

Blix, Mårten; Sellin, Peter

Working Paper

A Bivariate Distribution for Inflation and Output Forecasts

Sveriges Riksbank Working Paper Series, No. 102

Provided in Cooperation with:

Central Bank of Sweden, Stockholm

Suggested Citation: Blix, Mårten; Sellin, Peter (2000) : A Bivariate Distribution for Inflation and Output Forecasts, Sveriges Riksbank Working Paper Series, No. 102, Sveriges Riksbank, Stockholm

This Version is available at:

<http://hdl.handle.net/10419/82383>

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.

A Bivariate Distribution for Inflation and Output Forecasts

by

Mårten Blix and Peter Sellin*

Sveriges Riksbank

February 15, 2000

Abstract

The contribution of this paper is to derive a bivariate distribution for inflation and output uncertainty with a well-defined role for subjective judgements. The marginal distributions for inflation and output growth are derived from uncertainty in the macro variables that are deemed to be important for future inflation and output growth. The uncertainty in the 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. Given the marginal distributions for inflation and output growth we derive a bivariate distribution using the translation method. Having derived the bivariate distribution we are in a position to discuss inflation forecast uncertainty conditional on the growth of output (or vice versa). The analysis can readily be extended to the case of more than two variables.

JEL Classification: C19, C53, E39.

Mathematics Subject Classification (1991): 62E15, 62E17, 62P20.

Keywords: Inflation forecast, output forecast, conditional forecasts, two-piece normal distribution, translation method, Johnson system.

* Economics Department, Sveriges Riksbank, S-103 37 Stockholm, Sweden. Email: marten.blix@riksbank.se and peter.sellin@riksbank.se. We have benefitted from discussions with and comments from Tomas Lindström, Anders Vredin and Lars E.O. Svensson. We would also like to thank Hanna Widell for research assistance and Ann-Christine Högberg for secretarial assistance. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank. Author for correspondence: Mårten Blix.

1 Introduction

For policy purposes having uncertainty bands around the inflation forecast is useful for several reasons. First, they serve to illustrate that the inflation forecast is inherently uncertain. Second, the bands serve to present the central bank's view of the balance of risks to the public. In particular, it allows the central bank to communicate whether the risk (i.e. probability) is believed to be higher that inflation will be below or above the forecast. Third, the construction of the bands helps to focus internal discussion in the central bank about the sources of inflation uncertainty and about their quantitative importance.

These are some of the reasons why Sveriges Riksbank started publishing an inflation forecast with uncertainty bands around it since the December 1997 Inflation Report. In Blix and Sellin (1998) we describe how these uncertainty bands are derived. The main contribution of that paper was the aggregation of uncertainty with a well-defined role for subjective judgements. With one crucial but reasonable assumption we were able to relate the balance of risks in a set of macro variables deemed to be of importance for future inflation 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 Blix and Sellin (1998) the focus was exclusively on deriving forecast distributions for inflation. In this paper we extend the analysis by including output growth (or the output gap) as well. Given the forecast distributions for inflation and output growth we derive a bivariate distribution with these distributions as the marginal distributions. Having derived the bivariate distribution we are in a position to discuss inflation forecast uncertainty conditional on the growth of output (or vice versa).

Why do we not simply use an econometric model for deriving confidence bands? There are several reasons for this. First, no single model is used at the Riksbank for making inflation and output forecasts.¹ Second, it would 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. Likewise we make a distinction between output and the variables that affect it. The macro variables are directly adjusted for subjective uncertainty, as discussed in section 2. The uncertainty in the inflation and output forecasts, on the other hand, are *derived* from the uncertainty assessments on the macro variables. The different treatment of inflation and output relative to the other macro variables is a reflection of the central role these variables play in the monetary policy analysis. Thus, inflation and output uncertainty are *endogenously* determined from underlying assumptions.

¹ Uncertainty regarding the parameters of the various forecast models is not dealt with in our analysis. For a discussion of this type of uncertainty see Söderström (1999).

The paper is outlined as follows. In the next section, we discuss the subjective assessments and the distributional assumptions. Section 3 shows how potential skewness in the probability distributions of the macro variables can be linked to skewness in the inflation and output forecast distributions. Section 4 deals with the construction and fitting of a bivariate distribution given the marginal distributions of the inflation and output growth forecasts. We then discuss forecast distributions of inflation conditional on different assumptions regarding output growth. Section 5 discusses some issues for monetary policy. Section 6 concludes.

2 Uncertainty Assessment

In this section we will discuss the framework for the uncertainty assessment of the inflation and output forecasts.² The inflation and output forecasts themselves are of course the *sine qua non* input in this method, but will not be discussed explicitly. They are taken as *given* for the purposes of this paper.

2.1 Mode Forecast

One aspect of the inflation and output forecasts, however, needs to be addressed. They are *mode* forecasts rather than *mean* forecasts. The mode of a distribution is a different measure of central tendency, but will sometimes coincide with the mean or the median (as in the normal distribution). The *mode* is the most frequent observation in a distribution or the most likely outcome.³ It is not affected by the possibility of extreme events such as observations in the tails of the distribution.

² This section builds on Blix and Sellin (1998). The new feature is that the analysis now includes output forecasts as well as inflation forecasts.

³ This is strictly true only for a discrete distribution, where a single outcome can be associated with a positive probability measure. For the continuous distribution in this paper, the two piece normal, the *mode* is the most

The primary reason for associating point forecasts with the mode rather than with the mean is that it makes more intuitive sense to view the point forecast as the most likely outcome, i.e. the mode. This conjecture builds on two premises. The first is that the point forecasts are predominantly judgemental forecasts. Although econometric and other models are used in forecasting some of the macro variables there is always an element of judgement involved. The second premise is that judgemental forecasts are closer to the mode than to the mean, because forecasters consider the most likely scenario and do not weigh in all possible outliers with the appropriate (small) probabilities into their forecasts. We have found support for this conjecture in the practical implementation of our method at the Riksbank. During the emerging markets crises of 1997-98 judgemental forecasts of the balance of risks resulted in an inflation forecast distribution exhibiting downside risk. Had the point forecasts been mean forecasts the inflation forecast distribution – as we compute it – would instead have displayed upside risk, as the point forecasts of variables having a positive effect on inflation would have shifted down to below the mode (and the median).

The *mean* of the distribution is calculated implicitly in our method. The difference between the *mean* $\tilde{\mu}$ and the *mode* μ , 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 (upside and downside) risks in terms of a measure of skewness.

In section 5 we discuss monetary policy considerations arising from the choice of measure of central tendency and argue for using a framework of inflation forecast *distribution* targeting.

2.2 Distributional Assumption

Let us denote the macro variables that are deemed to influence the future level of inflation by $X_j^{(\pi)}(t)$ and those that affect output growth (or alternatively the output gap) by $X_j^{(y)}(t)$. Although there is much commonality between these two sets of macro variables, there are some that only affect inflation and some that only affect output. We summarize these sets of variables in $X_j^{(l)}(t)$ where $l = \{\pi, \Delta y\}$ and $j = 1, \dots, n_l$. We will assume that each of the $X_j^{(l)}$ (as well as inflation and output) 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*. For simplicity, the dependence on l, t and j is suppressed. 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 normal distribution with *mean* μ and with standard deviation σ_2 . The distribution has the property that it collapses to the normal distribution 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 $\sigma_1 < \sigma_2$.

We have chosen to work with the distribution in (1) for three main reasons. First, the two-piece normal distribution is easy to work with. It provides us with simple analytical expressions for the moments we need in terms of the three parameters that describe it. Second, asymmetries in either direction are treated in the same way unlike most other well known asymmetric distributions. Third, it has the normal distribution as a special case. Thus, in the absence of any evidence of skewness we will be working with the normal distribution, which is often implicitly assumed among forecasters.

The (three-parameter) two-piece normal distribution is discussed in great detail in John (1982), who has shown that the variance is

$$(2) \quad \sigma_x^2 = (1 - k^2)(\sigma_2 - \sigma_1)^2 + \sigma_1\sigma_2$$

and that the third central moment is given by

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

The sign of the third central moment is given by $k(\sigma_2 - \sigma_1)$ since $2k^2 - 1 > 0$.

Therefore,

$$(4) \quad \gamma \equiv \tilde{\mu} - \mu = k(\sigma_2 - \sigma_1)$$

is proportional to (3) and we will use it as a measure of skewness. The advantage of using (4) rather than the usual measure of skewness, which is (3) divided by the standard deviation cubed, stems from it being (exactly) the difference between the *mean* and the *mode* of the distribution. Moreover, from (4) the mean of the distribution is easily obtained by $\tilde{\mu} = \mu + k(\sigma_2 - \sigma_1)$.

2.3 Subjective Assessment of Uncertainty

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

$P_j^{(l)} = \text{pr}[X_j^{(l)} \leq \mu_j^{(l)}]$? The reference value is 50 percent, i.e. equal upside and

downside risk, which would be the natural value to assign in the absence of any specific information regarding the balance of risks.

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 $h_j^{(l)}$, 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 $h_j^{(l)} < 1$ implies that we are less uncertain than historically and conversely for $h_j^{(l)} > 1$.

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.⁴

How can we use the answers above in the forecast distribution specified in (1)? This is where our choice of distribution is useful. We can easily translate the answers to the straightforward questions above into the parameters of the two-piece normal distribution. Let the variance of $X_j^{(l)}$ that has been scaled with uncertainty parameter be denoted by

$$(5) \quad \omega_{j,j}^{(l)}(t) = \left(h_j^{(l)}(t) \sigma_j^{(l)}(t) \right)^2,$$

where $\sigma_j^{(l)}(t)$ is the historical standard deviation of $X_j^{(l)}$. In the appendix of Blix and Sellin (1998) it is shown that the expressions

$$(6) \quad \sigma_{1,j}^2(t,l) = \omega_{j,j}^{(l)}(t) \left[(1-k^2) \left(\frac{1-2P_j^{(l)}(t)}{1-P_j^{(l)}(t)} \right)^2 + \left(\frac{1-P_j^{(k)}(t)}{P_j^{(l)}(t)} \right) \right]^{-1},$$

$$(7) \quad \sigma_{2,j}^2(t,l) = \omega_{j,j}^{(l)}(t) \left[(1-k^2) \left(\frac{1-2P_j^{(l)}(t)}{1-P_j^{(l)}(t)} \right)^2 + \left(\frac{P_j^{(l)}(t)}{1-P_j^{(k)}(t)} \right) \right]^{-1}$$

are such that the variance is fixed by (5) and $P_j^{(l)}(t) = \Pr[X_j^{(l)}(t) \leq \mu_j^{(l)}(t)]$ as desired.

⁴ See for example Hamilton (1989).

The intuition for these expressions is simple. Note that (6) and (7) 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 and Output Forecast Distributions

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. The starting point is the assumption that the inflation and output forecasts are distributed as the two-piece normal in (1) with parameters μ_l , $\sigma_{1,l}$ and $\sigma_{2,l}$, where $l = \{\text{inflation, output}\}$. The forecast $\mu_l(t)$ as well as the variance of the forecast $\sigma_l^2(t)$ are taken as given. The forecast variance is given as a historical average of past forecasting errors.

The key question is how to relate the forecast distributions for the macro variables to the inflation (or output) forecast. As discussed in Blix and Sellin (1998) this is extremely complicated to do rigorously. Our approach is to take a shortcut. Let the skewness in inflation and output be denoted by γ_π and γ_y respectively. We make the following key assumptions about how the uncertainty in the macro variables $X_j^{(l)}$ is connected to future inflation:

$$(8) \quad \gamma_l(t) = \sum_{j=1}^{n_l} \beta_{j,l}(t) \gamma_{j,l}(t), \quad l = \{\pi, \Delta y\}$$

where $\gamma_{j,l}$ is the skewness of the forecast distribution for variable $X_j^{(l)}$. Equation (8) implies that the skewness from the macro variables $X_j^{(l)}$ affects the skewness of inflation&output with the weight $\beta_{j,l}$. The weights $\beta_{j,l}$ are the elasticities with respect to inflation&output obtained from considering a change in each $X_j^{(l)}$ in a macroeconomic model and deriving the effects on inflation&output one and two years ahead.

The skewness parameters on the RHS of (8) are immediately obtained by substituting (6) and (7) into (4),

$$(9) \quad \gamma_{j,l}(t) \equiv \tilde{\mu}_{j,l}(t) - \mu_{j,l}(t) = k(\sigma_{2,j}(t,l) - \sigma_{1,j}(t,l)).$$

Given the skewness $\gamma_l(t)$ from (8) and the standard deviation of past forecasting errors $\sigma_l(t)$, we can find $\sigma_{1,l}$ and $\sigma_{2,l}$ by solving the following equation system,


$$(10) \quad \sigma_l^2(t) = (1 - k^2) [\sigma_{2,l}(t) - \sigma_{1,l}(t)]^2 + \sigma_{1,l}(t) \sigma_{2,l}(t)$$

$$(11) \quad \gamma_l(t) = k(\sigma_{2,l}(t) - \sigma_{1,l}(t)).$$

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

$$(12) \quad \sigma_{1,l}^2(t) + b_l \sigma_{1,l}(t) + c_l = 0$$

where $b_l = (\gamma_l / k)$ and $c_l = -[(1 - 1/k^2)\gamma_l^2 + \sigma_l^2]$. There are two solutions to (12), but only one that will be relevant (the other solution is typically negative). This solution for $\sigma_{1,k}(t)$ can then be substituted into (11), which gives $\sigma_{2,k}(t)$. Figure 1 gives an example of a right skewed distribution of the two-piece normal family.

Figure 1. Marginal distribution for inflation forecast

4. A Bivariate Distribution with Given Marginal Distributions

In the first subsection below we present a method for deriving a bivariate distribution.

In the second subsection we use this method to derive a bivariate distribution, with given marginal distribution functions for inflation and output. The marginal

distributions belong to the two-piece normal family of distributions. In the third

subsection we discuss the fitting of the bivariate distribution using historical data and

the subjective judgements described in the previous two sections. Finally, we give an

example of a bivariate distribution of inflation and output, as well as a conditional inflation forecast.

4.1 The Translation Method

There are several methods that can be used to derive a bivariate distribution function (d.f.) with desired characteristics. We will use the translation method, which dates back to Edgeworth (1896). In general, the translation method is used to transform a complicated d.f. (that has been successfully fitted to the data) into a d.f. of a more convenient form (which allows the use of standard statistical tests to be performed). But it can also be used in reverse, as in this paper, in order to create a non-standard d.f. with certain pre-specified properties.

Perhaps the most successful translation system is the set of transformations investigated by Johnson (1949a). The basic translation is based on the principle of translating to normality, and takes the form

$$(13) \quad z = \alpha + \beta J(x),$$

where z is a standard normal variable, J is a monotone function of the x -variable that we are interested in, and α and β are parameters. The simplest transformation results from assuming that $J(x) = x$. Then in order for z to be standard normal, β will have to be the reciprocal of the standard deviation of x , $\beta = 1/\sigma_x$, and α will be minus the expected value of x divided by the standard deviation of x , $\alpha = -\mu_x/\sigma_x$. Thus,

$$(14) \quad z = \frac{x - \mu_x}{\sigma_x},$$

which is a transformation that is usually referred to as standardizing x . Johnson (1949a) considers a few more complicated J functions. The resulting d.f.'s are usually referred to as the Johnson family of distributions.

In Johnson (1949b) the Johnson family is extended to the bivariate case. Assume that the pairs of variables (z_1, x_1) and (z_2, x_2) both conform to (13). A bivariate d. f. can easily be derived by assuming that (z_1, z_2) have the joint standard normal bivariate probability density function (p.d.f.),

$$(15) \quad p(z_1, z_2) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left\{-\frac{1}{2(1-\rho^2)}(z_1^2 - 2\rho z_1 z_2 + z_2^2)\right\}.$$

This also determines the joint distribution of (x_1, x_2) , since by transforming the variables we obtain the bivariate p.d.f.

$$(16) \quad q(x_1, x_2) = \frac{1}{2\pi\sqrt{1-p^2}} \exp\left\{-\frac{1}{2(1-p^2)}[\tilde{J}_1^2(x_1) - 2\rho\tilde{J}_1(x_1)\tilde{J}_2(x_2) + \tilde{J}_2^2(x_2)]\right\},$$

where $\tilde{J}_i(x_i) = \alpha_i + \beta_i J_i(x_i)$.

4.2 Deriving a Bivariate Distribution with Given Marginal Distributions

We would like to create a bivariate distribution with marginal distribution functions (marginal d.f.'s) that are the ones we have used above for our inflation and output forecasts. The transformation used in this paper will thus have to take the marginal d.f.'s as given in deriving the bivariate distribution function. In other words, we will choose the function $J(x)$ such that the marginal d.f.'s will belong to the two-piece normal family of distributions.

Let $f(\pi)$ and $g(y)$ be the marginal p.d.f.'s, derived in the previous section, for inflation and output growth respectively. We can define a continuous bivariate distribution function with the marginal d.f.'s $F(\pi)$ and $G(y)$ by using the transformation

$$(17) \quad F(\boldsymbol{\pi}) = \Phi(\boldsymbol{\pi}_0), \quad G(y) = \Phi(y_0)$$

where $\Phi(\cdot)$ is the standard normal distribution function. Thus, the variables $\boldsymbol{\pi}_0$ and y_0 are normally distributed with mean zero and variance equal to one. Let J_π and J_y be the transformations imposed by the marginal distributions, i.e.,

$$(18) \quad \begin{aligned} \boldsymbol{\pi}_0 &= J_\pi(\boldsymbol{\pi}) \equiv \Phi^{-1}(F(\boldsymbol{\pi})), \\ y_0 &= J_y(y) \equiv \Phi^{-1}(G(y)). \end{aligned}$$

Hence, the transformation we use is of the general type discussed in the previous subsection.

Nataf (1962) has shown that any convenient multivariate d.f. can be used together with the transformation in (18) to obtain a multivariate d.f. with given marginal d.f.'s. Hence, we are not restricted to using only two variables (inflation and output)⁵ and we are not restricted to using the multivariate normal distribution. Instead of the bivariate normal we will use the joint p.d.f. given in Mardia (1970, p. 30):

$$(19) \quad h(\boldsymbol{\pi}, y) = \frac{f(\boldsymbol{\pi})g(y)}{\sqrt{1-\rho^2}} \exp\left\{-\frac{\rho}{2(1-\rho^2)} \left[\rho(J_\pi^2(\boldsymbol{\pi}) + J_y^2(y)) - 2J_\pi(\boldsymbol{\pi})J_y(y) \right]\right\}.$$

This p.d.f. has some nice properties. For example, it is easily seen that for $\rho = 0$ inflation and output would be independent, since then the joint p.d.f. would simply be given by the product of the marginal p.d.f.'s. This is of course not likely to be the case and we need to obtain an estimate of the parameter ρ . We will discuss the estimation of this parameter in the next subsection. Another nice property is the simple form of the conditional p.d.f.'s, as we will see later.

⁵ Another variable of special interest might be the instrument rate.

Mardia (1970) considers the sufficient conditions for h to be a probability density function. He presents proofs that these are met. Our interest is in making sure that the F and G functions we have used fulfill any conditions required in the proofs. The only conditions required for Mardia's proofs are that the functions be proper distribution functions of the continuous type. These conditions are indeed fulfilled by the d.f.'s belonging to the family of two-piece normal distributions, that we have used.

4.3 Fitting the Distribution

The bivariate distribution in (19) is fully specified by seven parameters:

$\mu_\pi, \sigma_{1,\pi}, \sigma_{2,\pi}, \mu_y, \sigma_{1,y}, \sigma_{2,y}, \rho$. The first six pertain to the marginal d.f.'s of inflation and output. In order to determine the sigma parameters we solve the equation system (10)-(11), which requires RHS values of $(\sigma_\pi^2, \gamma_\pi)$ and (σ_y^2, γ_y) respectively. The gamma parameters are easily computed from Eq. (8). It is more complicated to compute the average forecast error of inflation and output, respectively, for the relevant horizon. This is because in computing the forecast errors account has to be taken of the fact that the forecasts used by the Riksbank are made under the assumption of a constant repo rate.⁶

⁶ This is done mainly for pedagogical reasons. If the inflation forecast is above (below) the target, then this is an argument for raising (lowering) the repo rate to bring the inflation forecast back in line with the target. The fact that the forecasts are made under the assumption of an unchanged repo rate means that the historical forecasts will have to be adjusted, before they are compared with actual inflation and output growth for the relevant forecast periods.

Next we have to obtain an estimate of the parameter ρ . This parameter is the correlation coefficient between the transformed variables π_0 and y_0 . Hence, we must apply the transformations in (18) to the historical data of inflation and output growth before we can compute an estimate of ρ . The estimated value of ρ will in general be greater than the correlation between inflation and output. Lancaster (1957) has shown that

$$(20) \quad |\rho(\pi_0, y_0)| \geq |\rho(\pi, y)|,$$

i.e. the correlation between the transformed variables is greater than or equal to the correlation between the original variables.

4.4 An Example

We will consider an example where there is upside risk in both inflation and output. In Figure 2a we see what the distribution will look like under this scenario. The mode forecasts of output growth and inflation are 3 percent and 2 percent respectively, while the means are 3.4 and 2.16 respectively. This implies the positive skewness measures $\gamma_y = 0.4$ for output and $\gamma_\pi = 0.16$ for inflation. What value should we assign to the parameter ρ ? It may seem reasonable that the upside risk in output should be accompanied by an upside risk in inflation if demand shocks dominate. This type of question should be raised at meetings with the economists involved, which are primarily held to make sure that the overall picture of the balance of risks is internally consistent. An input to such a discussion would of course be an estimate of ρ , which is a measure of the average correlation in data. The discussion may then focus on to what extent the present situation warrants an adjustment of the historical average.

For purposes of illustration, the parameter ρ has been set to 0.5 for the d.f.

displayed in Figure 2a, which implies a correlation between inflation and output growth of at least 0.5. The positive correlation between inflation and output growth is clearly evident in the bivariate p.d.f. surface shown in figure 2b. It is also clear from this figure that it would be of considerable interest to consider inflation forecasts conditioned on output growth, since the inflation distribution varies with the given output growth.

Figure 2a. Bivariate distribution for output and inflation


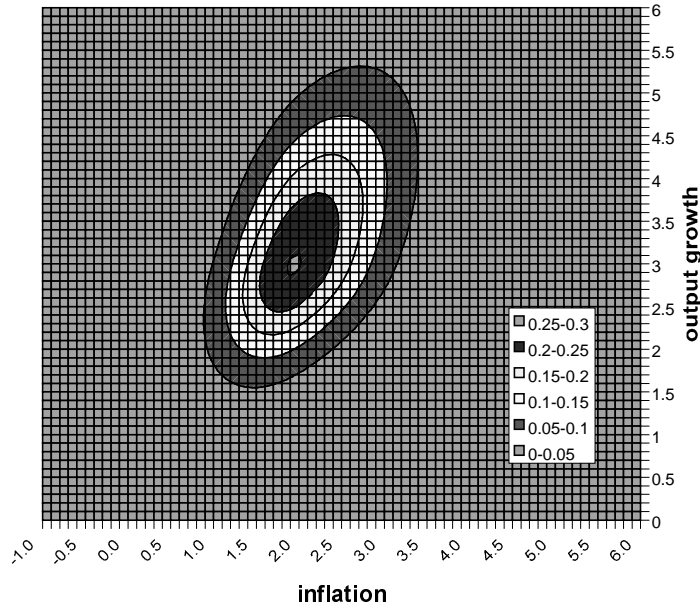


Figure 2b. Bivariate distribution surface

What is of special interest in this kind of bivariate analysis is the possibility of discussing what the inflation forecast uncertainty looks like under different assumptions regarding growth of output. This can easily be accomplished within our bivariate framework by computing the relevant conditional distributions,

$$\begin{aligned}
 (21) \quad f(\pi|y) &= \frac{h(\pi, y)}{g(y)} \\
 &= \frac{f(\pi)}{\sqrt{1-\rho^2}} \exp\left\{-\frac{\rho}{2(1-\rho^2)} \left[\rho(J_\pi^2(\pi) + J_y^2(y)) - 2J_\pi(\pi)J_y(y) \right]\right\}.
 \end{aligned}$$

In Figure 3a-b we compare the inflation forecast distribution conditional on growth in output of 3 percent with the distribution conditional on output growth of 4 percent. We see that, because of the positive correlation between inflation and output, the inflation mode forecast is higher conditional on output growth of 4 percent (2.3 percent inflation) compared to output growth of 3 percent (1.9 percent inflation). The conditional distribution of the inflation forecast has shifted to the right. And moreover, there is also more upside risk to the inflation forecast conditional on the higher growth of output.

Figure 3a. Inflation distribution conditional on 3% output growth





Figure 3b. Inflation distribution conditional on 4% output growth



5. Inflation Forecast Distribution Targeting

In this section we turn to the question of the rationale for the framework developed above in relation to monetary theory and in the practical implementation of monetary policy.

5.1 General issues

When a central bank takes asymmetric risks into account in the setting of monetary policy, it is implicitly deviating from a large body of models in the literature, see Bernanke et.al. (1999) for an overview. In many such models asymmetric risks are not included or do not arise. For example, Svensson (1999) has emphasised that in a linear model with additive uncertainty and a quadratic loss function, certainty equivalence holds. This implies, among other things, that monetary policy guided by inflation forecast targeting should consider the mean of future inflation.

However, this result hinges on a number of assumptions that may be violated. First of all, the quadratic loss function is merely an approximation, albeit a convenient one, of a general social welfare function for society. Depending on the form of the welfare function, it might not be enough to consider only a second order Taylor expansion with inflation and output as a proxy. For example, there might be other variables in the welfare function, or in the minds of policy makers.

Second, with small (not unreasonable) perturbations in the model, such as with multiplicative uncertainty, non-linear Phillips curves or with price/wage rigidities certainty equivalence no longer holds. As a result, it is no longer enough to consider only the mean of the loss function; the first order condition for an optimum depends on the whole distribution.

These are arguments that favour monetary policy being guided by inflation forecast *distribution* targeting (IFDT), a term coined in Svensson (1999). Simply put, it means that when the risks are asymmetric or the tails are larger than the Gaussian, policy makers should consider the whole forecast distribution when setting interest rates. It is not enough to simply use the mean, median or mode. Each measure of central tendency serves to provide different information about the forecast distribution in situations when it deviates from the Gaussian.

For example, assume that the forecast for inflation for the target horizon is slightly above the target but the downside risk is larger than the upside. Consequently, the median and the mean will be below the mode, perhaps differing only insignificantly from the target. In this example, the asymmetric risks thus provide one argument for keeping monetary policy unchanged. Of course, the reverse situation is also conceivable: the mode forecast is on target, but the upside risks are larger. Should monetary policy stay on hold or weigh in the fact that the median and mean are above target? This is a difficult question, but which is arguably very important for policymakers.

Another important question arises if the inflation forecast distribution significantly departs from the Gaussian, for example, by being bimodal. A bimodal distribution may represent two extreme but not implausible outcomes, such as the distribution for productivity growth under the so called “new” and “old” economic paradigms. In such a situation, most measures of central tendency are problematic (in the sense that the forecast uncertainty in each measure is large). Although IFDT does not “solve” the problem for monetary policy, it may more clearly focus policymakers on the risks represented by the forecast distribution.

It is difficult to generalise the role of uncertainty and skewness for monetary policy, but we can make a few remarks. First of all, the degree of skewness (i.e. how asymmetric the risks are) might play a role. If the difference between the measures is only a tenth of a percentage point, it would be nonsensical to let this have implications for policy. Second, the degree of uncertainty might be important. If the uncertainty is very large (the fan chart⁷ is very wide) this may imply that a *given* amount of skewness will tend to be less important for monetary policy. In other words, the larger the variance the less economic significance might be attributed to excess skewness as general forecast uncertainty becomes dominating. Finally, if the uncertainty is large this may in itself be an argument in favour of a more cautious policy, i.e. for leaving the instrument rate unchanged.

5.2 Skewness and uncertainty in practice

What guides policymakers in practice? In the appendix we have collected most of the statements concerning how uncertainty affects policy from past inflation reports (1997:4-1999:3) published by the Riksbank⁸ and a selection of those published by the Bank of England. The statements show that large uncertainty has been used as an argument for caution in setting interest rates and that skewness also can affect monetary policy, just as we discussed above. That uncertainty should make for more cautious policy would seem fairly uncontroversial ever since the work of Brainard (1967), but is it possible to say more about the skewness issue?

⁷ The fan chart shows the confidence bands around the median for different forecast horizons. Wallis (1999) has criticized the Bank of England for not having symmetric confidence bands around the median. The implications of this is that the Bank of England bands may have different probabilities in the tails of the distribution, which is a non-standard use of confidence intervals. We would like to point out that the Riksbank's confidence bands are symmetric around the median and have the standard interpretation.

⁸ Before 1997:4 no explicit inflation forecast was published.

As discussed in Blix and Sellin (1999), the economics department delivers a forecast and an uncertainty assessment to the Executive Board, which might be accepted or revised; in consequence, the assessment in the inflation report reflects the majority view of the board. Now, on several occasions it has transpired from published minutes⁹ that there has been dissent within the board. In such circumstances, skewness in the assessment can be used as a bargaining tool within the board. Suppose for instance that one member thinks a given forecast for some macroeconomic component contributing to future inflation is too high (low), but the others think the forecast is about right. Then it might be possible to reach a compromise by keeping the mode forecast unchanged, but introducing skewness in the desired direction. Using skewness to reach a compromise solution can thus be a tool whereby dissent is internalized when the views are not too far apart. In our analytical framework this is easy to incorporate by adjusting the skewness in the desired macro variable affecting inflation or output.

⁹ For example, minutes from Executive Board meetings on 1999/8/12, 1999/10/5, 1999/11/30 and 1999/12/8 (published on the Riksbank's web page www.riksbank.com).

Finally, data-revisions are an important source of uncertainty to policy makers. Recently revisions to Swedish GDP growth for the second quarter of 1999 lowered the estimate by almost half-a-percentage point. The recent US revisions of GDP and productivity are also illustrative of this problem: the data policy-makers have on the table at the decision moment may be quite different than the final revised data. Orphanides (1997), for instance, has compared the US output gap that was available at the time to the final figures used in econometric models today and found important differences. This is a source of uncertainty that is thus very real to policy-makers and which cannot be neglected. This form of uncertainty consequently makes it important to look at a wide variety of leading indicators to complement official statistics.

6 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 and output. 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 and output. Having made those subjective assessments we have then attempted to be as rigorous as possible in deriving the probability distribution for the inflation and output forecasts. A bivariate distribution for the inflation and output forecasts is then derived with given marginal distributions. The bivariate distribution makes it possible to discuss inflation forecasts conditional on specific assumptions regarding the growth rate of output.

Appendix: Quotes about uncertainty from inflation reports

Sveriges Riksbank:

"In addition to a main scenario, which is regarded as most probable, the Riksbank works on a number of alternative or risk scenarios; these are incorporated in the final assessment and accordingly also influence the construction of monetary policy." Inflation Report 1997:4, Sveriges Riksbank, page 35.

"Considering the elements of uncertainty in the assessment of inflation's future rate, monetary policy has to be conducted with caution." Inflation Report 1997:4, Sveriges Riksbank, page 36.

"The element of uncertainty in the inflation assessment can accordingly influence monetary policy's construction. A high degree of uncertainty can be a reason for giving policy a more cautious turn". Inflation Report 1998:1, Sveriges Riksbank, page 31.

"Monetary policy also has to consider the altered risk spectrum. The possibility of lower inflation, mainly on account of the Asian crisis, is still there. The crisis could worsen, in which case the repercussions on international activity and the Swedish economy might be more extensive and prolonged. The likelihood of higher inflation, generated by stronger wage increases and weak productivity growth, has decreased." Inflation Report 1998:1, Sveriges Riksbank, page 33.

"Inflation in the main scenario is now below the target in the perspective of twelve to twenty-four months that is most relevant for monetary policy. The reported assessment of uncertainties supports this. Against this background it is concluded that, at least for a time, the monetary conditions can be moved in a somewhat more stimulatory direction without risking fulfilment of the inflation target". Inflation Report 1998:2, Sveriges Riksbank, page 43.

"The spectrum of risks is also important for the construction of monetary policy. On this occasion the uncertainty in the inflation assessment is appreciably greater than usual on account of the financial market unrest and the consequences it and other factors may have for international economic development". Inflation Report 1998:3, Sveriges Riksbank, page 49.

"The spectrum of risks also has to be considered in the formation of monetary policy". Inflation Report 1998:4, Sveriges Riksbank, page 47.

"The assessments lead to the conclusion that, after adjustments for transitory effects from indirect taxes, subsidies and interest rates, the rate of inflation twelve to twenty-four months ahead will be somewhat below the Riksbank's target. In relation to the uncertainty associated with such assessments, however, the deviation from target is not sizeable". Inflation Report 1999:1, Sveriges Riksbank, page 45.

Bank of England:

“Inflation is a monetary phenomenon, and it is monetary policy which determines the rate of inflation. The lags between changes in monetary policy and changes in inflation are known only imprecisely, and will vary with the state of the economy. That is way monetary policy is set in relation not to the current rate of inflation but to inflationary trends over the next year or two.” Inflation Report, February 1993, Bank of England, page 5.

“Inflation is a monetary phenomenon: more rapid monetary growth will, other things being equal, lead to more rapid inflation. But the transmission from changes in monetary policy to changes in the rate of inflation is a complex process which is likely to change over time.” Inflation Report, May 1993, Bank of England, page 158.

“Monetary policy is based on an assessment of where inflation is headed some two years hence. It is this projection which is most relevant to setting monetary policy.” Inflation Report, February 1994, Bank of England, page 36.

“In the light of the central projection and the risks surrounding it, the Bank continues to see the need for a moderate tightening of policy.” Inflation Report, February 1997, Bank of England, page 54.

“ Given those uncertainties, the MPC concluded that monetary policy now reached a position at which it should be possible to pause in order to assess the direction in which the risks are likely to materialise.” Inflation Report, February 1997, Bank of England, page 50.

“There are considerable uncertainties surrounding these projections. Particular uncertainties exist in relation to the path of exchange rate, earnings, and price-cost margin. Alternative judgements about these issues led some Committee member to prefer a profile for inflation that would be higher or lower than that shown in Chart 2 – by $\frac{1}{4}\%$ to $\frac{1}{2}\%$ at the two-year forecast horizon.” Inflation Report, November 1999, Bank of England, page iii.

References

- Bernanke, B.S., T. Laubach, F.S. Mishkin, and A.S. Posen (1999), *Inflation Targeting — Lessons from the International Experience*, Princeton University Press, Princeton, New Jersey.
- Blix, Mårten and Peter Sellin (1998), “Uncertainty Bands for Inflation Forecasts,” Sveriges Riksbank Working Paper No. 65.
- Blix, Mårten and Peter Sellin (1999), “Inflation Forecasts with Uncertainty Intervals,” Sveriges Riksbank Quarterly Review No. 2, 12-28.
- Brainard, William (1967), “Uncertainty and the effectiveness of policy”, *American Economic Review* 57, Papers and proceedings, 411-425.
- 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.
- Edgeworth, F.Y. (1896), “The Compound Law of Error,” *Phil. Mag.* **41**, 207-215.
- Hamilton, J. (1989), “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica* 57 , 357-384.
- John, S. (1982), ”The two-piece normal family of distributions and its fitting”, *Communications in Statistical Theory and Methods*, **11**(8), 879-885.
- Johnson, N.L. (1949a), “Systems of Frequency Curves Generated by Methods of Translation,” *Biometrika* **36**, 149-176.
- Johnson, N.L. (1949b), “Bivariate Distributions Based on Simple Translation Systems,” *Biometrika* **36**, 297-304.
- Johnson, Kotz, and Balakrishnan (1994), *Continuous Univariate Distributions*, Vol. 1, P. 173.
- Lancaster, H.O. (1957), “Some Properties of the Bivariate Normal Distribution Considered in the Form of a Contingency Table,” *Biometrika* **44**, 289-92.
- Mardia, K.V. (1970), *Families of Bivariate Distributions*, Griffin, London.
- Nataf, A. (1962), Détermination des Distributions de Probabilités dont les Marges sont Données,” *C.R. Acad. Sci.*, Paris, **255**, 42-3.
- Orphanides, A. (1997), “Monetary Policy Rules Based on Real-Time Data,” Finance and Economics Discussion Series, 1998-03, Federal Reserve Board, December.

Svensson, Lars E.O. (1999), "Price stability as a target for monetary policy: defining and maintaining price stability", *mimeo*, Institute for International Economic Studies, Stockholm university.

Söderström, U. (1999), "Monetary policy with uncertain parameters," Sveriges Riksbank Working Paper No. 83.

Wallis, K.F. (1999), "Asymmetric Density Forecasts of Inflation and the Bank of England's Fan Chart," *National Institute Economic Review*, No. 1, 106-112.