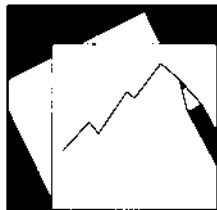


Working Paper

INTERNATIONAL MONETARY FUND



IMF Working Paper

Incorporating Market Information into the Construction of the Fan Chart

Selim Elekdag and Prakash Kannan

IMF Working Paper

Research Department

Incorporating Market Information into the Construction of the Fan Chart

Prepared by Selim Elekdag and Prakash Kannan¹

Authorized for distribution by Jörg Decressin

August 2009

Abstract

This Working Paper should not be reported as representing the views of the IMF.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

This paper develops a simple procedure for incorporating market-based information into the construction of fan charts. Using the International Monetary Fund (IMF)'s global growth forecast as a working example, the paper goes through the theoretical and practical considerations of this new approach. The resulting spreadsheet, which implements the approach, is available upon request from the authors.

JEL Classification Numbers: C82, E31, E37, E59

Keywords: Fan Chart, Balance of Risks, Option Prices, Consensus Forecasts

Author's E-Mail Address: Selim.Elekdag@tcmb.gov.tr, pkannan@imf.org

¹ The authors would thank Charles Collins, Jörg Decressin and Marco Terrones for insightful comments. Murad Omoev provided valuable research assistance. The views expressed in this paper are those of the authors and should not be associated with any of the institutions the authors are affiliated with.

I. Introduction	3
II. Outline of Theory	4
A. Characterizing the Distribution of Global Growth	5
B. Constructing Confidence Intervals	8
III. Using Survey- and Market-Based Information	9
A. Survey-based Measures	10
B. Market-based Measures	10
IV. An Example: Forecasting Global Growth	11
A. Choice of Risk Factors	11
B. Estimating the Weighting Parameters	12
C. Constructing the Fan Chart	13
D. Interpreting the Results	14
V. Conclusions	16
References	20
Box	
1. The Two-Piece Normal Distribution	6
Tables	
1. Estimated Elasticities and Skewness Coefficients	12
Figures	
1. Constructing Confidence Intervals	9
2. Dispersion of Forecasts for GDP and Selected Risk Factors	15
3. Fan Chart for Global Growth and Skewness of Risk Factors	18
4. Fan Chart for Global Growth Based on Direct Estimates of Variance and Skew	19

I. INTRODUCTION

Economic forecasts are inherently uncertain. While this is widely recognized, policymakers often face difficulties in communicating to the broader public the extent of the uncertainty involved. Often times, this difficulty stems from the fact that the forecasts are not based on a single model, but rather on a suite of models that are combined with elements of judgment. As such, these forecasts do not have standard confidence intervals associated with them.

The fan chart has emerged as a popular approach that primarily serves as a visual communication device that encourages the reader to recognize the degree of uncertainty surrounding a given baseline forecast. Fan charts have gained prominence through their use in inflation reports of central banks—of which the Bank of England features prominently as a pioneer—but could be readily adapted to other economic indicators as well (see Britton, Fisher, and Whitley, 1998).² Fan charts serve to address the following three questions:

What is the baseline forecast for the current and future years?

What level of uncertainty surrounds the forecast?

Where does the balance of risks lie?

In this paper, we present a modification to the standard methodology used to produce fan charts by incorporating market- and survey-based information regarding the uncertainty surrounding the inputs upon which forecasts are based. The motivation for this approach can be illustrated by considering the case of the International Monetary Fund (IMF)'s forecasts of global growth that are published semi-annually in the *World Economic Outlook* (WEO)—a case that we will continually use as an example throughout this paper. The baseline WEO forecast for global growth is based on an aggregation of individual country forecasts (with appropriate checks on global consistency), each of which relies on a variety of time-series and structural models. The outlook regarding several key global variables, however, serve as inputs to the individual forecasts. Examples of two such variables are oil prices and global interest rate developments. Variations in the outlook to these two variables can have significant impact on the forecasts made by individual country desks, which in turn lead to changes in the forecasts for global growth. The methodology presented in this paper allows us to incorporate market indicators on the balance of risks associated with these two variables (as well as others) into the construction of the fan chart, which then serves as a useful basis upon which the risks to the global growth forecast can be assessed.

The paper is structured as follows: Section II provides a self-contained overview of the theoretical foundation for the construction of the fan chart. Section III then presents our two main contributions: First, we describe how market- and survey-based information can be incorporated to better inform the balance of risks associated with a baseline forecast. The goal here is to encourage a more objective analysis at the outset of the forecasting process

² See Tay and Wallis (2000) for a survey of the use of fan charts.

before the incorporation of judgment associated with the forecasts. Second, we develop a user-friendly spreadsheet that constructs a fan chart from user-provided inputs. One benefit of this streamlined spreadsheet is the relative ease in which the user can understand the underlying assumptions that generate a non-symmetric fan chart.³ In Section IV, we provide an example that implements the new methodology drawing on the case of the WEO forecasts of global GDP growth, as mentioned earlier. This section discusses the inputs into the fan chart, its construction, and the interpretation of the results. The last section offers concluding remarks.

II. OUTLINE OF THEORY

We first begin by discussing the main elements of the underlying theory behind the construction of the fan chart. The framework is primarily based on John (1982) and Blix and Sellin (1997).

We will anchor the discussion around the example of the WEO global growth forecasting exercise purely for pedagogical reasons. Let's assume that global growth, Y , is a function of a potentially very large set of factors, Θ . In order to keep the analysis tractable, we focus our attention on a limited set of factors, $\theta \in \Theta$. We will call this set of variables "risk factors" for reasons that will be explained later.

In order to construct uncertainty bands around the baseline forecasts for Y , we first need to make assumptions regarding the underlying distribution for both global growth, Y , and the set of factors that we focus our attention on, θ . A convenient assumption is that both global growth and the factors θ are drawn from a two-piece normal distribution function (*TPN*).⁴ The two-piece normal distribution is widely used by central banks in the construction of fan charts for their inflation forecasts. It has the benefit of having a simple-to-compute density function along with the ability to incorporate asymmetries. Asymmetry is a feature that is necessary in order for the fan chart to convey to the reader where the balance of risks lies.

The two-piece normal distribution is a three-parameter distribution, which can be reparametrized in such a way that they capture the mode, variance, and skewness of the distribution. Specifically, each factor θ_i is assumed to have the following distribution:

$$\theta_i \sim TPN(\mu_i, c_{1,i}, c_{2,i}) \quad (1)$$

³ The spreadsheet is available upon request from the authors. Please direct your requests to the corresponding author, Prakash Kannan <pkannan@imf.org>.

⁴ Note that we need to make separate assumptions for the distribution of global growth and for the risk factors. Even if the function that maps elements in Θ to Y is linear, it is not necessarily the case that the process for global growth will follow the distribution of the individual elements of Θ .

where μ_i is the mode, and $\sigma_{1,i}$ and $\sigma_{2,i}$ are parameters that have a unique mapping to the variance and skew of the distribution. These three parameters convey views regarding the baseline forecast, the uncertainty surrounding the forecast, and the balance of risks, respectively. Box 1 presents an overview of the two-piece normal distribution.

The focus in the construction of the fan chart will be the variance and skew parameters of the respective factors. The mode, or central tendency, is, of course, the most important parameter. However, in most practical applications, the mode of the forecast for the main variable of interest should, in principle, already incorporate changes in the mode of the risk factors, θ . In the example discussed in the introduction, the baseline forecasts for oil prices and global interest rates (which we take as the mode of the distribution of these two factors) are already incorporated in the forecast for global growth. The construction of the fan chart, therefore, has to do with the higher moments of these individual factors.

A. Characterizing the Distribution of Global Growth

The process of determining the distribution of global growth involves two steps: (i) estimating the variance and skew parameters of the risk factors, and (ii) mapping these into the variance and skew of global growth. In Section III, we will show how one can obtain estimates of the variance and skew parameters of the distribution of the individual θ_i 's (which we label $\sigma^2_{\theta_i}$ and γ_{θ_i} , respectively) from both survey-based and market-based sources. In the remainder of this section, we will focus on the second step of the process, which is to map the variance and skew parameters of the risk factors to the equivalent counterparts in the distribution for global growth, which we label σ^2_Y and γ_Y , respectively.

We start with the skew parameter. In general, γ_Y will be a function of the individual skew coefficients as well as a vector of weights associated with each factor, which we label β , that captures the relative importance of the specific factor for global growth:

$$\gamma_Y = \varphi(\beta, \gamma_{\theta_i}) \quad (2)$$

Unfortunately, the two-piece normal distribution does not lend itself naturally to a multivariate characterization unlike the normal distribution. Even if we assume a linear relationship between global growth and the risk factors, the resulting distribution for global growth will not be *TPN*. To overcome this problem, we follow Blix and Sellin (1997) and assume a linear relationship between the skew coefficient for global growth and the risk factors where the weights β_i reflect the contribution of factor i to global growth:

$$\gamma_Y = \sum_i \beta_i \gamma_{\theta_i} \quad (3)$$

We now turn to the variance parameter. As is the case for the skew parameter for the distribution of global growth, the variance should also, in principal, be affected by the variance of the risk factors. However, we run into the same theoretical constraints as stated before requiring, once again, assumptions on the relationship between the variance-covariance matrix of the risk factors and the variance of global growth. We approach the

Box 1. The Two-Piece Normal Distribution

The two-piece normal distribution has been widely used in the literature on fan charts due to its asymmetric shape along with the relative ease in computing the cumulative density function. The early introduction and use of the distribution is briefly discussed in Johnson, Kotz, and Balakrishnan (1994). This box summarizes some of the key features of the distribution, drawing primarily from John (1982).

The density function of the two-piece normal can be thought of as a combination of two half-normal distributions, both with the same mean μ but with differing standard deviations, σ_1 and σ_2 respectively. The density function for the two-piece normal is therefore

$$f(x) = A \exp\left\{-\frac{(x-\mu)^2}{2\sigma_1^2}\right\} \quad \text{for } x \leq \mu \quad (1.1)$$

$$A \exp\left\{-\frac{(x-\mu)^2}{2\sigma_2^2}\right\} \quad \text{for } x > \mu$$

where $A = \sqrt{2}(\sigma_1 + \sigma_2)^{-1} / \sqrt{\pi}$ is a constant of proportionality introduced to ensure that the distribution is continuous and integrates to one. In the general case, when $\sigma_1 \neq \sigma_2$, the parameter μ will be the mode of the distribution. If instead $\sigma_1 = \sigma_2$, then the distribution collapses to the normal distribution with mean μ and standard deviation σ ($= \sigma_1 = \sigma_2$). The mean, variance and skew of the distribution are given by the following equations:

$$E(x) = \mu + k(c_2 - c_1)$$

$$V(x) = \sigma_1\sigma_2 + (1 - k^2)(\sigma_2 - \sigma_1)^2 \quad (1.2)$$

$$\gamma(x) = k(\sigma_2 - \sigma_1) \left[(2k^2 - 1)(\sigma_2 - \sigma_1)^2 + \sigma_1\sigma_2 \right]$$

where $E(x)$, $V(x)$, and $\gamma(x)$ represent the mean, variance, and skew of the distribution respectively, and $k = (2/\pi)^{1/2}$. Since the variance and skew of the distribution are uniquely identified by σ_1 and σ_2 , we can reparameterize the distribution using the mode, the variance, and skew. As in Blix and Sellin (1997), we will use

$$\gamma(x) = k(c_2 - c_1)$$

as a proxy for the measure of skew. This measure is positively related to the true measure of skew, is equal to the difference between the mean and the mode of the distribution, and, most importantly, can be combined with the expression for the variance to get closed-form expressions for σ_1 and σ_2 .

For our purposes, it is also useful to state the relationship between the area under the density between any two points, L_1 and L_2 , and the cumulative distribution function of the standard normal, $\Phi(\cdot)$:

$$\int_{L_1}^{L_2} f(x) dx = 2(\sigma_1 + \sigma_2)^{-1} \sigma \left[\Phi\left(\frac{L_2 - \mu}{\sigma}\right) - \Phi\left(\frac{L_1 - \mu}{\sigma}\right) \right] \quad (1.3)$$

where $\sigma = \sigma_1$ if $L_1 < L_2 < \mu$, and $\sigma = \sigma_2$ if $\mu < L_1 < L_2$.

computation of the variance parameter slightly differently by starting with the assumption that the variance is proportional to some baseline measure of forecast uncertainty:

$$\sigma_Y^2 = \phi \bar{\sigma}_Y^2 \quad (4)$$

A reasonable measure that can be used for $\bar{\sigma}_Y^2$ is the variance of historical forecast errors—in this case, of global growth. The coefficient ϕ then captures how the current assessment of the forecast variance compares with that of the historical forecast errors. We can use the information embedded in the distribution of the risk factors to compute ϕ :

$$\phi = \frac{\sum_i \sum_j \beta_i \beta_j \sigma_{ij}}{\sum_i \sum_j \beta_i \beta_j \bar{\sigma}_{ij}}$$

where σ_{ij} is the covariance between risk factor θ_i and θ_j (with σ_{ii} being the variance of factor i). $\bar{\sigma}_{ij}$ is some measure of historical covariance between these two risk factors. By defining ϕ in this way, the estimate of uncertainty associated with global growth can be higher or lower than the historical forecast error depending on whether the forecast variance of the linear combination of risk factors is higher or lower than its own historical level.

In some instances, we may have direct measures of the skew and variance of the forecast for the variable of interest either due to model-based simulations, traded securities, or a survey of forecasters. In this case, we can directly utilize the estimates of the skew and variance in characterizing the distribution of forecasts, and hence, constructing the fan chart. In Section IV, we provide examples using risk factors, as well as direct estimates of the skew and variance of global growth forecasts.

Putting it all together

Once we have the skew and variance parameters for global growth, equations (3) and (4), we can start characterizing its distribution in terms of the parameters of the two-piece normal, μ_i , $\sigma_{1,i}$ and $\sigma_{2,i}$. From the equations that determine the first three moments of the two-piece normal (see discussion in Box 1), we can show that $\sigma_{1,Y}$ is the solution to the following quadratic equation:

$$\sigma_{1,Y}^2 + b\sigma_{1,Y} + c = 0 \quad (5)$$

where $b = \gamma_Y / k$ and $c = -\left[\left(1 - 1/k^2\right)\gamma_Y^2 + \sigma_Y^2\right]$, with $k = \sqrt{(2/\pi)}$. Typically, one of the roots will be negative while the other will be positive. The general algorithm that one can use is to

select the highest real-valued solution to equation (5).⁵ Once we have $\sigma_{1,Y}$, we can directly compute $\sigma_{2,Y}$ from the simplified expression for the skew parameter (again, see discussion in Box 1).

$$\gamma_Y = k(c_{2,Y} - c_{1,Y}). \quad (6)$$

B. Constructing Confidence Intervals

With the three parameters of the two-piece normal distribution (the mode plus the two parameters that govern the dispersion of the distribution which we computed in the previous subsection), we have a full characterization of the distribution for global growth. Since the distribution is asymmetric, there isn't a unique way to compute confidence intervals around the central tendency of the distribution. One possible approach, is depicted graphically in Figure 1 (illustrating an arbitrary distribution that is skewed to the left). The approach requires solving for two values, z_1 and z_2 that satisfy the following two constraints:

$$\int_{z_1}^{\mu} f(Y)dY + \int_{\mu}^{z_2} f(Y)dY = q \quad (7)$$

and

$$f(z_1) = f(z_2) \quad (8)$$

where $f(\cdot)$ is the distribution function for global growth, and q is the appropriate confidence level required. The first constraint says that the areas to the left and right of the mode should sum up to the confidence level specified. The second constraint pins down the value of z_1 and z_2 , by requiring that the value of the distribution function evaluated at both points should be the same.

We can use the expression for the *TPN* distribution, as given by equation (1.1) in Box 1, to rewrite the second constraint as:

$$A \exp\left[-\frac{1}{2\sigma_{1,Y}^2}(z_1 - \mu)^2\right] = A \exp\left[-\frac{1}{2\sigma_{2,Y}^2}(z_2 - \mu)^2\right] \quad (9)$$

which implies that

$$z_1 = \mu - \frac{\sigma_1}{\sigma_2}(z_2 - \mu) \quad (10)$$

⁵ Note that the restriction to real-valued solutions implies that there is a restriction on a particular combination of values for the skew and variance of Y , namely $b^2 - 4c > 0$. In practice, this is typically satisfied, but the researcher should be wary of occasions where this does not hold.

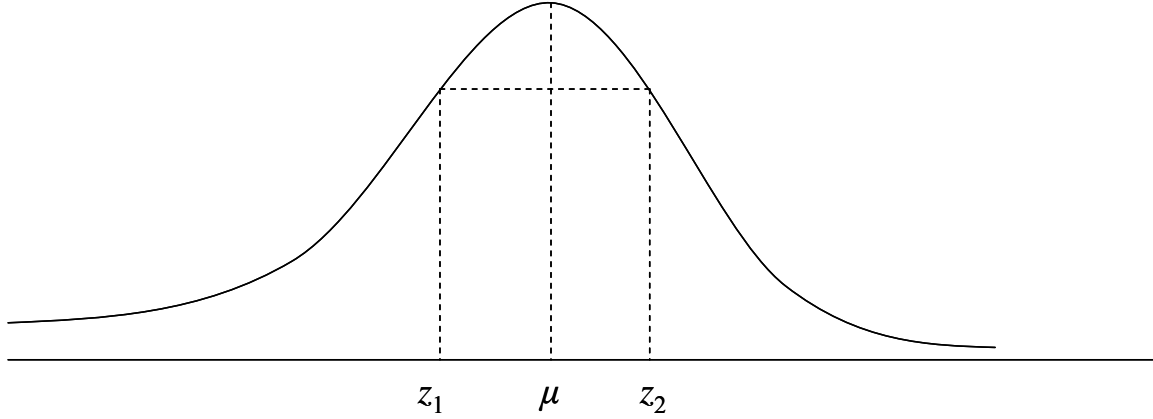


Figure 1: Constructing confidence intervals

Using equation (10), along with equation (1.3) in Box 1, we can rewrite equation (7) as:

$$\frac{2\sigma_{1,Y}}{\sigma_{1,Y} + \sigma_{2,Y}} \left[\frac{1}{2} - \Phi \left(\frac{\mu - z_2}{\sigma_{2,Y}} \right) \right] + \frac{2\sigma_{2,Y}}{\sigma_{1,Y} + \sigma_{2,Y}} \left[\Phi \left(\frac{z_2 - \mu}{\sigma_{2,Y}} \right) - \frac{1}{2} \right] = q \quad (11)$$

The expression above lends itself to a closed-form solution for the value of z_2 :

$$z_2 = \mu + \sigma_2 \Phi^{-1} \left(\frac{1+q}{2} \right) \quad (12)$$

where Φ^{-1} is the inverse of the standard normal distribution, which has a domain of $[0,1]$ and a range over the entire real line. Using the value of z_2 from equation (12) above, we can solve for z_1 using equation (10), and thus have the threshold values for the q -th percent confidence interval. The ability to obtain closed-form solutions in constructing the confidence intervals greatly reduces the computational burden, as approximations of a two-piece normal distribution function are not required.

III. USING SURVEY- AND MARKET-BASED INFORMATION

We now turn to methods related to the measurement of the skew and variance parameters of the risk factors. As one of the main contributions of this paper, we propose a couple of mutually reinforcing methods that provide information on the balance of risks to the baseline forecast arising from the risk factors. The first part of the section discusses the most straightforward method—using survey-based data—while the second part of the section shows how financial data—in particular, option prices—can be used to gauge the uncertainty associated with the forecast under consideration.

A. Survey-based Measures

The first method relies on survey-based data. Forecast surveys are regularly compiled by Consensus Economics and are also available on Bloomberg terminals. Consensus Economics surveys over 25 institutions each month for their forecasts of several key macroeconomic variables for the advanced economies and leading emerging markets over a two-year horizon. For most of these variables, the entire distribution of forecasts is published, which allows one to compute not just the mean, but also the higher moments associated with these forecasts. The sample variance provides an indicator of the level of uncertainty underlying the risk factor, while the sample skew gives an indication of the balance of risks associated with the forecast of the risk factor.

The use of survey-based measures are useful as some risk factors do not have active markets that are directly associated with them. In addition, such data is usually easy to obtain. It is important to note, however, that disagreement amongst survey respondents is a fundamentally different concept than the underlying uncertainty surrounding a particular variable. Their use can still be justified, however, insofar as the information signals from which analysts base their forecasts on are correlated with the underlying variable that they are trying to forecast. Abstracting from the precise cause, an increase in dispersion reflects heightened uncertainty which likely contains valuable information. For example, several recent studies, such as Prati and Sbracia (2002) and Kannan and Köhler-Geib (2009), show that the dispersion of analysts' forecasts is a significant predictor of financial crises.

B. Market-based Measures

Beyond the use of survey-based data, developments in financial econometrics provide tools that allow us to extract the risk-neutral density of an underlying variable of interest that is implicit in the market prices of options on this variable. The intuition behind the approach comes from the observation that options on the same underlying asset can be combined to form a portfolio whose returns are dependent on a particular realization of the "state" of nature. The price of this new portfolio will reflect the implied probability of such a state occurring. Therefore, by looking at the prices of call options at different strike prices, we can extract the entire distribution of the underlying state.⁶ There are several ways in which a researcher can extract the risk-neutral probability density function from option prices. Bahra (1997) is a good survey that covers both the theoretical basis for the methodologies along with some useful applications.

Options on a wide range of variables are traded in exchanges all around the world, with the markets for options on equities and currencies being particularly deep. Data on the prices at which options were traded are typically available from Bloomberg or some of the organized exchanges such as the Chicago Board Options Exchange (CBOE). One practical

⁶ The theoretical underpinning for most of these methods can be traced to Breeden and Litzenberger (1978) who show that the second derivative of the call price function with respect to the strike price is proportional to the risk-neutral density of the underlying asset.

complication that researchers will face is that different methodologies used to extract the underlying distribution function will often times yield different results. To a certain extent, this is unavoidable, as researchers may be restricted to using a particular approach due to data limitations or other computational considerations. However, the results using the same methodology over time should still provide some information on the evolution of the distribution, conditional on that particular methodology.

IV. AN EXAMPLE: FORECASTING GLOBAL GROWTH

With the conceptual framework in place, we now turn to implementation. As in the previous section, we use the IMF's WEO global growth forecasts as an example. Global growth forecasts are published alongside the WEO publication twice annually—once in April and once in October. In each issue, forecasts for global growth for the current and following year are published, though medium-term forecasts are also available. In what follows, we will base our example on a hypothetical construction of the October 2008 fan chart for global growth. The baseline projection for global growth was 3.9 percent for 2008 and 3.0 percent for 2009.⁷

A. Choice of Risk Factors

Three sets of macroeconomic variables are considered to represent key quantifiable risk factors associated with global growth prospects. Survey or options price data for these variables are then used to construct their one- and two-year ahead probability distributions. The variance and skew of these distributions together with the relationship between these variables and global real GDP growth are then used to build the confidence intervals around the WEO projections for global real GDP growth.

The three sets of variables that we have chosen as relevant risk factors cover: (i) financial conditions; (ii) oil price risk; and (iii) inflation risk. Financial conditions are proxied by the term spread (measured as the long-term minus the short-term interest rate) and the returns of the Standard and Poor's (S&P) 500 index. Financial market data are naturally forward-looking and so can convey useful information regarding growth prospects. Increased asset price volatility, for example, is a sign of heightened uncertainty, and will likely be associated with less favorable growth developments. The slope of the yield curve has been a reliable predictor of recessions as it embeds expectations of future monetary policy and inflation, which in turn are informative about future growth prospects (see Estrella and Mishkin, 1996). As a result, the risk of a decrease in the slope of the term spread is indicative of downside

⁷ The PPP-weighted average of the mean Consensus forecast made during the same period for GDP growth in the G7, Brazil, Mexico, China and India (which account for about 63 percent of global GDP) was 3.5 percent for 2008 and 2.5 for 2009.

Table 1. Estimated Elasticities and Skewness Coefficients

	<u>Elasticities</u>	<u>Skewness coefficient</u>			
		<u>Current</u>		<u>Next Year</u>	
		<u>2008</u>		<u>2009</u>	
		April	October A	April	October
β	γ	γ	γ	γ	
Term Spread	0.35	-0.22	-0.60	N.A.	-0.29
S&P 500	0.15	-0.46	0.26	0.04	-0.37
Inflation Risk	-0.40	0.43	0.16	0.08	-0.29
Oil Market Risks	-0.35	0.52	0.43	0.45	0.71
Global growth		0.04	-0.18	-0.33	-0.12

Source: Authors' calculations.

risk.⁸ Meanwhile, the oil price risk factor captures the risks associated with the baseline projection for oil prices, which serves as a key input to individual country growth projections. Finally, inflation risk is characterized by high or volatile price dynamics, which may trigger aggressive monetary tightening, thereby potentially depressing growth.

The inflation forecasts compiled by Consensus Economics for the United States, the euro area, Japan and several key emerging markets (Brazil, Russia, India, China, and Mexico) were used to provide information for global inflation risk. The calculations for the term spread and oil price risk factors are done in an analogous manner.⁹ In the case of the term spread, however, only data on the slope of the yield curves in the United States, the United Kingdom, Japan, and Germany (and thus, the euro area) are used.¹⁰ Finally, the balance of risks associated with the equity market risk factor are obtained by estimating the distribution function of equity returns implicit in call option data on the S&P 500 index.¹¹

⁸ Alternatively, a composite indicator of financial conditions could be used, such as the Financial Stress Index (FSI) proposed by Cardarelli, Elekdag, and Lall (2009).

⁹ The balance of risks associated with oil prices can also be obtained from options on WTI futures. Cheng (2009) develops a stable methodology using a mixture of lognormal distributions based on Bahra (1997).

¹⁰ While the rest of the forecasts were obtained from Consensus forecasts, the distribution of oil price forecasts was obtained from Bloomberg.

¹¹ The nonparametric constrained estimator introduced in Ait-Sahalia and Duarte (2003) was used to estimate the risk-neutral density of the S&P 500 returns. We use prices for December 2008 call options for Bloomberg for two cross-sections of the data: April and October. In order to obtain values for the 2009 distribution, we use the March 2009 call option prices for the April cross-section and the September 2009 call option prices for October. We did not use December 2009 options as these are thinly traded options with very minimal variation in the strike prices.

B. Estimating the Weighting Parameters

Now that the risk factors have been determined, we need to quantify how the chosen risk factors have historically affected global growth. As a first pass, we use annual data on global real GDP growth, oil prices (log deviations from a linear trend), world inflation, and S&P 500 returns from 1970 to 2007. After standardizing these series, we simply regress global growth against each of the risk factors (and lagged global growth). To assess sensitivity, a few different specifications are used, and in the end, we use the elasticities shown in the first column of Table 1.

These elasticities quantify how a risk factor would affect global growth. Naturally, these estimates can be refined. Estimated general equilibrium models, VARs, or panel specifications are all suitable candidate frameworks to refine these elasticities. At this stage, however, we use standard regressions for two main reasons. First, regression analysis is familiar to a wide audience, and the results are easy to produce and replicate. Second, the estimates used were robust to several specifications, and ultimately, judgment will be combined to determine the end result.

C. Constructing the Fan Chart

Using risk factors

The first step that needs to be taken is to estimate the variance and skew of the individual risk factors, which we have labeled $\sigma^2_{\theta_i}$ and γ_{θ_i} , respectively. When using survey-based forecasts, these will simply be the variance and skew of the distribution of forecasts. A convenient measure of the sample skew is given by a variant of Pearson's skewness coefficient, which is (3 times) the difference between the mean and the median divided by the standard deviation.¹² For market-based measures, the higher moments can be computed directly from the derived distribution. The resulting skewness coefficients for the April and October vintages of forecasts are shown in the first four rows of Table 1. As we assume a linear relationship between the skewness of global growth and the skewness of the risk factors, we then use equation (3) to compute the skew of global growth with the estimated elasticities serving as weights.

The next step is to construct estimates for the variance of global growth. The variance parameter is assumed to be a function of some baseline measure of uncertainty, with the variance-covariance matrix of the risk factors (and their respective weights) playing an amplification or dampening role (equation (4)). The baseline measure of uncertainty that we use is the variance of historical WEO forecast errors. There are two sets of forecasts of global growth published in the October WEO—current year and next year. The variance of the

¹² See Weisstein, Eric W. "Pearson's Skewness Coefficients." From *MathWorld*--A Wolfram Web Resource. <http://mathworld.wolfram.com/PearsonsSkewnessCoefficients.html>

forecast errors over the period 1990-2006 is 0.45 and 0.95 percentage points, respectively, for the current year estimate and the projection for the following year. From these forecasts error variances, we then compute an estimate of the variance of global growth forecasts using equations (4) and the expression for ϕ .

Once we have estimates of the variance and skew parameters of the distribution of global growth, we can proceed to construct the confidence intervals. Equations (5) and (6) are first used to back out the two parameters of the two-piece normal distribution (apart from the mode), $\sigma_{1,Y}$ and $\sigma_{2,Y}$. From there, the 50, 70, and 90 percent confidence intervals can be obtained in a straightforward manner from equation (12) and (10).

Using direct estimates of the variance and skew of global growth

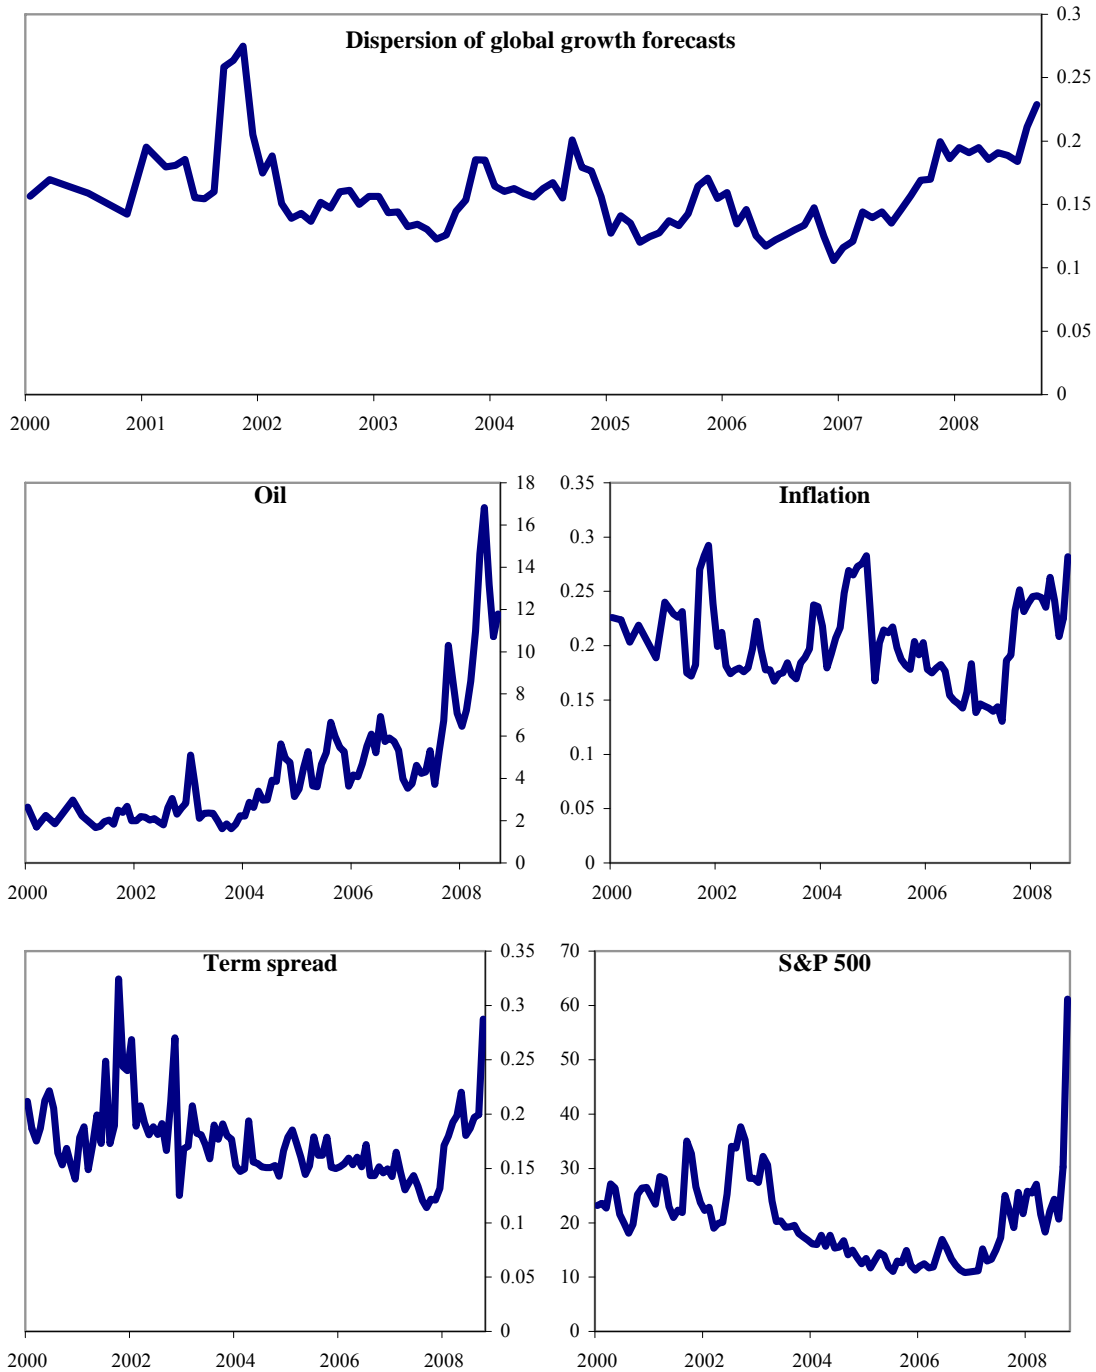
An alternative way of obtaining estimates of the variance and skew of global growth forecasts is through an aggregation of real GDP forecasts for individual countries. To do this, we compute the variance and skew of the forecasts for growth in the US, Japan, Germany, UK, France, Canada, Italy, Brazil, Mexico, China, and India—comprising 63 percent of world GDP in PPP terms—from Consensus Economics. Both these moments are then aggregated using PPP shares as the respective weights. The resulting skewness coefficients are reported in the last line of Table 1. As for the measure of variance, the deviation of the variance of the aggregate growth forecasts from their recent historical average is used as the amplification/dampening mechanism for the historical WEO forecasts error variance using equation (4). The time series for the aggregate variance is shown in Figure 2.

D. Interpreting the Results

The results of the example laid out in this section are shown graphically in Figures 2 through 4. What follows is our interpretation of what the survey- and market-based data reflect regarding the balance of risks to the central global growth projection. When interpreting the results, it is important to emphasize that these forecasts were based on available information before the release date of the WEO in October 2008.

The distribution of forecasts for GDP growth in key economies, as well as the identified risk factors, display much higher dispersion in October 2008 (the last observation used in this paper) relative to recent periods, indicating a larger degree of uncertainty associated with the baseline projection than has historically been the case (Figure 2). In the construction of the fan chart, the increase in the dispersion of growth forecasts and risk factors, relative to the average over the last ten years, is translated into a higher variance in the distribution of global growth projections by augmenting the historical current- and next-year forecast errors proportionately.

Figure 2. Dispersion of Forecasts for GDP and Selected Risk Factors 1/



Source: Consensus Economics, Bloomberg, CBOE and Staff estimates.

1/ The series for GDP and inflation measure the dispersion of GDP and inflation forecasts respectively for the G-7 economies, Brazil, India, China and Mexico. The series for term spread measures the dispersion of forecasts of the term spread for the United States, the United Kingdom, Germany and Japan. The oil price series measures the dispersion of one-year ahead oil forecasts. Finally, the series for equity risk is the VIX series which measures the implied volatility of the S&P 500.

Figure 3 shows the fan chart for global growth using information from the identified risk factors. The heightened uncertainty associated with the large dispersion of forecasts associated with the risk factors translate to a widening of the distribution of forecasts. If the distribution were solely based on historical WEO forecast errors (and is symmetric normal), the 90-percent confidence interval for 2009 global growth would have ranged from 1.4 percent to 4.6 percent. Instead, the heightened uncertainty increased that confidence interval to range from -1.8 percent to 7.8 percent (again, assuming symmetry).

The evolution of risk factors from April to October 2008 seems to be in line with the initial periods of a turning point in the business cycle. The survey- and market-based data show an increase in downside risks owing to financial conditions for 2009. Data from the S&P 500, however, show a small upside risk for the rest of 2008. Our interpretation of this result is that with equity prices dropping to very low levels during this period, the markets' expectations were for the risks to be only to the upside from there.

For inflation, the data indicate that risks for 2008 abated in October relative to April. Persistent inflationary pressures are a concern because they reduce the room to maneuver policy in response to downturns. However, the combination of rising slack and decreasing commodity prices seems to have contained the pace of price increases. As for oil market risks, with oil prices dropping from close to 150 dollars per barrel to under 70 dollars per barrel during this period, downside risks to global growth from this risk factor had surely abated. However, surveys regarding future oil prices see risks of increasing oil prices. The interpretation here is analogous to the dynamics of equity prices: because the market experienced such a sharp drop (which would have been incorporated into the baseline), survey participants are factoring in a rebound in the other direction.

Figure 4, instead, shows the fan chart for global growth using direct estimates of the skew and variance of the distribution. Compared to the dispersion of forecasts based on the risk factors, the forecast variance derived from the aggregation of growth forecasts feature a much tighter distribution. The 90-percent confidence interval for the forecast for global growth in 2009 ranges from 0.6 percent to 5.2 percent. The skewness coefficient associated with global growth shows some interesting patterns. In April, analysts were expecting some upside risks to growth in 2008, but a sharp downside risk for growth in 2009. This reading changed substantially in October. While analysts now saw downside risk for both years, the downside risk was more prominent for 2008 than for 2009.

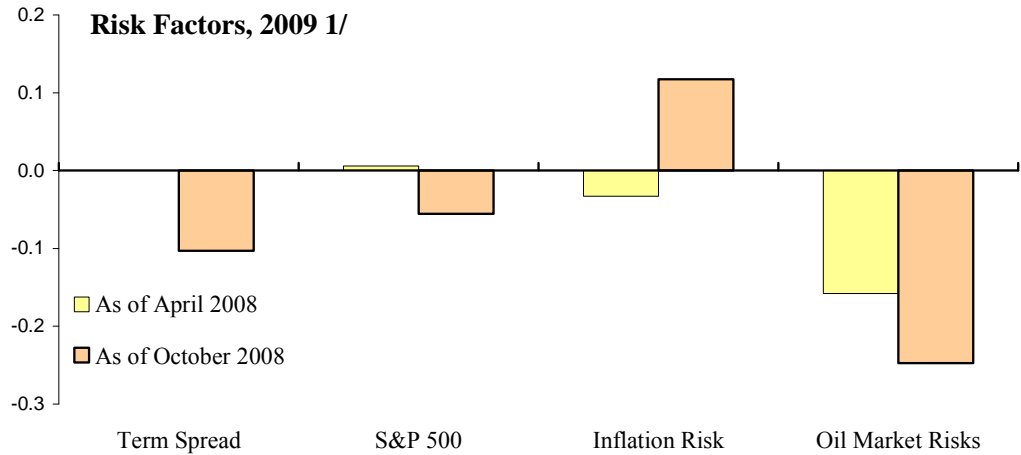
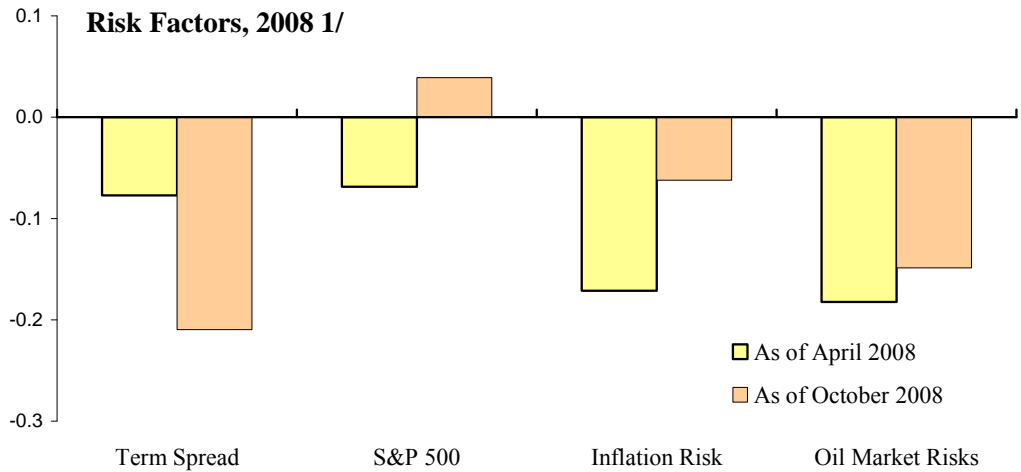
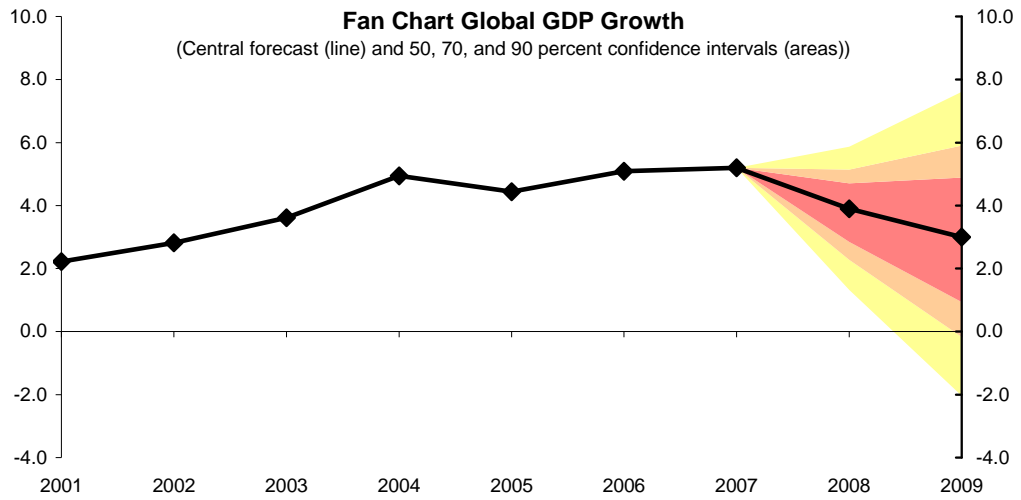
The choice between the two approaches used here to construct the fan chart—one using risk factors and the other based on direct estimates of the variance and skew—depends largely on the availability of the data. The risk factor approach is useful for the incorporation of market-based indicators, such as options, if such information is available. Additionally, if the forecast for the variable of interest depends vitally on a few key factors, then it may be worth getting more information about the higher moments of the forecasts for these factors to explicitly show the risks associated with them. The approach based on direct estimation of the skew and variance, on the other hand, is easily implementable and does not have large data requirements. For example, analysts who are interested in producing a fan chart for the growth forecasts of a single country can just use the variance of the historical forecast error

and the skew of the distribution of forecasts based on Consensus forecasts to produce a fan chart (assuming that the distribution takes the two-piece normal form).

V. CONCLUSION

The incorporation of market indicators into the construction of the fan chart represents a move towards having an objective analysis as a starting point to gauge the balance of risk and the level of uncertainty inherent in any forecasting exercise. The advantages of survey- and market-based data is that they are available at high frequencies and are inherently forward-looking, and could therefore inform policymakers on the evolution of risks as perceived by the markets. From this starting point, however, a layer of judgment can (and perhaps should) subsequently be introduced in order to incorporate other important risk factors that do not lend themselves to be easily quantified. At the end of the day, it is always useful to acknowledge that the fan chart primarily serves as a communication device. The methods introduced in this paper go some way towards achieving this purpose by showing how future uncertainty is related to several quantifiable risks.

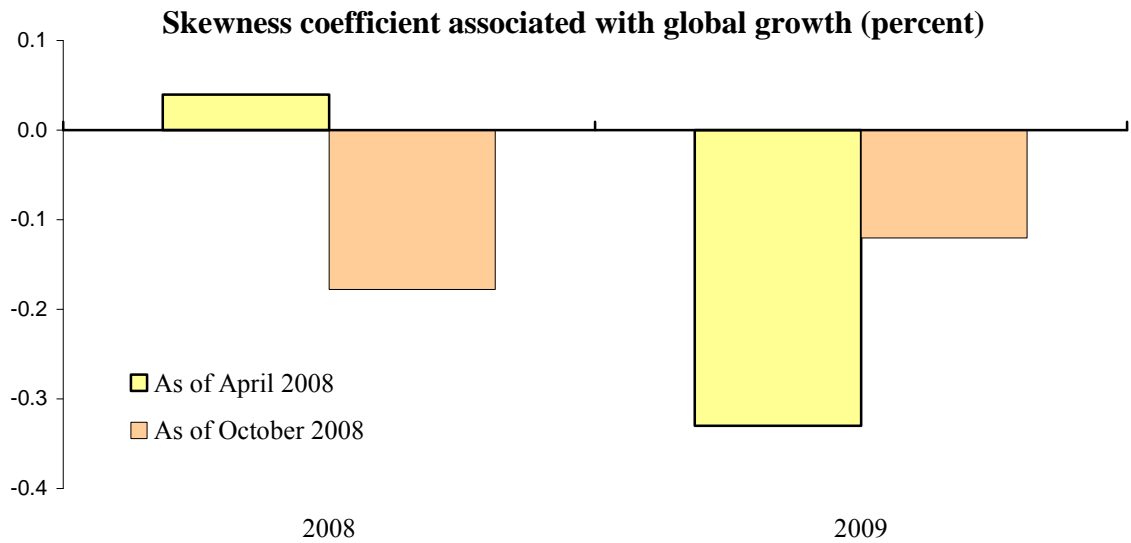
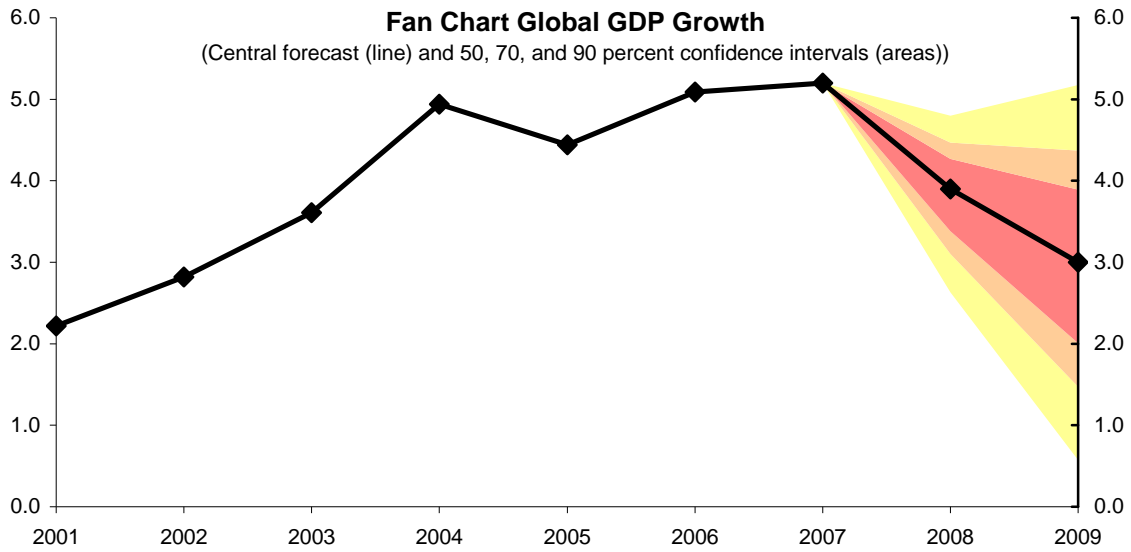
Figure 3. Fan Chart for Global Growth and Skewness of Risk Factors



Source: Consensus economics, Bloomberg and authors' estimates.

1/ Bars represent the percentage point impact of the risk factors on global growth (skewness coefficient multiplied by the estimated elasticity).

Figure 4. Fan Chart for Global Growth Based on Direct Estimates of Variance and Skew



Source: Consensus economics, Bloomberg and authors' estimates.

References

- Ait-Sahalia, Y. and J. Duarte, 2003, "Nonparametric Option Pricing Under Shape Restrictions," *Journal of Econometrics*, Vol. 116, pp. 9–47.
- Ait-Sahalia, Y. and A. Lo, 1998, "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices," *Journal of Finance*, Vol. 53, pp. 499–547.
- Bahra, Bhupinder, 1997, "Implied risk-neutral probability density functions from option prices: theory and application," Bank of England Working Paper, No. 66.
- Blix, M. and P. Sellin, 1998, "Uncertainty Bands for Inflation Forecasts," Sveriges Riksbank Working Paper No. 65 (Sveriges Riksbank: Sweden).
- Breeden, D. and R. Litzenberger, 1978, "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, Vol. 51, pp. 621–51.
- Britton, E., P. Fisher, and J. Whitley, 1998, *The Inflation Report Projections: Understanding the Fan Chart*, Bank of England Quarterly Bulletin (February).
- Cardarelli, R, S. Elekdag and S. Lall, 2009 "Financial Stress, Downturns, and Recoveries," forthcoming, IMF Working Paper..
- Cheng, Kevin, 2009, "A Stable and Flexible Method to Extract Information about Expected Future Price Prospects from Options Markets," (unpublished: International Monetary Fund).
- Estrella, Arturo, and Frederic S. Mishkin, 1996, "The Yield Curve as a Predictor of U.S. Recessions," *Current Issues in Economics and Finance*, Vol. 2, No. 7, pp. 1–6.
- International Monetary Fund (IMF), 2006, *World Economic Outlook: Globalization and Inflation* (Washington, April).
- John, S., 1982, The Three-Parameter Two-Piece Normal Family of Distributions and Its Fitting," *Communications in Statistics—Theory and Methods*, Vol. 11, pp. 879–85.
- Johnson, N. L., S. Kotz, and N. Balakrishnan, 1994, *Continuous Univariate Distributions* (John Wiley, New York)
- Kannan, Prakash and Frederike Köhler-Geib, 2009, "The Uncertainty Channel of Contagion," (unpublished: International Monetary Fund).
- King, T. B., A. T. Levin, and R. Perli, 2007, "Financial Market Perceptions of Recession Risk," Finance and Economics Discussion Series Working Paper No. 2007–57 (Board of Governors: Washington DC).

Prati, Alessandro and Massimo Sbracia, 2002, "Currency Crises and Uncertainty About Fundamentals," IMF Working Papers 02/3 (Washington: International Monetary Fund).

Tay, Anthony S. and Kenneth Wallis, 2000, "Density Forecasting: A Survey," *Journal of Forecasting*, Vol. 19, pp. 235–254.