

DISCUSSION PAPER SERIES

No. 3918

**THE CENTRAL BANKER AS A RISK
MANAGER: QUANTIFYING AND
FORECASTING INFLATION RISKS**

Lutz Kilian and Simone Manganeli

INTERNATIONAL MACROECONOMICS



Centre for Economic Policy Research

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP3918.asp

THE CENTRAL BANKER AS A RISK MANAGER: QUANTIFYING AND FORECASTING INFLATION RISKS

Lutz Kilian, University of Michigan, European Central Bank (ECB) and CEPR
Simone Manganelli, European Central Bank (ECB)

Discussion Paper No. 3918
June 2003

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Lutz Kilian and Simone Manganelli

June 2003

ABSTRACT

The Central Banker as a Risk Manager: Quantifying and Forecasting Inflation Risks*

In deciding a 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 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 show how to construct genuine real-time forecasts of year-on-year risks that may be used in policy-making. We demonstrate the usefulness of these risk forecasts in understanding the Fed's decision to tighten monetary policy in 1984, 1988 and 1994. Second, we forecast the risks of worldwide deflation for horizons of up to two years. Although recently fears of worldwide deflation have increased, we find that, as of September 2002, with the exception of Japan, there is no evidence of substantial deflation risks. We also put the estimates of deflation risk for the United States, Germany and Japan into historical perspective. We find that only for Japan is there evidence of deflation risks that are unusually high by historical standards.

JEL Classification: C22, E31, E37, E52 and E58

Keywords: deflation, forecast, inflation, monetary policy, price stability and risk

Lutz Kilian
European Central Bank
Eurotower
Kaiserstrasse 29
D-60311 Frankfurt a.M.
GERMANY
Tel: (49 69) 1344 8668
Fax: (49 69) 1344 8553
Email: lkilian@umich.edu

Simone Manganelli
European Central Bank
Eurotower
Kaiserstrasse 29
D-60311 Frankfurt a.M.
GERMANY
Tel: (49 69) 1344 7347
Fax: (49 69) 1344 8553
Email: simone.manganelli@ecb.int

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=132598

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=152569

*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, the Bank of Spain, the Deutsche Bundesbank, Barcelona Autònoma, Bocconi, Bonn, Erasmus University, Humboldt University and the Université de Montréal for useful comments. We also acknowledge helpful discussions with Robert Barsky, Ben Bernanke, Vítor Gaspar, Atsushi Inoue, Thorsten Koepl, Athanasios Orphanides, Adrian Pagan, Matthew Shapiro, Lars Svensson, Harald Uhlig, Ken Wallis and Jonathan Wright. We especially thank Athanasios Orphanides for providing access to the real-time dataset used in Orphanides (2003).

Submitted 24 April 2003

“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

“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

“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 preemptively 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

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

were selected to illustrate three requirements of risk measurement in the central banking context. First, the central bank's decision involves a trade-off between controlling upside risks and downside risks to price stability. Second, risks depend on the horizon of interest. The horizon relevant to central bankers is arguably at least one year, and perhaps even longer. Third, it is important from the central banker's point of view to be able to quantify these risks. 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 the event of inflation exceeding a certain threshold. Similarly, downside risks are commonly associated with inflation below a certain threshold. For example, Bernanke (2003) defines price stability as avoidance of deflation as well as excessive inflation. 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". Duisenberg (2003) elaborates that "this basically implies that, in practice, we are more inclined to act when inflation falls below 1% and we are also inclined to act when inflation threatens to exceed 2% in the medium run". Although not all central banks are as explicit about their objectives as the European Central Bank, most central banks today at least implicitly appear to pursue similar objectives.¹

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. 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

¹ See, e.g., Bernanke and Mishkin (1997) and Orphanides and Wieland (2000) for details. Bernanke (2002) notes that central banks with inflation targets invariably set their target for inflation above zero, generally between 1% and 3% per annum. 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.

probability distribution of a random variable of interest (such as the inflation rate) provides a complete and exhaustive description of its underlying uncertainty. Techniques to estimate and represent the probability distribution of the inflation rate are easily available. The “fan chart”, popularized 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 of future inflation. The mapping from the probability distribution of inflation forecasts to an assessment of the risks to price stability, however, is left to the eye of the beholder.

In this paper, our objective is to reduce the information contained in the probability distribution of inflation forecasts to indicators of risk that are easy to interpret and that at the same time effectively summarize the features of the predictive density that are most important to policy makers. We propose measures of risk that are directly related to the decision problem of the central banker. The proposed risk measures 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 in turn may serve as an input into monetary policy decision-making.

The remainder of the paper is organized as follows. In section 2 we propose a risk management model for central bankers. From this model we derive measures of the risk of deflation and the risk of excessive deflation. We establish formally conditions for the congruence of the risk management model with the expected utility model on the one hand and with the concept of stochastic dominance on the other. We also derive the implications of the risk management model for the construction of policy rules. We discuss conditions for the optimality of the policy rule and propose a measure of the balance of upside and downside risks to price stability. In section 3 we show how the proposed risk measures may be computed in practice based on an econometric model. In section 4 we construct genuine real-time forecasts of year-on-year 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 5, we formally assess the risks of worldwide 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 also put the deflation risks for these countries into historical perspective. We conclude in section 6.

2. A MODEL OF CENTRAL BANK RISK MANAGEMENT

A situation is said to involve risk if the randomness facing an economic agent can be expressed in terms of specific numerical probabilities (see Machina and Rothschild 1987). Therefore, any measure of risk must be related to the probability distribution of the underlying random variable – in the present context, the inflation rate. We will express this uncertainty in terms of a probability distribution function $F(\cdot)$ about possible inflation outcomes π .

As the discussion in the introduction illustrates, central bankers typically associate inflation risk with the failure to attain a given inflation target. This target may be defined as a point target, as in the case of an “inflation targeter”, or more commonly as a range, as in the case of an “inflation zone targeter” (see, e.g., Bernanke and Mishkin 1997; Orphanides and Wieland 2000). In the latter case, the target may be expressed as the objective of containing inflation within two threshold points, $\underline{\pi}$ and $\bar{\pi}$, where $\underline{\pi} < \bar{\pi}$. In the former case, we have $\underline{\pi} = \bar{\pi} = \pi^*$.

In section 2.1 we propose a general measure of inflation risk that is motivated by a risk management model. This model is explicitly designed to capture the decision problem of the central banker in that it involves a trade-off between upside and downside inflation risks. The model applies equally to inflation targeters and inflation zone targeters. Our analysis builds on the (α, t) -model of financial risk proposed by Fishburn (1977) to measure the trade-off between expected returns and downside risk relative to a target t .² Our model may be viewed as an $(\alpha, \underline{\pi}, \beta, \bar{\pi})$ -model of inflation risk, where α and β refer to measures of risk aversion. In section 2.2, we state conditions under which this risk-management model is consistent with the von Neumann-Morgenstern expected utility

² Holthausen (1981) proposed a generalization of this model, referred to as the (α, β, t) -model of financial risk.

model. In section 2.3, we show that the efficient set of the risk management model is a subset of the first-order stochastic dominance set for all $\alpha \geq 0$ and $\beta \geq 0$, and that it is a subset of the second-order stochastic dominance set for all $\alpha \geq 1$ and $\beta \geq 1$. In section 2.4, we discuss implications of the risk management model for the construction of policy rules. We derive conditions implied by the optimal policy rule. In section 2.5, we briefly discuss how our analysis can be extended to encompass multiple objectives.

2.1. The General Risk Management Model

The general model of risk that we propose may be viewed as a downside risk-upside risk model, in which risk is measured by a probability weighted function of the deviations of inflation from the threshold points $\underline{\pi}$ and $\bar{\pi}$:

Definition 1 [*Deflation risk (DR) and risk of excessive inflation (EIR)*]

$$\boxed{DR_{\alpha}(\underline{\pi}) \equiv - \int_{-\infty}^{\underline{\pi}} (\underline{\pi} - \pi)^{\alpha} dF(\pi), \quad \alpha > 0} \quad \boxed{EIR_{\beta}(\bar{\pi}) \equiv \int_{\bar{\pi}}^{\infty} (\pi - \bar{\pi})^{\beta} dF(\pi), \quad \beta > 0}$$

We adopt the convention of defining the risk of deflation as a negative number and the risk of excessive inflation as a positive number. Each risk is a scalar defined in terms of a two-parameter function involving a fixed target (or threshold), denoted by $\underline{\pi}$ and $\bar{\pi}$, and a fixed parameter governing risk aversion, denoted by α and β . The parameters α and β may be different.³ It is possible to compute the proposed measures of risk for any choice of α and β . The values of α and β can in principle be elicited from the central banker by means of a questionnaire that requires choices between pre-specified gambles (see Fishburn (1977) for details). We do not pursue this possibility here. Instead we will follow the academic literature in postulating $\alpha = \beta = 2$ for illustrative purposes (see e.g., Blinder 1997; Svensson 1997, 2002; Orphanides and Wieland 2000; Levin, Wieland and Williams 2002). We also consider the alternative assumption that $\alpha = \beta = 1$.

The idea that risk is measured by probability-weighted dispersions beyond given

³ A value of $\alpha = 1$ implies risk neutrality. Risk-seeking behavior is implied by $\alpha < 1$, whereas risk averse behavior follows from $\alpha > 1$. See Fishburn (1977) for a formal proof. Similar comments apply to β .

thresholds of policy interest is appealing in that it recognizes the central banker's desire to avoid inflation realizations outside the target zone. To the extent that this contention is correct, it casts serious doubt on the appropriateness of measures of inflation risk that assess the dispersion of inflation with respect to a parameter that changes from distribution to distribution. Such measures include for example the variance of inflation about the sample mean. They also include range or interval forecasts for inflation, and more generally quantile-based risk measures such as value at risk or the corresponding tail conditional expectations.

Risk management by definition always involves a trade-off. In the mean-risk model of returns proposed by Fishburn (1977), for example, the trade-off is between achieving a higher expected return and increased downside risk of a return below a given threshold. In the risk management model proposed below, in contrast, the trade-off is between the upside risk of inflation exceeding $\bar{\pi}$ and the downside risk of inflation falling below $\underline{\pi}$. As we will show, the quadratic loss function for inflation that is commonly postulated in the academic literature can be viewed as a special case of this risk management model.

Let $DR_{\alpha}(\underline{\pi}, F)$ denote the deflation risk computed under the distribution F and let $EIR_{\beta}(\bar{\pi}, F)$ denote the corresponding risk of excessive inflation. Then the risk management problem can be described as follows:

Definition 2 [Risk management problem]

Let F and G denote two alternative probability distributions of inflation. Then F is weakly preferred to G if $|DR_{\alpha}(\underline{\pi}, F)| \leq |DR_{\alpha}(\underline{\pi}, G)|$ and $EIR_{\beta}(\bar{\pi}, F) \leq EIR_{\beta}(\bar{\pi}, G)$. If this condition does not hold, the central banker faces a risk management problem.

The following proposition highlights some practical implications of Definition 2.

Proposition 1:

(a) *When $DR_{\alpha}(\underline{\pi}, F) = DR_{\alpha}(\underline{\pi}, G) = k$ where $k \leq 0$ is a constant, then F is weakly preferred to G if and only if $EIR_{\beta}(\bar{\pi}, F) \leq EIR_{\beta}(\bar{\pi}, G)$.*

(b) When $EIR_{\beta}(\bar{\pi}, F) = EIR_{\beta}(\bar{\pi}, G) = k$ where $k \geq 0$ is a constant, then F is weakly preferred to G if and only if $|DR_{\alpha}(\underline{\pi}, F)| \leq |DR_{\alpha}(\underline{\pi}, G)|$.

Proof: See Appendix.

Proposition 1 implies that whenever the risks of deflation (excessive inflation) are equal, then the distribution with lower excessive inflation (deflation) risk is preferred. This result is potentially useful in comparing inflation distributions across time or in comparing inflation distributions implied by alternative hypothetical policy paths for a point in time. For example, whenever deflation risks are zero over extended periods of time, Proposition 1 implies that for those periods the central banker may disregard DR_{α} and focus on EIR_{β} .

The practical usefulness of Proposition 1, however, is limited. Changes in upside and in downside risks are often offsetting, if only in part. In that case, there will be a trade-off between upside and downside risks that cannot be solved in the absence of further information about preferences. Solving this trade-off requires a risk management model. The defining characteristic of the central bank risk management model that we propose is that utility can be expressed as a function of DR_{α} and EIR_{β} .

Definition 3 [Risk management model]

We say that a central banker's preferences satisfy a risk management model if and only if there is a real valued function U in risks such that for all relevant distributions F and G :

$$F \text{ is preferred to } G \text{ if and only if } U(DR_{\alpha}(\underline{\pi}, F), EIR_{\beta}(\bar{\pi}, F)) > U(DR_{\alpha}(\underline{\pi}, G), EIR_{\beta}(\bar{\pi}, G))$$

Note that a central banker's preferences may satisfy a risk management model without also satisfying the von Neumann-Morgenstern axioms for expected utility. Since the expected utility model provides additional useful insights into the decision problem of the central banker, however, it will be useful to specify under what conditions such a utility function and the preference relations underlying the risk management model coincide. This is the question we turn to next.

2.2. Congruence with Expected Utility

The defining characteristic of the central bank risk management model is that utility U can be expressed as a function of DR_α and EIR_β . In contrast, in the expected utility model utility u is a non-monotonic function of π . Nevertheless, these two models will be congruent under certain conditions:

Definition 4 [Congruence between risk management model and expected utility model]

We say that the risk management model in Definition 3 is congruent with the expected utility model if

$$U(DR_\alpha(\underline{\pi}, F), EIR_\beta(\bar{\pi}, F)) > U(DR_\alpha(\underline{\pi}, G), EIR_\beta(\bar{\pi}, G)) \Leftrightarrow \int_{-\infty}^{\infty} u(\pi) dF(\pi) > \int_{-\infty}^{\infty} u(\pi) dG(\pi).$$

The following proposition shows that for an important class of utility functions that includes the leading examples discussed in the academic literature, the expected utility model is congruent with a risk management model that is linear in the risks.

Proposition 2:

If the central banker's preferences can be expressed as:

$$u(\pi) = \begin{cases} -a(\underline{\pi} - \pi)^\alpha & \text{if } \pi < \underline{\pi} \\ 0 & \text{if } \underline{\pi} \leq \pi \leq \bar{\pi} \\ -(1-a)(\pi - \bar{\pi})^\beta & \text{if } \pi > \bar{\pi} \end{cases},$$

then $E[u] = aDR_\alpha - (1-a)EIR_\beta$.

Proof: See Appendix.

It follows immediately that for this class of preferences the linear function $U(DR_\alpha, EIR_\beta) \equiv aDR_\alpha - (1-a)EIR_\beta$ is equal to $E[u]$ and hence satisfies congruence. Proposition 2 implies that DR_α and EIR_β are directly proportionate to the expected disutility associated with deflation and with excessive inflation. In this sense, DR_α and EIR_β quantify the expected consequences associated with the realization of the undesirable events of deflation and excessive inflation. The precise form of these expected consequences depends on the degree of risk aversion.

Figure 1 illustrates the effect of the parameters α and β on the shape of $u(\pi)$. These parameters reflect the central bank's attitude toward missing the threshold values by different amounts. For values of α and β below 1, the central banker's main concern is that he may fail to attain the target. The extent to which the target is missed is of lesser concern. In contrast, for values of α and β in excess of 1 the central banker considers smaller deviations from target as less of a concern than large deviations.

Interpretation of Risk in Some Interesting Special Cases

We now discuss the interpretation of leading examples of our risk measures that correspond to different degrees of risk reversion. First, consider the limiting case of $\alpha = \beta = 0$. In this case, we obtain:

$$DR_0 = - \int_{-\infty}^{\underline{\pi}} dF(\pi) = -\Pr(\pi < \underline{\pi}) \qquad EIR_0 = \int_{\bar{\pi}}^{\infty} dF(\pi) = \Pr(\pi > \bar{\pi})$$

This result is intuitive because for $\alpha = \beta = 0$ the central banker is only concerned about not missing the target zone, but, conditional on missing the target zone, does not care at all by how much inflation will exceed the target. This set of preferences would apply to a central banker concerned about avoiding at all cost the unpleasant task of having to report to parliament a failure to attain the inflation target. Although instructive, this limiting case is implausible in that central bankers in practice would be not be indifferent to whether inflation misses the target zone by a small amount or by a large amount.

Next, consider the case of $\alpha = \beta = 1$. In that case, we obtain:

$$DR_1 = - \int_{-\infty}^{\underline{\pi}} (\underline{\pi} - \pi) dF(\pi) \qquad EIR_1 = \int_{\bar{\pi}}^{\infty} (\pi - \bar{\pi}) dF(\pi)$$

By construction DR_1 is a measure of expected deflation (defined as inflation below $\underline{\pi}$), and EIR_1 a measure of expected excess inflation (defined as inflation above $\bar{\pi}$).⁴ A different

⁴ One must be careful in interpreting this number, since the expected excess relative to the threshold cannot be added to the threshold. For example, suppose that the inflation rate can be either 4% or -4% with probability 0.5 each. Then expected deflation will be $(1\% - 4\%) \cdot 0.5 = -1.5\%$. It would be wrong to conclude that expected deflation given that inflation falls below the threshold of 1% is $(1\% - 1.5\%) = -0.5\%$. Instead, the expected deflation given that inflation falls below the threshold of 1% would be -4%.

way of interpreting these measures is to disaggregate them into two components – a conditional expectation and a tail probability, both of which are time varying. For example, we may write: $DR_1 = E(\pi - \underline{\pi} | \pi < \underline{\pi}) \Pr(\pi < \underline{\pi})$. In other words, this measure of deflation risk is given by the product of the expected shortfall of inflation given that the inflation rate falls below the lower threshold, times the probability that this event occurs. A symmetric interpretation holds for the risk of excessive inflation.⁵

A third leading example is $\alpha = \beta = 2$. In that case, our general risk measure reduces to the target semi-variance:

$$DR_2 = - \int_{-\infty}^{\underline{\pi}} (\underline{\pi} - \pi)^2 dF(\pi) \qquad EIR_2 = \int_{\bar{\pi}}^{\infty} (\pi - \bar{\pi})^2 dF(\pi)$$

In this case, the central banker aims to minimize in expectation the squared deviations of inflation from the lower threshold and the squared deviations of inflation from the upper threshold. An important special case of this measure is obtained under quadratic utility for inflation.

$$u(\pi) = -0.5(\pi_t - \pi^*)^2$$

In that case, the optimal risk management strategy reduces to minimizing the variance about the target, as shown for example by Svensson (1997). This result is consistent with our analysis. Note that by Proposition 2 we have

$$E[u] = aDR_2 - (1 - a)EIR_2,$$

Under the additional assumption of symmetric preferences for the risks of deflation and excessive inflation ($a = 1/2$), and the further assumption that the central banker is an inflation targeter ($\underline{\pi} = \bar{\pi} = \pi^*$) it follows that:

$$E[u] = 0.5(DR_2 - EIR_2) = 0.5\left(- \int_{-\infty}^{\pi^*} (\pi^* - \pi)^2 dF(\pi) - \int_{\pi^*}^{\infty} (\pi - \pi^*)^2 dF(\pi)\right) = -0.5E(\pi - \pi^*)^2$$

From the last line it follows immediately that expected utility is maximized when

⁵ Interestingly, our measure of deflation risk, for $\alpha = 1$, is formally equivalent to the measure of (downside) risk recently proposed by Basak and Shapiro (2001) in the context of value-at-risk applications under the name of *weighted expected shortfall*.

minimizing the variance of inflation about the target π^* .⁶ Our analysis makes explicit that even for inflation targeters this strategy will be correct only under additional assumptions about preferences. More importantly, this strategy will be misleading when $\underline{\pi} < \bar{\pi}$.

2.3. Congruence with Stochastic Dominance Criteria

Stochastic dominance criteria are of interest because of their ability to deliver - under suitable conditions - preference rankings that are consistent with a wide range of unknown von Neumann-Morgenstern utility functions. For this reason it is of interest to study the relationship between our proposed model of risk management and stochastic dominance, whenever stochastic dominance can be shown to exist. The discussion below establishes that our proposed risk measures have the attractive feature of never ranking distribution F ahead of distribution G when stochastic dominance would suggest that we prefer G to F .

Being able to rank risks requires us to restrict the class of preferences. We begin by putting minimal structure on central bank preferences. Although no specific central banker's utility function is publicly available, we do know that central bankers typically aim at keeping inflation in a band between $\underline{\pi}$ and $\bar{\pi}$, and that they dislike inflation in excess of the upper threshold ("excessive inflation") as well as inflation below the lower threshold ("deflation"). These basic assumptions are formalized in Definition 1 below.

Definition 5 [*Central banker's preferences*]

Let $\underline{\pi} \leq \bar{\pi}$ represent a lower and upward threshold for inflation. A central banker's preferences are defined by a utility function $u(\pi)$ that satisfies the following conditions:

- $u'(\pi) > 0$ for $\pi < \underline{\pi}$
- $u'(\pi) = 0$ for $\underline{\pi} < \pi < \bar{\pi}$
- $u'(\pi) < 0$ for $\pi > \bar{\pi}$

$\bar{\pi}$ and $\underline{\pi}$ are known to the central banker. For $\bar{\pi} = \underline{\pi}$ the central banker may be viewed as

⁶ Note that the variance about the target in general will differ from the variance of inflation about the sample mean.

an inflation targeter, whereas for $\bar{\pi} > \underline{\pi}$ the central banker may be viewed as an inflation-zone targeter. In general, $\bar{\pi}$ and $\underline{\pi}$ could be time-varying.

If one is willing to put more structure on the preferences of the central banker, a natural assumption to make is that central bankers are risk averse. With this additional assumption preferences may be summarized as:

Definition 6 [*Risk-averse central banker's preferences*]

Let $\underline{\pi} \leq \bar{\pi}$ represent a lower and upward threshold for inflation. A central banker's preferences are defined by a utility function $u(\pi)$ that satisfies the following conditions:

- $u'(\pi) > 0$ and $u''(\pi) \leq 0$ for $\pi < \underline{\pi}$
- $u'(\pi) = 0$ for $\underline{\pi} < \pi < \bar{\pi}$
- $u'(\pi) < 0$ and $u''(\pi) \leq 0$ for $\pi > \bar{\pi}$

Given these definitions of preferences, we can define first-order and second-order stochastic dominance in the central banking context. Let $F(\pi)$ and $G(\pi)$ be two alternative probability distribution functions for the inflation rate.

Definition 7 [*Stochastic dominance*]

(a) We say that distribution F first-order stochastically dominates (FOSD) distribution G if every central banker prefers distribution F to distribution G.

(b) We say that distribution F second-order stochastically dominates (SOSD) distribution G if every risk-averse central banker prefers distribution F to distribution G.

To make this definition operational, we need to be able to characterize the relative risks in terms of the probability distributions without reference to preferences. For the class of monotonically increasing utility functions, this reformulation is possible, provided the distributions in question satisfy certain conditions (see Hadar and Russell 1969; Hanoch and Levy 1969; Rothschild and Stiglitz 1970). It is natural to adapt these results to the central banking context. A modification is needed to account for the fact that central bank utility functions are not monotonically increasing in the underlying random variable. The

following proposition shows when this extension is possible. The proofs are omitted because they are straightforward extensions of the proofs provided in the original papers.

Proposition 3:

(a) F FOSD G if and only if conditions 1 and 2 below hold:

1. $\int_{-\infty}^{\underline{\pi}} f(h)dh \leq \int_{-\infty}^{\underline{\pi}} g(h)dh$ for every $\underline{\pi} \in (-\infty, \underline{\pi})$

2. $\int_{\bar{\pi}}^{\infty} f(h)dh \leq \int_{\bar{\pi}}^{\infty} g(h)dh$ for every $\bar{\pi} \in (\bar{\pi}, \infty)$.

(b) F SOSD G if and only if conditions 3 and 4 below hold:

3. $\int_{-\infty}^{\underline{\pi}} F(h)dh \leq \int_{-\infty}^{\underline{\pi}} G(h)dh$ for every $\underline{\pi} \in (-\infty, \underline{\pi})$

4. $\int_{\bar{\pi}}^{\infty} (1-F(h))dh \leq \int_{\bar{\pi}}^{\infty} (1-G(h))dh$ for every $\bar{\pi} \in (\bar{\pi}, \infty)$

The next proposition states the main result of this sub-section. It says that whenever F stochastically dominates G , the linear risk management model selects F over G .

Proposition 4:

(a) If F FOSD G , then $aDR_{\alpha}(\underline{\pi}, F) + (1-a)EIR_{\beta}(\bar{\pi}, F) < aDR_{\alpha}(\underline{\pi}, G) + (1-a)EIR_{\beta}(\bar{\pi}, G)$, for all $\alpha > 0$ and $\beta > 0$.

(b) If F SOSD G , then $aDR_{\alpha}(\underline{\pi}, F) + (1-a)EIR_{\beta}(\bar{\pi}, F) < aDR_{\alpha}(\underline{\pi}, G) + (1-a)EIR_{\beta}(\bar{\pi}, G)$, for all $\alpha > 1$ and $\beta > 1$. If $\alpha=1$ and $\beta=1$, this result continues to hold, except when $DR_{\alpha}(\underline{\pi}, F) = DR_{\alpha}(\underline{\pi}, G)$ and $EIR_{\beta}(\bar{\pi}, F) = EIR_{\beta}(\bar{\pi}, G)$.

Proof: See Appendix.

Proposition 4 shows that whenever first-order stochastic dominance holds, using the linear representation of the risk management model will produce the same rankings. The same result holds for second-order stochastic dominance provided that risks are computed based on $\alpha > 1$ and $\beta > 1$. This choice is natural, given that second-order stochastic dominance is a concept designed for risk averters. Proposition 4(b) shows that a risk measure based on risk neutrality ($\alpha = 1, \beta = 1$) nevertheless will provide the correct ranking under risk

aversion, except in rare cases. Even in those rare cases, however, it will not rank a second-order stochastically dominated distribution strictly higher than the alternative.

2.4. Implications of the Risk Management Model for Policy Rules

Suppose the central bank aims at containing inflation between $\underline{\pi}$ and $\bar{\pi}$ over a given horizon by setting in advance the policy instrument, say, the discount rate, denoted by i . Let $F_{\pi}(i)$ denote the distribution of inflation as a function of the policy instrument i . Then the general one-period decision problem of the central banker reduces to choosing the interest rate that maximizes U :

$$(1) \quad i^* = \operatorname{argmax}_i \{U[DR_{\alpha}(\underline{\pi}, F_{\pi}(i)), EIR_{\beta}(\bar{\pi}, F_{\pi}(i))]\}.$$

In contrast, in the conventional expected utility model, the one-period decision problem of the central banker reduces to choosing the interest rate that maximizes $E(u)$:

$$(2) \quad i^* = \operatorname{argmax}_i E\{u(\pi(i))\}.$$

Proposition 2 establishes a simple link between (1) and (2). It allows us to rewrite the stylized decision problem underlying expression (2) equivalently in terms of a trade-off between inflation risks and deflation risks:

$$(3) \quad i^* = \operatorname{argmax}_i \{aDR_{\alpha} - (1-a)EIR_{\beta}\}$$

Note that we do not presume to have knowledge of the parameter a . Nevertheless, our risk measures allow policymakers who are willing to take a stand on a to assess the desirability of alternative policy choices. Specifically, we may postulate alternative hypothetical paths for the interest rate on a grid with increments of 25 basis points. Then, for a given a , these choices may be ranked by the approximate expected utility they generate. Our risk measures also allow in principle an investigation of the effect of alternative choices of a on the optimal setting of the policy instrument.

It is common for central bankers to stress the need to balance the upside risks and the downside risks to price stability. As we will show next, under plausible assumptions there is indeed a close link between the balance of risks and the optimal choice for the policy

instrument. Given our risk measures, the balance of risks can be formally defined as follows:

Definition 8 [Balance of Risks]

$BR_{\alpha,\beta} = \omega DR_{\alpha} + \nu EIR_{\beta}$ where ω and ν are positive real numbers.

When central bank preferences satisfy risk aversion (see Definition 6), it is possible to derive conditions for the optimality of the choice of i in terms of the balance of risks.

Proposition 5:

Let $\pi = g(i) + \varepsilon$, where $\varepsilon \sim F$ and $g(i)$ is twice continuously differentiable. When $\alpha > 1$ and $\beta > 1$, the optimal choice i^* in problem (3) satisfies the following necessary and sufficient condition:

$$BR_{\alpha-1,\beta-1}(i^*) = 0 \text{ where } \omega = \alpha\alpha \text{ and } \nu = (1-\alpha)\beta.$$

Proof: See Appendix.

This means that under risk aversion the central banker should aim for $BR_{\alpha-1,\beta-1} = 0$. Note that the coefficients ω and ν are scale invariant, which mirrors the fact that utility may be rescaled. In the special case of $\alpha = \beta = 2$ and $a = 0.5$, as discussed, for example, by Blinder (1997), we obtain $BR_{1,1} = DR_1 + EIR_1$.

2.5. Extensions to Multiple Policy Objectives

Our analysis can also be extended to situations involving multiple objectives. We will not pursue this possibility in this paper, but briefly outline how the analysis would change. For concreteness let us postulate that the central bank pursues inflation and output objectives. Let $F_{\pi}(i)$ denote the distribution of inflation as a function of the policy instrument i , and $F_y(i)$ the corresponding distribution of the output gap, denoted by y . The general solution to the central bank's decision problem can be written in terms of four risks:

$$(1') \quad i^* = \underset{i}{\operatorname{argmax}} \{U[DR_{\alpha}(\underline{\pi}, F_{\pi}(i)), EIR_{\beta}(\bar{\pi}, F_{\pi}(i)), NGR_{\gamma}(\underline{y}, F_y(i)), PGR_{\delta}(\bar{y}, F_y(i))]\}.$$

Here NGR_t and PGR_t stand for the risk of a negative gap and a positive gap, respectively, and α , β , γ , and δ refer to the respective degrees of risk aversion. A natural choice for \underline{y} and \bar{y} would be $\bar{y} = \underline{y} = 0$. If we follow the literature in assuming separability of the utility function in the objectives, then, given a generalization of Proposition 2, the implied decision rule can equivalently be expressed as:

$$(3') \quad i^* = \arg \max_i \{aDR_\alpha + bEIR_\beta + cNGR_\gamma + dPGR_\delta\},$$

where a , b , c and d refer to the weights that sum to 1.

Note that the existence of additional objectives does not render the measures of deflation risk and of excessive inflation risk that we proposed earlier invalid. Under additive separability, one can simply view DR_α and EIR_β as part of a larger set of inputs into the decision-making process of the central banker. Moreover, these measures may be computed even in the absence of data on the output gap.

Our analysis again includes as a special case the loss functions currently under consideration in the academic literature for multiple objectives. Consider for example the period utility function $u(\pi, y) = -(\pi_t - \pi^*)^2 - \lambda y^2$, discussed by Svensson (2002). In that case, we obtain a policy reaction function of the form: $i^* = \arg \max_{\{i\}} \{E(\pi - \pi^*)^2 + \lambda E(y)^2\}$, where $a = b = 1$ and $c = d = \lambda$, consistent with Svensson's result.

3. COMPUTING MEASURES OF INFLATION RISK IN PRACTICE

We illustrate the estimation of the inflation risk measures defined in section 2 in the context of a parametric model of inflation dynamics. In the presence of possible model misspecification, it is natural to fit the model directly at the horizon of interest. This approach also avoids the well-known problems of time aggregation of volatility dynamics. Let π_τ , $\tau = 1, 2, \dots, T$, denote nonoverlapping observations for the inflation rate measured at the horizon of interest. Then:

$$(4) \quad \begin{aligned} \pi_{\tau+1} &= \mu_{\tau+1} + u_{\tau+1}, & u_{\tau+1} &= \sqrt{h_{\tau+1}} \varepsilon_{\tau+1}, & \varepsilon_{\tau+1} | \Omega_\tau &\sim i.i.d.(0,1) \\ h_{\tau+1} &= a u_\tau^2 + b h_\tau + c \end{aligned}$$

where Ω_τ denotes the information set at date τ . The use of the GARCH framework for modeling the conditional variance of inflation, here denoted by h_τ , was originally suggested by Engle (1982) and Bollerslev (1986). For $a = 0$ and $b = 0$, this model reduces to the homoskedastic case with $h_{\tau+1} = c \forall \tau$.

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 i.i.d. white noise.

A natural choice for the conditional mean, $\mu_{\tau+1}$, is an autoregressive model for inflation:

$$\mu_{\tau+1} = \phi_0 + \sum_{i=1}^p \phi_i \pi_{\tau-i+1},$$

where p denotes the autoregressive lag order (see, e.g., 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:

$$\mu_{\tau+1} = \phi_0 + \sum_{i=1}^p \phi_i \pi_{\tau-i+1} + \sum_{k=1}^s \sum_{j=1}^{q_k} \phi_{k,j} x_{k,\tau-j+1},$$

where the additional predictors, $x_{k,\tau}$, $k = 1, \dots, s$, may enter with potentially different lag orders q_k for each predictor (see, e.g., Engle 1982). In our empirical work we will consider alternative specifications of the conditional mean and use information criteria to choose between them.

The conditional mean model may be estimated consistently by ordinary least squares (OLS). Consistent estimates of the GARCH parameters may be obtained by maximizing the quasi-log-likelihood function (see Bollerslev and Wooldridge 1992). Given

parameter estimates of model (4), we proceed as follows. In general, all proposed risk measures may be computed by bootstrap simulation. Consider $EIR_{\beta, \tau+1}(\bar{\pi}) =$

$\int_{\bar{\pi}}^{\infty} (\pi_{\tau+1} - \bar{\pi})^{\beta} dF(\pi_{\tau+1})$, for example. Denote the estimated standardized residuals of model (4) by $\hat{\varepsilon}_{\tau+1} \equiv (\pi_{\tau+1} - \hat{\mu}_{\tau+1}) / \sqrt{\hat{h}_{\tau+1}}$. Given the empirical distribution of $\hat{\varepsilon}_{\tau+1}$, we first generate the bootstrap predictive density $\{\pi_{\tau+1}^*\}_{\tau=1}^T$, conditional on the parameter estimates of model (4).⁷ Then we make use of the fact that

$$EIR_{\beta, \tau+1}(\bar{\pi}) = \Pr(\pi_{\tau+1}^* > \bar{\pi}) E[(\pi_{\tau+1}^* - \bar{\pi})^{\beta} | \pi_{\tau+1}^* > \bar{\pi}],$$

where the probability $\Pr(\pi_{\tau+1}^* > \bar{\pi})$ can be computed as the fraction of observations $\pi_{\tau+1}^*$ that exceed $\bar{\pi}$, and where the expectation $E[(\pi_{\tau+1}^* - \bar{\pi})^{\beta} | \pi_{\tau+1}^* > \bar{\pi}]$ can be estimated by the sample moment $\frac{1}{\#(\pi_{\tau+1}^* > \bar{\pi})} \sum_{\pi_{\tau+1}^* > \bar{\pi}} (\pi_{\tau+1}^* - \bar{\pi})^{\beta}$.

When $\beta = 1$ or $\beta = 2$, the risk measures can be computed directly without resorting to bootstrap simulation. For example, in computing the risk measure $EIR_{1, \tau+1}(\bar{\pi}) = \Pr(\pi_{\tau+1} > \bar{\pi}) E[(\pi_{\tau+1} - \bar{\pi}) | \pi_{\tau+1} > \bar{\pi}]$ simply note that

$$E_{\tau}[\pi_{\tau+1} - \bar{\pi} | \pi_{\tau+1} > \bar{\pi}] = \mu_{\tau+1} - \bar{\pi} + \sqrt{h_{\tau+1}} E[\varepsilon_{\tau+1} | \varepsilon_{\tau+1} > (\bar{\pi} - \mu_{\tau+1}) / \sqrt{h_{\tau+1}}],$$

where the last expectation can be computed by taking averages of the standardized residuals that satisfy the inequality condition. The probability $\Pr(\pi_{\tau+1} > \bar{\pi})$ can be computed, as discussed before. Similarly, for $EIR_{2, \tau+1}(\bar{\pi}) = \Pr(\pi_{\tau+1} > \bar{\pi}) E[(\pi_{\tau+1} - \bar{\pi})^2 | \pi_{\tau+1} > \bar{\pi}]$ we have

$$E_{\tau}[(\pi_{\tau+1} - \bar{\pi})^2 | \pi_{\tau+1} > \bar{\pi}] = (\mu_{\tau+1} - \bar{\pi})^2 + h_{\tau+1} E[\varepsilon_{\tau+1}^2 | \varepsilon_{\tau+1} > (\bar{\pi} - \mu_{\tau+1}) / \sqrt{h_{\tau+1}}] + 2(\mu_{\tau+1} - \bar{\pi}) \sqrt{h_{\tau+1}} E[\varepsilon_{\tau+1} | \varepsilon_{\tau+1} > (\bar{\pi} - \mu_{\tau+1}) / \sqrt{h_{\tau+1}}],$$

which may be computed analogously.

Throughout this paper, 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 distribution $\pi_{\tau+1}^* | \Omega_{\tau}$ and the associated measures of risk, given the full-sample

⁷ The estimation uncertainty in the model parameters will be asymptotically negligible (see, e.g., Cao et al. 1997).

parameter estimates. In contrast, we compute *forecasts* of risks by recursively re-estimating the forecast model at each point in time, before computing the conditional distribution of $\pi_{\tau+1}^* | \Omega_\tau$ and the associated measures of risk, given the most recent recursive parameter estimates.

We conclude this section with an illustration of how to compute risks of year-on-year inflation based on monthly data for the CPI and West Texas Intermediate (WTI) oil prices. Let y_t^h denote the percent change in variable y from h months 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} and the corresponding percent change in oil prices, Δwti_t^{12} , where $t = 1+12, 2+12, \dots, T$. Given these data, the model for year-on-year inflation may be written as $\pi_{t+12}^{12} = \mu_{t+12}^{12} + u_{t+12}^{12}$. We postulate that $u_{t+12}^{12} = \sqrt{h_{t+12}^{12}} \varepsilon_{t+12}^{12}$, where $\varepsilon_{t+12}^{12} | \Omega_t \sim i.i.d.(0,1)$ and h_{t+12}^{12} is governed by a GARCH(1,1) process. In practice, we test for the absence of serial correlation in $(u_{t+12}^{12})^2$. If the null cannot be rejected, we treat u_{t+12}^{12} as i.i.d. Given $\pi_i^{12}, \dots, \pi_{i-12(p-1)}^{12}$ and $\Delta wti_i^{12}, \dots, \Delta wti_{i-12(q-1)}^{12}$, the conditional mean model is $\mu_{t+12}^{12} = \phi_0 + \sum_{j=0}^{p-1} \phi_{j+1} \pi_{t-12j}^{12} + \sum_{j=0}^{q-1} \phi_{k,j+1} \Delta wti_{t-12j}^{12}$, for $t = i, i+12, \dots, T-12+i$, where $i = 1, 2, \dots, 12$ refers to the i th month in the sample. This means that we fit the model separately for each month. For each point in time, we generate predictive densities for π_{t+12}^{12} based on the appropriate model estimate and compute the associated risk measures. Finally, we assemble the risk estimates into a monthly time series.

4. FORECASTING INFLATION RISKS IN REAL TIME: A STEP BEYOND THE NARRATIVE APPROACH TO MONETARY POLICY

Risk analysis based on full-sample parameter estimates for ex post-revised data provides the most accurate picture of inflation risks that prevailed at the time, but it is not suitable in general 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 illustrate the feasibility of our methodology in a true real time setting. We will construct forecasts of U.S. inflation and deflation risks period-by-period. In constructing these risk forecasts, we will recursively re-estimate the forecasting model for each period, using only the incompletely revised data actually available to the Fed in real 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. We will show that there is a close match between the Fed's own assessment of inflation risks at the time and the quantitative measures of inflation risk.

4.1. Real-Time Estimation Methodology

The analysis is 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 Orphanides (2001, 2003).

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. For expository purposes we will focus on forecasts of year-on-year inflation risks. Specifically, let $\pi_{r|s}^h$ 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 $\pi_{r|s}^1$ where $r = s - 1, \dots, s - 4$, and $s = 1, \dots, T$. Let N denote the length of the sample used in generating recursive real time forecasts. By definition,

$N = R, R + 1, R + 2, \dots, T$, where R is the length of the initial recursive sample. Furthermore,

define $\pi_{t-|t}^h \equiv \sum_{r=t-h}^{t-1} \pi_{r|t}^1 / h$ for $t = i + h, i + 2h, \dots, N - h + i$ and

$\pi_{t+h-|t+2h-1}^h \equiv \sum_{r=t}^{t+h-1} \pi_{r|t+h}^1 / h$ for $t = i + h, i + 2h, \dots, N - 2h + i + 1$, where $i = 1, \dots, h$. Here

$i = 1$ refers to the first quarter of the sample, $i = 2$ to the second quarter, etc., and $h = 4$. Given these data, we compute the horizon- h risks as described in section 3.

For given h , the forecast model can be written as $\pi_{t+h-1}^h = \mu_{t+h-1|t}^h + u_{t+h-1}^h$, where π_{t+h-1}^h denotes the ex-post revised inflation rate (which in practice may be proxied by $\pi_{t+h-1|t+2h-1}^h$). 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 $\widehat{gap}_{t|t}$.⁸ Specifically, the conditional mean model is $\mu_{t+h-1|t}^h = \phi_0 + \sum_{j=0}^{p-1} \phi_{j+1} \pi_{t-jh-1|t}^h + \phi_1 \widehat{gap}_{t|t}$, where the lag order p may change with the horizon. We set $p = 2$. The parameters of the conditional mean model are estimated by running the regression:

$$\pi_{t+h-1|t+2h-1}^h = \phi_0 + \sum_{j=0}^{p-1} \phi_{j+1} \pi_{t-jh-1|t}^h + \phi_1 \widehat{gap}_{t|t} + u_{t+h-1}^h$$

for $t = i + h, i + 2h, \dots, N - 2h + i + 1$ and $i = 1, \dots, h$. This means that we estimate h forecast models, one for each i . Each of these forecast models is estimated excluding the last $2h - 1$ 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 compute the real-time forecasts of the conditional mean of $\pi_{N+h-1|N}^h$ 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 discussed in section 3, whenever the Ljung-Box test indicates the presence of conditional heteroskedasticity. In the specific example discussed below, we typically found no evidence of GARCH.

⁸ 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 and the limitations of the data set.

Given these conventions and definitions, our procedure can be summarized as follows: Step 1) Recursively estimate the models of the conditional mean and of the conditional variance of $\pi_{t+h-1|t}^h$ for $t = i + h, i + 2h, \dots, N - 2h + i + 1$ and $i = 1, \dots, h$. Step 2) Compute the recursive conditional distribution of $\pi_{N+h-1|N}^h$ at each point in time. Step 3) Given the real-time forecast distributions, compute the implied risk measures.

4.2. Real-Time Forecasts of Inflation Risks at the One-Year Horizon

We are ultimately interested in relating the one-year ahead measures of inflation risk to actual policy discussions. For expository purposes, we postulate that the Fed in the Volcker-Greenspan period was risk averse in the form of piecewise-quadratic preferences with $\alpha = 2$ and $\beta = 2$. We also postulate that the Fed was intent on containing inflation between 1% and 3% in the medium run. Although there is no independent evidence to support these choices, we will show that the resulting risk estimates closely mimic the Fed's own perception of inflation risks.⁹

Under these assumptions, it follows from Proposition 5 that in the absence of an additional output objective for monetary policy the Fed should have aimed for

$$(5) \quad BR_{1,1}(i^*) = aDR_1 + (1 - a)EIR_1 = 0 .$$

Is this a plausible description of Fed behavior? Arguably, there is no output-inflation trade-off in the long run, so when the forecast horizon is chosen long enough expression (5) is likely to be a reasonable description of the disinflation policy under Volcker and Greenspan. In practice, however, we will work with one-year forecast horizons, given the data constraints. At these horizons a monetary tightening can be expected to have a cost in terms of output. Thus, the speed, with which the Fed will aim to achieve balance will in general depend on the extent to which it pursues an additional output objective. Even if the Fed could lower the inflation risk to zero in one period, it would in general choose not to do so, given the negative output risk this would imply. Instead, the Fed will aim to drive down inflation risks smoothly. The degree of inter-temporal smoothing may also reflect changes in the relative weights of the inflation and the output objective over time. As we will show

⁹ Our qualitative results are not sensitive to the choice of $\bar{\pi}$. The choice of α and β affects the sensitivity of the risk measure, but leaves the trend and turning points unchanged.

below, there were periods in which the Fed, by its own accounts, largely discarded the output objective and aggressively combated inflation risks, and there were other periods, in which the Fed postponed measures to lower inflation risks with an explicit reference to the output objective.

The year-on-year real-time risk forecasts of DR_t and EIR_t based on the data of Orphanides (2003) are plotted in Figure 2. The time line refers to the date, at which the forecasts were made. Figure 2 shows that real-time measures of DR_t were essentially zero throughout this period. As a result, the real-time inflation risk, EIR_t , will be proportionate to the balance of risks, $BR_{t,1}$, and, in what follows, we may focus on the risk of excessive inflation without much loss of information. Given that DR_t is effectively zero throughout the sample period, by Proposition 1a, the Fed's decision problem simplifies to raising the interest rate to reduce the risk of excessive inflation. This is indeed consistent with the Fed's decision to raise interest rates sharply in 1980. In response to this policy decision, Figure 2 shows a steady decline of inflation risks at the beginning of the sample, which in turn was followed by a decline in the Fed Funds rate. Our analysis suggests that for the remainder of the sample, the objective of achieving balance would have been well served, if the Fed had simply raised the interest rate, whenever the process of disinflation threatened to stall or to reverse itself. In the next subsection we will study how closely actual Fed decisions followed this stylized description of policy.

4.3. Policy Analysis Using Real-Time Forecasts of Risks

We follow the existing literature in treating the Fed Funds rate as the policy instrument of the Fed since August 1982. Figure 3 plots the Federal Funds rate and quarterly real-time forecasts of the risk of inflation for the year to come. Given that concerns about the level of output and unemployment may substantially weaken the empirical relationship that our stylized model leads us to expect, our discussion focuses on selected episodes in our sample, during which the Fed - by its own account - explicitly acted only in pursuit of the inflation objective. In identifying such episodes 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 3 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 year-on-year inflation risks, as measured by our indicator in Figure 3. This resurgence coincided with the gradual recovery of the economy from the trough of the recession in November 1982. As early as 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 monetary tightening of the Fed in 1984 is fully consistent with the predictions of our stylized policy rule for inflation zone targeters. Given that the process of disinflation had stalled and was threatening to unravel, in the absence of an output objective, the only possible response for an inflation zone targeter would have been to raise interest rates to prevent the balance from rising and - if possible – to force the balance closer to zero.

The policy intervention of 1984 appears to have been successful in that year-on-year 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.

Both of these motivations are consistent with the view that a central banker should strive to move the balance of risks toward zero by raising the interest rate. 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. Figure 3 shows that indeed the risks of inflation had steadily increased since 1987, even as the Fed Funds rate had been gradually adjusted upward. Starting in early 1988, forecasts of year-on-year inflation risks appeared to increase further, signalling increased dangers to price stability. From March 1987 to March 1988 alone, inflation risk had increased considerably. Between June 1988 and June 1989, one-year forecasts of upward risks increased even further. By mid-1989 forecasts of inflation risk were at their highest level since 1981. Figure 3 shows that the risks of inflation during this episode were much higher than during 1984-85.

The evidence in Figure 3 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, year-on-year 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) and that actual policy decisions closely adhered to the predictions of our stylized model. Again this close match is not surprising, given that the FOMC during this sub-period 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 (see Romer and Romer 1994, p. 82). In other words, the FOMC subordinated the output objective to that of lowering inflation (also see Romer and Romer 2002, p. 43).

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 3. It began in early 1994 and continued until early 1995.

Unlike in the other two episodes, there is no apparent increase in inflation risks in 1994 leading up to the monetary tightening. Rather forecasts of inflation risks, although noisy, show no trend at all between 1992 and 1994. 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. The downward trend in inflation risks ended in late 1992. For one year, inflation risks remained almost unchanged. Given this evidence that the process of disinflation had stalled, according to our policy rule for inflation zone targeters, one would have expected a monetary tightening as early as 1992 or 1993. This tightening did not happen. Rather the interest rate remained flat. Apparently the Fed decided in response to its output and employment objective to keep interest rates unchanged. 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)

Interest rates remained unchanged for almost a year. Only in early 1994, the Fed decided to raise the Fed Funds rate to levels more consistent with inflation risks. Rather than being prompted by an unforeseen increase in inflation risk, that response reflected an explicit shift in priorities away from output and back toward the primary objective of price stability. When the recovery from the 1991 recession was in full swing, the Fed returned to its objective of fighting inflation and finally responded to the evidence of positive and persistent inflation risks, consistent with our stylized policy rule. 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 began to decline further, as one would expect, following the increase in the Fed Funds rate.

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 3 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.¹⁰

¹⁰ A subtle point in comparing inflation risk forecasts to the Fed’s own assessment of inflation risks is that the Fed – rather than forecasting inflation – arguably based its policy decisions on projections of inflation conditional on the assumption that the interest rate remains unchanged. Investigating this possibility requires adding lags of the interest rate to the conditional mean model and generating forecasts of inflation conditional

5. 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 warned that “the risk of falling prices is greater than at any time since the 1930s”. Similarly, in November 2002 the Goldman-Sachs Global Economics Weekly concluded that “investors continue to worry about the spectre of global deflation” (see Wilson 2002).¹¹ These concerns in turn have elicited a response by central bankers. For example, Bernanke (2002) discussed measures the “central bank can take to reduce the risk of falling into deflation” and makes the case that “the chance of a significant deflation in the United States in the foreseeable future is extremely small”. At about the same time, Issing (2002b) made 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.

A Snapshot of the Risks to Price Stability

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, when the issue of deflation risks was first raised.¹² 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

on the interest rate remaining unchanged. Although this modification is conceptually straightforward, adding more predictors to the conditional mean model rapidly depletes the degrees of freedom available for this real-time data set, further shortening the sample available for prediction. Thus, we are only able to compute conditional risk forecasts for episodes 2 and 3. These results are substantively identical to those in Figure 3.

¹¹ 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).

¹² Note that the risk of deflation is conceptually different from the risk that the zero bound on interest rates becomes binding. The latter question can only be addressed within the context of a specific macroeconomic model. We do not pursue this latter question, because it is ultimately a question about the appropriate choice of the policy instrument, whereas our focus is on quantifying the risks for any possible choice of policy instrument.

construction of risk forecasts for the horizons of interest here.¹³

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 data source for the U.S. data is the Bureau of Labor Statistics. 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.¹⁴ 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 their balance for horizons of one year and of two years, as described earlier. Clearly, these results should be viewed with some caution, as the effective sample size is small, especially for the two-year horizon. Nevertheless, it is of interest to obtain at least a preliminary and tentative assessment of the risks.

We consider three alternative specifications of the conditional mean in model (4): one involving only lagged inflation ($\Delta\pi_t$) and the other models including in addition lagged percentage changes of the oil price (Δwti_t) or lagged growth rates of money (Δm_t):

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

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

$$\text{Model 3:} \quad \mu_{t+1} = c + \sum_{i=1}^p \phi_i \pi_{t-i+1} + \sum_{j=1}^r \zeta_{k,j} \Delta m_{t-j+1}$$

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

¹³ 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. There is no statistical evidence of a structural break in the German inflation rate process due to unification or due to the introduction of the Euro.

¹⁴ 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.

did not investigate the use of additional predictors, given the large number of parameters involved.¹⁵

Table 1 presents the model diagnostics by country. We first use the SIC to select the optimal number of lags for each model.¹⁶ We then choose between the alternative models based on the SIC values of the models evaluated at these lag orders. We select Model 1 for the United States, Model 3 for Japan and Model 1 or Model 3, depending on the horizon, for Germany on the grounds that these forecast models have the lowest SIC value among the models considered (see Table 1 for details). Further diagnostic tests confirm that the preferred models fit 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 in all cases. We also test for conditional heteroskedasticity in the model residuals. The Ljung-Box test shows no evidence of conditional heteroskedasticity.¹⁷

The results of the forecast exercise are summarized in Table 2. Table 2 provides a snapshot of the risks of deflation and the risks of excessive inflation as of September 2002. It shows the projected risks by geographic area for a horizon of one year and, alternatively, of two years. We also include the conditional mean forecast for comparison. We note that for all countries there is some evidence of deflation risk, although the magnitudes differ greatly. The U.S. estimate of -0.02 at the one-year horizon is likely to be negligible. The corresponding estimate of -2.19 for Japan appears large, whereas the practical significance of the estimate of -0.29 for Germany is unclear.

¹⁵ The monetary aggregates are M2 for the Japan, M1 for the United States and M3 for Germany. The oil price is the price per barrel of West Texas Intermediate (converted into local currency). The data sources are the OECD Main Economic Indicators, FRED, the European Central Bank and the Deutsche Bundesbank.

¹⁶ See Inoue and Kilian (2003) for a discussion of alternative criteria for forecast model selection. For simplicity we restrict the number of lags to be the same across predictors. The upper bound is three lags.

¹⁷ 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 the models selected provide 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.

One way of assessing how much of a concern these forecasts of deflation risk should be is to pursue the counterfactual question of how an inflation zone targeter would view this evidence. For expository purposes, let us presume that central bankers in these three countries have symmetric, piecewise-quadratic preferences for inflation and that their sole objective is to contain inflation between $\underline{\pi} = 1\%$ and $\bar{\pi} = 3\%$. In that case, as shown in section 2, optimal policy would aim for $BR_{1,t} = DR_t + EIR_t = 0$. This result suggests that it would be a mistake to focus on the forecasts of deflation risk alone, without accounting for the simultaneous presence of inflation risks. Table 2 shows the balance of risk forecasts obtained under symmetric piecewise-quadratic preferences. It is clear that the United States at the one-year horizon are well within the inflationary region on balance, and the existence of minor deflation risks is inconsequential. In contrast, Japan is clearly in the deflationary region. Finally, Germany is only slightly in the deflationary region on balance. For Germany and for the United States forecasts for the two-year horizon indicate a slight improvement of the balance of risks.

This analysis, of course, provides at best a benchmark, since we do not know the actual preferences of central bankers in these countries. The observed deviations from zero balance may be rationalized in a number of ways. We already referred to the possibility of intertemporal smoothing due to an additional output objective. Another possibility is the existence of asymmetric preferences for inflation and deflation. Our results may also be affected by the choice of thresholds. For example, the use of an upper threshold of 2%, as practiced by the European Central Bank, would have shifted the balance for Germany away from the deflationary region. The analysis of Germany is further complicated by the fact that the policy of the European Central Bank is appropriately geared toward the Euro area, as opposed to a single member country.

Putting Deflation Risks into Historical Perspectives

A different approach to assessing how serious the evidence of deflation risk is, is to put the numbers in Table 3 into historical perspective. Figure 4 shows the evolution of the risks of deflation and of excessive inflation since the 1960s. The plots allow us to address, at least in part, the recent claim in the The Economist that “the risk of falling prices is greater than at any time since the 1930s”. We find that deflation risks for the United States today are

not unlike those in the 1960s, for example. For Germany, deflation risks were much higher in the late 1980s than they are today. Only for Japan there is evidence that deflation risks have reached unprecedented levels by post-war historical standards. Moreover, the risks of deflation in Japan are highly persistent. Table 3 presents the same evidence as historical averages by decade together with the average balance of risks. Only for Japan the risks are tilted toward deflation and this tendency has existed since the 1990s. For the United States there is no evidence that the balance of risks today is much different from the 1960s or the 1990s. For Germany the balance has been closer to zero since the year 2000 than for any previous decade, but on average remains slightly above zero.

6. CONCLUSION

In deciding the monetary policy stance, central bankers need to evaluate carefully the risks the current economic situation poses to price stability. We proposed to regard the central banker as a risk manager who aims to contain inflation within pre-specified bounds. We developed formal tools of risk management that may be used to quantify and forecast the risks of failing to attain that objective. We illustrated the use of these risk measures in practice, first, by showing how to construct genuine real time forecasts of year-on-year risks that may be used in policy-making. We demonstrated the usefulness of these risk forecasts in understanding the Fed's decision to tighten monetary policy in 1984, 1988, and 1994. Second, we forecast the risks of worldwide deflation for horizons of up to two years, as of September 2002. Notwithstanding recent concerns about renewed global deflation, we found that, with the exception of Japan, there is no evidence of substantial deflation risks. We also found that only for Japan the estimated deflation risks are high by historical standards and persistent.

There are several interesting extensions of the methodology proposed in this paper. First, it would be important to generalize the simple rules proposed in this paper 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. Such an extension would be

consistent with recent interpretations of United States monetary policy.¹⁸ In this context, it would also be interesting to compare the fit of risk-based policy rules for the United States to more traditional Taylor type rules.

Second, as noted in the introduction, the relevant risk measure must be tailored to the medium-term horizons of interest to the policy maker. Medium term here refers to the horizon over which monetary policy can be expected to have effects. Blinder (1997), for example, notes that these effects are likely to last for several years. It is difficult to compute risks for such long horizons. In the current paper we have restricted ourselves to the analysis of inflation forecasts for at most two years. The reason is that our methodology is based on one-step-ahead forecasts. To compute risks at the two-year horizon, we effectively work with data at bi-annual frequency, which greatly reduces the sample size. This makes it difficult to estimate directly the risks at longer horizons that are of interest to central bankers. We therefore are currently developing an alternative procedure of computing medium-term risks based on multi-step ahead forecasts of risks from models fit at a higher frequency than the horizon of interest.

The third extension relates to the fact that standard forecast models of inflation risk are not suitable for policy analysis. The risk forecasts proposed in this paper implicitly reflect expectations of the future path of policy variables, but such expectations form at best a benchmark for policy analysis. A central banker, in contrast, often wishes to assess how alternative paths of the policy variable may affect the risks of inflation or of deflation. We therefore are currently working on a method of computing forecasts of risks that are conditional on the hypothesized path of the policy instrument. This method will allow us to study how risks change as a function of the path of the policy instrument. It also will allow us to analyze how the central bank's policy decisions are likely to affect future risks to price stability in the context of specific policy decisions.

¹⁸ 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 in the foreseeable future".

APPENDIX

Proof of Proposition 1

Proposition 1 follows from Definition 2. We will prove only part (a), as part (b) can be proved in the same way. By assumption we have $DR_\alpha(\underline{\pi}, F) = DR_\alpha(\underline{\pi}, G) = k$ and therefore, by Definition 2, F is weakly preferred to G if $EIR_\beta(\bar{\pi}, F) \leq EIR_\beta(\bar{\pi}, G)$. This proves the sufficient condition. To prove the necessary condition, suppose that the deflation risks are equal and that F is weakly preferred to G . Further suppose that $EIR_\beta(\bar{\pi}, F) > EIR_\beta(\bar{\pi}, G)$. Then by definition 2, G should be weakly preferred to F , which implies a contradiction. ■

Proof of Proposition 2

The result follows from integrating the utility function with respect to the distribution function of π . ■

Proof of Proposition 4

For $\alpha > 0$ and $\beta > 0$, $aDR_\alpha(\underline{\pi}) + (1-a)EIR_\beta(\bar{\pi})$ is the expected utility of the central banker for any utility function that satisfies Proposition 2. Since this utility function for $\alpha > 0$ and $\beta > 0$ satisfies Definition 5, the result follows. Analogous reasoning can be used to prove the second part of the proposition. Simply note that for $\alpha > 1$ and $\beta > 1$ the utility function satisfies Definition 6. When $\alpha = 1$ and $\beta = 1$, the result follows, except when $DR_\alpha(\underline{\pi}, F) = DR_\alpha(\underline{\pi}, G)$ and $EIR_\beta(\bar{\pi}, F) = EIR_\beta(\bar{\pi}, G)$. ■

Proof of Proposition 5

Under the conditions of Proposition 2, substitute $g(i) + \varepsilon$ for π in the definition of the utility function U . We obtain:

$$U(DR_\alpha, EIR_\beta) \equiv aDR_\alpha - (1-a)EIR_\beta = -a \int_{-\infty}^{\underline{\varepsilon}} (\underline{\pi} - g(i) - \varepsilon)^\alpha dF(\varepsilon) - (1-a) \int_{\bar{\varepsilon}}^{\infty} (g(i) + \varepsilon - \bar{\pi})^\beta dF(\varepsilon)$$

where $\underline{\varepsilon} \equiv \underline{\pi} - g(i)$ and $\bar{\varepsilon} \equiv \bar{\pi} - g(i)$. This function is twice continuously differentiable in i . Therefore a necessary and sufficient condition for a maximum is that the first-order derivative is zero and that the second-order derivative is negative:

$$\frac{\partial U}{\partial i} \Big|_{i=i^*} = a \int_{-\infty}^{\underline{\varepsilon}} \alpha(\underline{\pi} - g(i^*) - \varepsilon)^{\alpha-1} g'(i^*) dF(\varepsilon) - (1-a) \int_{\underline{\varepsilon}}^{\infty} \beta(g(i^*) + \varepsilon - \bar{\pi})^{\beta-1} g'(i^*) dF(\varepsilon) = 0$$

$$\Rightarrow a\alpha DR_{\alpha-1} + (1-a)\beta EIR_{\beta-1} = BR_{\alpha-1, \beta-1}(i^*) = 0$$

Given that $\alpha > 1$ and $\beta > 1$, we have

$$\frac{\partial^2 U}{\partial i^2} \Big|_{i=i^*} = g''(i^*) BR_{\alpha-1, \beta-1}(i^*) - [g'(i^*)]^2 a \int_{-\infty}^{\underline{\varepsilon}} \alpha(\alpha-1)(\underline{\pi} - g(i^*) - \varepsilon)^{\alpha-2} dF(\varepsilon) -$$

$$- [g'(i^*)]^2 (1-a) \int_{\underline{\varepsilon}}^{\infty} \beta(\beta-1)(g(i^*) + \varepsilon - \bar{\pi})^{\beta-2} dF(\varepsilon) < 0$$

which ensures a maximum. ■

References

- Basak, S., and A. Shapiro (2001), "Value-at-Risk Based Management: Optimal Policies and Asset Prices," *Review of Financial Studies*, 14, 407-405.
- Bernanke, B.S. (2002), "Deflation: Making Sure "It" Doesn't Happen Here", Remarks by Governor Ben S. Bernanke, Federal Reserve System, before the National Economists Club, Washington, DC, November 21, 2002.
- Bernanke, B.S. (2003), "Constrained Discretion and Monetary Policy", Remarks by Governor Ben S. Bernanke, Federal Reserve System, before the Money Market Marketeers of New York University, New York, February 3, 2003.
- Bernanke, B.S., and F.S. Mishkin (1997), "Inflation Targeting: A New Framework for Monetary Policy?" *Journal of Economic Perspectives*, 11, 97-116.
- Blinder, A. (1997), "What Central Bankers Could Learn from Academics – and vice versa," *Journal of Economic Perspectives*, 11, 3-19.
- Board of Governors of the Federal Reserve System (1983, 1984, 1992, 1994), *Annual Report*, Washington, DC: Federal Reserve System.
- 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, 143-172.
- Cao, R., M. Febrero-Bande, González-Manteiga, W., Prada-Sánchez, J.M., and I. García-Jurado, (1997), "Saving Computer Time in Constructing Consistent Bootstrap Prediction Intervals for Autoregressive Processes," *Communications in Statistics, Simulation and Computation*, 26, 961-978.
- Drost, F.C., and T.E. Nijman (1993), "Temporal Aggregation of GARCH Processes," *Econometrica*, 61, 909-927.
- Duisenberg, W. (2003), "Testimony before the Committee on Economic and Monetary Affairs of the European Parliament," Brussels, February 17, 2003.

- The Economist (2002a), "Dial D for Deflation," September 14.
- The Economist (2002b), "Of Debt, Deflation and Denial," October 10.
- 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.
- Fishburn, P.C. (1977), "Mean-Risk Analysis with Risk Associated with Below-Target Returns," *American Economic Review*, 67, 116-126.
- Hadar, J., and W. Russell (1969), "Rules for Ordering Uncertain Prospects," *American Economic Review*, 59, 24-34.
- Hanoch, G., and H. Levy (1969), "The Efficiency Analysis of Choices Involving Risk," *Review of Economic Studies*, 36, 335-346.
- Holthausen, D.M. (1981), "A Risk-Return Model with Risk and Return Measured as Deviations from Target Return," *American Economic Review*, 71, 182-188.
- Inoue, A., and L. Kilian (2003), "On the Selection of Forecasting Models," Working Paper No. 214, European Central Bank.
- 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.
- Levin, A., Wieland, V., and J.C. Williams (2002), "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty," forthcoming: *American Economic Review*.
- Machina, M.J., and M. Rothschild (1987), "Risk," in: Eatwell, J., Millgate, M., and P. Newman (eds.), *The New Palgrave Dictionary of Economics*. London: MacMillan, 203-205.
- Orphanides, A. (2001), "Monetary Policy Rules Based on Real-Time Data," *American Economic Review*, 91, 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*.
- Orphanides, A., and V. Wieland (2000), "Inflation Zone Targeting," *European Economic Review*, 44, 1351-1387.
- 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.

- Rothschild, M., and J. Stiglitz (1969), "Increasing Risk I: A Definition," *Journal of Economic Theory*, 2, 225-243.
- 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.
- Svensson, L.E.O (1997), "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," *European Economic Review*, 41, 1111-1146.
- Svensson, L.E.O. (2002), "Inflation Targeting: Should it be Modeled as an Instrument Rule or a Targeting Rule?" *European Economic Review*, 46, 771-780.
- Wilson, D. (2002), "Deflation – An Inflated Concern?", *Global Economics Weekly*, 02(40), Goldman-Sachs.

Table 1 – Forecast Model Diagnostics by Country

	Predictors	<u>Inflation</u>		<u>Inflation & Oil Prices</u>		<u>Inflation & Money</u>	
		1 Year	2 Year	1 Year	2 Year	1 Year	2 Year
United States	\hat{p}	3	1	1	1	1	1
	\hat{q}	-	-	1	1	-	-
	\hat{r}	-	-	-	-	1	1
	<i>SIC</i>	1.37	1.91	1.47	2.03	1.44	1.97
	<i>LB(15) : u_t</i>	0.47	0.74	7.49	0.40	6.09	0.38
	(<i>p-value</i>)	(0.99)	(0.98)	(0.19)	(1.00)	(0.30)	(1.00)
	<i>LB(15) : u_t²</i>	2.58	4.61	4.06	4.26	4.03	5.55
(<i>p-value</i>)	(0.76)	(0.47)	(0.54)	(0.51)	(0.55)	(0.35)	
Germany	\hat{p}	1	3	1	1	1	3
	\hat{q}	-	-	1	1	-	-
	\hat{r}	-	-	-	-	1	3
	<i>SIC</i>	0.58	0.94	0.66	1.26	0.61	0.35
	<i>LB(15) : u_t</i>	3.51	1.00	3.89	9.36	3.42	0.77
	(<i>p-value</i>)	(0.62)	(0.96)	(0.57)	(0.10)	(0.64)	(0.98)
	<i>LB(15) : u_t²</i>	6.78	6.58	6.51	1.41	5.77	1.68
(<i>p-value</i>)	(0.24)	(0.25)	(0.26)	(0.92)	(0.33)	(0.89)	
Japan	\hat{p}	1	1	1	1	2	1
	\hat{q}	-	-	1	1	-	-
	\hat{r}	-	-	-	-	2	1
	<i>SIC</i>	2.22	2.40	2.32	2.56	1.97	1.42
	<i>LB(15) : u_t</i>	3.14	0.42	3.21	0.46	2.36	2.18
	(<i>p-value</i>)	(0.68)	(0.99)	(0.67)	(0.99)	(0.80)	(0.82)
	<i>LB(15) : u_t²</i>	13.84	0.65	14.04	0.66	2.41	0.69
(<i>p-value</i>)	(0.02)	(0.99)	(0.02)	(0.99)	(0.79)	(0.98)	

NOTE: Estimates based on monthly data for 1960.1-2002.9. The models and data are described in the text. Boldface indicates the preferred model for each country.

Table 2 – Forecasts of the Risks to Price Stability as of September 2002

		Horizon of 1 Year	Horizon of 2 Years
Inflation risk forecast	<i>United States</i>	1.03	0.82
	<i>Germany</i>	0.07	0.00
	<i>Japan</i>	0.10	0.10
Balance of Risks forecast	<i>United States</i>	1.00	0.78
	<i>Germany</i>	-0.22	-0.18
	<i>Japan</i>	-2.09	-2.09
Deflation risk forecast	<i>United States</i>	-0.02	-0.04
	<i>Germany</i>	-0.29	-0.18
	<i>Japan</i>	-2.19	-2.19
Conditional mean forecast	<i>United States</i>	3.68	3.06
	<i>Germany</i>	1.45	1.20
	<i>Japan</i>	-0.92	-0.99


NOTE: See Table 1 for models. The risk measures are described in the text.

Table 3 – Average Historical Year-on-Year Risks by Decade and Country

	1960s	1970s	1980s	1990s	2000s
Inflation Risk					
<i>United States</i>	0.75	3.00	2.91	1.02	0.84
<i>Germany</i>	0.40	1.42	0.80	0.55	0.21
<i>Japan</i>	3.92	4.83	0.79	0.34	0.11
Deflation Risk					
<i>United States</i>	-0.13	0.00	-0.04	-0.06	-0.04
<i>Germany</i>	-0.06	-0.02	-0.15	-0.07	-0.18
<i>Japan</i>	0.00	0.00	-0.21	-1.27	-2.01
Balance of Risks					
<i>United States</i>	0.62	3.00	2.87	0.95	0.80
<i>Germany</i>	0.34	1.40	0.65	0.47	0.03
<i>Japan</i>	3.92	4.83	0.58	-0.92	-1.90


NOTE: See Table 2.

Figure 1 – Alternative Central Bank Preferences



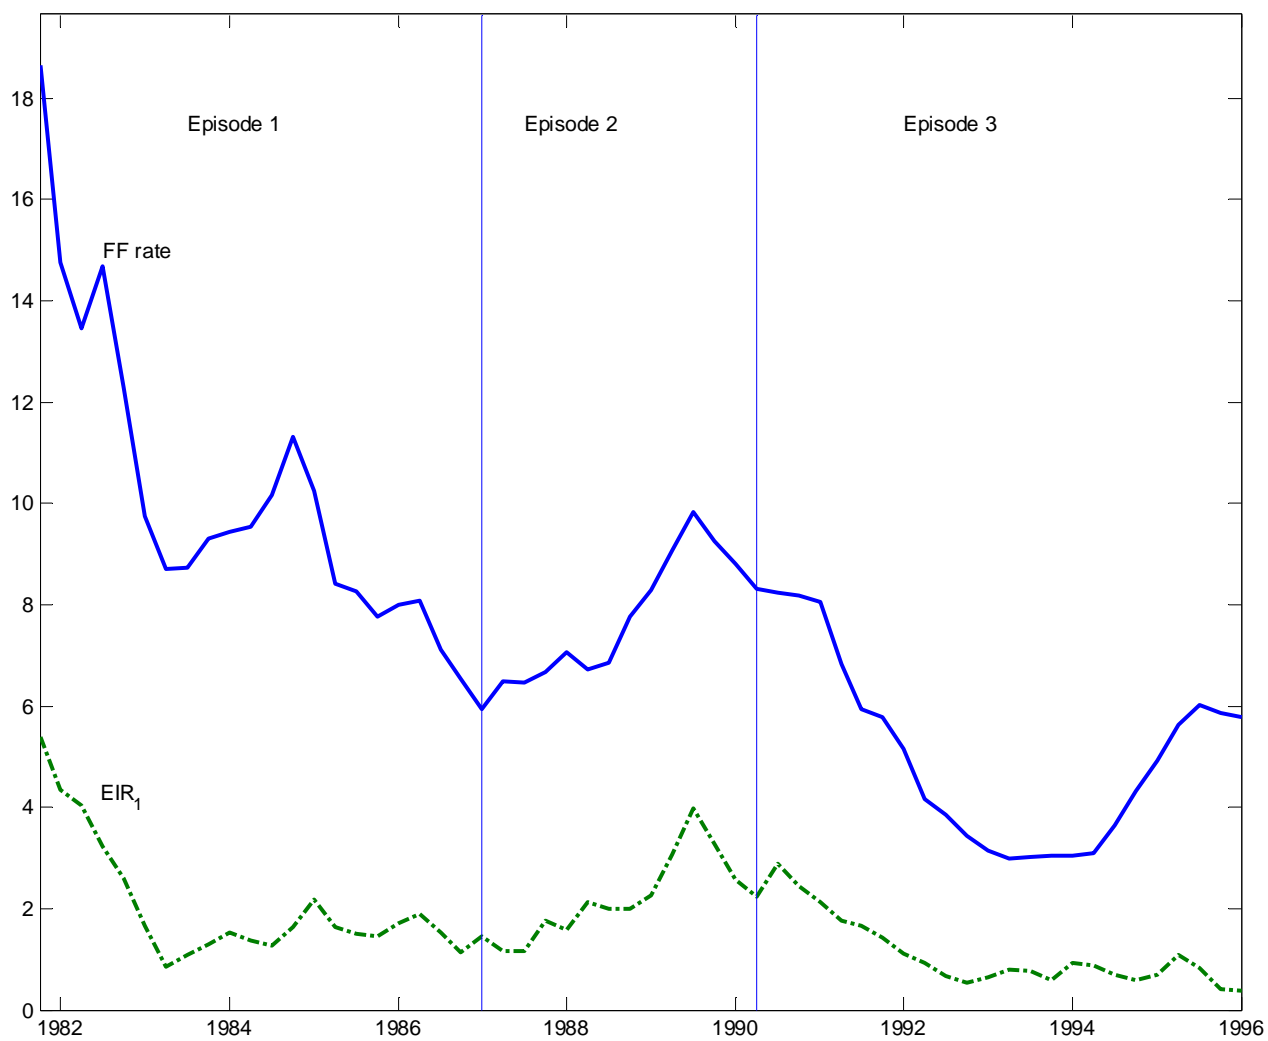
NOTE: For expository purposes we set $\underline{\pi} = 1\%$ and $\bar{\pi} = 3\%$ and impose symmetry.

Figure 2 – Quarterly Real-Time Forecasts of U.S. Year-on-Year Risks to Price Stability




NOTE: Based on real-time data in Orphanides (2003). The forecast model and risk measures are described in the text.

Figure 3 – Quarterly Real-Time Forecasts of U.S. Year-on-Year Inflation Risks and the Fed Funds Rate



NOTE: Based on real-time data in Orphanides (2003). The Fed Funds rate data are from FRED. The forecast model and risk measure are described in the text.

Figure 4 – Historical Evolution of Year-on-Year Risks by Country



NOTE: See Table 2. The different starting dates reflect different lag order choices.