

Blix, Mårten; Sellin, Peter

Working Paper

Uncertainty bands for inflation forecasts

Sveriges Riksbank Working Paper Series, No. 65

Provided in Cooperation with:

Central Bank of Sweden, Stockholm

Suggested Citation: Blix, Mårten; Sellin, Peter (1998) : Uncertainty bands for inflation forecasts, Sveriges Riksbank Working Paper Series, No. 65, Sveriges Riksbank, Stockholm

This Version is available at:

<https://hdl.handle.net/10419/83036>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Uncertainty Bands for Inflation Forecasts

by

Mårten Blix
and
Peter Sellin*

1998-09-09.

Abstract

The contribution of this paper is to show how the balance of risk for various macro variables can be linked to inflation uncertainty. Inflation uncertainty is *derived* from uncertainty in the macro variables that are deemed to be important for future inflation. The paper focuses on the technical derivation of inflation forecast skewness from uncertainty in such macro variables. The uncertainty in these macro variables is based on their historical standard deviations, but we allow these to be subjectively adjusted if there is reason to be more or less uncertain than historically. We also allow for a subjective assessment of the balance of risk, i.e. whether the distributions are symmetric or not. The baseline case is that the distributions of the macro-variables are symmetric and Gaussian with standard deviations based on historical data; any departures from the baseline can typically be justified from some indicator correlated with future inflation.

JEL Classification: C19, C53, E39.

Keywords: Inflation forecast, inflation uncertainty, two-piece normal distribution, skewness.

* Economics Department, Sveriges Riksbank, S-103 37 Stockholm, Sweden. Email: marten.blix@riksbank.se and peter.sellin@riksbank.se. We have benefitted from a discussion with John Whitley and Erik Britton at the Bank of England. We would also like to thank Hans Dillén, Sten Hansen, Hans Linberg, Hans Lindblad, Marianne Nessén and Anders Vredin for valuable comments. The opinions expressed in the paper are those of the authors and need not reflect those of the Riksbank. Author for correspondence: Mårten Blix.

1 Introduction

In January 1993 the Governing Board of Sveriges Riksbank (the Swedish central bank) adopted an explicit inflation target, stating that the annual increase in the consumer price index in 1995 and onwards should be limited to 2 percent with a tolerance band of ± 1 percentage point. Since monetary policy influences inflation with a lag of 1-2 years the Riksbank must base monetary policy on an assessment of future inflation. This is done with an inflation forecast, which is conditional on all information considered relevant. It is also conditional on the assumption that the current repo rate is *unchanged* over the forecast horizon. The Riksbank communicates its view of the inflation outlook to the public in its Inflation Report, which is published on a quarterly basis. Since December 1997, the Inflation Report contains the inflation forecast with uncertainty bands around the forecast.

For policy purposes, having uncertainty bands around the inflation forecast is useful for several reasons. First and foremost, they serve to illustrate that the inflation forecast is inherently uncertain. The uncertainty is both about the shocks that will affect the economy as well as uncertainty about both the qualitative and quantitative nature of the transmission mechanism.

Second, the bands serve to present the Riksbank's view of the balance of risks to the public and to market participants. In particular, it allows the Riksbank to communicate with a minimum of ambiguity whether the risk is believed to be higher that inflation will be *below* the forecast than that it will be *above*, as was the case in Inflation Report 1998:2. In other circumstances, the situation may be the reverse, which could then also be communicated.

Third, the construction of the bands helps to focus internal discussion in the Riksbank about the sources of inflation uncertainty and about their quantitative

importance.

In this paper we propose a new method for constructing the uncertainty bands around the forecast. There are some similarities to the Bank of England (BoE) method. In particular, we are using the setup from Britton, Cunningham and Whitley (1997) in posing the questions for the subjective assessments (in section 2.3) and we have the same distributional assumptions.

Although there are some similarities, the methods differ in important ways. Our method starts with an assessment of uncertainty in various macro variables and then aggregates the implications of that uncertainty for the inflation forecast. The aggregation of uncertainty with a well-defined role for subjective judgements is the main contribution of this paper. With one reasonable assumption we can relate the balance of risks in the macro variables to the balance of risks for inflation. In other words, the subjective assessments of the macro variables determine the balance of risks for the inflation forecast.

In the BoE method by contrast, as presented in Britton, Fisher and Whitley (1998), the forecast distribution of inflation originates from the monetary policy committee (MPC) and might in that sense be characterized as a “Top-Down” approach. In our method, the initial assessments and the aggregation is instead done at the Economics Department, and then filtered upwards. Indeed, the method as discussed in this paper is structured to work in symbiosis with the inflation forecast. This is one of its key advantages.

The standard statistical approach in deriving forecast error bands would start with estimating an econometric model for making inflation forecasts. In a linear multivariate model exogenous shocks would typically be assumed to be normally distributed, which would imply that the endogenous variables (inflation among them)

would be normally distributed as well. Deriving forecast error bands in such a setting is a well known and understood statistical problem. In a nonlinear multivariate model, simulation exercises would have to be performed in order to derive the forecast error bands. This could be more time-consuming but fairly straightforward.

We do not use the standard approach for a number of reasons. First, no single model is used at the Riksbank for making inflation forecasts. Second, the standard approach does not allow specific information relevant to the particular forecast period to be used. Third, for the relatively short forecast horizons we are interested in (up to two years) subjective judgements have proved to be important in making good forecasts. We would therefore prefer to use an approach that explicitly, and as rigorously as possible, takes subjective judgements about uncertainty into account. Judgements about upside or downside risks as well as judgements regarding whether uncertainty is greater or smaller than in the recent past are of interest.

It is worth emphasizing that we will make a distinction between the macro variables that are deemed to affect inflation and inflation itself. The macro variables are directly adjusted for subjective uncertainty, as discussed in section 2. The uncertainty in the inflation forecast, on the other hand, is *derived* from the uncertainty assessments on the macro variables after making one key assumption, as discussed in section 3. The different treatments of inflation and other macro variables are only a reflection of the Riksbank's inflation target. Since it is the inflation forecast and the inflation uncertainty that matters for policy decisions, it is desirable to have inflation uncertainty *endogenously* determined from underlying assumptions.

The paper is outlined as follows. In the next section, we discuss the subjective assessments and the distributional assumptions. Section 3 discusses how the subjective assessments can be aggregated. In particular, it shows how potential skewness in the

probability distributions of the macro variables can be linked to skewness in the inflation forecast distribution. The variance of the forecast distribution is initially taken as fixed, determined from the standard deviation of historical forecast errors. We further show how this assumption can be relaxed, allowing subjective uncertainty in the macro variables to also affect the variance of the inflation forecast distribution. Section 4 contains an example to illustrate the forecast distribution. Finally, section 5 contains some concluding remarks.

2 Uncertainty Assessment

In this section we will discuss the framework for the uncertainty assessment of the inflation forecast. The inflation forecast itself is of course the *sine qua non* input in this method, but will not be discussed explicitly. It is taken as given for the purposes of this paper; the inflation forecast is discussed in the Riksbank's quarterly inflation report.

2.1 Mode Forecast

One aspect of the inflation forecast, however, needs to be addressed. Similar to BoE, it is a *mode* forecast, denoted by μ , rather than a *mean* forecast, denoted by $\tilde{\mu}$. The mode of a distribution is a different measure of central tendency, but will sometimes coincide with the *mean* or the *median*, for example in the standard Gaussian distribution. The *mode* is the most frequent observation in a distribution and in that sense also the most likely¹ outcome; it is not affected by the possibility of extreme events such as observations in the tails of the distribution.

This property is both the strength and weakness of the *mode*: it uses less information about the distribution and is therefore also less sensitive to unlikely outcomes. For example, the *mode* might be misleading if the distribution were multi-modal, since it

¹ For the distribution in this paper, the two piece normal, the *mode* is the most likely outcome in the sense that it maximizes the probability density function $f(x)$.

would select one of the peaks in the distribution and disregard the others. This is not a serious concern as long as the distribution is single-peaked and not too flat.

The *mean* of the distribution is calculated implicitly in our method. This does not imply that it is unimportant. Indeed, the difference between the *mean* and the *mode*, denoted by $\gamma \equiv \tilde{\mu} - \mu$, plays a central role in our analysis. The parameter γ , as discussed in the next section, can be viewed as a measure of skewness for the distribution. When γ is negative the distribution is skewed to the left, or in other words, there is more downside risk than upside. Formally, this can be expressed as $\text{pr}[X \leq \mu] > 0.5$. Conversely, if γ is positive this implies that there is more upside risk than downside risk. When the distribution is symmetric there is no skewness ($\gamma = 0$).

The parameter γ thus summarizes the balance of risks in terms of skewness. Suppose the distribution is very skewed so that γ is large. This would then be reason to re-examine the assumptions behind the inflation forecast and perhaps iterate on the forecasting round. Whether or not the possibility of an extreme observation, such as a severe deterioration in the Asian crisis, should motivate a revision of the forecast is a matter of judgement.

Finally, note that the *mode* does not have the same asymptotic justification as the *mean* as measure of central tendency. Under quite general conditions, allowing for heteroskedastic and non-iid processes we can still have mean squared consistency as the sample size goes to infinity. This might be viewed as a disadvantage for the *mode* in that its use cannot readily be justified from large-sample theory. However, the main use of asymptotic theory is in allowing inferences when the finite-sample distribution is *unknown*. When the distribution is assumed to be known this is not an issue². Instead

² The choice of distribution is discussed in the next section.

the question is how to best estimate the parameters of the distribution.

2.2 Distributional Assumption

Let us denote the macro variables that are deemed to influence the future level of inflation by $X_j(t)$ where $j = 1, \dots, n$, while inflation is denoted by $\pi(t)$. We will assume that each of the X_j (as well as inflation) is drawn from the univariate distribution given by

$$(1) \quad f(x; \mu, \sigma_1, \sigma_2) = \begin{cases} C \exp\left\{-\frac{1}{2\sigma_1^2}(x-\mu)^2\right\} & x \leq \mu \\ C \exp\left\{-\frac{1}{2\sigma_2^2}(x-\mu)^2\right\} & x > \mu, \end{cases}$$

where $C = k(\sigma_1 + \sigma_2)^{-1}$, $k = \sqrt{2/\pi}$ and μ is the *mode*. This distribution is known in the statistical literature as the "two-piece normal", see Johnson, Kotz, and Balakrishnan (1994). Three parameters, the *mode* and two measures of standard deviation define it. To the left of the *mode*, it is proportional to a standard Gaussian with *mean* μ and standard deviation σ_1 ; to the right of the *mode*, to a standard Gaussian with *mean* μ and with standard deviation σ_2 . The distribution has the property that it collapses to the standard Gaussian when $\sigma_1 = \sigma_2$. When $\sigma_1 > \sigma_2$ it is skewed to the left, i.e. $\text{pr}[X \leq \mu] > 0.5$ and conversely when $\sigma_1 < \sigma_2$.

This distribution is discussed in John (1982), who has shown that $\text{pr}[L_1 \leq x \leq L_2]$ is given by

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

where $\Phi(\cdot)$ is the standard normal cumulative distribution function and

$$(3) \quad \begin{cases} \sigma = \sigma_1 & \text{if } L_1 \leq L_2 \leq \mu \\ \sigma = \sigma_2 & \text{if } \mu \leq L_1 \leq L_2. \end{cases}$$

Note that if we are interested in $L_1 \leq \mu \leq L_2$, the integral must be split as

$\int_{L_1}^{\mu} f(x)dx + \int_{\mu}^{L_2} f(x)dx$. Moreover he has shown that the variance is

$$(4) \quad \text{var}(x) = (1 - k^2)(\sigma_2 - \sigma_1)^2 + \sigma_1\sigma_2$$

and that the third central moment (skewness) is given by

$$(5) \quad \mathbb{E}[(x - \mu)^3] = k(\sigma_2 - \sigma_1)[(2k^2 - 1)(\sigma_2 - \sigma_1)^2 + \sigma_1\sigma_2],$$

which is proportional to $k(\sigma_2 - \sigma_1)$ since $2k^2 - 1 > 0$. Therefore, we will use

$$(6) \quad \gamma \equiv \tilde{\mu} - \mu = k(\sigma_2 - \sigma_1)$$

as measure of skewness. The advantage of using (6) rather than (5) stems from it being (exactly) the difference between the *mean* and the *mode* of the distribution. Note that both (5) and (6) are zero when there is no skewness and the variance in (4) reproduces the standard formula. In other words, with no skewness the distribution collapses to the standard Gaussian. Moreover, from (6) the mean of the distribution is easily obtained by $\tilde{\mu} = \mu + k(\sigma_2 - \sigma_1)$.

This distributional assumption is important and convenient. Since it is closely related to the standard Gaussian, central limit theorems can be used as the basic rationale; departures from the standard Gaussian can most often be justified in terms of some specific event or some particular indicator (discussed more below).

2.3 Uncertainty Assessment

As input in the method we need uncertainty assessments of the macro variables. The assessment is partly subjective, but takes as starting point the historical data. This we formalize by posing two questions for each variable X_j given the *mode* forecast.

1. What is the chance that the outcome will be lower than the *mode* forecast μ ? In other words, what is the downside risk? Or more formally, what is

$P_j = \text{pr}[X_j \leq \mu_j]$? Unless there is some specific information available, the answer to this question will be 50%, which is the reference value.

2. How large is the uncertainty of the forecast compared to the historical uncertainty as measured by the standard deviation? The answer is given as h_j , a multiplicative factor on the standard deviation. The reference value is one unless there is some specific information available that would give us reason to be more or less uncertain. A value of $h_j < 1$ implies that we are less uncertain than historically and conversely for $h_j > 1$.

As an example of answers to the questions above, consider imports. Due to the crisis in Asia, the forecast might be seen as having both downside risk as well as being more uncertain than historically, say $P_j = 0.6$ and $h_j = 1.3$.

This way of quantifying uncertainty is useful for several reasons. While it is true that there is little chance of having a serious debate about whether $h_j = 1.29$ or $h_j = 1.30$, the answers provide a reference point for discussion. The difference between being 10% more uncertain than historically can be contrasted to being for example 30% more uncertain. The uncertainty can also be compared to the last forecasting round with a minimum of ambiguity. Being quantitative in this way focuses the discussion on the issues behind the uncertainty and on the underlying assumptions.

On what grounds might an assessment be more uncertain than historically? For example, if we believe that the economy is approaching a turning point in the business cycle this might justify being more uncertain, since turning points are notoriously hard to forecast. Another example is if we are in an election year with a possible shift in

policies.

But there may also be less uncertainty than historically. For example, once the wage bargaining round is completed there is less reason to be uncertain about future wage increases. Another example is if all indicators point in the same direction.

How can we use the answers above in the forecast distribution specified in (1)? Let the variance of X_j that has been scaled with uncertainty parameter be denoted by

$$(7) \quad \omega_{j,j}(t) = (h_j(t) \sigma_j(t))^2$$

where $\sigma_j(t)$ is the historical standard deviation of X_j . In the appendix it is shown that the expressions

$$(8) \quad \sigma_{1,j}^2(t; \omega_{j,j}, P_j) = \omega_{j,j}(t) \left[(1-k^2) \left(\frac{1-2P_j(t)}{P_j(t)} \right)^2 + \left(\frac{1-P_j(t)}{P_j(t)} \right) \right]^{-1},$$

$$(9) \quad \sigma_{2,j}^2(t; \omega_{j,j}, P_j) = \omega_{j,j}(t) \left[(1-k^2) \left(\frac{1-2P_j(t)}{1-P_j(t)} \right)^2 + \left(\frac{P_j(t)}{1-P_j(t)} \right) \right]^{-1}$$

are such that the variance is fixed by (7) and $P_j(t) = \Pr[X_j(t) \leq \mu_j(t)]$ as desired.

Hereafter the dependence in (8) and (9) on the parameters $P_j(t)$ and $\omega_{j,j}(t)$ will be suppressed to keep the notation simple.

The intuition for these expressions is simple. Note that (8) and (9) are proportional to $\sigma_1^2 \cong h^2 \sigma^2 P / (1-P)$ and $\sigma_2^2 \cong h^2 \sigma^2 (1-P) / P$. The factor h thus has the effect of scaling the measures of standard deviation so that a large h will increase both σ_1 and σ_2 , and conversely for a small h . The effect of P is perhaps best seen with the example above in which there was some downward risk – given by $P = 0.6$. We then have that $P / (1-P) > 1$ and $(1-P) / P < 1$, and consequently σ_1 will be scaled upwards and σ_2 downwards. A larger σ_1 than σ_2 is of course the same as having more

probability mass to the left of the mode, or in other words, more downside risk.

3. Inflation Forecast Distribution

In the previous section we allowed subjective judgements to play a well-defined role. In this section we will discuss how these assessments can be aggregated.

3.1 Inflation Forecast Skewness

The starting point is the assumption that the inflation forecast is also distributed as the "two-piece normal" in (1) with parameters μ_π , $\sigma_{1,\pi}$ and $\sigma_{2,\pi}$. The inflation forecast $\mu_\pi(t)$ as well as the variance of the inflation forecast are taken as given. The remaining problem is thus to connect the uncertainty assessments of the X_j 's with inflation and thereby deriving estimates of $\sigma_{1,\pi}$ and $\sigma_{2,\pi}$.

The key question is how to relate the forecast distributions for the macro variables to the inflation forecast. If we were to assume a linear relationship between the X_j variables and inflation, we could in principle derive the forecast distribution for inflation. This approach, however, appears infeasible. Given n variables X_j , it requires n -dimensional integration, which even for small values of n becomes hopelessly complicated.³ Moreover, the resulting distribution would certainly *not* be a two piece-normal distribution or any other known distribution that could be summarised with a few parameters.

Our approach is instead to make a key assumption about how the uncertainty in the macro variables X_j is connected to future inflation,

³ Even using the moment-generating function we would not obtain a tractable parametric distribution for the inflation forecast.

$$(10) \quad \gamma_{\pi}(t) = \sum_{j=1}^n \beta_j(t) \gamma_j(t)$$

where γ_{π} is the skewness of inflation and γ_j is the skewness of variable X_j . Equation (10) implies that the skewness from the macro variables X_j affects the skewness of inflation with the weight β_j .

Although (10) is a statistical approximation, its basic rationale comes from economic arguments. These are perhaps best illustrated by some simple examples, which show that assumption (10) is *qualitatively* sensible. First, let us consider the benchmark case where there is no skewness in any of the macro variables. Equation (10) will imply that there will not be any skewness in inflation either. This we believe is a sensible property of the assumption, at least under linearity. Second, if there is negative skewness in for example consumption, this will result in negative skewness for inflation since the weight for consumption is typically non-negative. Finally, suppose some other variable, such as the wage rate, shows positive skewness. Whether or not the sum of consumption skewness and wage skewness should result in positive or negative skewness will then depend on their relative importance for future inflation, as captured by the weights β_j . This last example can also be used to judge the extent to which the results are also *quantitatively* sensible.

How are the weights derived? The weights β_j are the elasticities with respect to inflation obtained from considering a change in each X_j in a macroeconomic model and deriving the effects on inflation one and two years ahead.

The skewness parameters on the RHS of (10) are immediately obtained by substituting (8) and (9) into (6),

$$(11) \quad \gamma_j(t) \equiv \tilde{\mu}_j(t) - \mu_j(t) = k(\sigma_{2,j}(t) - \sigma_{1,j}(t)).$$

Given the skewness of inflation $\gamma_\pi(t)$ from (10) and the standard deviation of past forecasting errors $\sigma_\pi(t)$, how⁴ can we find $\sigma_{1,\pi}$ and $\sigma_{2,\pi}$? Recall from (4) and (6) that $\sigma_\pi(t)$ and $\gamma_\pi(t)$ are defined from

$$(12) \quad \sigma_\pi^2(t) = (1 - k^2)[\sigma_{2,\pi}(t) - \sigma_{1,\pi}(t)]^2 + \sigma_{1,\pi}(t)\sigma_{2,\pi}(t)$$

$$(13) \quad \gamma_\pi(t) = k(\sigma_{2,\pi}(t) - \sigma_{1,\pi}(t)).$$

This gives us two equations and two unknowns, which can be reduced to the equation

$$(14) \quad \sigma_{1,\pi}^2(t) + b\sigma_{1,\pi}(t) + c = 0$$

where $b = (\gamma_\pi / k)$ and $c = -[(1 - 1/k^2)\gamma_\pi^2 + \sigma_\pi^2]$. There are two solutions to (14), but only one that will be relevant (the other solution is typically negative). This solution for $\sigma_{1,\pi}(t)$ can then be substituted into (13), which gives $\sigma_{2,\pi}(t)$.

3.2 Inflation Forecast Variance

In the previous section the variance of the inflation forecast was taken as given, calculated from past forecast errors. With some further assumptions we can let this variance be affected by the assessments on the X_j 's. In particular, if we assume that

$$(15) \quad \sigma_\pi^2(t; \sigma_\varepsilon^2) = \beta'(t)\Sigma(t; \sigma_\varepsilon^2)\beta(t), \quad t = 1, \dots, T$$

where $\beta' = [1 \quad \beta_1 \quad \dots \quad \beta_n]$,

⁴ The method used to find $\sigma_{1,j}(t)$ and $\sigma_{2,j}(t)$ discussed in the appendix cannot be used, since we do not yet know $P_\pi(t) = \text{pr}[\pi(t) \leq \mu_\pi(t)]$.

$$(16) \quad \Sigma(t; \sigma_\varepsilon^2) = \begin{bmatrix} \sigma_\varepsilon^2 & 0 & \cdots & 0 \\ 0 & \omega_{1,1} & \cdots & \omega_{1,n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \omega_{n,1} & \cdots & \omega_{n,n} \end{bmatrix}$$

and

$$(17) \quad \omega_{i,j} = \begin{cases} \text{var}[X_j(t)] & \text{for } i = j \\ \text{cov}[X_i(t), X_j(t)] & \text{for } i \neq j. \end{cases}$$

The variance in (17) is given in (7), while the covariance can be obtained from

$$(18) \quad \text{cov}[X_i(t), X_j(t)] = \rho_{i,j} \sqrt{\omega_{i,i}(t) \omega_{j,j}(t)},$$

where $\rho_{i,j}$ is the estimated data correlation between X_i and X_j (assumed to not depend on t). The covariance matrix Σ can be viewed as a scaled version of the empirical covariance matrix, where any difference between the two reflects subjective judgements. The term $\sigma_\varepsilon^2(t)$ is the variance of an inflation shock independent of the X_j 's and might be interpreted as the part of the forecast uncertainty which increases with horizon due to forecasts further into the future being inherently more uncertain.

The role of σ_ε^2 is perhaps best illustrated with the following calibration exercise.

- I. Set $h_j(t) = 1 \forall j$. Choose $\sigma_\varepsilon^2(t)$ such that (15) is equal to the standard deviation of the historical forecast errors t years ahead.

The calibration in (I) is computed for $t = 1, \dots, T$, in our case one and two years ahead. We can then return to the subjectively adjusted variances in (7) and use them together with the calibrated $\sigma_\varepsilon^2(t)$ in (15). This yields a subjectively adjusted $\sigma_\pi^2(t)$, which can then be used in (12) with the same methods as in that section applied.

What has this accomplished? In the baseline case with the same uncertainty as historically, the empirical standard deviation of the forecast errors is re-produced by

construction. When there is more uncertainty in some given variable X_j this will tend to increase $\sigma_\pi^2(t; \sigma_\varepsilon^2)$ by the weight of this variable β_j ; the converse occurs when there is less uncertainty than historically.

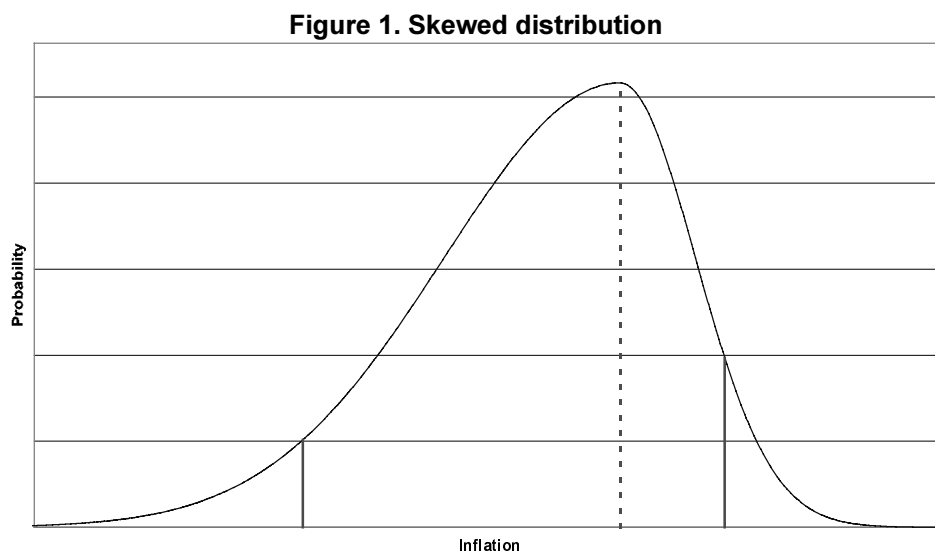
This approach differs from the standard regression methods. The more conventional way of relating unconditional variances is to assume some (typically linear) relation of the form $\pi_t = \theta' X_t + \eta_t$, whence $\sigma_\pi^2 = \theta' \Omega_X \theta + \sigma_\eta^2$ where Ω_X is the covariance matrix of X_t and $\sigma_\eta^2 = \text{var}[\eta_t]$ (note that $\eta_t \neq \varepsilon_t$).⁵ This is of course the mathematically correct way of relating variances with standard OLS assumptions and could be used instead of the method outlined above. The problem with this approach is that the standard regression assumptions are probably violated, as discussed in Blix (1998).

4. An example

In Inflation Report 1998:2 the forecast distribution was slightly skewed to the left so that the difference between a the standard Gaussian and the two-piece normal was small. To illustrate the properties of the two-piece normal, it will be more convenient to consider an *hypothetical* example where there is strong downside risk around the inflation forecast.

Suppose the aggregation of risks results in parameters for the inflation forecast distribution such that $\text{pr}[\pi \leq \mu_\pi] = 0.7$. The resulting distribution is illustrated in figure 1.

⁵ It would also be possible to derive skewness (the third central moment) of the inflation forecast distribution as a weighted average of the various components' co-skewnesses with inflation (see Diacogiannis (1994) for an application to asset pricing theory).



The dotted line shows the mode to the left of which there is 70% of the probability mass. The straight line to the left of the mode is such that 5% of the probability mass is to its left and conversely for the one to the right. Thus, the two lines represent a 90% confidence interval.

The figure serves to illustrate how the two-piece normal is affected by strong skewness. First, the two areas outside the confidence interval are of different shape. Second, most of the confidence interval lies to the left of the mode.

Another way to illustrate the properties of the distribution is to consider the probability that the outcome will lie in a given interval, as was done in the Inflation Report 1998:2. For example, we can calculate the probability that inflation is less than one, between one and two, and so on. An interval of particular interest is that between one and three, $\int_1^3 f(\pi) d\pi$, the Riksbank's tolerance interval for inflation.

5. Concluding Remarks

In this paper we show how the balance of risks for various macro components, i.e. the skewness of the distributions, can be linked to the balance of risk for inflation. The assessment of risk for the macro variables is partly subjective but also based on historical data. In the baseline case the uncertainty is the same as the historical and the risks are symmetric.

The aim of the paper is to provide a well-defined role for subjective assessments of the macro variables that are deemed to influence inflation. Having made those subjective assessments we have then attempted to be as rigorous as possible in deriving the probability distribution for the inflation forecast.

Appendix

Here we outline a procedure for finding the standard deviation parameters in (8) and

(9). Consider a given answer to questions one and two in section 2.3, i.e. $P_j(t)$ and

$\omega_{j,j}(t)$ are fixed. How do we choose $\sigma_{1,j}(t)$ and $\sigma_{2,j}(t)$? Note that the formula in (2)

implies that

$$(A.1) \quad \int_{-\infty}^{\mu} f(x) dx = \frac{\sigma_1}{\sigma_1 + \sigma_2}, \quad \forall \mu.$$

Thus, we choose $\sigma_{1,j}(t)$ and $\sigma_{2,j}(t)$ such that

$$(A.2) \quad \frac{\sigma_{1,j}(t)}{(\sigma_{1,j}(t) + \sigma_{2,j}(t))} = P_j(t),$$

and

$$(A.3) \quad (1 - k^2)(\sigma_{2,j}(t) - \sigma_{1,j}(t))^2 + \sigma_{1,j}(t)\sigma_{2,j}(t) = \omega_{j,j}(t),$$

which is implied by the expression for the variance in (4). This gives the solutions

$$(A.4) \quad \sigma_{1,j}^2(t) = \omega_{j,j}(t) \left[(1 - k^2) \left(\frac{1 - 2P_j(t)}{P_j(t)} \right)^2 + \left(\frac{1 - P_j(t)}{P_j(t)} \right) \right]^{-1},$$

$$(A.5) \quad \sigma_{2,j}^2(t) = \omega_{j,j}(t) \left[(1 - k^2) \left(\frac{1 - 2P_j(t)}{1 - P_j(t)} \right)^2 + \left(\frac{P_j(t)}{1 - P_j(t)} \right) \right]^{-1}$$

as claimed in the text.

References

Blix, Mårten (1998), "Forecasting Swedish Inflation with a Markov Switching VAR", *mimeo*, Sveriges Riksbank.

Britton, Erik, Paul Fisher och John Whitley (1998), "The Inflation Report Projections: understanding the Fan Chart", *Bank of England Quarterly Bulletin*.

Britton, Erik, Alastair Cunningham och John Whitley (1997), "Asymmetry, risks and a probability distribution of inflation", *mimeo*, Bank of England.

Diacogiannis, G.P. (1994), "Three Parameter Asset Pricing", *Managerial and Decision Economics* 15, 149-158.

John, S. (1982), "The three-parameter two-piece normal family of distributions and its fitting", *Communications in Statistical Theory and Methods*, 11 (8), 879-885.

Johnson, Kotz, and Balakrishnan (1994), *Continuous Univariate Distributions*, Vol. 1, P. 173.