \lambda_{t+h,h+1} + \phi_{i,h+1} + \varepsilon_{i,t+h,h+1} \right) = \\ &\beta_h \left(A_t - \left(A_t - \lambda_{t,1} + \phi_{i,1} + \varepsilon_{i,t,1} \right) \right) + \eta_{ith} \end{aligned} \quad (0.8)$$

It can be shown by example that $\lambda_{t,1} = \lambda_{t+h,h+1} - \lambda_{t+h,h}$. Combining the idiosyncratic error terms and η_{ith} into ξ_{ith} reduces (0.8) to

$$\phi_{i,h+1} - \phi_{ih} = \beta_h \phi_{i,1} + \left(1 - \beta_h\right)\left(\lambda_{t+h,h+1} - \lambda_{t+h,h} \right) + \xi_{ith} \quad (0.9)$$

For a forecaster who fully incorporates past forecast errors into current forecasts, $\beta_h = 1$, we have

$$\phi_{i,h+1} - \phi_{ih} = \phi_{i,1} + \xi_{ith} \quad (0.10)$$

Since (0.10) holds for all h and the expected value of ξ_{ith} is zero,

$$\mathrm{E}\left(\phi_{ih}\right) = h\mathrm{E}\left(\phi_{i,1}\right) \quad (0.11)$$

This suggests that a forecaster who fully adapts under the adaptive expectations model is equivalent to a forecaster whose expected forecast bias increases linearly with the forecast horizon. For the forecaster who incorporates none of his past errors into his current forecast, $\beta_h = 0$, we have

$$\phi_{i,h+1} - \phi_{ih} = \lambda_{t+h,h+1} - \lambda_{t+h,h} + \xi_{ith} . \qquad (0.12)$$

The non-adaptive forecaster in the adaptive expectations model is equivalent to a forecaster whose bias change, ignoring the idiosyncratic error, exactly matches the shocks that occurred in the most recent period. In the general case of a partially adaptive forecaster, $0 < \beta_h < 1$, the change in the forecaster's bias is a weighted average of the most recent shocks and the shortest-horizon bias.

The Davies-Lahiri framework also suggests a test for the presence of private information. Under the assumption of rationality, shocks should be uncorrelated with idiosyncratic errors. The rational forecaster (i.e., the forecaster who is unbiased and who correctly processes all available information) generates forecast $\tilde{F}_{th}$. Combining (0.3) and (0.4), the rational forecaster's forecast error will be

$$A_t - \tilde{F}_{th} = \lambda_{th} . \qquad (0.13)$$

Compare forecaster *i*'s error shown in (0.6) to the rational forecaster's error shown in (0.13). Suppose that forecaster *i* is unbiased so that $\phi_{ih} = 0$. Dropping the subscripts, the rational forecaster's error variance is σ_λ^2, and forecaster *i*'s error variance is $\sigma_\lambda^2 + \sigma_\varepsilon^2$. Suppose that forecaster *i*'s idiosyncratic error variance were correlated with the shocks such that $\operatorname{cov}(\lambda_{th}, \varepsilon_{ith}) = \rho$. Given this correlation, forecaster *i*'s error variance would then be $\sigma_\lambda^2 + \sigma_\varepsilon^2 - 2\rho$. Now, if $\rho > \sigma_\varepsilon^2 / 2$, then forecaster *i*'s error variance would be <u>less</u> than the rational forecaster's error variance. Given that the rational forecaster has correctly incorporated all publicly available information, the only way for an unbiased forecaster *i* to obtain a forecast error variance less than that of the rational forecaster is for forecaster *i* to have access to private information.

Clements, et al. (2007) employ this framework in testing Federal Reserve forecasts of inflation, real GDP growth, and unemployment for rationality. They test the Fed's Greenbook forecasts for each forecast horizon separately and pooling all the horizons together. Interestingly, they find the forecasts to be unbiased when each horizon is tested separately, but find that the forecasts are biased when pooling the horizons and allowing biases to vary across horizons. They find that forecast revisions are correlated across horizons, implying that the Fed does not fully adjust its forecasts. The authors suggest an explanation that amounts to rational irrationality in which the Fed is motivated both to attain accuracy and to maintain the credibility. The latter can be called into question if the Fed reverses previous forecasts that were based on early data in light of data revisions. By smoothing forecast revisions, the Fed is able to avoid reversing earlier forecasts at the cost of only partially adjusting forecasts in light of the latest data.

Measuring Shocks, Volatilities, and Anticipated Changes

A multi-dimensional forecast panel provides the means to distinguish between anticipated and unanticipated changes in the forecast target as well as volatilities associated with the anticipated and unanticipated changes. This is also important in determining the correct expression for aggregate forecast uncertainty based on such panel of forecasts. Davies (2006) describes three types of shocks: cumulative shocks, cross-sectional shocks, and discrete shocks. The shocks are distinguished by when they occur and when they impact the target being forecast.

Cumulative shocks, λ_{th}, are the total unanticipated changes in the actual that occur and impact the actual over the span starting from h periods prior to the realization of the actual. Cross-sectional shocks, u_{th}, are the shocks that occur in the single period that is h periods prior to the realization of the actual and that impact the actual at any point up to the realization of the actual. Discrete shocks, v_{th}, occur in the single period that is h periods prior to the realization of the actual and impact the actual in the single period at the end of which the actual is realized. These definitions are depicted in **Figure 1** where a time line depicts quarters 4 through 10. For a forecaster standing at the beginning of quarter 6, the horizontal bracket labeled $\lambda_{9,4}$ is the span of time over which cumulative shocks ($\lambda_{9,4}$) can occur that will impact the realization of the forecast target, A_9. For a forecaster standing at the beginning of period 7, the horizontal bracket labeled $\lambda_{9,3}$ is the span of time over which cumulative shocks ($\lambda_{9,3}$) can occur that will impact the realization of the forecast target, A_9. The difference in the two, $u_{9,4}$, is the set of cross-sectional shocks occurring in quarter 6 that impact the realization of the forecast target, A_9. Notice that, there is a second measure of cross-sectional shocks occurring in quarter 6, $u_{8,3}$. These cross-sectional shocks, while *occurring* in the same period as $u_{9,4}$, impact the realization of the actual, A_8. Thus, the difference in these two cross-sectional shocks ($u_{9,4} - u_{8,3}$) represents information that *occurs* in quarter 6 but *impacts* the target in quarter 9. This difference is the set of discrete shocks, $v_{9,4}$.
[Insert Figure 1 here]

The model parameters can be estimated by assuming that the idiosyncratic errors are white noise over all three dimensions and that shocks are white noise over the two dimensions. For T target periods, averaging (0.6) over t yields

$$-\hat{\phi}_{ih} = \frac{1}{T}\sum_{t=1}^{T}\left(A_t - F_{ith}\right) \quad (0.14)$$

The first difference across the horizon dimension of (0.8) is

$$F_{ith} - F_{i,t,h-1} = -\lambda_{th} + \lambda_{t,h-1} + \phi_{ih} - \phi_{i,h-1} + \varepsilon_{ith} - \varepsilon_{i,t,h-1}\ . \tag{0.15}$$

Substituting the estimates in (0.14) into equation (0.15) and averaging over *i* yields estimates of the changes in cumulative shocks over horizons where, for *N* forecasters,

$$\hat{\lambda}_{th} - \hat{\lambda}_{t,h-1} = \frac{1}{N}\sum_{i=1}^{N}\left(-F_{ith} + F_{i,t,h-1} + \hat{\phi}_{ih} - \hat{\phi}_{i,h-1}\right)\ . \tag{0.16}$$

The differences in the cumulative shocks over horizons,

$$\hat{u}_{th} = \hat{\lambda}_{th} - \hat{\lambda}_{t,h-1} \tag{0.17}$$

are the cross-sectional shocks impacting the economy over the single period beginning *h* periods prior to the end of period *t*. From (0.17), we estimate the discrete shocks as

$$\hat{v}_{th} = \hat{u}_{th} - \hat{u}_{t-1,h-1}\ . \tag{0.18}$$

Similarly, discrete anticipated changes can be derived from cumulative anticipated changes, γ_{th}. Taking the appropriate difference in (0.5) and we have:

$$F_{ith} - F_{i,t-1,h-1} = A^{*}_{th} - A^{*}_{t-1,h-1} + \gamma_{th} - \gamma_{t-1,h-1} + \phi_{ih} - \phi_{i,h-1} + \varepsilon_{ith} - \varepsilon_{i,t-1,h-1}\ . \tag{0.19}$$

Provided that the horizon index is measured in the same units as the target index, "*h* periods prior to the end of period *t*" is the same point in time as "$h - j$ periods prior to the end of period $t - j$."[3] This means that $A^{*}_{th} = A^{*}_{t-j,h-j}\ \forall j$. Incorporating this into (0.19) causes the actuals to cancel and we have:

$$F_{ith} - F_{i,t-1,h-1} = \gamma_{th} - \gamma_{t-1,h-1} + \phi_{ih} - \phi_{i,h-1} + \varepsilon_{ith} - \varepsilon_{i,t-1,h-1}\ . \tag{0.20}$$

Estimating the forecaster biases as in (0.14) and averaging (0.20) over *i* yields estimates of the difference in cumulative anticipated changes over horizons:

[3] For example, while the Survey of Professional Forecasters measures both the horizon and target indices in quarters, the Blue Chip Economic Indicators measures the target index in years but the horizon index in months.

$$\hat{\gamma}_{th} - \hat{\gamma}_{t-1,h-1} = \frac{1}{N}\sum_{i=1}^{N}\left(F_{ith} - F_{i,t-1,h-1} - \hat{\phi}_{ih} + \hat{\phi}_{i,h-1}\right) \tag{0.21}$$

where the cumulative anticipated change, γ_{th}, is the sum of changes the rational forecaster anticipates occurring starting *h* periods prior to the end of period *t*. The first difference in the cumulative anticipated changes, the discrete anticipated change,

$$\hat{a}_{th} = \hat{\gamma}_{th} - \hat{\gamma}_{t-1,h-1} \tag{0.22}$$

is the change in the actual anticipated, from a horizon of *h* periods, to occur in period *t*.

Each of the shock measures implies a corresponding volatility measure. From the definition for discrete shocks, we have

$$\hat{v}_{th} = \frac{1}{N}\sum_{i=1}^{N}\left(-F_{ith} + F_{i,t,h-1} + F_{i,t-1,h-1} - F_{i,t-1,h-2} + \hat{\phi}_{ih} - 2\hat{\phi}_{i,h-1} + \hat{\phi}_{i,h-2}\right) \tag{0.23}$$

and

$$\hat{\text{var}}\left(v_{th}\right) = \frac{1}{N-1}\sum_{i=1}^{N}\left(-F_{ith} + F_{i,t,h-1} + F_{i,t-1,h-1} - F_{i,t-1,h-2} + \hat{\phi}_{ih} - 2\hat{\phi}_{i,h-1} + \hat{\phi}_{i,h-2} - \hat{v}_{th}\right)^2 . \tag{0.24}$$

As with the discrete shocks, we distinguish between when the volatility *occurred* and when the volatility *impacted* the target variable. Equation (0.24) shows the volatility of shocks that occurred in the single period *h* periods prior to the end of period *t* and that impact the target variable only in period *t*. Similar calculations yield the volatilities of anticipated changes,

$$\hat{\text{var}}\left(a_{th}\right) = \frac{1}{N-1}\sum_{i=1}^{N}\left(F_{ith} - F_{i,t-1,h-1} - \hat{\phi}_{ih} + \hat{\phi}_{i,h-1}\right)^2 . \tag{0.25}$$

In the past, researchers have used the variance of forecast errors as proxies for shocks. Such an approach assumes that all changes in the target variable are unanticipated. Consistent with approaches to modeling monetary shocks (Bernanke and Mihov, 1998, and Christiano et al., 1997) and trade shocks (Chang and Velasco, 2001), the Davies-Lahiri framework demonstrates that changes in the target variable might be either anticipated or unanticipated, and describes a method for separating shocks from anticipated changes.

Measuring Forecast Uncertainty

Multi-dimensional panel data sets also provide information necessary to distinguish between forecast uncertainty and disagreement. Earlier research (Levi and Makin 1979; Bomberger and Frazer 1981; Makin 1983) used the Livingston Survey in an attempt to measure uncertainty about future inflation. As a proxy for uncertainty, these studies used the dispersion of individual forecasts for a given target. The justification for this proxy is the belief that there is a high correlation between the dispersion of point forecasts across individuals and the level of market uncertainty at the same moment in time. Zarnowitz and Lambros (1987) point out that this proxy is not so much a measure of market uncertainty as it is a measure of disagreement among forecasters about expected inflation. They define the dispersion of point forecasts across forecasters as *disagreement*, and the average diffuseness of the forecasters' probability distributions about their point forecasts as *uncertainty*.

Using the ASA-NBER probability forecast data set, Zarnowitz and Lambros (1987) directly compute the forecast uncertainty for each forecaster at each point in time. Let F_{ithp} be individual i's forecast for target t made at horizon h and to which the forecast assigns probability p. The uncertainty associated with forecaster i's forecast for target t at horizon h, $\tilde{s}^2_{ith}$, is

$$\tilde{s}^2_{ith} = \sum_p p\left(F_{ithp} - F_{ith}\right)^2 \qquad (0.26)$$

where F_{ith} is mean of forecaster i's probability forecasts. They define disagreement among forecasts for target t at horizon h, s^2_{th}, as

$$s_{th}^2 = \frac{1}{N}\sum_{i=1}^{N}\left(F_{ith} - F_{\bullet th}\right)^2 \quad (0.27)$$

where $\overline{F}_{\bullet th}$ is the mean of the individual forecasts. It is the addition of the fourth dimension to the data set, the probabilities, that makes the distinction between $\tilde{s}_{ith}^2$ and s_{th}^2 possible. They find that the dispersion of point forecasts and uncertainty are correlated, but that the dispersion measure understates true uncertainty. In a broader sense, their study is noteworthy as an example of how adding an additional dimension to a data set (in their case, the additional dimension was the probabilities associated with each forecast) allows researchers to describe phenomenon with a clarity impossible to achieve without the dimension. In this sense, the additional dimension represents not merely more data, but qualitatively different data.

Based on the SPF density forecasts data, Giordani and Söderlind (2003) compare uncertainty as estimated in (0.26) to disagreement as estimated in (0.27). They find that the two measures are highly correlated so that they claim that disagreement can be used as a reasonable proxy for uncertainty. As pointed out by Bomberger (1996) and Giordani and Söderlind (2003), however, disagreement remained a theoretically unfounded measure of uncertainty.

Lahiri and Sheng (2010a) demonstrate that the Davies-Lahiri framework suggests a simple way of identifying the relationship between forecast uncertainty and disagreement. It shows that the perceived volatility of anticipated change in the target variable mediates the direct relationship between the two. Following Engle (1983), we can decompose the average squared individual forecast errors as

$$\frac{1}{N}\sum_{i=1}^{N}\left(A_t - F_{ith}\right)^2 = \left(A_t - F_{\bullet th}\right)^2 + \left(1 - \frac{1}{N}\right)d_{th}, \quad (0.28)$$

where the observed disagreement that is observable at the time forecasts are made can be written as:

$$d_{th} \equiv \frac{1}{N-1}\sum_{i=1}^{N}\left(F_{ith} - F_{\bullet th}\right)^2, \tag{0.29}$$

and $F_{\bullet th} = \frac{1}{N}\sum_{i=1}^{N} F_{ith}$. Taking expectations on both sides given all available information at time *t-h* including F_{ith} and d_{th}, we get the following conditional relationship between aggregate uncertainty, the variance of aggregate forecast errors and observed disagreement

$$U_{th} = \mathrm{E}\left(A_t - F_{\bullet th}\right)^2 + \left(1 - \frac{1}{N}\right) d_{th}. \tag{0.30}$$

The first term on the right-hand side of (0.30) can alternatively be written as:

$$\mathrm{E}\left(A_t - F_{\bullet th}\right)^2 = \frac{1}{N^2}\mathrm{E}\left[\sum_{i=1}^{N}\left(A_t - F_{ith}\right)^2\right] + \frac{1}{N^2}\mathrm{E}\left[\sum_{i=1}^{N}\sum_{j\neq i}^{N}\left(A_t - F_{ith}\right)\left(A_t - F_{jth}\right)\right]. \tag{0.31}$$

Given our framework, (0.31) can be expressed as

$$\mathrm{E}\left(A_t - F_{\bullet th}\right)^2 = \sigma_{\lambda_{th}}^2 + \frac{1}{N^2}\sum_{i=1}^{N}\sigma_{\varepsilon_{ith}}^2 . \tag{0.32}$$

Substituting (0.32) into (0.30), we obtain

$$U_{th} = \sigma_{\lambda_{th}}^2 + \frac{1}{N^2}\sum_{i=1}^{N}\sigma_{\varepsilon_{ith}}^2 + \left(1 - \frac{1}{N}\right) d_{th} . \tag{0.33}$$

For large values of *N*, the second term on the right-hand side of (0.35) will be very close to zero and can be ignored. Thus, the wedge between uncertainty and disagreement will be determined partly by the size of the forecast horizon over which the aggregate shocks accumulate – the longer is the forecast horizon the bigger will be the difference on average. It also suggests that the robustness of the relationship between the two will depend on the variability of aggregate shocks over time. In relatively stable time periods where the perceived variability of the aggregate shocks is small, whether the perceptions are correct or not, disagreement will be a good proxy for the unobservable aggregate uncertainty. In periods where the perceived volatility of the aggregate shocks is high, disagreement can become a tenuous proxy for uncertainty. This finding has important implications on how to estimate forecast uncertainty in real time and how to construct a measure of average historical uncertainty. We address each of the implications below.

To form a measure of forecast uncertainty in real time, Lahiri and Sheng (2010a) suggest that one should use the observed disagreement from the survey, d_{th} and the variance of aggregate shocks generated conditionally by a suitably specified GARCH-type model, $\hat{\sigma}^2_{\lambda_{th}}$ to estimate U_{th} as

$$\hat{U}_{th} = \hat{\sigma}^2_{\lambda_{th}} + \left(1 - \frac{1}{N}\right) d_{th}. \tag{0.36}$$

The justification is as follows. Uncertainty comes from two sources: the error components in common information and in private information. The $\hat{\sigma}^2_{\lambda_{th}}$ term captures the imprecision in common information, and d_{th} reflects the imprecision in forecasters' idiosyncratic information and diversity in forecasting models. The measure of uncertainty in (0.36) avoids the drawback of the inability to capture the heterogeneity of forecasting models in using GARCH measure of uncertainty alone. Their suggestion is supported by the findings in Batchelor and Dua (1993) and Bomberger (1996); in a comparison of ARCH and survey measures of uncertainty, these two studies concluded that the former tends to be lower than the latter, and more importantly the former is less variable over time than the latter. Thus, if one accepts survey measures as valid, the ARCH measure alone underestimates the level and the variation in uncertainty over time. Using the SPF density forecasts, Lahiri and Sheng (2010a) find that, compared to the uncertainty constructed using the squared error in the mean forecast, the uncertainty measure in (0.36) is less volatile and matches better the survey measure of uncertainty. This underscores the important point that *ex ante* uncertainty has to be generated conditionally based on the information known to survey respondents when making their forecasts, which is exactly what GARCH-type models do.

Since November 2007, the Federal Open Market Committee has released a summary of participants' views about how the current level of uncertainty compares with that seen on average in the past. This calls for the construction of an appropriate historical benchmark uncertainty. Using squared forecast errors of a number of private and government forecasters averaged over 1986-2006, Reifschneider and Tulip (2007) proposed such a measure of past forecast uncertainty. They first calculated the individual root mean squared error (RMSE) over the period and then took the average across forecasters of the individual RMSEs to obtain

$$RMSE_{RT}^{h} = \frac{1}{N}\sum_{i=1}^{N}\sqrt{\frac{1}{T}\sum_{t=1}^{T}(A_t - F_{ith})^2} \quad (0.37)$$

Note that the above measure is different from the one suggested in Lahiri and Sheng (2010a). According to Lahiri and Sheng (2010a), one should use

$$RMSE_{LS}^{h} = \sqrt{\frac{1}{TN}\sum_{t=1}^{T}\sum_{i=1}^{N}(A_t - F_{ith})^2} \quad (0.38)$$

to estimate the typical uncertainty of a randomly drawn forecaster from the sample. It is clear that the Reifschneider-Tulip measure in (0.37), like Lahiri-Sheng measure in (0.38), will have the disagreement and the squared consensus forecast error as components of uncertainty. Also, because of the averaging of squared mean forecast errors over the last twenty years, the Reifschneider-Tulip measure may not be very sensitive to occasional large forecast errors, and thus, may be a reasonable measure for the average historical uncertainty. However, by Jensen inequality, $RMSE_{RT}^{h} \le RMSE_{LS}^{h}$, where the equality takes place when there is no individual heterogeneity with respect to idiosyncratic error variance, that is, $\sigma_{\varepsilon_i}^2 = \sigma_{\varepsilon}^2$ for all *i*. This hypothesis has been overwhelmingly rejected in the studies of inflation forecasts, cf. Davies and Lahiri (1999) and Boero, et al. (2008). Thus, the uncertainty measure constructed according to (0.37) will necessarily underestimate the "true" *ex post* uncertainty.

Rationality Tests

We have shown that with a multi-dimensional forecast panel it is possible to extract estimates of shocks, anticipated changes, and volatilities. These estimates can be analyzed directly or used in constructing error covariance matrices for use in conducting rationality tests.

Bonham and Cohen (2001) show that forecast rationality tests of panel data will falsely accept unbiasedness when microhomogeneity does not hold. That is, when regression coefficients are not constant across forecasters, it is possible for individual biases to cancel each other out leaving a panel that appears unbiased in the aggregate despite being biased in the individual. The authors show that microhomogeneity does not hold for the majority of SPF forecasts and so conclude that tests for unbiasedness should only be carried out for the forecasters individually or for the panel of forecasters using seemingly unrelated regression. There is no reason to assume that, similarly, microhomogeneity holds for other panel data sets of survey forecasts. Therefore, Bonham and Cohen's results underline the need to avoid both collapsing the individuals dimension by using consensus forecasts, and constraining regression parameters to be constant across individuals in panel data sets.

Keane and Runkle's (1990) attempt to analyze the SPF data set is noteworthy for their use of generalized method of moments. Using SPF data, they estimate the rational expectations model:

$$A_t = \alpha + \beta F_{it} + \gamma X_{it} + \varepsilon_{it} \tag{0.39}$$

where the i^{th} individual's forecast for the target date t is F_{it}, and the actual at time t is A_t. X_{it} is information available to forecaster i at the time he made his forecast, and ε_{it} is noise. Under the rational expectations hypothesis, the forecasts are unbiased (i.e., $\alpha = 0$, $\beta = 1$) and efficient (i.e., $\gamma = 0$). To reduce the three-dimensional SPF data set to two-dimensions, Keane and Runkle used only the single forecast horizon that was closest to the realization of the actual.

In a departure from forecast rationality research up to that time, Keane and Runkle claimed to have found no evidence of irrationality, but their results are suspect for several reasons in addition to their choice to evaluate only the nearest forecast horizon. Bonham and Cohen (1995) point out that the results of their analysis are invalid due to unaddressed non-stationarity. Further, since information available to the forecaster, X_{it}, is predetermined but not strictly exogenous, a GMM (or GLS) estimation of (0.39) yields inconsistent parameter estimates when there are individual specific dummy variables on the right hand side because the regression becomes equivalent to a regression on demeaned variables. The demeaned $\bar{X}_{i\bullet}$ are functions of future values of X_{it} and the demeaned errors likewise are functions of future errors. Because past innovations can affect future information, the error and the regressor in the demeaned regression will be contemporaneously correlated. Keane and Runkle attempt to sidestep this problem by assuming a common bias for all forecasters thereby avoiding the use of individual-specific dummies. However, including a constant term causes the same cotemporaneous correlation between the error and the regressor as does including individual specific dummies, therefore Keane and Runkle's model suffers from the very problem they attempt to avoid. In addition, the assumption of a common bias can mask individual forecaster biases. If some forecasters exhibit positive biases while others exhibit negative biases, assuming a common bias can cause the forecasters to appear to be unbiased in the aggregate despite the fact that they are biased in the individual.

Although they did not analyze their data in three-dimensions, Keane and Runkle did describe a rudimentary error covariance matrix for a three-dimensional analysis. Lacking an underlying model describing how the actuals and forecasts are generated and how the effects of shocks accumulate over horizons, their error covariance matrix was neither complete nor reduced to the minimal number of parameters. However, it did provide the first glimpse into the complexity of forecast evaluation in multi-dimensional data.

Given that stochastic components appear in both the actual and the forecasts, the correct formulation for a rationality test is

$$A_t - F_{ith} = -\phi_{ih} + \lambda_{th} - \varepsilon_{ith} \tag{0.40}$$

where the ϕ_{ih} are fixed effects to be tested, and the λ_{th} and ε_{ith} are components of the error term. From the definition of λ_{th} in (0.18), the assumption that cross-sectional shocks are independent over both dimensions, and that the idiosyncratic shocks are independent over all three dimensions, it can be shown that the error covariance matrix, $\mathbf{\Sigma}$, takes the form

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{B} & \mathbf{B} & \dots & \mathbf{B} & \mathbf{B} \\ \mathbf{B} & \mathbf{A}_2 & \mathbf{B} & \dots & \mathbf{B} & \mathbf{B} \\ \vdots & & & & & \\ \mathbf{B} & \mathbf{B} & \mathbf{B} & \dots & \mathbf{B} & \mathbf{A}_N \end{bmatrix}, \ \mathbf{A}_i = \sigma_{\varepsilon_i}^2 \mathbf{I} + \mathbf{B} \tag{0.41}$$

where $\sigma_{\varepsilon_i}^2$ is the variance of the idiosyncratic error for forecaster i, and $\mathbf{I}$ is a TH x TH identity matrix. The matrix $\mathbf{A}_i$ contains the covariance of error terms across targets and horizons for forecaster i, and the matrix $\mathbf{B}$ contains the covariance of error terms across targets and horizons and any two forecasters. Matrix $\mathbf{B}$ is comprised of component matrices, $\mathbf{b}_{t_1,t_2}^j$, where each H x H

component matrix represents the error covariance for different forecasters, across targets t_1 and t_2 (where $j = |t_1 - t_2|$), and across the horizons.

$$\mathbf{B} = \begin{bmatrix} \mathbf{b}_1^0 & \mathbf{b}_{1,2}^1 & \mathbf{b}_{1,3}^2 & \mathbf{b}_{1,4}^3 & \cdots & \mathbf{b}_{1,T}^{T-1} \\ \mathbf{b}_{1,2}^{1\,'} & \mathbf{b}_2^0 & \mathbf{b}_{2,3}^1 & \mathbf{b}_{2,4}^2 & \cdots & \mathbf{b}_{2,T}^{T-1} \\ \mathbf{b}_{1,3}^{2\,'} & \mathbf{b}_{2,3}^{1\,'} & \mathbf{b}_3^0 & \mathbf{b}_{3,4}^1 & \cdots & \mathbf{b}_{3,T}^{T-1} \\ \mathbf{b}_{1,4}^{3\,'} & \mathbf{b}_{2,4}^{2\,'} & \mathbf{b}_{3,4}^{1\,'} & \mathbf{b}_4^0 & \cdots & \mathbf{b}_{4,T}^{T-1} \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ \mathbf{b}_{1,T}^{T-1'} & \mathbf{b}_{2,T}^{T-1'} & \mathbf{b}_{3,T}^{T-1'} & \mathbf{b}_{4,T}^{T-1'} & \cdots & \mathbf{b}_T^0 \end{bmatrix} \tag{0.42}$$

The pattern in the elements of $\mathbf{b}_{t_1,t_2}^j$ is determined by the forecast panel being analyzed. For the Survey of Professional Forecasters (SPF) data set (**Figure 2**), in each quarter, individuals generate forecasts for the last quarter, the current quarter, and each of the next four quarters.[4] For the Blue Chip Economic Indicators (BCEI) data set (**Figure 3**), in each month, individuals generate forecasts for the current year and the next year. In both figures, the arrows indicate the points in time at which the indicated forecasts are made. The horizontal brackets show the ranges of time over which cumulative shocks occur that affect the various forecasts.

[Insert Figure 2]

[Insert Figure 3]

Decomposing the cumulative shocks into cross-sectional shocks, for the SPF panel, we have:

$$\operatorname{cov}\left(\lambda_{t_1,h_1}, \lambda_{t_2,h_2}\right) = \operatorname{cov}\left(\sum_{j_1=-1}^{h_1} u_{t_1,j_1}, \sum_{j_2=-1}^{h_2} u_{t_2,j_2} \right) \tag{0.43}$$

$$\operatorname{cov}\left(u_{t_1,h_1}, u_{t_2,h_2}\right) = \begin{cases} \sigma_{u_{t_1,h_1}}^2 = \sigma_{u_{t_2,h_2}}^2 & \forall\ t_1 - h_1 = t_2 - h_2 \\ 0 \text{ otherwise} \end{cases} \tag{0.44}$$

For the BCEI panel, we have:

[4] See Lahiri and Liu (2005, 2006) for the detailed analysis of the SPF density forecast data.

$$\mathrm{cov}\left(\lambda_{t_1,h_1}, \lambda_{t_2,h_2}\right) = \mathrm{cov}\left(\sum_{j_1=1}^{h_1} u_{t_1,j_1},\ \sum_{j_2=1}^{h_2} u_{t_2,j_2}\right) \tag{0.45}$$

$$\mathrm{cov}\left(u_{t_1,h_1}, u_{t_2,h_2}\right) = \begin{cases} \sigma^2_{u_{t_1,h_1}} = \sigma^2_{u_{t_2,h_2}} & \forall\ t_2 = t_1 + 1 \text{ and } h_2 = h_1 + 12 \\ 0 \text{ otherwise} \end{cases} . \tag{0.46}$$

Holding the target period constant, as the horizon increases, the variances of the cross-sectional shocks accumulate. We can use this fact to construct the H x H covariance matrix of forecast errors for different forecasters, the same target and across the horizons. Let m be the shortest forecast horizon and M be the longest forecast horizon such that $M - m + 1 = H$. For example, for the SPF $m = -1$ and $M = 4$; for the BCEI, $m = 1$ and $M = 24$. For both the SPF and BCEI data sets,

$$\mathbf{b}_t^0 = \begin{bmatrix} \sum_{h=m}^{M} \sigma^2_{u_{t,h}} & \sum_{h=m}^{M-1} \sigma^2_{u_{t,h}} & \sum_{h=m}^{M-2} \sigma^2_{u_{t,h}} & \cdots & \sigma^2_{u_{t,m}} \\ \sum_{h=m}^{M-1} \sigma^2_{u_{t,h}} & \sum_{h=m}^{M-1} \sigma^2_{u_{t,h}} & \sum_{h=m}^{M-2} \sigma^2_{u_{t,h}} & \cdots & \sigma^2_{u_{t,m}} \\ \sum_{h=m}^{M-2} \sigma^2_{u_{t,h}} & \sum_{h=m}^{M-2} \sigma^2_{u_{t,h}} & \sum_{h=m}^{M-2} \sigma^2_{u_{t,h}} & \cdots & \sigma^2_{u_{t,m}} \\ \vdots & \vdots & \vdots & & \vdots \\ \sigma^2_{u_{t,m}} & \sigma^2_{u_{t,m}} & \sigma^2_{u_{t,m}} & \cdots & \sigma^2_{u_{t,m}} \end{bmatrix} \tag{0.47}$$

The covariance of forecast errors can be different depending on the panel, and are determined by examining the structure of the forecasts as shown in **Figure 2** and **Figure 3**. Based on the structure of the BCEI forecast panel, we know that forecast errors will be correlated (depending on the forecast horizons) for targets separated by up to two periods.[5] For the BCEI panel, we have the following H x H covariance matrix describing the covariance of forecast errors for different forecasters, across adjacent target periods, and across the horizons:

[5] For the BCEI panel, targets are measured in years while horizons are measured in months.

$$\mathbf{b}_{t_1,t_2}^{1}=\begin{bmatrix}
\sum_{h=1}^{12}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{11}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{10}\sigma_{u_{s,h}}^2 & \cdots & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 \\
\sum_{h=1}^{12}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{11}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{10}\sigma_{u_{s,h}}^2 & \cdots & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 \\
\vdots & \vdots & \vdots & & \vdots \\
\sum_{h=1}^{12}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{11}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{10}\sigma_{u_{s,h}}^2 & \cdots & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 \\
\sum_{h=1}^{11}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{11}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{10}\sigma_{u_{s,h}}^2 & \cdots & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 \\
\sum_{h=1}^{10}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{10}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{10}\sigma_{u_{s,h}}^2 & \cdots & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 \\
\vdots & \vdots & \vdots & & \\
\sum_{h=1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 & \cdots & \sum_{h=1}^{1}\sigma_{u_{s,h}}^2 \\
0 & 0 & 0 & & 0 \\
\vdots & \vdots & \vdots & & \vdots \\
0 & 0 & 0 & & 0
\end{bmatrix},\ s=\min\left(t_1,t_2\right)\text{ and } t_2-t_1=1 \qquad (0.48)$$

For the BCEI panel, there is no error covariance (under rationality) when targets are separated by more than one period. Therefore, $\mathbf{b}_{t_1,t_2}^{j}=\mathbf{0}\ \forall\, j>1$.

Based on the structure of the SPF forecast panel, we know that forecast errors will be correlated (depending on the forecast horizons) for targets separated by up to five quarters. Corresponding to each of the five degrees of separation, we have the following *H* x *H* covariance matrices:

$$\mathbf{b}_{t_1,t_2}^{1}=\begin{bmatrix}
\sum_{h=-1}^{3}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{2}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 \\
\sum_{h=-1}^{3}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{2}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 \\
\sum_{h=-1}^{2}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{2}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 \\
\sum_{h=-1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 \\
\sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 \\
\sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0
\end{bmatrix},\ s=\min\left(t_1,t_2\right)\text{ and } \left|t_1-t_2\right|=1 \qquad (0.49)$$

$$\mathbf{b}_{t_1,t_2}^{2}=\begin{bmatrix} \sum_{h=-1}^{2}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 \\ \sum_{h=-1}^{2}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 \\ \sum_{h=-1}^{2}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 \\ \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 \\ \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 \end{bmatrix},\ s=\min\left(t_1,t_2\right) \text{ and } \left|t_1-t_2\right|=2 \qquad (0.50)$$

$$\mathbf{b}_{t_1,t_2}^{3}=\begin{bmatrix} \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 & 0 \\ \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 & 0 \\ \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 & 0 \\ \sum_{h=-1}^{1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 & 0 \\ \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{0}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^{2} & 0 & 0 & 0 \end{bmatrix},\ s=\min\left(t_1,t_2\right) \text{ and } \left|t_1-t_2\right|=3 \qquad (0.51)$$

$$\mathbf{b}_{t_1,t_2}^{4} = \begin{bmatrix} \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{0}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad s = \min\left(t_1, t_2\right) \text{ and } \left|t_1 - t_2\right| = 4 \qquad (0.52)$$

$$\mathbf{b}_{t_1,t_2}^{5} = \begin{bmatrix} \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 & 0 \\ \sum_{h=-1}^{-1}\sigma_{u_{s,h}}^2 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad s = \min\left(t_1, t_2\right) \text{ and } \left|t_1 - t_2\right| = 5 \qquad (0.53)$$

By making use of the structure of the forecast panel, the *NTH* x *NTH* error covariance matrix in (0.41) can be constructed from *N* + *TH* parameter estimates. These covariance matrices then can be used to test rationality in a Generalized Method of Moments framework.

Conclusion

Now there exist a number of very rich panel data sets that record forecasts made by professional forecasters collected at alternative frequencies, for multiple horizons, and with rolling and fixed targets. The forecasts are typically for a wide array of macroeconomic and financial variables. For example, in the U.S. the Livingston data have been available since 1946, SPF data since 1968 and the Blue Chip surveys since 1976. The European Central Bank is conducting a SPF-type survey since early 1990s. More interestingly, Consensus Economics Inc. has been collecting macroeconomic forecasts on a large number of countries since October 1989.[6] With the proliferation of quality multi-dimensional surveys, it becomes increasingly important for researchers to employ an econometric framework in which these data can be properly analyzed and put to their maximum use.

In this chapter we have summarized such a framework developed in Davies and Lahiri (1995, 1999), and illustrated some of the uses of these multi-dimensional panel data. In particular, we have characterized the adaptive expectations mechanism in the context of broader rational and implicit expectations hypotheses, and suggested ways of testing one hypothesis over the others. We find that, under the adaptive expectations model, a forecaster who fully adapts to new information is equivalent to a forecaster whose forecast bias increases linearly with the forecast horizon. A multi-dimensional forecast panel also provides the means to distinguish between anticipated and unanticipated changes in the forecast target as well as volatilities associated with the anticipated and unanticipated changes. We show that a proper identification of anticipated changes and their perceived volatilities are critical to the correct understanding and estimation of forecast uncertainty. In the absence of such rich forecast data, researchers have typically used the variance of forecast errors as proxies for shocks. It is the perceived volatility of the anticipated change and not the (subsequently-observed) volatility of the target variable or the unanticipated change that should condition forecast uncertainty. This is because forecast uncertainty is formed when a forecast is made, and hence anything that was unknown to the

[6] See Isiklar et al. (2006) for an extension of the Davies-Lahiri framework to a multi-country context.

forecaster when the forecast was made should not be a factor in determining forecast uncertainty. This finding has important implications on how to estimate forecast uncertainty in real time and how to construct a measure of average historical uncertainty, cf. Lahiri and Sheng (2010a). Finally, we show how the Rational Expectations hypothesis should be tested by constructing an appropriate variance-covariance matrix of the forecast errors when a specific type of multi-dimensional panel data is available.

References

Batchelor, R. and P. Dua, 1993. Survey vs. ARCH measures of inflation uncertainty. *Oxford Bulletin of Economics and Statistics* 55: 341-353.

Batchelor, R. and P. Dua, 1991. Blue Chip rationality tests. *Journal of Money, Credit, and Banking* 23: 692-705.

Bernanke, B. and I. Mihov, 1998. Measuring monetary policy. *Quarterly Journal of Economics* 113: 869-902.

Boero, G., J. Smith, and K.F. Wallis, 2008. Evaluating a three-dimensional panel of point forecasts: the Bank of England. *International Journal of Forecasting* 24: 354-367.

Bomberger W.A., 1996. Disagreement as a measure of uncertainty. *Journal of Money, Credit and Banking* 28: 381-392.

Bomberger, W. and W. Frazer, 1981. Interest rates, uncertainty, and the Livingston data. *Journal of Finance* 36: 661-675.

Bonham, C. and R. Cohen, 1995. Testing the rationality of price forecasts: comment. *American Economic Review* 85: 284-289.

Bonham, C. and R. Cohen, 2001. To aggregate, pool, or neither: testing the rational expectations hypothesis using survey data. *Journal of Business and Economic Statistics* 19: 278-291.

Chang, R. and A. Velasco, 2001. Monetary policy in a dollarized economy where balance sheets matter. *Journal of Development Economics* 812: 445-464.

Christiano, L., M. Eichenbaum, and C. Evans, 1997. Sticky price and limited participation models of money: a comparison. *European Economic Review* 41: 1201-1249.

Clements, M.P., F. Joutz, and H.O. Stekler, 2007. An evaluation of the forecasts of the Federal Reserve: a pooled approach. *Journal of Applied Econometrics* 22: 121-136.

Davies, A., 2006. A framework for decomposing shocks and measuring volatilities derived from multi-dimensional panel data of survey forecasts. *International Journal of Forecasting* 22: 373-393.

Davies, A. and K. Lahiri, 1995. A new framework for testing rationality and measuring aggregate shocks using panel data. *Journal of Econometrics* 68: 205-227.

Davies, A. and K. Lahiri, 1999. Re-examining the rational expectations hypothesis using panel data on multi-period forecasts. In Hsaio, C., K. Lahiri, L.F. Lee, and M.H. Pesaran (eds.), *Analysis of Panels and Limited Dependent Variable Models*, 226-254. Cambridge University Press.

De Bont, W.F.M. and M.M. Bange, 1992. Inflation forecast errors and time variation in term premia. *Journal of Financial and Quantitative Analysis* 27: 479-496.

Engle, R.F., 1983. Estimates of the variance of U.S. inflation based upon the ARCH model. *Journal of Money, Credit and Banking* 15: 286-301.

Figlewski, S.F. and P. Wachtel, 1981. The formation of inflationary expectations. *Review of Economics and Statistics* 63: 1-10.

Fildes, R. and H. Stekler, 2002. The state of macroeconomic forecasting, *Journal of Macroeconomics* 24: 435-468.

Giordani, P. and P. Söderlind, 2003. Inflation forecast uncertainty. *European Economic Review* 74: 1037-1060.

Isiklar, G., K. Lahiri, and P. Loungani, 2006. How quickly do forecasters incorporate news? Evidence from cross-country surveys. *Journal of Applied Econometrics* 6: 703-725.

Keane, M.P. and D.E. Runkle, 1990. Testing the rationality of price forecasts: new evidence from panel data. *American Economic Review* 80: 714-735.

Lahiri, K. and F. Liu, 2005. ARCH models for multi-period forecast uncertainty: a reality check using a panel of density forecasts. *Advances in Econometrics* 20: 321-363.

Lahiri, K. and F. Liu, 2006. Modeling multi-period inflation uncertainty using a panel of density forecasts. *Journal of Applied Econometrics* 21: 1199-1219.

Lahiri, K. and X. Sheng, 2008. Evolution of forecast disagreement in a Bayesian learning model. *Journal of Econometrics* 144: 325-340.

Lahiri, K. and X. Sheng, 2010a. Measuring forecast uncertainty by disagreement: The missing link. *Journal of Applied Econometrics*, forthcoming.

Lahiri, K., and X. Sheng, 2010b. Learning and heterogeneity in GDP and inflation forecasts. *International Journal of Forecasting* 26: 265-292.

Levi, M. and J. Makin, 1979. Fisher, Phillips, Friedman and the measured impact of inflation on interest. *Journal of Finance* 34: 35-52.

Maddala, G.S., 1990. Survey data on expectations: what have we learnt? In M. Nerlove (ed.), *Issues in Contemporary Economics, Macroeconomics and Econometrics*, 2: 319-344, New York.

Makin, J., 1983. Real interest, money surprises, anticipated inflation and fiscal deficits. *Review of Economics and Statistics* 65: 374-384.

Mills, E.S., 1957. The theory of inventory decisions. *Econometrica* 25: 222-238.

Mincer, J.A., 1969. Models of adaptive forecasting. In J. Mincer (ed.), *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, 6: 83-111, New York.

Murphy, A.H., 1972. Scalar and vector partitions of the probability score: part I. two-state situation. *Journal of Applied Meteorology* 11: 273-282.

Murphy, A.H. and R.L. Winkler, 1992. Diagnostic verification of probability forecasts. *International Journal of Forecasting* 7: 435-435.

Muth, J.F., 1961. Rational expectations and the theory of price movements. *Econometrica* 29: 315-355.

Nordhaus, W.D., 1987. Forecasting efficiency: concepts and applications. *Review of Economics and Statistics* 69: 667-674.

Pesaran, M.H. and M. Weale, 2006. Survey expectations. In Elliot, G., C.W.J. Granger, and A. Timmermann (eds.), *Handbook of Economic Forecasting*: 715-776, North-Holland.

Reifschneider, D. and P. Tulip, 2007. Gauging the uncertainty of the economic outlook from historical forecasting errors. *Federal Reserve Board, Finance and Economics, Discussion Series*: 60.

Swindler, S., and D. Ketcher, 1990. Economic Forecasts, rationality, and the processing of new information over time. *Journal of Money, Credit, and Banking* 22: 65-76.

Valchev, R. and A. Davies, 2009. Transparency, performance, and agency budgets: a rational expectations modeling approach. *George Washington University Research Program on Forecasting Working Paper*, 2009-004.

Yates, J.F., 1982. External correspondence: decompositions of the mean probability score. *Organizational Behavior and Human Performance* 30: 132-156.

Zarnowitz, V. and P. Braun, 1993. Twenty-two years of the NBER-ASA quarterly economic outlook surveys: aspects and comparisons of forecast performance. In Stock, J.H. and M.W. Watson (eds.), *Business Cycles, Indicators, and Forecasting*: 11-84, University of Chicago Press.

Zarnowitz, V. and L.A. Lambros, 1987. Consensus and uncertainty in economic prediction. *Journal of Political Economy* 95: 591-621.

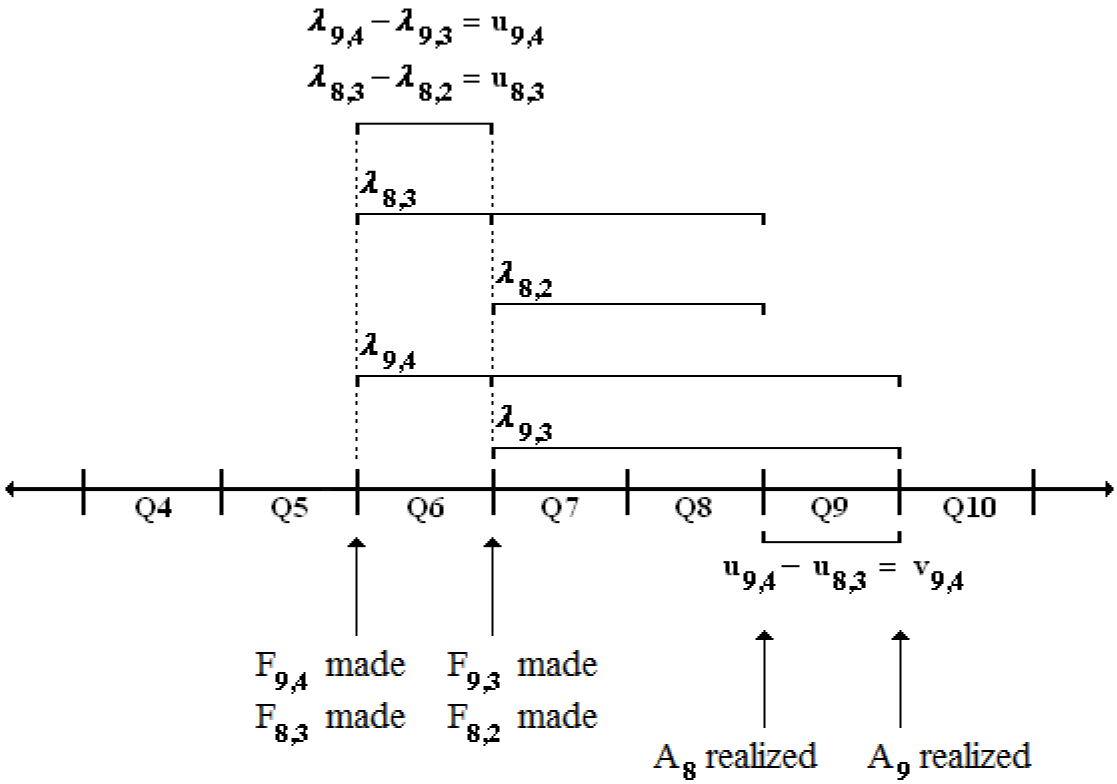

Figure 1. Definitions of Shocks

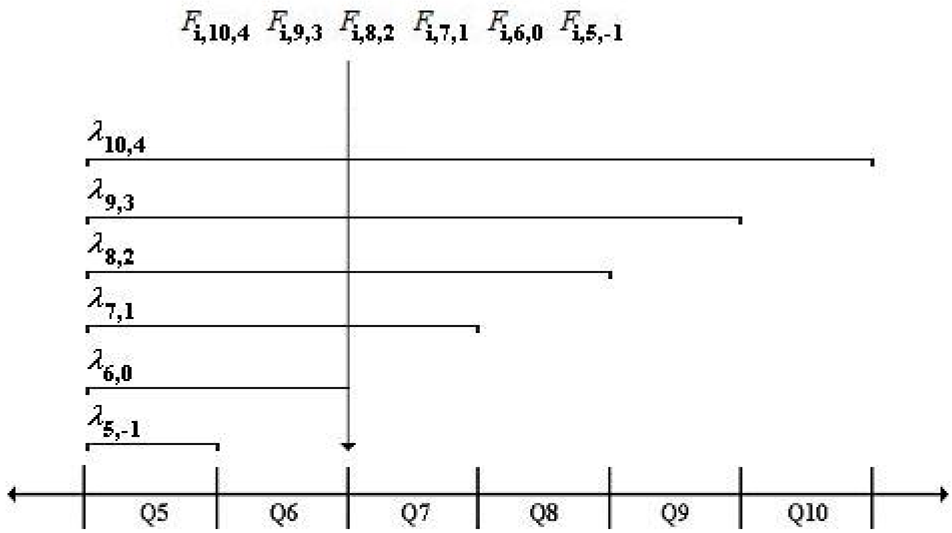

Figure 2. Structure of the Survey of Professional Forecasters Panel

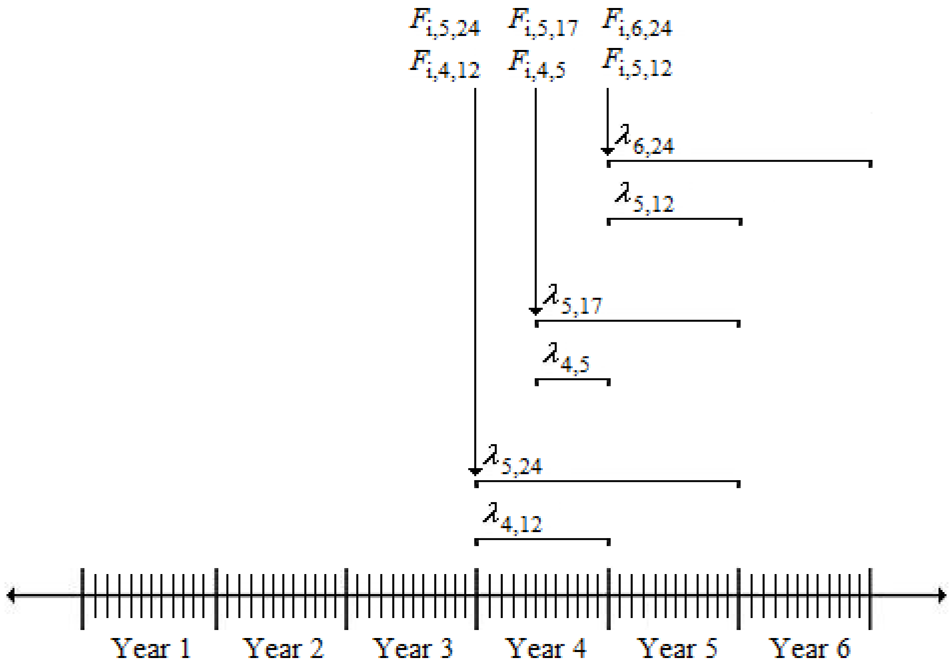

Figure 3. Structure of the Blue Chip Economic Indicators Panel