THE CENTRAL BANKER AS A RISK MANAGER:

QUANTIFYING AND FORECASTING INFLATION RISKS

PRELIMINARY AND INCOMPLETE. DO NOT QUOTE!

December 23, 2002

Lutz Kilian

University of Michigan, European Central Bank and CEPR

Simone Manganelli

European Central Bank

Abstract: In deciding the monetary policy stance, central bankers need to evaluate carefully the risks the current economic situation poses to price stability. We propose to regard the central banker as a risk manager who aims to contain medium-term inflation within pre-specified bounds. We develop formal tools of risk management that may be used to quantify and forecast the risks of failing to attain that objective. We illustrate the use of these risk measures in practice. First, we quantify and compare the historical evolution of inflation and deflation risks in the United States, Germany and Japan. Second, we show how to construct genuine real time forecasts of short- and medium term risks that may be used in policy-making. We demonstrate the practical usefulness of these forecasts in understanding the Fed's decision to tighten monetary policy in 1984, 1988, and 1994. Third, we forecast the risks of world-wide deflation for horizons of up to two years. Although recently fears of world-wide deflation have increased, we find that, as of September 2002, with the exception of Japan there is no evidence of substantial deflation risks. Fourth, and last, we discuss implications of our analysis for the specification of monetary policy rules.

JEL classification: E31, E37, E52, E58, C22.

Keywords: Inflation; Deflation; Risk; Forecast; Monetary policy.

Acknowledgements: The views expressed in this paper are our own and do not necessarily reflect the views of the European Central Bank (ECB) or its staff. We thank seminar participants at the ECB for useful comments. We also acknowledge helpful discussions with Atsushi Inoue, Athanasios Orphanides, Adrian Pagan, Matthew Shapiro, and Ken Wallis. We especially thank Athanasios Orphanides for providing access to the real-time data set used in Orphanides (2003).

"Because monetary policy works with a lag, it is not the conditions prevailing today that are critical but rather those likely to prevail six or twelve months, or even longer, from now. Hence, as difficult as it is, we must arrive at some judgement about the most probable direction of the economy and the distribution of risks about that expectation."

Alan Greenspan

Chairman of the Federal Reserve
Testimony before the Senate Committee on the Budget
January 21, 1997

"Let me repeat that … we analyse risks of deflation as well as inflation and will act accordingly to prevent both phenomena in case we detect risks in one or the other direction."

Otmar Issing
Member of the Executive Board, European Central Bank
Speech at the European Finance Convention
December 2, 2002

"Suppose a central bank adopts dramatic easing actions way beyond "standard" or rule-based monetary policy at an early stage. The intention would be to create an accelerated inflation pre-emptively to avoid the future risk of a deflation … The central bank pursuing such a strategy would have to be fully convinced, substantiated by quantitative analyses, and strongly concerned about the risk of deflation a few years into the future."

Yukata Yamaguchi

Deputy Governor of the Bank of Japan
Remarks at a Symposium sponsored by the Federal Reserve Bank of Kansas City
August 30, 2002

1. Introduction

The quotes reported above demonstrate that risk management is an integral part of the art of central banking. In deciding the monetary policy stance, central bankers need to evaluate carefully the risks that the current economic situation poses to price stability.

From this perspective, it is somewhat surprising that closer links between modern risk management techniques and monetary policy decisions have not been established.

What are central bankers referring to when they talk about upside or downside risks to price stability? There is a general consensus that upside risks refer to inflation in excess of a certain threshold. Similarly, downside risks are commonly expressed as inflation below a certain threshold. For the European Central Bank, Issing (2002a) notes that the "objective of maintaining price stability in the euro area … was clearly intended from the start to mean *neither* prolonged inflation *nor* prolonged deflation." He defines the objective of price stability as the objective of containing "medium term inflation … below 2%" in the harmonized CPI, while "maintaining flexibility at the lower bound", defined as inflation of 0% or slightly in excess of 0%.[1] Although not all central banks are as explicit about their objectives as the European Central Bank, most central banks at least implicitly appear to pursue similar objectives.[2]

Thus, it is natural to frame the problem of managing the risks to price stability in terms of keeping the inflation rate within a well-defined band. This task is made more difficult by the great uncertainty about future levels of inflation when making monetary policy decisions. Among practitioners, there has been a gradual shift in recent years toward the explicit recognition that inflation uncertainty is pervasive. From a purely statistical standpoint, knowledge of the probability distribution of a random variable of interest (such as the inflation rate) provides a complete and exhaustive description of its underlying uncertainty. Econometric techniques to estimate and represent the probability distribution of the inflation rate are easily available. The "fan chart", popularised by the Bank of England, is one such example. The fan chart graphs the central 10% prediction interval as a dark band and graphs successively wider intervals in ever lighter shades of red. The selective shading of the intervals is intended to draw attention to the uncertainty in future inflation. The mapping from the fan chart to an assessment of the risks to price stability, however, are left to the eye of the beholder.

---

[1] For example, Issing (2002a) notes that "inflation and interest rate levels excessively close to zero entail risks of deflation". Specifically, he argues that "most of the problems associated with the lower bound on nominal interest rates could be avoided for rates of inflation as low as 1%."

[2] For example, the Bank of New Zealand targets a range of 1%-3% (formerly 0%-3%); the Bank of Canada aims at 1%-3% of inflation.

The exercise of extracting a precise indication of the risk of inflation or the risk of deflation from the predictive density of inflation is nontrivial. Part of the problem is that graphs or tables of the probability distribution contain more information that can be processed easily. In practice, decision-makers are primarily interested in certain aspects of the distribution of inflation. In this paper, we will propose techniques that will allow the reader to construct indicators of the risk of excessive inflation and of the risk of deflation. We also propose statistical measures of the inflation rates that may be expected, in the event that excessive inflation does occur (or, alternatively, in the event that deflation does occur). Our objective is to reduce the information contained in the probability distribution to a few key indicators of risk that are easy to interpret and that at the same time effectively summarise the information relevant for policy makers. Our statistical measures of risk may be used to shed light on historical episodes. They may also be used to compute real-time forecasts of the risks of excessive inflation and the risks of deflation. These forecasts may serve as an input into monetary policy decision making.

The remainder of the paper is organized as follows. In section 2 we propose measures of risk that are specifically suited for measuring the risks of excessive inflation and of deflation. In section 3 we show how to compute these risk measures and illustrate their interpretation for U.S. CPI inflation data. Based on this discussion, we propose an index of the risk of excessive inflation and a corresponding index of the risk of deflation, as well as a measure of the balance of upside and downside risks to price stability. Section 4 contains a comparison of historical inflation and deflation risks in the United States, Japan and Germany. We also discuss some possible determinants of historical inflation risks. In section 5 we construct genuine real-time forecasts of medium term risks to price stability and illustrate their usefulness in understanding policy decisions. We focus on the example of the 1984, 1988 and 1994 decisions of the Fed to raise the Federal Funds rate. In section 6, we formally assess the risks of world-wide deflation as of September 2002, which have recently received much attention among policymakers and the general public. We forecast the risks of inflation and of deflation for a horizon of up to two years, based on data for Germany, Japan and the United States. We conclude in section 7 with a discussion of the implications of our analysis for the specification of monetary policy rules.

2. MEASURES OF INFLATION RISK

2.1. Motivation

Notwithstanding the important role that central bankers assign to managing the risks to price stability, it appears that in practice the assessment of inflation risks is largely judgmental. No formal tools are used to quantify and forecast inflation risks. This practice contrasts sharply with the best practice in many areas of the private sector. Formal risk management tools are used on a routine basis in many areas of financial management. In designing measures of inflation risk it is therefore useful to start with a review of the measures of risk currently in use in finance. A general measure of the downside risk of the random variable x at a given point in time is given by

$$(1) \qquad \underline{\rho}_{\alpha,t,\eta,k} = k \int_{-\infty}^{t} |x-\eta|^{\alpha} \, dF(x), \qquad \alpha \geq 0,$$

where $F(x)$ is the probability of obtaining a realization that does not exceed x given the current information set, η is the reference point, relative to which deviations are measured, α is a measure of the relative impact of small and of large deviations, $\underline{t}$ is a range parameter that specifies what values are to be included in the risk measure, and k is a normalizing constant.[3] The empirical measures of risk discussed in the finance literature may be understood as special cases of this general measure of risk. An analogous measure of upside risk, denoted by $\overline{\rho}_{\alpha,t,\eta,k}$ is obtained by integrating from $\bar{t}$ to ∞ instead. We will discuss these measures and their potential limitations one by one, and we will show how they may be modified for the purpose of assessing inflation risks.

Since Markovitz (1952), portfolio choice problems have been posed as a trade-off between expected returns and risk. For example, in Markovitz' original mean-variance model risk is associated with the variance of returns. Specifically, it is assumed that investors have a preference for expected returns and an aversion to variable returns. This *variance* measure of risk is obtained from (1) by setting $\alpha = 2$, $\eta = E(x)$, $\underline{t} = \infty$ and $k = 1$.

---

[3] Expression (1) is a slight modification of the risk measure proposed by Stone (1973).

(2) $$\rho = \int_{-\infty}^{\infty} |x - E(x)|^2 \, dF(x)$$

The use of the variance of returns as a measure of risk is consistent with an expected utility framework only to the extent that agents' preferences assume a quadratic utility form, or that returns are multivariate normally distributed. Despite this drawback, the mean-variance portfolio model has become extremely popular, essentially because of its analytical tractability and its rich empirical implications. As a measure of risk, however, the variance concept has been subject to severe criticism. Its main weakness is that, by construction, the variance weights equally upside and downside risks. Most people would agree that portfolio managers care more about the risk of a downside movement in returns than the risk of an upside movement.

An alternative risk measure, which is robust to this criticism, is the *semi-variance*, introduced by Markovitz (1959). The semi-variance measure of risk is obtained from (1) by setting $\alpha = 2$, $\underline{t} = \eta = E(x)$ and $k = 1$.

(3) $$\underline{\rho} = \int_{-\infty}^{E(x)} |x - E(x)|^2 \, dF(x)$$

By focusing on only one side of the distribution, the semi-variance concept improves upon the variance as a measure of risk, but, in practice, risk managers are often interested in deviations from a pre-specified bound rather than deviations from the population mean. This is particularly true for central bankers. For example, the European Central Bank has defined price stability as an inflation rate bounded between close to 0% and 2%. If the semi-variance is taken as a measure of inflation risk, a situation where inflation oscillates between 0% and 2% would be considered as risky as one, in which inflation oscillates between 2% and 4%. Most people would agree that the second scenario poses greater threats to price stability than the first one. This ambiguity of the semi-variance concept is addressed by the introduction of the *target semi-variance*, which is computed by squaring the distance of each observation from a certain threshold, $\underline{t}$ or $\bar{t}$, that is of economic interest. If we are concerned with a lower bound, the target semi-variance measure of downside risk is obtained from (1) by setting $\alpha = 2$, $\eta = \underline{t}$ and $k = 1$.

(4) $$\underline{\rho} = \int_{-\infty}^{\underline{t}} |x - t|^2 \, dF(x)$$

Alternatively, if we are concerned with an upper bound, as in the inflation example, we obtain:

(4′) $$\overline{\rho} = \int_{\overline{t}}^{\infty} |x - t|^2 \, dF(x)$$

The measure of upside risk in (4′) assigns zero penalty to observations that are at or below the threshold. The penalty assigned to observations above the threshold increases disproportionately with their distance from the threshold. In the inflation example above, setting the threshold, $\overline{t}$, to 2%, the target semi-variance would be able to identify the second scenario (where inflation oscillates between 2% and 4%) as riskier than the first one (where inflation moves between 0% and 2%).

Alternatively, it has been proposed to compute the probability of the random variable not remaining within a pre-specified range of values. Such a measure may be motivated in the context of the problem of maximizing expected returns subject to the constraint of keeping low the probability of a catastrophic event - defined as returns falling below a certain threshold $\underline{t}$ (see, e.g., Roy 1952). This probability measure of risk can be obtained from (1) by setting $\alpha = 0$ and $k = 1$.

(5) $$\underline{\rho} = \int_{-\infty}^{\underline{t}} dF(x) = \Pr(x < t) = \underline{\theta}$$

In the case of portfolio risk management, the primary concern is downside risk. More generally, risk may occur at the lower end as well as the upper end of the distribution of the random variable of interest. In the latter case, we have

(5′) $$\overline{\rho} = \int_{\overline{t}}^{\infty} dF(x) = \Pr(x > t) = \overline{\theta}$$

In the inflation example, the central banker may compute the probability of deflation based on expression (5) with $\underline{t} = 0\%$ and the probability of excessive inflation from (5′) with $\overline{t} = 2\%$.

Although intuitively appealing, the probability measure of risk in (5) and ($5'$) is not widely used in practice. One reason is that this measure – similarly to statistical measures that simply express the risk in terms of standard deviations from a threshold value - does not allow us to quantify risks in the units of ultimate interest to the decision maker. For example, a risk manager in the financial sector finds it easier to relate to risks when they can be expressed in terms of foregone profits. Similarly, a central banker is naturally interested in quantifying inflation risks in terms of additional percentage points of inflation. In response to this criticism, the financial industry more recently has developed an alternative measure of risk, commonly known as *Value at Risk* (*VaR*).[4]

The value at risk of a portfolio describes the monetary loss (or change in net income) that may occur over a given period, at a given confidence level $\underline{\theta}$, due to adverse movements of asset prices. Statistically, the *VaR* measure of risk is given by the quantile of the distribution of net income that corresponds to the pre-specified tail probability $\underline{\theta}$. This quantile in practice is obtained from the corresponding $\underline{\theta}$-quantile of the distribution of returns by weighting the return by the value of the underlying portfolio. Formally, the *VaR* measure of downside risk is obtained by inverting the probability measure in (5):

(6) $$\underline{\rho} = t \quad s.t. \int_{-\infty}^{\underline{t}} dF(x) = \underline{\theta}$$

where the threshold $\underline{t}$ will depend on the tail probability $\underline{\theta}$. Similarly, we could compute the upside risk

($6'$) $$\overline{\rho} = t \quad s.t. \int_{\overline{t}}^{\infty} dF(x) = \overline{\theta}$$

The basic idea is that instead of solving for the probability for a pre-specified threshold $\underline{t}$ or $\overline{t}$ as in (5) and ($5'$), in (6) and ($6'$) we solve for the quantile corresponding to a pre-specified probability $\underline{\theta}$ or $\overline{\theta}$.

Although the *VaR* approach has been developed for asset returns in finance, it is

[4] The theoretical foundations of *VaR* can be traced back to the contributions of Roy (1952), Kataoka (1963), Arzac and Bawa (1977) and Fishburn (1977).

equally straightforward to apply this approach to the problem of managing the value of the domestic currency. The value at risk at a given point in time can simply be identified with the quantile of the distribution of inflation rates corresponding to a pre-specified tail probability, $\underline{\theta}$ or $\overline{\theta}$. The resulting risk measure is already expressed in the unit of ultimate interest. This fact suggests that the *VaR* concept would be a natural choice for central bankers interested in managing inflation risks.

The *VaR* approach solves the problem of expressing risk in units that have a natural economic interpretation, but – similarly to the probability measure in (5) and ($5'$) - it is subject to another important criticism: For any given probability $\underline{\theta}$, the value at risk corresponds to the financial loss that will be exceeded with probability $\underline{\theta}$, but the *VaR* measure is silent about the amount by which this loss will be exceeded. It does not take into account the severity of the downside risk. Similarly, the probability measure gives no indication by how much the threshold will be missed, in the event of a catastrophic event.

This shortcoming may be addressed by computing a different measure of risk, the *expected shortfall* or *tail conditional expectation*. This measure of risk was recently introduced by Artzner et al. (1999). The *tail conditional expectation* is defined as the expected value of the random variable of interest, given that it exceeds a pre-specified quantile of its distribution. This quantile may be chosen either based on a fixed tail probability, $\underline{\theta}$ or $\overline{\theta}$, as in (6) and ($6'$), or based on a pre-specified threshold, $\underline{t}$ or $\overline{t}$, as in (5) and ($5'$). By setting $k = sign(x)/\Pr(x < \underline{t})$ for the lower tail and $k = sign(x)/\Pr(x > \overline{t})$ for the upper tail, respectively, and by setting $\alpha = 1$ and $\eta = 0$, we obtain:

$$(7) \qquad \underline{\rho} = \frac{\int_{-\infty}^{\underline{t}} x dF(x)}{\int_{-\infty}^{\underline{t}} dF(x)} = E(x \mid x < \underline{t}) \quad \text{and} \quad \overline{\rho} = \frac{\int_{\overline{t}}^{\infty} x dF(x)}{\int_{\overline{t}}^{\infty} dF(x)} = E(x \mid x > \overline{t})$$

In the inflation context, this measure corresponds to the expected value of inflation, conditional on inflation falling below the level associated with the $\underline{\theta}$-quantile of interest (or alternatively inflation exceeding the level associated with the $\overline{\theta}$-quantile of interest). Computing such a tail conditional expectation is important in practice. Suppose that inflation will exceed a level of 2% with probability $\overline{\theta}$, then surely a central banker must be able to distinguish between a situation in which we can expect inflation to be 2.1 % with probability $\overline{\theta}$, and one in which we expect inflation to be 4% with probability $\overline{\theta}$. It is exactly this question that the tail conditional expectation is designed to answer.[5]

2.2. Definitions

In subsection 2.1., we have considered measures that allow the assessment of inflation risk at a point in time. In that case, the probability risk measure (5) and the *VaR* risk measure (6) are formally equivalent and by implication so are the tail conditional expectations associated with these two risk measures. This formal equivalence breaks down when we consider sequences of inflation risks computed on historical data or sequences of real-time forecasts of inflation risks. In that case, the probability measure of inflation risk is based on the notion of a time-varying tail probability and a time-invariant threshold, whereas the *VaR* measure of risk is based on the notion of a time-invariant tail probability and a time-varying threshold. Thus, there may be real differences of interpretation when these measures of risk are applied to actual data, and we need to consider the merits of these measures separately.

We now proceed by defining formally the analogues of the *probability* measure (5), the *VaR* measure (6), and the corresponding *tail conditional expectations* (7) in the context of assessing inflation risks. These definitions are required to lay the groundwork for section 3, in which we will contrast the properties and interpretation of these alternative measures. Based on this comparison, in section 3 we will define an *Index of Inflation Risk* (*IIR*) that corresponds to our preferred measure of (upward) inflation risk. We also will propose an analogous *Index of Deflation Risk* (*IDR*). For expository purposes, we identify excess inflation risk with inflation rates in excess of 2% and deflation risk with inflation rates below 0%. We will focus on one-step ahead measures of risk throughout this paper. This convention involves no loss of generality because, as we

---

[5] A related measure proposed by Fishburn (1977) and recently reintroduced by Basak and Shapiro (2001) is the *weighted expected shortfall*, defined by $\rho = \int_{-\infty}^{\underline{t}} x dF(x) = E(x \mid x < \underline{t}) \Pr(x < \underline{t})$. It is obtained from (1) by setting $\alpha = 1$, $\eta = 0$, and $k = sign(x)$. Unlike the expected shortfall this alternative measure of risk does not express risk in terms of the units of interest and hence will not be pursued here.

will show later, we may compute one-step ahead risks for any horizon of interest by fitting a regression model to suitably aggregated inflation data.

Definition 1a (Probability of Excess Inflation, Probability of Deflation)

We begin with the probability measures in (5) and ($5'$). Formally, for each period *t*, we define the level of *Probability of Excess Inflation* and the *Probability of Deflation* as:

(8) $PEI_{t+1} \equiv \Pr(\pi_{t+1} > 2\% \mid \Omega_t)$

and

($8'$) $PD_{t+1} \equiv \Pr(\pi_{t+1} < 0\% \mid \Omega_t)$,

respectively, where Ω_t denotes the information set.

The benchmark value of 2% for the upper bound is chosen for expository purposes only. Depending on the economic context, other benchmark values may be more appropriate (see footnote 2). Similarly, the benchmark value of 0% is chosen for expository purposes only. For example, it may be reasonable to alternatively set the benchmark equal to 1% (see footnotes 1 and 2).

The probability measure answers the question of what the probability is that inflation will exceed a pre-specified level of 2% next period. As noted earlier, this probability measure does not convey by how much the threshold is likely to be exceeded. The expected excess inflation is measured by the corresponding tail conditional expectation. Clearly, neither measure alone would allow a proper assessment of the risks to price stability. Thus, it is natural to view the probability measure of risk and the corresponding *tail conditional expectation* as conjugates that only in conjunction allow a proper assessment of the risk.

Definition 1b (Expected Excess Inflation, Expected Deflation)

We define the measures of *Expected Excess Inflation* (EEI_{t+1}^{PEI}) and *Expected Deflation* (ED_{t+1}^{PD}) associated with (8) and ($8'$) as:

(9) $EEI_{t+1}^{PEI} \equiv E(\pi_{t+1} \mid \pi_{t+1} > 2\%)$

and

$$(9') \qquad ED_{t+1}^{PD} \equiv E(\pi_{t+1} \mid \pi_{t+1} < 0\%)$$

respectively, where the dependence of the expectation on Ω_t has been suppressed.

A key feature of Definitions 1a and 1b is that the probability, with which inflation exceeds a given benchmark, is time-varying, whereas the benchmark itself is time-invariant. Alternatively, we may fix the tail probability ex ante and study how the threshold corresponding to that time invariant probability evolves over time. This is the approach underlying the *VaR* measure in (6) and ($6'$). Definitions 2a and 2b below are analogues of the risk measurement tools used in finance. Specifically, we are asking what level of inflation will be exceeded with a pre-specified probability of $\overline{\theta}$. We are also asking, how high we can expect inflation to be, conditional on this $\overline{\theta}$-probability event.

Definition 2a (Inflation at Risk, Deflation at Risk)

Formally, for each period *t*, we define the level of *Inflation at Risk* (IaR_{t+1}) in analogy to the *VaR* literature as the solution to

$$(10) \qquad \Pr(\pi_{t+1} > IaR_{t+1} \mid \Omega_t) = \overline{\theta},$$

where Ω_t denotes the information set.

Here $\overline{\theta}$ is the pre-specified probability with which the inflation rate next period, π_{t+1}, will exceed the level IaR_{t+1}. For example, if the value of IaR_{t+1} associated with $\overline{\theta}$ = 25% is estimated to be 2.2%, then the central banker knows that inflation over the next period (that could be next month, quarter or year) will exceed a level of 2.2% with probability 25%. In practice, $\overline{\theta}$ must be chosen by the central banker. Clearly, $\overline{\theta}$ must be reasonably high to capture the risk of a rare event (as opposed to a likely event), but not too high, lest the event becomes too rare to be of much interest to policy makers.

We similarly define the corresponding measure for the downside risk. Let *Deflation at Risk* (DaR_{t+1}) refer to the solution of

$$(10') \qquad \Pr(\pi_{t+1} < DaR_{t+1} \mid \Omega_t) = \underline{\theta},$$

where $\underline{\theta}$ is the pre-specified probability with which inflation next period will be lower than the level DaR_{t+1}.

Definition 2b (Expected Excess Inflation, Expected Deflation)

We define the measures of *Expected Excess Inflation* (EEI_{t+1}^{IaR}) and *Expected Deflation* (ED_{t+1}^{DaR}) associated with (10) and (10′) as:

(11) $$EEI_{t+1}^{IaR} \equiv \max[E(\pi_{t+1} \mid \pi_{t+1} > IaR_{t+1}), 2\%]$$

and

(11′) $$ED_{t+1}^{DaR} \equiv \min[E(\pi_{t+1} \mid \pi_{t+1} < DaR_{t+1}), 0\%]$$

respectively, where the dependence of the expectation on Ω_t has been suppressed.

Note that by construction the measures of expected excess inflation and expected deflation in Definition 1b and Definition 2b will differ in practice, because they condition on different events.

3. COMPUTING MEASURES OF INFLATION RISK IN PRACTICE

3.1. Estimation methodology

We illustrate the estimation of the inflation risk measures defined in section 2.2. in the context of a simple parametric model of inflation dynamics:

(12) $$\pi_{t+1} = \mu_{t+1} + \mathrm{u}_{t+1}, \qquad \mathrm{u}_{t+1} = \sqrt{h_{t+1}}\varepsilon_{t+1}, \qquad \varepsilon_{t+1} \mid \Omega_t \sim i.i.d.(0,1)$$
$$h_{t+1} = \varpi + \alpha\varepsilon_t^2 + \beta h_t$$

where Ω_t denotes the information set at date *t*. The use of the GARCH framework for modeling the conditional variance of inflation, here denoted by h_t, was originally suggested by Engle (1982) and Bollerslev (1986). A natural choice for the conditional mean, μ_t, is an autoregressive model for inflation:

$$\mu_{t+1} = c + \sum_{i=1}^{p} \phi_i \pi_{t-i+1} ,$$

where *p* denotes the autoregressive lag order (see Engle 1982, Bollerslev 1986). Alternatively, one may wish to model explicitly various other variables that may affect future inflation, such as the percent change of oil prices or monetary aggregates. In that case, the model for μ_t would take the form:

$$\mu_{t+1} = c + \sum_{i=1}^{p} \phi_i \pi_{t-i+1} + \sum_{k=1}^{s} \sum_{j=1}^{q_k} \varphi_{k,j} x_{k,t-j+1}$$

where the additional predictors, x_k, may enter with potentially different lag orders for each predictor. In our empirical work we will consider alternative specifications of the conditional mean and use statistical criteria to choose between them.

We propose a three-step procedure for estimating the inflation risk measures based on model (12). In step 1, we estimate the first and second moments of inflation. In step 2, we compute the θ-quantile of $\pi_{t+1} \mid \Omega_t$, where $\theta \in [0,1]$. In step 3, we compute the corresponding tail conditional expectation. Given that the number of parameters can become quite large in practice, estimation of model (12) by full maximum likelihood methods tends to be numerically challenging. Therefore, in step 1, we first estimate the conditional mean equation by ordinary least squares (OLS), and then fit the GARCH(1,1) model to the OLS residuals. This estimation procedure, although inefficient, is consistent and greatly reduces the computational burden. Consistent estimates of the GARCH parameters are obtained by maximizing the quasi-log-likelihood function (see Bollerslev and Wooldridge 1992).

Given the estimate of model (12), we are in a position to compute the quantile of interest of the distribution of $\pi_{t+1} \mid \Omega_t$. Denote the estimated standardized residual by $\hat{\varepsilon}_{t+1} \equiv (\pi_{t+1} - \hat{\mu}_{t+1}) / \sqrt{\hat{h}_{t+1}}$. By assumption, these standardized residuals are i.i.d. with zero mean and variance 1. Various options are available for estimating the θ-quantile of an i.i.d distribution.[6] Here we use the empirical quantile method. We sort the standardized residuals in ascending order. Then for a sample of size *T*, the θ-quantile estimator of this residual distribution is given by the θ*Tth* residual.[7] Given this estimate, denoted by $\hat{q}^{\varepsilon}_{\theta,t+1}$, and the estimates of $\hat{h}_{t+1}$ and $\hat{\mu}_{t+1}$, we can compute in step 2, for any probability level *θ*, the *θ*-quantile of π_{t+1}, denoted by $\hat{q}^{\pi}_{\theta,t+1}$:

---

[6] For example, one may use kernel estimators, linear interpolation methods or flexible parametric approximations.

[7] When θ*T* is not an integer, we interpolate between the closest values.

(13) $$\hat{q}^{\pi}_{\theta,t+1} = \hat{\mu}_{t+1} + \sqrt{\hat{h}_{t+1}}\hat{q}^{\varepsilon}_{\theta,t+1}.$$

Step 3 of our procedure is the estimation of the tail conditional expectation. In practice, we rely on the regression method proposed by Manganelli and Engle (2001) that exploits the i.i.d. property of the standardized residuals. Suppose we are interested in the upper tail conditional expectation. Observe that:

$$\begin{aligned} E_{t+1}(\pi_{t+1} \mid \pi_{t+1} > q^{\pi}_{\theta,t+1}) &= \mu_{t+1} + E_{t+1}(\sqrt{h_{t+1}}\varepsilon_{t+1} \mid \pi_{t+1} > q^{\pi}_{\theta,t+1}) \\ &= \mu_{t+1} + \sqrt{h_{t+1}}E(\varepsilon_{t+1} \mid \pi_{t+1} > q^{\pi}_{\theta,t+1}) \end{aligned}$$

where the first and second equalities follow from the application of the law of iterated expectations. Since the expected value term in the second equality is constant by the i.i.d. assumption of model (12), we may drop the $t+1$ subscript from the expectation. This expectation $E(\varepsilon_{t+1} \mid \pi_{t+1} > q^{\pi}_{\theta,t+1}) \equiv \gamma$ may be estimated by the sample mean of $\hat{\varepsilon}_{t+1}$ for all observations that satisfy $\pi_t > \hat{q}^{\pi}_{\theta,t+1}$. Denoting by $\hat{\gamma}$ the estimator of γ, the upper tail conditional expectation can be written as:

$$\hat{E}_{t+1}(\pi_{t+1} \mid \pi_{t+1} > \hat{q}^{\pi}_{\theta,t+1}) = \hat{\mu}_{t+1} + \hat{\gamma}\sqrt{\hat{h}_{t+1}}$$

The lower tail conditional expectation may be obtained analogously.[8]

3.2. Model Estimation Results

We now implement the estimation procedure discussed in subsection 3.1. for U.S. data. For expository purposes, we will focus on risks at the one-month horizon. The computation of risks for longer horizons involves added complexity and will be addressed in sections 4 and 5. Our objective in this section is simply to illustrate the interpretation of risk measures in the simplest possible setting. The data used are the seasonally adjusted CPI for all urban consumers, as released by the Bureau of Labor Statistics, and the West Texas Intermediate spot oil price series. The sample runs from January 1947 through December 2001. Consumer price inflation rates and the percent change in oil prices were computed as the month-to-month log difference of the respective price index and subsequently annualized.

---

[8] For extreme quantiles, i.e. when θ is close to 0 or close to 1, this regression method becomes unreliable, as very few observations are available in the extreme tail. An obvious alternative in these cases is to appeal to extreme value theory (see Manganelli and Engle 2001). In practice, however, values of θ very close to 0 or to 1 are unlikely to be of interest to central bankers, so for our application standard regression methods are adequate.

We consider four alternative specifications of the conditional mean in model (12): one involving only lagged inflation ($\Delta\pi_t$) and the other three models including in addition lagged percentage changes of the oil price (Δwti_t), lagged growth rates of M3 ($\Delta m3_t$) or both.

Model 1: $\mu_{t+1} = c + \sum_{i=1}^{p} \phi_i \pi_{t-i+1}$,

Model 2: $\mu_{t+1} = c + \sum_{i=1}^{p} \phi_i \pi_{t-i+1} + \sum_{j=1}^{q} \varphi_{k,j} \Delta wti_{t-j+1}$

Model 3: $\mu_{t+1} = c + \sum_{i=1}^{p} \phi_i \pi_{t-i+1} + \sum_{j=1}^{q} \varphi_{k,j} \Delta wti_{t-j+1} + \sum_{j=1}^{r} \zeta_{k,j} \Delta m3_{t-j+1}$

Model 4: $\mu_{t+1} = c + \sum_{i=1}^{p} \phi_i \pi_{t-i+1} + \sum_{j=1}^{r} \zeta_{k,j} \Delta m3_{t-j+1}$

The use of oil prices and of monetary aggregates as potential additional predictors is a natural choice. Although theory does not restrict the set of conditioning information, we did not investigate the use of additional predictors, given the large number of parameters involved.

We first use the AIC to select the optimal number of lags for each model. For Model 1 the AIC selects 13 lags. For Model 2 and Model 4 the AIC selects 13 lags each. For Model 3 it selects 9 lags each. We then choose between these four models based on the AIC values of the models evaluated at these lag orders. Model 2 is preferred on the grounds that it has a slightly lower AIC value than the other models (see Table 1).

Further diagnostic tests reveal that Model 2 fits the data reasonably well. The Ljung-Box (LB) statistic for the first 15 lags does not reject the null of no serial correlation in the residuals, $\hat{u}_t$, suggesting that our AR model does a good job in describing the dynamics of the mean equation. The squared model residuals, $\hat{u}_t^2$, are highly persistent. In contrast, the Ljung-Box statistics for the first 15 lags shows no evidence of serial correlation in the time series of the squared standardized residuals, $\hat{\varepsilon}_t^2$, suggesting that the GARCH(1,1) model successfully removes all conditional heteroskedasticity in the residuals.

A different way of judging the fit of the model is to focus on the relative frequency, with which the realizations of the inflation rate fall outside the bounds given by the conditional θ quantiles. If the model is well specified, one would expect that the empirical probability of a miss would be close to the theoretical probability of θ. One also would expect that the misses are not clustered around a few dates. We found that the probability of inflation falling below the lower quartile was 25.0% and that of exceeding the upper quartile was 25.0%, with misses being evenly spread across the sample. This result provides further support for the selected model specification.[9]

In addition to evaluating the fit of our model, we also examined a variety of different model specifications and verified that the qualitative results are robust to these alterations, even when statistical tests are able to discriminate between these models. Furthermore, we estimated the same models using a variety of semi-parametric techniques based on quantile regression, as proposed by Engle and Manganelli (1999), again finding no evidence that the qualitative results are driven by the parametric specification.[10] Finally, we also repeated the analysis based on quarterly data. The results again were qualitatively similar.

3.3. Interpreting the Estimated Measures of Inflation Risk

We compute all historical measures of inflation risk by first fitting Model 2 to the full-sample of post-war data. We then recursively compute the conditional quantiles of the distribution of $\pi_{t+1} \mid \Omega_t$ and the associated measures of risk. A first crude impression of the risks to price stability in the United States is conveyed in Figure 1 by the estimated probabilities of excess inflation (upper panel) and by the estimated probabilities of

---

[9] An important concern in modelling inflation is the possibility of structural change. Stock (2002) presents evidence that the persistence of U.S. inflation has been constant and that its autocorrelations have remained stable over the post-war period. In closely related work, Sims (2002) has suggested that apparent time variation in U.S. inflation data may be accounted for as shifts in the variances of the structural disturbances. Our model allows for explicit time variation in the conditional variance. The model diagnostics suggest that this form of time variation provides a good approximation to the U.S. data. Thus, we do not need to allow for additional structural breaks. We note, however, that more generally, to the extent that structural shifts may be estimated consistently, they can be imposed in modelling the inflation process that underlies the risk estimates.

[10] Estimates of conditional quantiles are usually built by modelling first the conditional mean function, and then the higher moments of the residuals. This second step is generally carried out using mixtures of normal distributions, GARCH models, or option implied volatilities (see Tay and Wallis 2000 for a survey). This is the approach we follow in this paper. The key assumption behind this modelling strategy is that the distribution of the innovations is invariant to the values of the vector of covariates. In reality, however, this assumption may be violated, thus invalidating the parametric model in expression (12). For example, an oil price increase may increase the risk of inflation without affecting the risk of deflation. The opposite conclusion may be true in the event of a sharp reduction of oil prices, which increases the downside risks of deflation, but leaves unchanged the upside risks. The fact that our qualitative results are not affected by the imposition of a parametric model suggests that these asymmetries, although potentially important, are not an important feature of our data set.

deflation (middle panel). The probability of inflation exceeding the upper threshold of 2% was computed as the fraction of standardized residuals that exceed $\hat{q}^{\varepsilon}_{\theta,t+1} \equiv (2-\hat{\mu}_{t+1})/\sqrt{\hat{h}_{t+1}}$. This cut-off point is obtained by inverting equation (13) after substituting $q^{\pi}_{\theta,t+1} \equiv 2$. An analogous approach was used to compute the probability that inflation will fall below 0%.

The upper panel shows that throughout much of post-war U.S. history the probability of inflation exceeding 2% has been rather high. On average, the probability is 68%. Only rarely this probability falls below 30% and in some cases in the 1970s and early 1980s it reaches virtually 100%. The corresponding probability of deflation in the middle panel has historically been much lower with an average of 13.6%. Moderate risks of deflation persisted throughout the 1950s and early 1960s. Deflationary risks appear negligible from the late 1960s until the early 1980s as well as in the 1990s. A sharp, but short-lived, increase in deflationary risk is associated with the breakdown of OPEC in 1986 and two smaller spikes are plausibly related to the Volcker recession and to the collapse of oil prices in 1998. There also is a noticeable increase of deflationary risks since 2000. We postpone a detailed discussion of the historical evidence until section 4 and focus instead on the interpretation of these statistical measures.

It may be tempting to take the probability estimates in Figure 1 as prima facie evidence of the pervasive risks to price stability, but such an interpretation would be ill-advised. For example, a probability of excess inflation of near 100% does not necessarily constitute a serious risk to price stability. As noted earlier, even a very high probability of inflation exceeding 2% does not convey by how much inflation is likely to exceed the 2% threshold. For the same probability, we may be in a situation, in which inflation can be expected to barely miss the threshold or one in which inflation can be expected to surpass the threshold by a wide margin. This latter, closely related question is answered in the last panel of Figure 1 that plots the expected excess inflation and expected deflation associated with the probability measures in the upper and middle panels. For example, the probability of excess inflation in 1970 and in 1974 was quite similar, yet the expected excess inflation was much higher in 1974.

Similarly, it would not be enough to study the last panel because it conditions on events with time-varying probabilities. For example, the expected deflation in 1978 and in

1986 is about the same, yet the probability of excess deflation is close to 0% in 1978 and close to 80% in 1986. Thus, neither measure alone is informative; they must be viewed simultaneously. This fact makes the use of Definitions 1a and 1b awkward in practice.

The latter problem of interpretation can be avoided altogether by focusing on the *VaR* based measures of inflation risk and deflation risk. Since in this case, by construction, the probability of obtaining a rare event is time-invariant and known, it is possible to focus directly on the plot of tail conditional expectations, without looking at a probability plot at the same time. The same holds true for the deflation at risk measure and the corresponding expected deflation. This feature allows us to compress the information into a single index by defining the tail conditional expectation in (11) and (11′) as a unique measure of inflation risk and deflation risk, respectively. This fact greatly facilitates the interpretation of the risks to price stability in practice.

Definition 3 (Index of Inflation Risk, Index of Deflation Risk)

We formally define the *Index of Inflation Risk* (IIR_{t+1}) and the *Index of Deflation Risk* (IDR_{t+1}) as:

(14) $$IIR_{t+1} \equiv \max[E(\pi_{t+1} \mid \pi_{t+1} > IaR_{t+1}),\ 2\%]$$

and

(14′) $$IDR_{t+1} \equiv \min[E(\pi_{t+1} \mid \pi_{t+1} < DaR_{t+1}),\ 0\%]$$

where the dependence of IaR_{t+1} on $\overline{\theta}$ and of DaR_{t+1} on $\underline{\theta}$ has been suppressed.[11]

These two indices are shown in the upper panel of Figure 2. We set $\overline{\theta} = 0.75$ and $\underline{\theta} = 0.25$ throughout this paper.[12] The interpretation of these indices is straightforward: For example the *Index of Inflation Risk* indicates, at each point in time, what level of inflation we can expect, conditional on the 25% probability event that inflation will

[11] In defining the indices, we choose to suppress values of inflation risk below 2% and values of deflation risk above 0%. This convention facilitates the graphical presentation and hence the interpretation of the results. Without these bounds, Definition 3 reduces to Definition 2b. We will use Definition 3 when representing time series of risks and Definition 2b when discussing risks evaluated at a point in time.

[12] We also experimented with other choices. The qualitative results are not very sensitive over a range of reasonable values of $\overline{\theta}$ and $\underline{\theta}$.

exceed the IaR_{t+1} level. Put more succinctly, the *Index of Inflation Risk* indicates that there is a 25% chance that expected inflation will turn out as high as the value given by the index. By analogy, the *Index of Deflation Risk* indicates that there is a 25% chance that expected inflation will turn out as low as the value given by the index.

Although the risk measures in Figure 1 and Figure 2 are based on the same information, the interpretation of Figure 2 is much more straightforward in practice. For example, both Figures 1 and 2 show a significant increase in inflation risk in the 1970s, but only Figure 2 makes it readily apparent that the deflation risk in the 1970s was zero, as one would expect, given the high levels of actual inflation during that period. Also note that Figure 2 shows a significant decrease in inflation risk towards the end of the sample that is not readily apparent from Figure 1.

Figure 2 shows that inflation risks were quite high even during the 1960s, at a time when inflation was relatively low. For example, in November 1960 the Index in Figure 2 shows that there was a 25% chance that inflation in December 1960 would exceed the *IaR* level and actually turn out to be 6.9% in expectation. In contrast, the actual inflation rate in November 1960 was 0%. Even the average rate of inflation for 1960 only reached about 1.4%. Nevertheless, there was a substantial upside risk to price stability, as suggested by the tail behavior of the conditional distribution. Note that one also would have underestimated that risk by focusing on the conditional mean of inflation, which in November 1960 stood at 2.9%. By construction, the conditional mean neither accounts for differences in variance over time nor for asymmetries in the conditional distribution of inflation that alter the upside and downside risks to price stability.

This same example also is instructive because it illustrates that the existence of high inflation risks by itself is not necessarily an indication of poor monetary policy. It stands to reason that inflation risks (as well as deflation risks) often will be triggered by events outside of the control of central bankers. It is the central banker's job to respond to these surprise increases in risks and to balance the upward and downward risks to price stability. It would be unreasonable to expect the central banker to always eliminate inflation risks, if at the same time the central banker has to care about the risk of deflation. In the short run, the central banker may be able to fight one risk at the expense of the other, but he cannot eliminate both risks at the same time. Indeed, in this specific example, the Index of Deflation Risk suggests a non-negligible simultaneous risk of deflation throughout 1960 that may have prompted the Fed not to steer a more contractionary course. Thus, a sensible measure of how successful the central bank has been at balancing upward and downward risks to price stability is provided by the average of the tail conditional expectations. Perfect balance will translate into an average of zero. We introduce such an index as a natural complement to the indices in Definition 3:

Definition 4 (Index of the Balance of Risks)

We formally define the *Index of the Balance of Risks* (IBR_{t+1}) as:

$$IBR_{t+1} \equiv [E(\pi_{t+1} \mid \pi_{t+1} > IaR_{t+1}) + E(\pi_{t+1} \mid \pi_{t+1} < DaR_{t+1})\ ]/2 \qquad (15)$$

Note that this measure of balance is descriptive, not normative. Specifically, when central bank preferences incorporate a trade-off between higher output and lower inflation, the central bank may aim for a slightly positive balance rather than zero balance.[13]

The *Index of the Balance of Risks* is plotted in the lower panel of Figure 2. It is immediately apparent that, although there has been some inflationary bias throughout most of the post-war history of the United States, including the early 1960s, the risks have become severely unbalanced mainly during four periods: the Korean war, the periods leading up to and following the 1974 and 1979 OPEC oil crises, and the Gulf war episode of 1991. In contrast, periods in which the balance has swung toward deflation are rare and short-lived. They include the immediate post-war period (possibly as a consequence of deregulation) and a negative spike following the collapse of OPEC in 1986. Most recently, the balance seems to have shifted slightly toward deflation, but not more so than in 1998. How persistent this shift is, is too soon to tell. We will examine the outlook for U.S. inflation over the next years in more detail in section 6.

We conclude that the early 1960s indeed were a golden age of monetary policy by this measure, as the collective memory would suggest - not because the Fed succeeded in eliminating inflation risks – but because it came close to achieving balance between upside and downside risks to price stability. Moreover, save for a few outliers, the degree of balance since the mid-1980s has been similar to that in the 1950s and 1960s. The next section will explore more systematically the risks of inflation and deflation in historical perspective. We will contrast the experience of the United States, Germany and Japan based on the indices for inflation risk, deflation risk and balance of risks proposed in this section. Unlike in this section, we will focus on risks at medium-term horizons.

4. Inflation risks in historical perspective

The starting date for our comparative analysis is January 1960, given the availability of international data. The German data were truncated in 1999 because of the introduction of the Euro.[14] The Japanese data and U.S. data extend to December 2001. We discarded the data for 2002 on the grounds that they are still preliminary. The data source for the German and Japanese CPI data are the OECD Main Economic Indicators. Unlike the U.S. CPI data, these data are seasonally unadjusted. We used the X-12 procedure in Eviews to remove the seasonal variation in the data.[15]

Since historical analysis aims to uncover broader trends in the data, it is natural to focus on medium-term horizons, as opposed to horizons of only one month. We will focus on a horizon of one year in the empirical analysis. Although this is not the only horizon one may consider, a one-year horizon is often referred to by policy makers (see the earlier quote by Greenspan; also see Issing 2002b). The analysis of year-on-year changes in inflation requires some notational changes and modifications of the data, although the basic framework of analysis is unchanged. In essence, we need to define the analogue of the forecast model in section 3 for year-on-year inflation. Let y_t^h denote the percent change in variable *y* from *h* periods ago to date *t*. Let *T* denote the number of observations for π_t^1. We first compute the monthly observations for year-on-year inflation, π_t^{12}, where $t = 1+12, 2+12, ..., T$. The forecast model for year-on-year inflation may be written as $\pi_{t+12}^{12} = \mu_{t+12}^{12} + \mathrm{u}_{t+12}^{12}$. Analogous to section 3, we postulate that $\mathrm{u}_{t+12}^{12} = \sqrt{h_{t+12}^{12}}\varepsilon_{t+12}^{12}$, where $\varepsilon_{t+12}^{12} \mid \Omega_t \sim i.i.d.(0,1)$ and h_{t+12}^{12} is governed by a GARCH(1,1)

---

[13] When the conditional distribution is perfectly symmetric, this measure of balance reduces to the conditional mean of inflation, provided that $\overline{\theta} = 1 - \underline{\theta}$. For the historical data in sections 3 and 4 we find that there is a moderate degree of asymmetry in the conditional distribution.

[14] A potential concern is the effect of German unification on the structural stability of the German inflation process after October 1989. We also computed the results German inflation data after October 1989. Since the results were qualitatively unchanged, we chose to present the results for the longer sample.

[15] The implementation of our procedures does not require the use of seasonally adjusted data, but in practice the plots of the risk measures are smoother and hence easier to read for seasonally adjusted data.

process. In practice, we test for the absence of serial correlation in $(\mathrm{u}_{t+12}^{12})^2$. If the null cannot be rejected, we treat u_{t+12}^{12} as i.i.d.[16] The conditional mean model is $\mu_{t+12}^{12} = c + \sum_{j=0}^{p-1} \phi_{j+1} \pi_{t-12j}^{12} + \sum_{j=0}^{p-1} \varphi_{k,j+1} \Delta wti_{t-12j}^{12}$, given $\pi_i, ..., \pi_{i-12(p-1)}$ and $\Delta wti_i, ..., \Delta wti_{i-12(p-1)}$, for $t = i, i+12, ..., T-12+i$, where $i = 1, 2, ..., 12$ refers to the *i*th month in the sample. This means that we fit a different forecast model for each month. Based on the AIC we select $p = 3$ for the United States and for Germany, and $p = 1$ for Japan. Although we include lagged percent changes in oil prices in the forecast model, we note that very similar empirical results would be obtained without the oil price variable. As in section 3.3, we compute all historical measures of a country's inflation risk by first fitting the forecast model to the full sample. We then recursively compute the conditional quantiles of the distribution of $\pi_{t+12}^{12} \mid \Omega_t$ and the associated measures of risk.

In Table 2 we show the average risks by decade and country. This allows us to address the question of how large these risks have been on average historically, whether there are any systematic trends over time and whether the risks have evolved similarly across countries. The first panel shows that historically average inflation risks have moved in the range of 3.6-11.1% with an average of 6%. Although all countries historically experienced the highest inflation risks in the 1970s, there are also striking differences across countries. For example, inflation risks in Japan were on average almost twice as high as in Germany during that period. German inflation risks were on average somewhat higher in the 1960s than in the United States, but clearly lower than in the United States in the 1970s, 1980s and 1990s. Inflation risks in Germany were much lower than in Japan in the 1960s and 1970s, but the differences have been steadily declining since then.

The second panel shows that in Germany and in the United States the risks of deflation have been negligible over the post-war period. Among the three countries, Japan has had the lowest risks of deflation in the 1960s and 1970s, but the highest risks of deflation since the 1980s. Japan's deflation risk has been steadily growing since the

[16] Note that even in the absence of GARCH dynamics the error distribution may be asymmetric. Our risk measures are designed to capture that asymmetry.

1980s, but have been especially high since 2000.

The last panel of Table 2 reveals that the balance of risks has been increasing from the 1960s to the 1970s in all three countries, with an overall tendency to decline since then. The one exception to this pattern is a slight increase in the balance from 3.0% to 3.1% for Germany in the 1990s, driven by an increase in inflation risk that is likely to be related to German reunification. Despite this increase, Germany's balance of risks in the 1990s was still lower than for the United States. Notwithstanding the overall similarity of trends, there are important differences across countries. For example, the degree of balance achieved by the Bundesbank has remained relatively stable over time, whereas that of the Bank of Japan has undergone the most dramatic shifts. Japan not only experienced the greatest bias toward inflation in the 1970s, but Japan in the 1990s and 2000s also provides the only example of a country, in which the balance of risk has shifted into the deflationary region. Moreover, although in all countries the balance of risks turned strongly inflationary in the 1970s, in Japan this increase started from an already very high level.

To summarize, Table 2 suggests that, despite some dramatic differences in national experiences, inflation risks in the United States, Germany and Japan followed similar long run trends. An interesting question thus is to what extent this evolution of risks was driven by common global factors as opposed to country-specific factors. Figure 3 plots the indices of inflation risk and of deflation risk for the three countries under consideration. The corresponding balance of risks is shown in Figure 4. We first focus on the United States. There are two major increases in inflation risk since 1963. Both episodes appear to be associated with oil events: OPEC I in 1973/74 and OPEC II in 1979/80. This pattern is suggestive at first sight of a close relationship and perhaps even a causal role of oil prices for year-on-year CPI inflation risks. A more careful inspection of the data, however, casts doubt on such a link. First, not all major oil price increases since 1963 were associated with increases in inflation risk. The 1999-2000 tripling of oil prices is a case in point. There is no indication in Figure 3 of a major increase in inflation risk during that period.[17] Similarly, there is no systematic increase in inflation risk during Gulf war of 1990/91.

A second reason that casts doubt on the notion of a tight link between oil price shocks and important movements in inflation risk is that the evolution of inflation risks is quite different across countries during these oil dates. Countries like Germany or Japan with a similarly high dependence on imported oil presumably should experience an increase in inflation risk of similar magnitude, following an oil price increase. Figure 3, however, shows that Germany in 1973/74 or 1979/80 was subject to a much smaller increase in inflation risk than the United States. In contrast, Japan shows an even greater increase in inflation risk than the United States in 1973/74, reaching a level of risk almost twice as high as the United States; yet the increase in inflation risk in 1979/80 is much smaller than in 1973/74 (or for that matter the concurrent increase in the United States). This heterogeneity across countries and across time in the evolution of inflation risks strongly suggests that the explanation of the two major spikes in U.S. inflation risk is country-specific.

One plausible country-specific explanation of U.S. inflation risks at least for 1973/74 and 1979/80 are changes in the overall stance of monetary policy adopted by the Federal Reserve Board. For example, the index of the monetary policy stance proposed by Bernanke and Mihov (1998) shows both major increases in inflation risk in the 1970s were preceded by periods of expansionary policy stance and, similarly, that the subsequent reductions in inflation risk were preceded by moves toward a more contractionary stance (also see Barsky and Kilian 2002).

Monetary policy also is a good candidate for understanding the relative magnitudes of inflation risks in Germany and Japan in the 1970s. The fact that German inflation risks remained well contained in the 1970s is consistent with the view that the Bundesbank remained relatively hawkish on inflation throughout that period compared with the Bank of Japan and the Fed. Similarly, the differences in Japanese inflation risks between the 1973/74 and 1979/80 episodes appear to be linked to the monetary policy stance. As

---

[17] It may be objected that this latest increase in oil prices occurred only gradually and hence perhaps is not a true oil price shock, but the same is true for the 1979/80 increase in oil prices, which does appear to be associated with a dramatic increase in inflation risk.

noted by Bohi (1989, pp. 78-79) the high Japanese inflation rates of 1973/74 were largely home-made and resulted from deliberate monetary policy choices in the period leading up to the first oil shock: Unlike in the United States, "… inflation was considered a less important problem than maintaining a stable foreign exchange rate for the yen. Indeed, in response to President Richard M. Nixon's 1971 decision to let the U.S. dollar float in international markets, the Bank of Japan purposely inflated the economy to prevent the yen from appreciating." According to Bohi, this inflationary policy prompted a sharp rise in excess liquidity and aggregate demand, which in turn fuelled inflation into 1974. This argument suggests that the extraordinary increase in inflation risk in Japan is not related to the oil crisis at all. In sharp contrast to the Bank of Japan, the Bundesbank responded to the collapse of the Bretton Woods system by switching to a more restrictive monetary policy and letting the mark appreciate (also see Bruno and Sachs 1985, p. 234). This difference in monetary policy stance not only helps to understand the difference in inflation risks between Germany and Japan during 1973/1974, but it also helps to explain the striking disparity in inflation risks between the 1973/74 and the 1979/80 episode in Japan.

Monetary policy of course is not the only source of country-specific risks. For example, it is widely believed that German reunification created renewed risks of inflation in Germany. Figure 3 shows that inflation risks in Germany indeed increased after 1989. The peak level of risk in 1993 is comparable in magnitude to the highest levels of risk in the 1970s. Such major events, however, are rare. We conclude that historically the evolution of inflation risks and of the balance of risks is likely to have reflected at least in part shifts in the attitudes of central banks toward inflation and implicitly toward the trade-off between containing inflation and pursuing secondary objectives. An interesting question for future research would be to investigate the link, if any, between monetary policy, oil price shocks, and real wage behavior in determining the medium-term risk of inflation.

5. Forecasting Inflation Risks in Real Time: A Step Beyond the Narrative Approach to Monetary Policy

Historical analyses, as in section 4, provide the most accurate picture of inflation risks that prevailed at the time, but they are not suitable for assessing the rationale of specific policy decisions. The reason is that the information set available to a central banker at the time of decision-making tends to differ from that of the economist who reviews the historical evidence using the ex-post revised data. The important role of data revisions for understanding monetary policy decisions has recently been highlighted by a number of papers, including Orphanides (2001, 2002, 2003). A serious analysis of policy decisions therefore has to incorporate genuine real-time data.

In this section, we will construct forecasts of U.S. inflation and deflation risks period-by-period, given the data available to the Fed in real time. In constructing these risk forecasts, we will recursively re-estimate the forecasting model for each period, using only the data available at each point in time. These risk forecasts will allow us to discuss the rationale for policy decisions taken by the Fed in the 1980s and early 1990s. This time period is especially interesting in that it includes the transition from a period of high-inflation risk to one of moderate risk.

5.1. Real-Time Estimation Methodology

The analysis will be based on the real-time data set for U.S. deflator inflation used in Orphanides (2003). The data are for the GDP or GNP implicit price deflator, as appropriate for each period. The data set contains vintages of quarterly real time inflation data for the period from 1966.I-1995.IV. Each vintage contains data for the most recent four periods. The data set also contains real time within-quarter forecasts of the output gap. For further discussion of the construction of the real-time data see, e.g., Orphanides (2001, 2003).

We avoid the well-known problems of time aggregation of volatility dynamics by fitting different models for each forecast horizon *h*. Each forecast model is fit initially on data for 1966.I-1981.II and then recursively re-estimated, as new vintages of data become available each quarter. We start by defining the inflation rate for the horizon of interest. Given that there are only four lags of inflation in the Orphanides (2003) data set, we cannot consider forecast horizons in excess of four quarters. For expository purposes we will consider horizons of one quarter and four quarters. Specifically, let $\pi^{h}_{r|s}$ denote the inflation rate from period $r-h+1$ to period *r*, given information as of period *s*. Orphanides provides inflation data in the format $\pi^{1}_{r|s}$ where $r = s-1, ..., s-4$, and

$s = 1,...,T$. Let N denote the length of the sample used in generating recursive real time forecasts. We have $N = R, R+1, R+2,...,T,$ where R is the length of the initial recursive sample. Furthermore, define $\pi^h_{t-1|t} \equiv \sum_{r=t-h}^{t-1} \pi^1_{r|t} \big/ h$ for $t = i+h, i+2h,..., N-h+i$ and $\pi^h_{t+h-1|t+2h-1} \equiv \sum_{r=t}^{t+h-1} \pi^1_{r|r+h} \big/ h$ for $t = i+h, i+2h,...,\ N-2h+i+1,$ where $i = 1,...,h$. Here $i = 1$ refers to the first quarter of the sample, $i = 2$ to the second quarter, etc. Given these data, we compute the horizon-*h* risks analogous to the horizon-1 risks in section 3.

For given h, the forecast model can be written as $\pi^h_{t+h-1} = \mu^h_{t+h-1|t} + u^h_{t+h-1}$, where π^h_{t+h-1} denotes the ex-post revised inflation rate (which in practice may be proxied by $\pi^h_{t+h-1|t+2h-1}$). The dependent variable is the ex-post revised inflation rate, because we are interested in generating forecasts for the actual inflation rate, not the preliminary data releases. Our preferred model for the conditional mean includes lagged inflation and the Fed's within-quarter forecast of the output gap. Since the data set of Orphanides (2003) does not include lags of the real time output gap, we are unable to compute the average gap for *h* periods. Instead the forecast model includes only the most recent quarter of output gap data, denoted by $\hat{gap}_{t|t}$.[18] Specifically, the conditional mean model is $\mu^h_{t+h-1|t} = c + \sum_{j=0}^{p-1} \phi_{j+1} \pi^h_{t-jh-1|t} + \varphi_1 \hat{gap}_{t|t}$, where the lag order p may change with the horizon. We use $p = 4$ for the one-quarter forecast model and $p = 2$ for the one-year forecast model. The parameters of the conditional mean model may be estimated by running the regression:

$$\pi^h_{t+h-1|t+2h-1} = c + \sum_{j=0}^{p-1} \phi_{j+1} \pi^h_{t-jh-1|t} + \varphi_1 \hat{gap}_{t|t} + u^h_{t+h-1}$$

for $t = i+h, i+2h,..., N-2h+i+1$ and $i = 1,...,h.$ This means that we estimate *h* forecast models, one for each *i*. Each of these forecast models is estimated excluding the last

[18] There are no established procedures for choosing between alternative real-time forecast models of inflation risk. Our choice of variables and lag orders is guided by a priori considerations. Arguably, lagged inflation and the within-quarter output gap will be included in any inflation risk assessment by the Fed. We also experimented with adding lagged percent changes in WTI oil prices. Since there are no revisions in spot price series, the oil price data are measured in real time as well. The qualitative results are similar. The main difference is that the collapse of the OPEC oil price in 1986 and the invasion of Kuwait in 1990 are followed by more pronounced spikes in the risk forecasts than for the benchmark model.

$2h-1$ quarters of the N observations in the recursive sample, since in real time no ex-post revised data are available for the most recent time period. Given the resulting parameter estimates, we then compute the real-time forecasts of the conditional mean of $\pi^h_{N+h-1|N}$ based on the full recursive sample, including the last $2h-1$ observations.

The corresponding conditional variance of inflation for each model is estimated by a GARCH(1,1) process as in section 3. Standard results on the time aggregation of GARCH models imply that the form of conditional heteroskedasticity will be affected by the choice of horizon (see e.g., Drost and Nijman 1993). We view the fitted GARCH models, as they evolve with the choice of forecast horizon, as convenient approximations. Note that in theory, as we lengthen the forecast horizon, the GARCH dynamics will ultimately vanish. In the limit, the conditional variance will equal the unconditional variance, and we may compute the risk measures from the unconditional distribution of inflation. In practice, we test for the existence of GARCH for each model and time period based on the Ljung-Box test. If there is no statistical evidence of GARCH, we model the residuals as white noise. In the specific example discussed below, we found that GARCH effects were typically quantitatively important only for $h=1$.

Given these conventions and definitions, our procedure can be summarized as follows: Step 1) Estimate the models of the conditional mean and of the conditional variance of $\pi^h_{t+h-1|t}$ for $t=i+h, i+2h, ..., N-2h+i+1$ and $i=1,...,h$. Step 2) Estimate $\hat{\gamma}$ as described in section 3.1. Step 3) Forecast the conditional mean and the conditional variance of $\pi^h_{N+h-1|N}$. Step 4) Given these forecasts and the parameter estimate $\hat{\gamma}$, compute the forecasts of the risk measures, as proposed in section 3.

5.2. Real-Time Forecasts of Short-Term and of Medium-Term Inflation Risks

The one-year real-time risk forecasts based on the data of Orphanides (2003) are plotted in Figure 5a. The time line refers to the date, at which the forecasts were made. Figure 5a shows that real-time deflation risks were essentially non-existent throughout this period. The same conclusion holds for the one-quarter horizon forecasts, which are not shown. Since the inflation risk and the balance of risks largely move in parallel, in what follows, we may focus on the risk of inflation without much loss of information.

In the remainder of section 5, we will focus on one-year forecasts of risks because that horizon is more relevant for policy decisions. It is nevertheless instructive to compare briefly short-term and medium-terms measures of inflation risks. Figure 5b plots the real-time forecasts of U.S. inflation risk computed for the one-quarter ("short-term") and the one-year ("medium term") horizon. The plot suggests several interesting conclusions. First, with the exception of the early 1980s, the short-term risk of inflation is consistently lower than the medium term risk of inflation.

Second, there are striking differences in the evolution of short-term risks and medium-term risks. As the quote by Alan Greenspan in the preface makes it clear, the Fed focuses on medium-term horizons rather than short-term horizons in assessing the risks to price stability. This distinction also appears important in implementing our risk analysis. For example, during 1988 and 1994 the medium term inflation risks increased dramatically, without much change in the short-term risks. This evidence confirms what central bankers have intuitively known for a long time, namely that forecasts of short-term inflation risks may be very misleading.

Third, oil price shocks – at least for the United States - tend to affect inflation risks only in the short-run, but do not appear to have important consequences for medium-term inflation risks. For example, the collapse of OPEC oil prices in 1986/87 sharply reduced the short-term risks of inflation without any apparent effect on the medium-term risk. Similarly, the 1990/91 Gulf war caused a sharp spike in short-term inflation risks in 1991, but very little change in medium-term risks. This finding is of independent interest in thinking about policy responses to future oil price shocks.

5.3. Policy Analysis Using Real-Time Forecasts of Risks

We are ultimately interested in relating the one-year ahead measures of inflation risk to actual policy decisions. We follow the existing literature in treating the Fed Funds as the policy instrument of the Fed since August 1982. Figure 6 plots the Federal Funds rate and quarterly real-time forecasts of the risk of inflation for the year to come. To the extent that our quantitative measures of risk mimic the Fed's assessment of future risks to price stability, we would expect to see common trends in risk forecasts and in the policy instrument. Moreover, we would expect major interest rate increases to be preceded by increases in inflation risk. Indeed, for most sub-periods there is a substantial degree of co-movement, but the relationship is not perfect. One complication is that the Federal Reserve Board over the relevant period at times pursued other objectives than merely to contain inflation. Notably, concerns about the level of output and unemployment may substantially weaken this co-movement. We therefore focus on selected sub-periods of our sample, during which the Fed - by its own account - explicitly acted only in pursuit of the inflation objective. In identifying such sub-periods we follow the analysis of Romer and Romer (1989, 1994, 2002), supplemented by various issues of the Annual Reports of the Federal Reserve Board. For our sample, we are able to identify three episodes of policy interventions that occurred explicitly in response to inflation risks. These include the monetary tightening of 1984, that of 1988 and finally the tightening of 1994. We divided the plot in Figure 6 into three sub-periods, corresponding to these three different episodes of monetary tightening.

Episode 1: Real Time Risk Forecasts During the 1984 Monetary Tightening

Between late 1981 and early 1983 both inflation risks and the Fed Funds rate fell sharply. This decline in inflation risk reflects in large part the process of disinflation launched by Paul Volcker in October 1979, but was undoubtedly helped by the 1981/82 recession and the strength of the dollar. In early 1983, however, there were clear signs of a resurgence of medium-term inflation risks, as measured by our indicator in Figure 6. This resurgence coincided with the gradual recovery of the economy from the trough of the recession in November 1982. As early as February 1983, the Fed concluded that "although inflation moderated substantially in 1982, many potential investors, scarred by the experience of the 1970s, remained cautious about the longer-range outlook – and about the government's commitment to maintain forceful anti-inflationary policies" (Annual Report, February 16, 1983, p. 30). In July 1983, the Fed warned of the risk that "some of the progress against inflation could be reversed as the private economy strengthens" (Annual Report, July 20, 1983, p. 49). In response to the renewed risk of inflation, the Fed almost immediately raised interest rates. The increase continued throughout 1984. In fact, in March 1984, some committee members felt that "inflationary expectations appeared to be worsening", as "capacity utilization rates … were approaching levels that had been associated with rising rates of inflation in previous periods of economic expansion". (Annual Report, March 26-27, 1984, p. 96).

The magnitude of the interest rate response in Figure 6 suggests that perhaps the Fed was concerned about re-establishing its long-run credibility and intended to signal that it would fight inflation at any cost. For example, as early as February 1983, the Fed warned that "overcoming the still deep skepticism about the anti-inflation effort is crucial […] Periods of slowing inflation in the past two decades have proved to be temporary. […] Unless the commitment to see the present effort through is made fully credible, there will be a danger that as markets improve with recovery we will see a reversion to aggressive patterns of wage and price behavior" (Annual Report, February 16, 1983, p. 43). Throughout 1983 and 1984 the Fed cautioned that "the persistence of inflationary expectations is evident both in recent surveys of private opinion and in the behavior of financial markets." (Annual Report, July 20, 1983, p. 49) and warned that there may be a temptation "to revert to the pricing and wage bargaining patterns of earlier years of rapid inflation (e.g., Annual Report, February 7, 1984, p. 49).

The policy intervention of 1984 appears to have been successful in that medium-term inflation risks after 1983 never returned to their level before the 1981/1982 recession, even as interest rates fell by 5 percentage points between late 1984 and late 1986. Despite a booming economy, inflation risks increased only slightly over the four years following the recession. We conclude that our risk forecasts for this episode match up quite well with the actual policy decisions. This is not surprising because the Fed, as noted by Romer and Romer (2002, p. 43), attached very little importance to the output objective during the early stage of disinflation (see Record of Policy Actions, July 6-7, 1981, p. 116, February 1-2, 1982, p. 89).

Episode 2: Real Time Risk Forecasts During the 1988 Monetary Tightening

Romer and Romer (1989, 1994, pp. 81-82) identify a second policy intervention in response to inflation risks in 1988. They note that in March 1988, most members of the FOMC agreed that the risks of more inflation had increased. According to their narrative account, over the next six months policy discussions centered around "the potential for higher rates of inflation", the need "to counter the risks of rising inflation", "risks that inflationary pressures would intensify" and in general "risks of higher inflation". In response to these concerns the Fed decided to increase the Fed Funds rate from 6.6% to 8.4% in November 1988. In late 1988, Romer and Romer (1994) document a partial shift in goals. Reducing the level of inflation became an important additional consideration. Notably, in December 1988, the FOMC tightened policy significantly "not just to forestall a pickup in inflation" but also "to permit progress to be made in reducing inflation over time". As a result, the Fed Funds rate was gradually raised to almost 9.9% in early 1989.

To the extent that containing inflation was the only objective of Fed Policy during 1988, we would expect persistent increases in inflation risk forecasts to be good predictors of a monetary tightening. This assumption seems valid for this sub-period in that the FOMC, as noted by Romer and Romer (1994, p. 82), explicitly acknowledged that the "risks of a downturn … needed to be accepted" and expressed their willingness to accept "a less robust economy" in the interest of disinflation. In other words, the FOMC subordinated the output objective to that of lowering inflation (also see Romer and Romer 2002, p. 43). Thus, the high degree of co-movement in the second sub-period in Figure 6 is not surprising.

Figure 6 shows that indeed the risks of inflation had steadily increased since the end of 1987, even as the Fed Funds rate had been gradually adjusted upward. Starting in early 1988, forecasts of medium-term inflation risks appeared to increase further, signalling increased dangers to price stability. From December 1987 to March 1988 alone, inflation risk jumped by almost one percentage point. The timing coincides exactly with the March date identified by Romer and Romer.

Between June 1988 and June 1989, one-year forecasts of upward risks increased by an additional two percentage points. By mid-1989 forecasts of inflation risk were at their highest level since 1981. Indeed, Figure 6 shows that the risks of inflation during this episode were much higher than during 1984-85. The evidence in Figure 6 lends credence to the FOMC's continued concern about the upside risks to inflation and helps rationalize the observed increase in the Fed Funds rate in 1988 and 1989. Indeed, medium term risks declined sharply in late 1989, following the increase in the Fed Funds rate. We conclude that our risk forecasts for this episode match up quite well with the narrative evidence in Romer and Romer (1994).

Episode 3: Real Time Risk Forecasts During the 1994 Monetary Tightening

A final example that illustrates the usefulness of our methodology is provided by the policy decisions taken by the Federal Reserve Board in 1994. As noted by Romer and Romer (2002), the Fed tightened moderately in 1994 in response to risks of inflation that had yet to materialize. In support of this view, they cite the Record of Policy Actions, February 3-4, 1994, pp. 131, 134, 137 and March 25, 1997, pp. 118-121. This 1994 tightening is clearly visible in Figure 6. It began in early 1994 and continued until early 1995.

Our real-time methodology again allows us to quantify the inflation risks that the Fed was responding to. Unlike in the other two episodes, Figure 6 shows no close co-movement between the Fed Funds rate and inflation risk forecasts for this sub-period. Moreover, there is no apparent increase in inflation risks in 1994. Rather forecasts of inflation risks, although noisy, show no trend at all between 1992 and 1995. This pattern may appear puzzling at first. To understand this evidence it is important to keep in mind that between 1991 and 1994 the Fed had subordinated its inflation objectives to the output and employment objective. As of 1990, both interest rates and forecasts of inflation risk trended down at about the same rate, but in March 1991 the plot shows a sudden sharp decrease in the Fed Funds rate not related to any change in inflation risk forecasts. This change in slope coincides almost exactly with the trough of the 1991 recession identified by the NBER. Apparently the Fed decided in response to its output and employment objective to lower interest rates over and above what would have been called for, had the Fed responded only to the risks of inflation. This interpretation is consistent with the Fed's own accounts. For example, the Annual Report of February 19, 1992 (p. 59) states that "the principal objective of monetary policy [in 1991] has been to help lay the groundwork for a sustainable expansion without sacrificing the progress against inflation that had already been set in motion" and elaborates that "in formulating its objectives for monetary policy for 1992, the FOMC has sought to promote a sustainable upturn in economic activity while continuing to build upon the hard won gains against inflation that have already been made" (p. 45)

The downward trend in the Fed Funds rate ends in early 1993 and for almost a year interest rates remain almost unchanged. Only in early 1994, the Fed decided to raise the

Fed Funds rate to levels more consistent with inflation risks. The data in Figure 6 suggest that indeed the Fed had been responding to inflation risks, as stated in the Record of Policy Actions, but - rather than being prompted by an unforeseen increase in inflation risk - that response had only become necessary, given the Fed's earlier departure from the primary objective of price stability. When the recovery from the 1991 recession was in full swing, the Fed returned to its primary objective of fighting inflation. This view is again consistent with the Fed's own account. According to the Annual Report of 1994, "Federal Reserve policy had remained very accommodative in 1993 in order to offset factors that had been inhibiting economic growth. By early 1994, however, the expansion clearly had gathered momentum, and the maintenance of the prevailing stance of policy would eventually have led to rising inflation" (p. 3). "Consequently, the FOMC, at its meeting in early February 1994, agreed that policy should be moved to a less stimulative stance." (p. 26) Indeed, in 1995 inflation risks declined by more than 1 percentage point, as one would expect, following the increase in the Fed Funds rate. We conclude that the risk forecast data support the Fed's decision to tighten monetary policy pre-emptively in 1994.

These three examples illustrate the usefulness of our proposed measures of risk in understanding policy decisions in a real-time context. The close match between the narrative evidence and our quantitative indicators of risk is encouraging. We demonstrated that our approach is useful in helping researchers understand historical policy decisions. Real-time risk forecasts complement the narrative approach in that they provide additional information about the timing and magnitude of risks that are not readily apparent from narrative evidence. In addition, our real-time approach also is designed to help central bankers make ongoing policy decisions. Although the information in Figure 6 is not the only information that central bankers will consider in practice (and containing inflation may not be the only policy objective), this example illustrates that forecasts of inflation risks may provide a useful input into the decision-making process of central bankers.

6. IS THE WORLD HEADING TOWARD DEFLATION?

Recently, the concern has been increasing that the risks of price stability have tilted noticeably toward deflation. This concern has been expressed for example by The Economist (2002a,b,c) in a series of articles published between September and November 2002. Notably, in October 2002 The Economist warns that "the risk of falling prices is greater than at any time since the 1930s".[19] These concerns in turn have elicited a response by central bankers. For example, Issing (2002b) makes the case that based on current data as well as conditional mean forecasts of inflation for next year there are no apparent risks of deflation in the Euro area or for that matter in Germany. There has been no formal analysis of these risks, however, for any of the major OECD countries.

In this section, we will use data for Japan, for the United States and for Germany to quantify the risks of deflation, as they existed in September of 2002. We focus on German, as opposed to European, inflation data for two reasons. First, Germany is often perceived to be the country most exposed - among European countries - to the risks of deflation (see, e.g., Issing 2002b). Second, there are no data for Euro area wide inflation that extend far enough back in time to allow the construction of risk forecasts for the horizons of interest here.[20]

We use inflation data, as they are currently publicly available. We are unable to use real-time data for this exercise because there are at this point no real time data sets for Japan and for Germany. Nor are there publicly available updates for the data set used by Orphanides (2003), given that the Board of Governors imposes a five calendar year restriction on the publication of its Greenbook data. Our estimation sample starts in January 1960 and ends in September 2002. For each country, we compute the risk of inflation, the risk of deflation and of their balance for horizons of one year and of two years, as described earlier. Clearly, these results should be viewed with some caution, as

[19] Similar concerns have been expressed in Business Week, the Wall Street Journal and the Financial Times, among others (see Issing (2002b), p. 4, for additional references).

[20] For the Euro area the short time span of inflation data since 1999 makes it impossible to estimate reliably econometric models. Although one could rely on synthetic Euro data, these data are only available back to the 1970s and become increasingly unreliable, as one extrapolates back in time. In contrast, the German data only need to be extrapolated forward for a few years. Although it is possible that the structural stability of the process that generates German inflation data was affected by the introduction of the Euro, we will abstract from that possibility.

the effective sample size is small, especially for horizons exceeding one year. Nevertheless, it is of interest to obtain at least a preliminary and tentative assessment of the risks.

Table 3 represents a snapshot of the profile of risks as of September 2002. It shows the projected risks by geographic area for a horizon of one year and, alternatively, of two years. Consistent with the rest of the paper, all risk measures are based on 25% tail probabilities. The rates have been annualized. The predictors and lag structure of the underlying forecast models were selected by the SIC.[21] The preferred forecast models include between one and three lags of the variable to be predicted. If we take deflation to mean a persistent absolute decline in the price level, then of the three countries in question only Japan appears to be subject to substantial deflation risk. Notably, there is a 25% chance that expected inflation rates in Japan will reach levels as low as –2.5% or –2.2%, depending on the horizon. This risk assessment is tempered, however, by the fact that these deflation risks appear to be more than offset by even higher inflation risks. On balance, the Japanese data appear to be slightly inside the inflationary region. For Germany, we forecast a minor risk of deflation for the coming year at the rate of –0.2%. The risk of deflation vanishes as we extend the forecast horizon by one more year, and even for the coming year it is more than offset by the risk of inflation. On balance, the risks of inflation far outweigh those of deflation. In contrast to Japan and Germany, the United States do not suffer from any foreseeable risk of deflation. In fact, there are considerable upward risks to price stability, especially for the United States, where the balance of risks is close to 3.8% for next year.

Alternatively, it has been suggested that it may be necessary to maintain inflation rates above a threshold of 1% to avoid the zero bound on interest rates, to account for upward measurement bias in CPI inflation, and perhaps to facilitate structural change in the economy (see, e.g., Issing 2002a,b). By this more stringent criterion not just Japan, but also Germany would be subject to the risk of deflation at the two-year horizon. Specifically, with 25% probability, expected inflation in Germany over the next two years will be as low as 0.2%. Clearly, however, a central bank intent on fighting equally the dangers of deflation and of inflation would not be prompted into action by this risk of

[21] The SIC in all cases rejects the use of lagged percent changes in oil prices as an additional predictor.

deflation, given the simultaneous and even greater risk of inflation. The predicted balance of risks for Germany remains slightly above the threshold, at a level of 1.8% for the next two years. In contrast, by this more stringent definition of deflation, Japan is on balance slightly inside the deflationary region with no sign of future improvements.

We conclude this section with a reminder that standard forecasts of the conditional mean of inflation do not in general adequately account for the risks of inflation or deflation. Nor do they properly account for the balance of risks. For example, for Japan and for the United States the conditional mean forecast of inflation tends to be as much as 0.6 percentage points lower than the balance of risks, suggesting a more deflationary stance than is actually the case.

7. CONCLUSION

The view of policy behavior implicit in our analysis is consistent with the policy discussions by central bankers, but it differs remarkably from standard descriptions of policy making implicit in the specification of standard Taylor-type rules. Ignoring idiosyncratic policy errors, Taylor rules take the form

$$i_t = \alpha + \beta E[\pi_{t+1} \mid \Omega_t] + \gamma E[gap_{t+1} \mid \Omega_t],$$

where i_t refers to the interest rate set by the central bank at date t (see, e.g., Orphanides 2001, 2003) The dependence of interest rate decisions on the output gap is optional and will be suppressed in what follows. What is of interest to us here, is that, according to the Taylor rule, policy decisions by construction depend only on the conditional mean of inflation. This fact is surprising, considering the emphasis central bankers put on assessing upside and downside risks to price stability in justifying policy decisions. The conditional mean may be viewed as a measure of risk only under quite unrealistic assumptions. It emerges as a special case of the risk measure in (1) for $\alpha = 1$, $\eta = 0$, $k = sign(x)/\Pr(x < t)$ and $t = \infty$:

$$\rho = \int_{-\infty}^{\infty} x dF(x) = E(x).$$

Note that this measure fails to account for the difference between upside and downside risks and in general will differ from the tail conditional expectation. The conditional expectation for horizon *h* will coincide with EEI_{t+h}^{IaR} only when $\overline{\theta} = 1$ and it will coincide with ED_{t+h}^{DaR} only when $\underline{\theta} = 0$. Thus, the conditional mean implies an extreme and rather implausible choice of $\overline{\theta}$ and $\underline{\theta}$. Furthermore, as noted earlier, the conditional mean also differs from the balance of risks, unless $\underline{\theta} = 1 - \overline{\theta}$ and in addition the conditional distribution is symmetric. Even if it turns out that the conditional distribution is symmetric at some point in time, however, the only way of verifying this fact is to compute explicitly the balance of risks. Finally, even in the unlikely case that symmetry holds exactly, the use of the conditional mean implies that the central banker is indifferent to an increase in the magnitude of both inflation and deflation risks (measured by the variance of $\pi_{t+h} \mid \Omega_t$) for the same conditional mean. This implication again seems implausible.

If we take central bankers by their word, then adjustments to the policy instrument must be related explicitly to the evolution of upside and downside risks. From the central bank perspective, the relevant risk is that inflation over a given time horizon will exceed an upper threshold or a lower threshold. For expository purposes, we will assume that the interest rate is the only policy instrument available and that containing inflation within the bounds given by [0%,2%] is the only objective of the central bank. Then a policy rule that is consistent with the stated objectives of central bankers might take the form:

$$\Delta i_t = \lambda(\widehat{EEI}_{t+1}^{IaR}) \max\{\widehat{EEI}_{t+1}^{IaR} - 2\%, 0\} + \eta(\widehat{ED}_{t+1}^{DaR}) \min\{0\% - \widehat{ED}_{t+1}^{DaR}, 0\},$$

where the adjustment parameters λ and η are increasing functions of the degree of risk, the hats denote forecasts conditional on Ω_t, and EEI_{t+1}^{IaR} and ED_{t+1}^{DaR} refer to the expectations over the forecast horizon relevant for decision-making. Such a policy rule would incorporate the following reaction patterns:

- If $\widehat{EEI}_{t+1}^{IaR} > 2\%$ and $\widehat{ED}_{t+1}^{DaR} > 0\%$, raise interest rates
- If $ED_{t+1}^{DaR} < 0\%$ and $\widehat{EEI}_{t+1}^{IaR} < 2\%$, lower interest rates
- If $EEI_{t+1}^{IaR} < 2\%$ and $\widehat{ED}_{t+1}^{DaR} > 0\%$, leave interest rates unchanged

- If $EEI_{t+1}^{IaR} > 2\%$ and $\hat{ED}_{t+1}^{DaR} < 0\%$, the decision will depend on the weights λ and η

The degree of risk aversion of the central banker towards inflation or deflation will be reflected in the functional form for the weights λ and η and in the choice of thresholds. The tighter the bounds, the higher the likelihood of central bank intervention. The more sensitive these parameters are to the degree of risk, the greater the interest rate adjustment, all else equal. Such a rule could be viewed at least as a crude approximation to the policy decisions of central banks such as the ECB, which view containing inflation as their only objective. Note that this policy rule is descriptive rather than normative. It is designed to reflect actual policies, not necessarily socially optimal policies.

There are four key differences between the proposed policy rule and standard Taylor rules, apart from the omission of the output gap. The first difference is that our rule makes policy decisions explicitly dependent on risks. The second difference is that our policy rule is stated in terms of changes to the policy instrument, rather than its levels. Thus, unlike the classic Taylor rule, it does not require knowledge of the appropriate level of the equilibrium real interest rate. As noted by Orphanides (2003), the dependence of Taylor rules on the equilibrium real interest rate is problematic in that this rate is notoriously difficult to estimate and may well vary over time. Our proposed rule does not require the central banker to know this rate and reflects more closely actual policy discussions. A third difference is that the Taylor rule is built around the notion of an implicit time-invariant inflation target. For example, Taylor assumes a target of 2%. There is ample evidence that if such a target exists at all, it is likely to be time-varying (see, e.g., Romer and Romer 2002). More importantly, many central banks do not appear to target a specific rate of inflation at all; rather they target a range of inflation rates. Our proposed rule is consistent with that observation. A fourth difference is that our rule allows for asymmetric preferences for inflation risk and deflation risk.

The cost of developing a policy rule that resembles more closely actual policy behavior is that at present this rule is ad hoc. One interesting topic for future research will be the development of micro-foundations for policy rules of this type. It also would be important to generalize this rule to allow for secondary policy objectives (such as ensuring high levels of employment and output). One could imagine conducting a similar risk assessment for the output gap for example and incorporating the result into the policy reaction function.[22] A second important avenue for research will be the estimation of such policy rules based on real time data. This will involve first finding a suitable parametric representation for λ and η and then developing a formal test of the estimated policy rule. Such extensions are beyond the scope of this paper.

References

Arzac, E.R., and V.S. Bawa (1977), "Portfolio Choice and Equilibrium in Capital Markets with Safety-First Investors," *Journal of Financial Economics*, 4, 277-288.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999), "Coherent Measures of Risk," *Mathematical Finance*, 9(3), 203-228.

Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in B. Bernanke and K. Rogoff (eds): *NBER Macroeconomics Annual*, 16, 137-183.

Basak. S., and A. Shapiro (2001), "Value-at-Risk Based Management: Optimal Policies and Asset Prices," *Review of Financial Studies*, 14(2), 407-405.

Bernanke, B.S., and I. Mihov (1998), "Measuring Monetary Policy," *Quarterly Journal of Economics*, 113(3), 869-902.

Board of Governors of the Federal Reserve System (1983, 1984, 1992, 1994), *Annual Report*, Washington, DC: Federal Reserve System.

Bohi, D.R. (1989), *Energy Price Shocks and Macroeconomic Performance*. Washington, DC: Resources for the Future.

Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307-328.

Bollerslev, T., and J.M. Wooldridge (1992), "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances," *Econometric Reviews*, 11(2), 143-172.

Bruno, M., and J. Sachs (1985), *Economics of Worldwide Stagflation*. Cambridge, MA: Harvard University Press.

Drost, F.C., and T.E. Nijman (1993), "Temporal Aggregation of GARCH Processes," *Econometrica*, 61(3), 909-927.

The Economist (2002a), "Dial D for Deflation," September 14.

The Economist (2002b), "Of Debt, Deflation and Denial," October 10.

---

[22] Such an extension would be consistent with recent interpretations of United States monetary policy. In a press release dated January 19, 2000, the Federal Reserve Board announced a change in its language describing the FOMC's assessment of future developments. The new language "will describe the FOMC's consensus about the balance of risks to the attainment of its long-run goals of price stability and sustainable economic growth […] More specifically, the announcement will indicate how the Committee assesses the risks of heightened inflation or economic weakness weakness in the foreseeable future".

The Economist (2002c), “Words that Can Harm You,” November 23.

Engle, R.F. (1982), “Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, 50, 987-1007.

Engle, R.F., and S. Manganelli (1999), “CAViaR: Conditional Value at Risk by Quantile Regression,” Working Paper No. 7341, National Bureau of Economic Research.

Fishburn, P.C. (1977), “Mean-Risk Analysis with Risk Associated with Below-Target Returns,” *American Economic Review*, 67(2), 116-126.

Issing, O. (2002a), “Monetary Policy and the Role of the Price Stability Definition,” panel speech at the ECB Watchers Conference in Milan on June 10, 2002.

Issing, O. (2002b), “The Euro after Four Years: Is There a Risk of Deflation?”, speech at the 16th European Finance Convention in London on December 2, 2002.

Kataoka. S. (1963), “A Stochastic Programming Model,” *Econometrica*, 31, 181-196.

Manganelli, S., and R.F. Engle (2001), “Value-at-Risk Models in Finance,” Working Paper No. 75, European Central Bank.

Markowitz, H. (1952) “Portfolio Selection,” *Journal of Finance*, 7, 77-91.

Markowitz, H. (1959), *Portfolio Selection*. New York.

Orphanides, A. (2001), “Monetary Policy Rules Based on Real-Time Data,” *American Economic Review*, 91(4), 964-985.

Orphanides, A. (2002), “Historical Monetary Policy Analysis and the Taylor Rule,” mimeo, prepared for: *Carnegie Rochester Conference on Public Policy*.

Orphanides, A. (2003), “Monetary Policy Rules, Macroeconomic Stability and Inflation: A View from the Trenches, ”Working Paper No. 115, European Central Bank, forthcoming: *Journal of Money, Credit and Banking*.

Romer, C.D., and D.H. Romer (1989), “Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz,” in O.-J. Blanchard and S. Fischer (eds.): *NBER Macroeconomics Annual*, 4, 121-170.

Romer, C.D., and D.H. Romer (1994), “Monetary Policy Matters,” *Journal of Monetary Economics*, 34, 75-88.

Romer, C.D., and D.H. Romer (2002), “The Evolution of Economic Understanding and Post-war Stabilization Policy,” mimeo, Department of Economics, UC Berkeley.

Roy, A.D. (1952), “Safety First and the Holding of Assets,” *Econometrica*, 20, 431-449.

Sims, C.A. (2002), “Comment on Cogley and Sargent: Evolving Post-World War II U.S. Inflation Dynamics,” in B. Bernanke and K. Rogoff (eds), *NBER Macroeconomics Annual 2001*, 137-183.

Stock, J.H. (2002), “Comment on Cogley and Sargent: Evolving Post-World War II U.S. Inflation Dynamics,” in B. Bernanke and K. Rogoff (eds), *NBER Macroeconomics Annual 2001*, 137-183.

Stone, B.K. (1973), “A General Class of Three Parameter Risk Measures,” *Journal of Finance*, 28, 675-685.

Tay, A.S., and K.F. Wallis (2000), “Density Forecasting a Survey,” *Journal of Forecasting*, 19(4), 235-254.

Table 1 – Model Diagnostics for U.S. Models

Conditional Mean	**Model 1**	**Model 2**	**Model 3**	**Model 4**
	Inflation	*Inflation & oil prices*	*Inflation & oil prices & money*	Inflation & money
$\hat{p}$	13	13	9	13
$\hat{q}$	-	13	9	-
$\hat{r}$	-	-	9	13
$AIC(\hat{p},\hat{q},\hat{r})$	1.847	1.796	1.828	1.883
$LB(15) : u_t$	0.38	1.02	1.05	0.33
(p-value)	(1.00)	(1.00)	(1.00)	(1.00)
$LB(15) : u_t^2$	28.21	26.11	26.92	29.10
(p-value)	(0.01)	(0.02)	(0.01)	(0.01)

Model 2			
Conditional Variance	*Coefficient*	*Robust s.e.*	*t-statistic*
ω	0.16	0.09	1.88
α	0.11	0.05	2.23
β	0.87	0.05	17.14
$LB(15)$: ε_t	9.41		
(p-value)	(0.74)		
$LB(15)$: ε_t^2	11.05		
(p-value)	(0.61)		

NOTE: The lag orders are selected by the AIC subject to the constraint that all lag order estimates are equal. LB denotes the Ljung-Box test.

Table 2 – Average Historical Risks by Decade

	1963-1969	1970-1979	1980-1989	1990-1999	2000-2001
Inflation Risk					
United States	5.12	8.06	7.94	5.67	5.05
Germany	4.09	5.59	4.35	4.49	-
Japan	8.68	11.12	6.53	5.54	3.61
Deflation Risk					
United States	1.09	4.03	3.91	1.65	1.03
Germany	1.31	2.82	1.57	1.71	-
Japan	2.54	4.98	0.39	-0.60	-2.53
Balance of Risks					
United States	3.10	6.04	5.93	3.66	3.04
Germany	2.70	4.20	2.96	3.10	-
Japan	5.61	8.05	3.46	2.47	0.54

NOTE: Estimates based on Model 2 for 1960.1-2002.2. The German and Japanese CPI data are from the OECD Main Economic Indicators.

Table 3 – Forecasts of the Profile of Risks as of September 2002

		1 year	2 years
Inflation risk forecast	US	5.86	6.03
	GE	3.11	3.40
	JP	4.19	3.15
Balance of Risks forecast	US	3.82	3.60
	GE	1.47	1.78
	JP	0.85	0.48
Deflation risk forecast	US	1.78	1.18
	GE	-0.17	0.16
	JP	-2.49	-2.19
Conditional mean forecast	US	3.68	3.06
	GE	1.45	1.78
	JP	0.22	-0.06

NOTE: Estimates based on monthly data for 1960.1-2009.2. The German and Japanese CPI data are from the OECD Main Economic Indicators. The U.S. data are from FRED. The predictors and lag structure of the forecast model were selected by the SIC. The preferred forecast models include between one and three lags of the variable to be predicted.

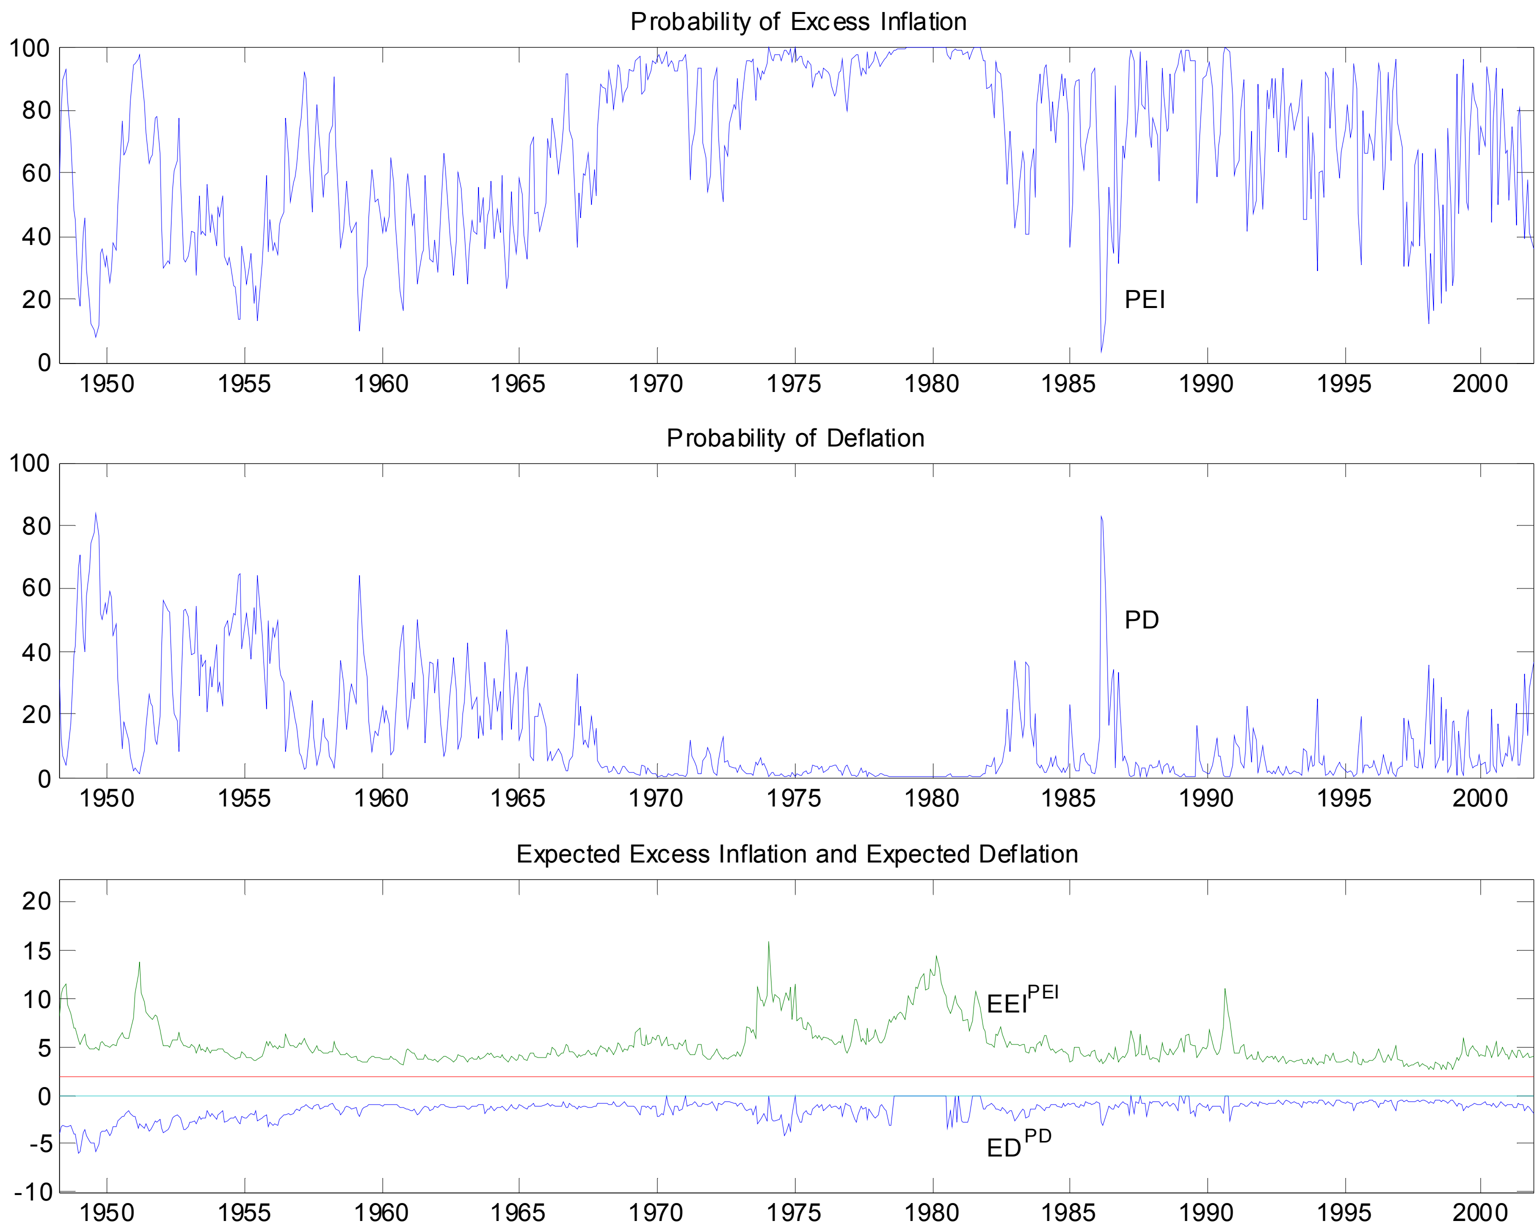

Figure 1 – Probability-based measures of risk (time-varying probability, time-invariant quantile)

United States CPI Inflation 1947.1-2001.12

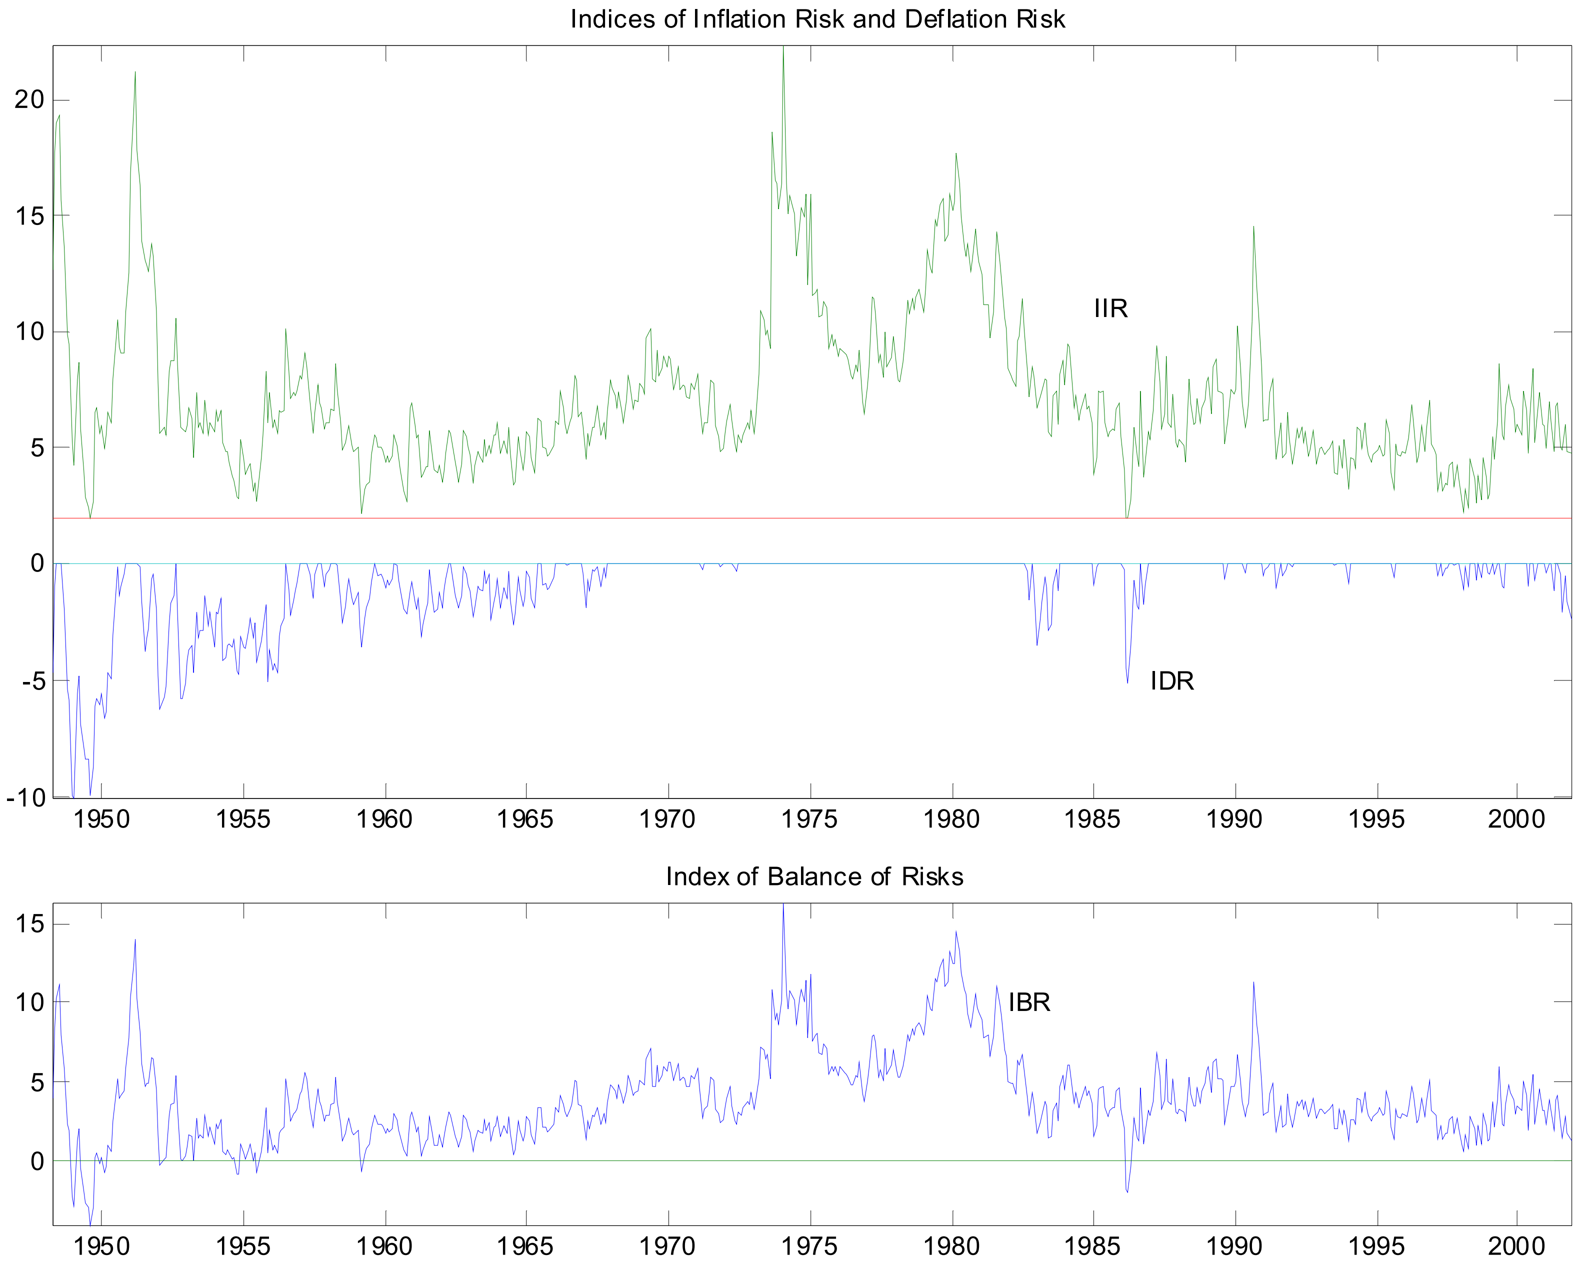

Figure 2 – Quantile-based measures of risk (time-invariant probability, time-varying quantile)

United States CPI Inflation 1947.1-2001.12

Figure 3 – Inflation Risks and Deflation Risks in Historical Perspective.

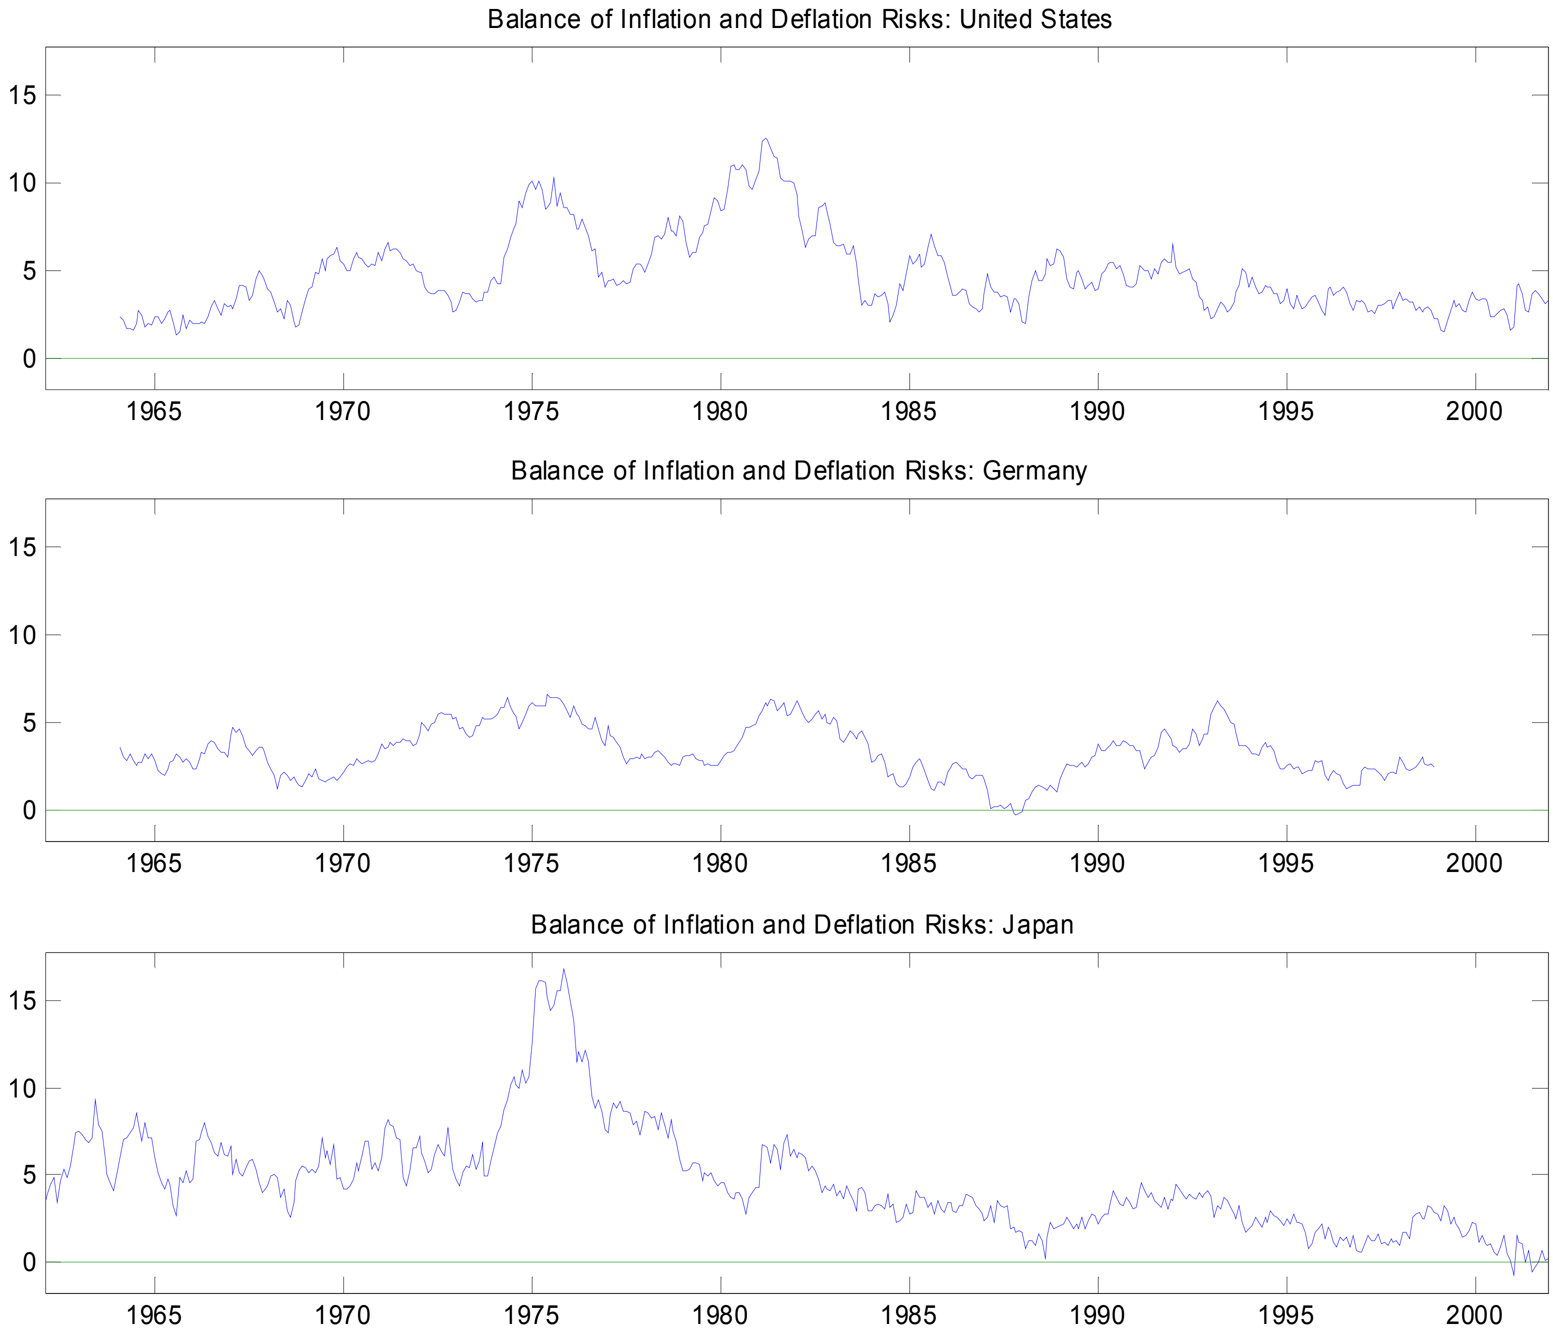

Figure 4 – Balance of Risk in Historical Perspective.

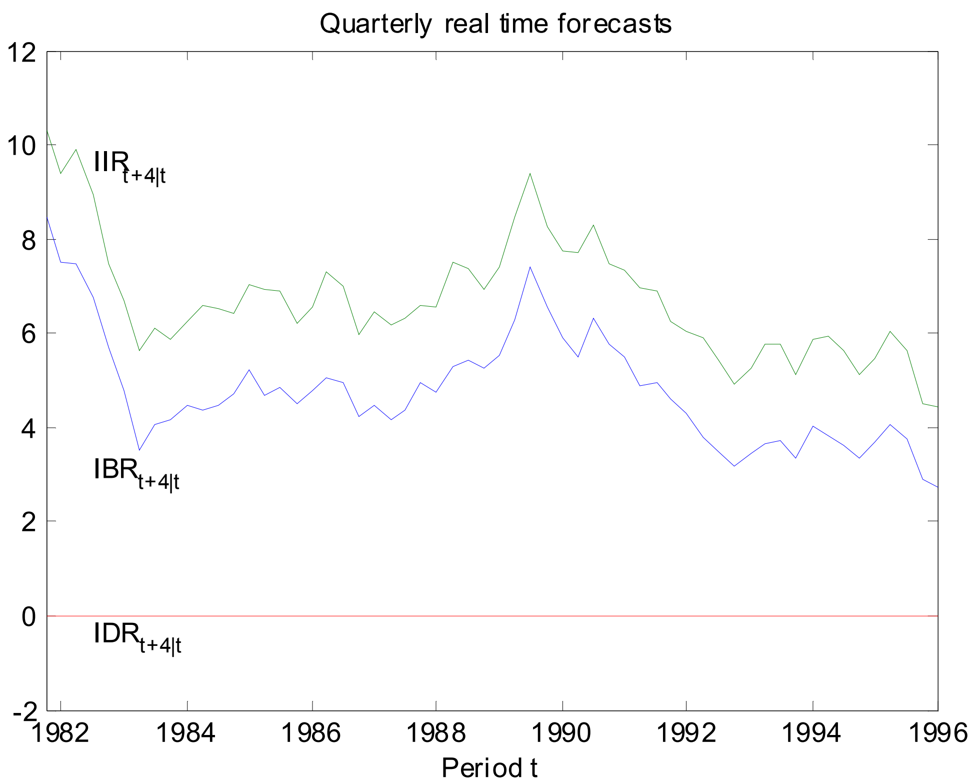

Figure 5a –Real-time Forecasts of Risk Measures One Year from Period *t*

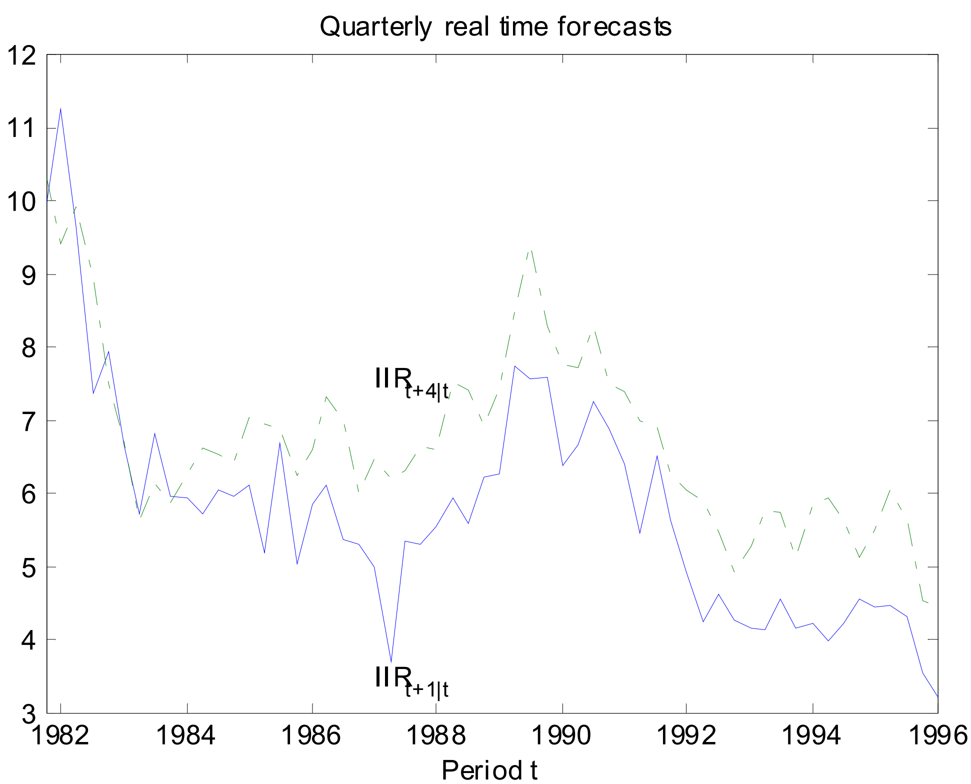

Figure 5b – Real-time Forecasts of Inflation Risk One Quarter and One Year from Period *t*

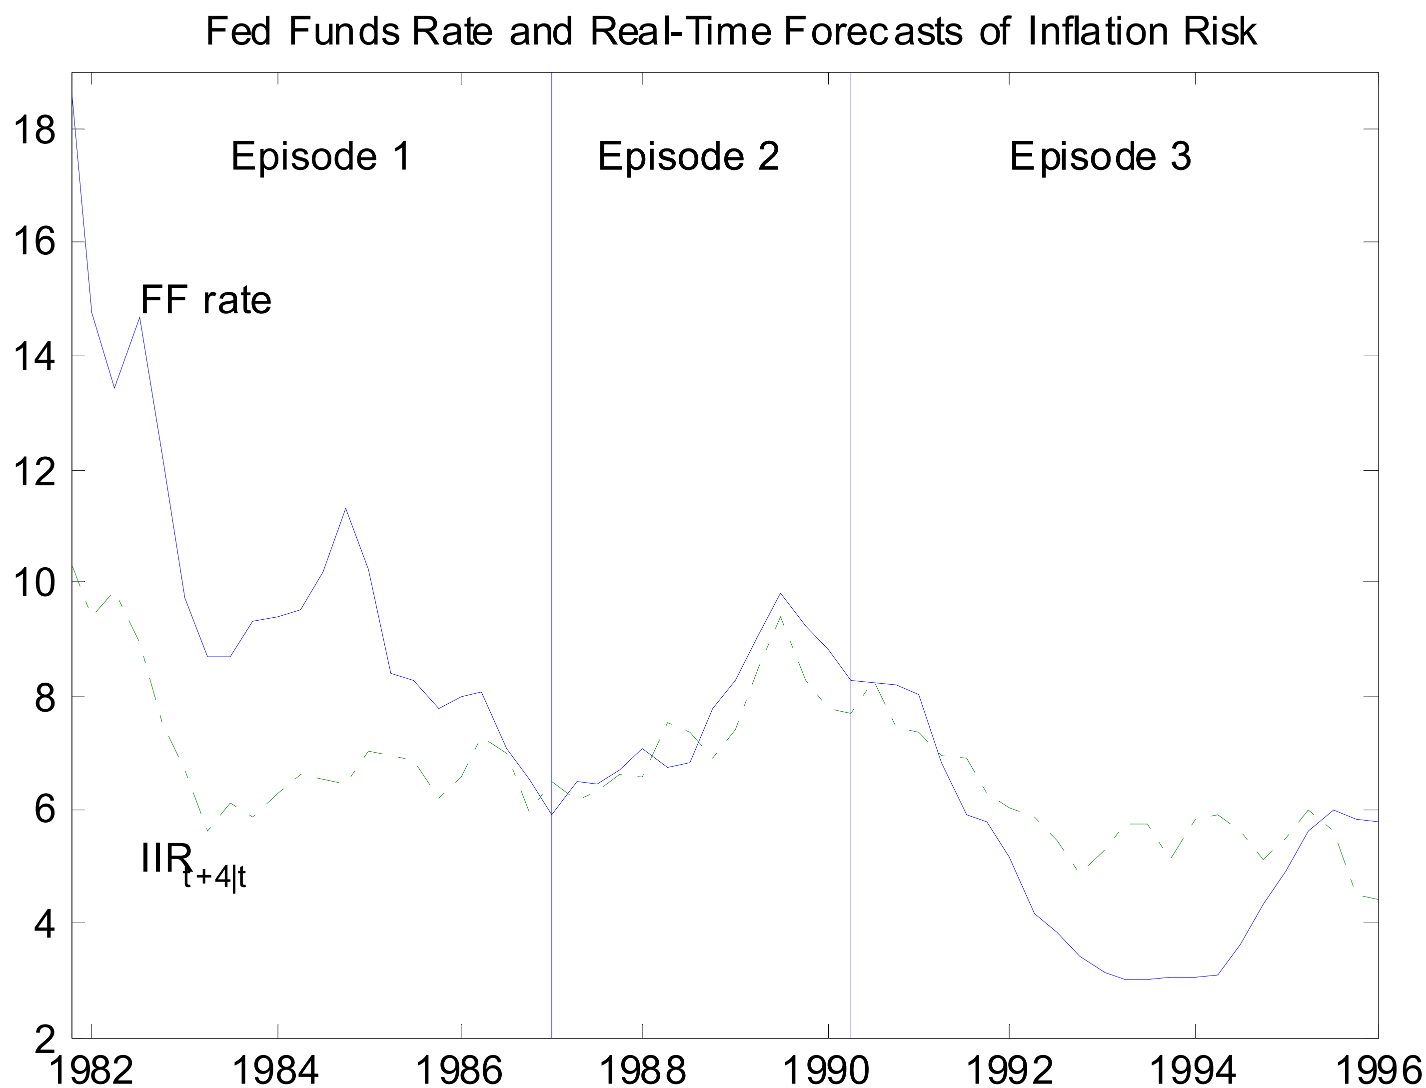

Figure 6 – Federal Funds Rate and Quarterly Real Time Forecasts of Inflation Risk One Year from Period *t*