QUANTIFYING THE RISK OF DEFLATION

Lutz Kilian
University of Michigan and CEPR

Simone Manganelli
European Central Bank

February 13, 2004

Abstract: We propose formal and quantitative measures of the risk that future inflation will be excessively high or low relative to the range preferred by the private sector agent. Unlike alternative measures of risk, our measures are designed to be consistent with the expected utility framework and make explicit the dependence of risk measures on the private sector agent's preferences with respect to inflation. We illustrate our methodology by estimating the risks of deflation for the United States, Germany and Japan for horizons of up to two years. The question of how large these risks are has been subject to considerable public debate. We find that, as of September 2002 when this question first arose, there was no evidence of substantial deflation risks for the United States and for Germany, contrary to some conjectures at the time. In contrast, there was evidence of substantial deflation risks in Japan.

JEL classification: C5, D81, E31, E37.

Keywords: Inflation; Deflation; Price stability; Risk; Expected utility framework; Asymmetric preferences.

Acknowledgements: The views expressed in this paper are our own and do not necessarily reflect the views of the European Central Bank (ECB) or its staff. We have benefited from comments at the 2003 CFS Conference on New Directions in Financial Risk Management, the Joint UCLA-Caltech-USC Workshop on Forecasting, the 2003 European Summer Symposium in International Macroeconomics of the CEPR and the 2003 RSQE Forecasting Conference. We also thank seminar participants at the European Central Bank, the Bank of Spain, the Deutsche Bundesbank, Barcelona Autonoma, Bocconi, Bonn, Erasmus, Humboldt, Montréal, Michigan, Brown and Virginia. We especially acknowledge helpful discussions with Robert Barsky, Vítor Gaspar, Chris House, Atsushi Inoue, Emre Ozdenorem, and Harald Uhlig.

1. INTRODUCTION

Starting in late 2002, there has been much public debate about whether the risks to price stability are tilted toward deflation in many OECD countries. For example, in October 2002 (and again in May 2003), The Economist announced that "the risk of falling prices is greater than at any time since the 1930s". In February 2003, The Economist elaborated that "in America deflation looks a bigger risk than inflation", and in July 2003, it warned that Japan has lived with deflation since the mid-1990s, but that "America faces the risk of deflation too. So, more markedly, does Germany."[1]

In response to these concerns in the financial press central bankers such as Bernanke (2002, 2003a,b) discussed measures the "central bank can take to reduce the risk of falling into deflation" and made the case that "the chance of a significant deflation in the United States in the foreseeable future is extremely small". Similarly, Issing (2002a) made the case that based on current data as well as conditional mean forecasts of inflation there are no apparent risks of deflation in the Euro area or for that matter in Germany. Greenspan (2003) further elaborated on the need to balance the risk of deflation against the risk of excessively high inflation. As late as October 2003, the FOMC issued a statement that "the risk of inflation becoming undesirably low remains the predominant concern for the foreseeable future" (see Board of Governors 2003a). Only the December 2003 FOMC statement signalled that "the probability of an unwelcome fall in inflation has diminished over recent months" (see Board of Governors 2003b).

The risk of deflation is also on the minds of consumers, investors, and businesses. The mere mention of the word "deflation" appears to be enough to stir markets. For example, in July 2003 Business Week noted that "simply by alluding to the risks of deflation, the Federal Reserve convinced investors that buying bonds was a sure thing because interest rates were headed lower". There also is concern that consumers, fearful of

---

[1] Similar concerns have been expressed in Business Week, the Wall Street Journal, and the Financial Times, among others (see Issing (2002a) for details). Also see The Economist (2002a,c) and Wilson (2002). Interest in this question continued throughout much of 2003. For example, in May 2003, the International Monetary Fund issued its own warnings that deflation in Japan may grow worse and that there was a "considerable" risk of deflation in Germany, but a "low" risk in the United States (see Kumar et al. 2003). The Financial Times in the same month warned of the increasing "risk that the world's largest and third-largest economies would follow Japan, the second-largest, into deflation". Business Week in June 2003 warned that the risk of deflation was rising in Germany and represented a threat to its neighbors in the euro zone as well.

deflation, may curtail their spending, further aggravating countries' economic stagnation (see, e.g., Krugman 2003). In this paper we propose formal measures of the risks to price stability that can be used by private sector agents to assess the risk of deflation. There is no established methodology for quantifying these risks. Even more importantly, it is unclear what the term "risk" refers to in this context and what the basis is for judging whether the risk of deflation is high or low. In fact, there has been no formal analysis of these risks for any of the major OECD countries. Existing attempts to estimate the risks of deflation have been ad hoc at best.[2]

Our formal measure of the risk of deflation is consistent with a utility-based approach to risk measurement. Our premise is that every private sector agent can express his or her preferences for inflation in the form of a utility function subject to weak restrictions. Given this parameterization of preferences, we can uniquely define measures of the risk that inflation is excessively high or excessively low relative to the preferred range of values. The proposed risk measures can be easily computed in practice. We also propose a measure of the balance of those risks. Our approach helps to reconcile seemingly contradictory interpretations of deflation risk in the financial press and provides a theoretically sound basis for further analysis. We show that the degree of risk will in general depend on the likelihood, the severity and the duration of the event of deflation. Thus, risk measures will be affected by the choice of forecast horizon and by the degree of risk aversion of the user.

The remainder of the paper is organized as follows. In section 2, we outline our utility-based approach to risk measurement. In section 3 we show how measures of risk may be computed by private sector agents in practice. In section 4, we apply our methodology to the problem of estimating the risk of deflation for the three largest OECD countries as of September 2002, when the concern about deflation first emerged. We show

---

[2] Note that the risk of deflation is fundamentally different from the likelihood of a "deflationary spiral" (see, e.g., Bernanke 2003b), which may or may not occur following a prolonged period of deflation. A deflationary spiral arises if economic slack leads to more aggressive wage and price cuts, which in turn deepens the economic slack. Such a spiral may be caused by the ineffectiveness of the short-term interest rate as a policy instrument in the presence of a "liquidity trap". If the nominal short-term interest rate cannot be reduced further, price cuts would raise the real short-term interest rate and further reduce real demand. Thus, the danger of a "deflationary spiral" is ultimately a question about the choice of policy instrument, not about the risk of deflation per se. While both questions are important, here we only focus on the risk of deflation in the proper sense.

that, as of September 2002, only for Japan is there empirical evidence of serious deflation risks and that the current deflation risks in Germany and in the United States are by no means high by post-war standards. We also illustrate the potential impact of the user's preferences with respect to inflation on the risk estimates. We conclude in section 5.

2. THE UTILITY-BASED APPROACH TO RISK MEASUREMENT

There are two basic requirements of any measure of risk. The first requirement is that the measure of risk must be related to the probability distribution of the underlying random variable (see Machina and Rothschild 1987). In the present context, the random variable of interest is the inflation rate over the horizon of interest, denoted by π. The probability distribution function of future inflation outcomes will be denoted by $F(\cdot)$. This distribution function may be estimated by the empirical distribution of inflation forecasts.

Second, any measure of risk must be linked to the preferences of the economic agent. As emphasized by the literature on the theory of forecasting, agents do not operate in a vacuum, but have well-defined preferences over the realizations of the random variable of interest (see, e.g., Granger and Newbold 1986). For the purpose of this paper – and in line with the discussion of the introduction – we take the perspective of a representative agent with a preference for price stability.[3]

Price stability in practice is usually interpreted as low inflation. We define *low inflation* as inflation in the range $[\underline{\pi}, \overline{\pi}]$, where $\underline{\pi} \leq \overline{\pi}$. We define *excessive inflation* as inflation in excess of $\overline{\pi}$ and *deflation* as inflation below $\underline{\pi}$.[4] Our analysis allows for the possibility that $\underline{\pi} = \overline{\pi} \equiv \pi^*$. Users of inflation forecasts tend to associate risk with the failure of inflation to remain inside the range $[\underline{\pi}, \overline{\pi}]$. There is an upside risk of inflation exceeding $\overline{\pi}$ as well as a downside risk that inflation falls below $\underline{\pi}$.

---

[3] See Issing (2000b) for a review of the theoretical and empirical arguments in favor of price stability.

[4] As noted by Bernanke (2003b), very low, but positive inflation and negative inflation pose qualitatively similar problems, and no sharp discontinuity exists at the point that measured inflation changes from positive to negative values. We therefore disregard this distinction for the purpose of this paper and refer to inflation below the lower threshold indiscriminately as "deflation".

The magnitude of these risks will depend on the probability distribution $F(\cdot)$ of inflation forecasts, the definition of price stability (or equivalently the choice of $\underline{\pi}$ and $\overline{\pi}$), and the economic agent's attitude towards departures from price stability. The agent's preferences can be expressed as a function of inflation. A simple, yet flexible parameterization of the preference function of the economic agent is given by:

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

where π denotes the inflation rate over the horizon of interest, $0 \le a \le 1$ represents the weight given to the objective of avoiding deflation and the parameters $\alpha \ge 0$ and $\beta \ge 0$ measure the degree of risk aversion of the economic agent. The existence of such a utility function may be viewed as a convenient short-hand expression for the costs imposed by the absence of price stability. While we do not claim that this parameterization is without loss of generality, it is likely to be a good approximation to most utility functions of interest in practice. It also provides the basis for an intuitive and economically appealing definition of risk, and it includes as a special case preference functions commonly used in macroeconomics (see, e.g., Svensson 1997, 2002).

A particular example of this preference function is shown in Figure 1. The plot in Figure 1 can be thought of as an index of the degree of satisfaction that the agent receives, as the inflation rate varies. For expository purposes we set $\underline{\pi} = 1\%$ and $\overline{\pi} = 3\%$ and impose symmetry. Figure 1 shows that the highest possible satisfaction (or utility) is achieved when inflation remains in the range consistent with price stability. We normalize this level of satisfaction to zero. As inflation exceeds the thresholds (whether it is too high or too low), the agent experiences dissatisfaction (negative utility), represented by the downward-sloping lines in the graph. The degree of dissatisfaction increases with the extent to which inflation exceeds the threshold. A key feature of the economic agent's preferences is the pair of parameters α and β that govern his attitude toward violations of price stability. The larger these parameters are, the more quickly the user becomes dissatisfied, when inflation exceeds the threshold by a given amount. We will illustrate the role and interpretation of these parameters shortly. Note that by choosing alternative values of α

and β we are able to approximate a wide range of attitudes toward inflation. We also are able to allow for asymmetry in α and β.

Given a distribution of inflation forecasts $F(\cdot)$, how can we quantify the risks that inflation may be too high or too low relative to the economic agent's preferences? Figure 2 provides a graphical presentation of the parts of the distribution of future inflation that are relevant to risk management. By construction, risks will not be related to the center of the distribution, but to the tails of the distribution, defined as the areas in which inflation exceeds the thresholds. Thus, a risk analyst would not be interested in the most likely (or modal) forecast value, nor would a risk analyst normally focus on the conditional mean forecast of inflation typically reported by forecasters. Rather he would be interested in the event that realizations of inflation fall into the tails of the distribution, defined as the areas exceeding the upper bound ($\overline{\pi}$) and falling below the lower bound ($\underline{\pi}$). The key question is how to quantify these risks in a manner that is consistent with the preference function of the economic agent. Once we understand the nature of this risk, we may estimate it, and we may characterize the trade-off between upside and downside risks to inflation.

A Proposal for Measures of Inflation Risk

In light of the above discussion, we propose to measure the risks of departing from price stability by a probability weighted function of the deviations of inflation from the threshold points $\underline{\pi}$ and $\overline{\pi}$:

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

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

We adopt the convention of defining the risk of deflation as a negative number and the risk of excessive inflation as a positive number. This convention only serves to facilitate the graphical representation of risks. It does not affect the analysis in any way. Each risk is a scalar defined in terms of a two-parameter function involving a fixed threshold, denoted by

---

[6] A value of $\alpha = 1$ implies risk neutrality in the relevant range. Risk-seeking behavior in that range 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.

$\underline{\pi}$ and $\overline{\pi}$, and a fixed parameter governing risk aversion, denoted by α and β. The parameters α and β may be different.[6] It is possible to compute the proposed measures of risk for any choice of α and β. Measures of risk of this type were first proposed by Fishburn (1977) in the context of portfolio allocation in the presence of downside risk. Similar integral-based measures of risk have also been used by Holthausen (1981) and Basak and Shapiro (2001), for example. As we will show below, many measures of risk that have been proposed in other contexts can be derived as special cases of Definition 1.

The idea that risk is measured by probability-weighted dispersions beyond the thresholds that define price stability is appealing in that it recognizes the economic agent's desire to avoid inflation realizations outside the preferred range. To the extent that this contention is correct, it casts serious doubt on the appropriateness of alternative 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 or the skewness of the distribution $F(.)$. 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.

More specifically, the measures of risk proposed in Definition 1 are attractive for several reasons:

- First, as shown below in Propositions 1 and 2, they are consistent with the expected utility framework. This is important for otherwise one could not make the case that the economic agent with preferences as given in (1) ought to use these specific risk statistics.
- Second, the proposed measures of risk are fully operational and can be computed easily in practice. They can be used in conjunction with any econometric forecasting model as well as in conjunction with judgmental inflation forecasts, as long as these forecasts can be expressed in the form of a distribution $F(\cdot)$.
- Third, the proposed risk measures have an intuitive economic interpretation in that they are proportionate to the expected disutility associated with deflation and with excessive inflation. This allows the construction of risk indices.

- Fourth, our analysis suggests a natural way of resolving the trade-off between the risk of deflation and the risk of excessively high inflation. We propose a specific weighted average of the upside and downside risks to price stability that may be viewed as a measure of the *balance of risks*. This measure also has an interesting economic interpretation. It turns out that – if risks are defined as in Definition 1 – under plausible assumptions a balance of risks of zero is equivalent to utility maximization of the economic agent.

- Fifth, the proposed measures of risk provide a unifying framework for several measures of risk proposed in the literature in other contexts. In special cases, they can be shown to reduce to the probability of exceeding a threshold, the weighted expected shortfall, the target semi-variance and the variance about target.

- Sixth, the measures of risk in Definition 1 explicitly incorporate preference parameters such as the degree of risk aversion.

- Seventh, although we focus on the case of a known preference function for inflation, we note that the proposed risk measures have the ability to deliver – under suitable conditions – preference rankings that are consistent with a wide range of unknown von Neumann-Morgenstern utility functions. Specifically, whenever one distribution $F(\cdot)$ first-order stochastically dominates an alternative distribution, our risk measures will produce the same ranking. The same result holds for second-order stochastic dominance, provided that $\alpha > 1$ and $\beta > 1$ (see Kilian and Manganelli (2003) for further discussion).

On the Interpretation of the Proposed Risk Measures

Proposition 1 establishes a close link between the risk measures of Definition 1 and the expected utility of the economic agent:

Proposition 1:

Let $DR_\alpha(\underline{\pi}, F)$ denote the deflation risk computed under the distribution F and let $EIR_\beta(\overline{\pi}, F)$ denote the corresponding risk of excessive inflation, as defined in Definition 1. Then

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

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

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

Proof: Immediate.

Proposition 1 implies that DR_{α} and EIR_{β} are proportionate to the expected disutility associated with deflation and with excessive inflation. Thus, they have a direct economic interpretation. Specifically, they measure the degree to which the agent dislikes economic scenarios where inflation exceeds the upper threshold or falls short of the lower threshold. In this sense, DR_{α} and EIR_{β} are a measure of 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. Next we will discuss some leading examples.

Interpretation of Risk in Some Interesting Special Cases

As it turns out, for appropriate choices of the risk aversion parameters α and β, our general measure of risk reduces to measures of risk proposed in the literature in other contexts. First, consider the limiting case of $\alpha = \beta = 0$. In this case, we obtain:

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

Thus, the general measures of risk simply reduce to the probabilities of exceeding the target range at either end.[7] This result is intuitive because for $\alpha = \beta = 0$ the agent is only concerned about not missing the price stability range, but, conditional on violating price stability, does not care at all by how much inflation will exceed the threshold.

---

[7] This interpretation of risk has been used in the press. For example, the Financial Times (2003b) implies in May 2003 that the risk of deflation in the U.K. is low because the Bank of England's "forecasts imply only a one-in-20 chance that inflation will be below 1.3 percent in two years' time, let alone below zero". Similarly, Kumar et al.'s (2003) IMF study refers to "a non-zero, but still low, probability of a … decline in prices" as evidence against deflation.

Although instructive, this limiting case is implausible in that economic agents in practice would not be indifferent to whether inflation misses the target zone by a small amount or by a large amount. There are at least three ways of making this point. One is by direct observation. For example, Bernanke (2003b) notes that "very low inflation and deflation pose qualitatively similar economic problems, though the magnitude of the associated costs can be expected to increase sharply as deflationary pressures intensify". This statement clearly implies that $\alpha > 1$ in his preference function.

Second, a simple counterexample illustrates that few economic agents would be indifferent to whether inflation misses the thresholds by a small amount or by a large amount. Suppose that an agent faces the choice between two situations: (a) 2.001% inflation with probability 100%; and (b) 10% inflation with probability 20% and inflation below 2% with probability 80%. If we go by the probabilities of missing the threshold, situation (a) is clearly worse. In practice, most agents would prefer (a) over (b), however, suggesting that their preferences are inconsistent with $\alpha = \beta = 0$.

A third argument against this specification is the language used by many economic agents in describing risks. For example, Greenspan (2003) discussed risks in terms of "the product of a low-probability event and a severe outcome, should it occur". In doing so, he ruled out $\alpha = \beta = 0$, because in that case it would have been sufficient to express risks in terms of probabilities alone. Although our focus in this paper is not on central bankers, their statements provide examples that are of more general relevance to how economic agents interpret risk. Greenspan's speech is but one example. The same careful distinction between the likelihood of the event of deflation and its magnitude is implicit in statements such as "there is a considerable risk of mild deflation" (Kumar et al. 2003) or "the minor probability of an unwelcome substantial fall in inflation" (Board of Governors 2003c). Again such statements rule out $\alpha = \beta = 0$. This is not to say that even the same institutions always use language consistent with the same definition of risk, but it highlights the importance of being precise about what notion of risk we have in mind.

We conclude that characterizations of risk merely in terms of the probability of missing a threshold are misleading. Note that Greenspan's language is actually much closer to what our general risk measure would imply for $\alpha = \beta = 1$. In that case, we obtain:

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

By construction DR_1 is a measure of expected deflation, and EIR_1 a measure of expected excess inflation. A different way of writing these measures is to interpret them as the product of a conditional expectation and a tail probability. For example, we may write:

$$DR_1 = E(\pi - \underline{\pi} \mid \pi < \underline{\pi}) \Pr(\pi < \underline{\pi}).$$

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.[8] This language closely mimics that used by Greenspan. In practice, the interpretation of this risk measure is best illustrated by an example. Let the upper threshold of inflation be 2%. Suppose that the inflation rate can be either 4% with probability 1/2 or 0% with probability 1/2. Then the *expected excess inflation* would be (4%-2%) 1/2 = 1%.[9]

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

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

a concept familiar from finance. Here the agent is best off when he or she minimizes in expectation the squared deviations of inflation from the lower threshold and the squared deviations of inflation from the upper threshold. An interesting special case of this measure is obtained under quadratic symmetric utility for inflation when $\underline{\pi} = \overline{\pi} \equiv \pi^*$:

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

In that case, expected utility is maximized when the variance about π^* is minimized, as shown for example by Svensson (1997).

---

[8] 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*.

The Balance of Risks

It is common in discussions of inflation risks to stress the need to balance the upside risks and the downside risks to price stability. This balance can be thought of as a weighted average of the risks in Definition 1. One way of providing economic content to the notion of a balance of risks is the following thought experiment: Suppose that the inflation process can be decomposed into a conditional mean component (μ_{t+1}) and a conditional innovation term (u_{t+1}), whose distribution is assumed to be location-invariant:

(2) $$\pi_{t+1} = \mu_{t+1} + u_{t+1} \qquad u_{t+1} \sim F^u_{t+1}(\cdot)$$

Further suppose that, for given $F^u_{t+1}(\cdot)$, the agent could choose the level of expected inflation μ_{t+1} that maximizes his expected utility. Then under the assumptions of Proposition 2 below the utility-maximizing level of expected inflation, μ^*_{t+1}, will correspond to a balance of risks equal to zero, if the balance of risks is defined as a suitable weighted average of the risk of excessive inflation and the risk of deflation with weights depending on the preference parameters of the economic agent.

Proposition 2:

Given the set of preferences in (1) and the inflation process (2), if α and β are greater than 1, the economic agent's expected utility is maximized for the level of expected inflation μ^*_{t+1} that solves the following equation:

$$BR_{\alpha-1,\beta-1}(\mu^*_{t+1}) \equiv \omega DR_{\alpha-1}(\mu^*_{t+1}) + \nu EIR_{\beta-1}(\mu^*_{t+1}) = 0$$

where $\omega = a\alpha$, $\nu = (1-a)\beta$ and $BR_{\alpha-1,\beta-1}$ denotes the balance of risks.

Proof: The result follows immediately from the first order condition of the expected utility maximization problem.

---

[9] It may be tempting to conclude that the expected inflation conditional on missing the target would be 3% (=2%+1%) therefore. That conclusion would be wrong, because the expected inflation rate conditional on exceeding the 2% target is 4%.

Proposition 2 shows that – under plausible assumptions – there is a close link between the balance of risks and the optimal level of inflation from the point of view of the economic agent. The coefficients ω and ν are scale invariant, which mirrors the fact that utility may be rescaled. The balance of risks may be computed easily in practice. In the special case of $\alpha = \beta = 2$ and $a = 0.5$, for example, we obtain $BR_{1,1} = DR_1 + EIR_1$.

Note that the balance of risks is not an indicator of the overall extent of risk (which would be measured by the expected utility expression in Proposition 1), but rather an indicator of the optimality of the distribution of risks. The balance of risks measure provides a convenient tool for judging departures from price stability. Specifically, if the balance of risks is negative, we say that the balance of risks is tilted toward deflation; if it is positive, the risks are tilted toward excessive inflation. Note that a balance of risk of zero does not mean that the economic agent would not prefer a reduction in both upside and downside risk, if given the choice. It means that the economic agent would not want the central banker to reduce one risk, say that of deflation, if doing so means increasing the other risk, namely that of excessive inflation.

The balance of risks differs in general from measures of the central tendency of inflation forecasts such as the conditional mean forecast. We will illustrate this point in section 4. The only situation, in which there is a formal link between the conditional mean forecast and the balance of risks is the case of economic agents with quadratic symmetric preferences and a preference for inflation rates of exactly π^*. For example, an economic agent may view all departures from an inflation rate of 2 percent as undesirable. In that case, certainty equivalence holds and the balance of risks equals the deviation of the conditional mean forecast of inflation from π^*:

$$BR_{1,1} = DR_{1,1}(\pi^*) + EIR_{1,1}(\pi^*) = -\int_{-\infty}^{\pi^*} (\pi^* - \pi) dF(\pi) + \int_{\pi^*}^{\infty} (\pi - \pi^*) dF(\pi) =$$
$$\int_{-\infty}^{\infty} \pi dF(\pi) - \pi^* = E(\pi) - \pi^*$$

Of course, in this case, the balance of risks is not a measure of how much risks are tilted toward deflation, but rather of how much risks are tilted toward inflation rates below π^*.

Implementation Issues

The implementation of our risk measures requires the user to specify the parameters a, α and β, $\underline{\pi}$ and $\overline{\pi}$, and the forecast horizon of interest. In practice, these design parameters will be chosen by the economic agent.

The choice of the thresholds $\underline{\pi}$ and $\overline{\pi}$ by the economic agent is straightforward. A reasonable starting point is to presume that the economic agent's definition of price stability is the same as that used by many central bankers. For example, Bernanke (2003b) suggests that core inflation between 1 and 2 percent per year is probably the de facto equivalent of price stability, taking into account factors such as measurement bias in inflation indices. An essentially identical interpretation has been adopted by the European Central Bank (see Issing 2002b, Duisenberg 2003). Other central banks such as the Bank of England, the Bank of New Zealand, the Swedish Riksbank, and the Bank of Canada prefer a somewhat wider range of 1 to 3 percent. Throughout this paper we will adopt the latter convention for expository purposes.

The values of α and β can in principle be elicited from the economic agent 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 $\alpha = \beta = 2$ as our benchmark case (see e.g., Blinder 1997; Svensson 1997, 2002). We also will consider larger values for α and β, and we will allow for asymmetry in α and β. The parameter a will be set to 0.5 for expository purposes. This choice is consistent with the macroeconomic literature and there is no compelling reason to depart from it.

For the purpose of this paper we take as given that inflation is defined as the percent change of the consumer price index over the relevant forecast period. The choice of forecast horizon thus affects the definition of π. In practice, most users of inflation forecasts are likely to be interested in medium-term rather than short-term horizons. This is especially true in the context of assessing the risk of deflation. There is little interest in risk assessments for periods shorter than a year. For expository purposes we will focus on horizons of one year and of two years. Our choice of horizon is consistent with the practice of central bankers and many private sector forecasters. For example, the Fed's forecasts typically cover a horizon of six to eight quarters (see Reifschneider, Stockton and Wilcox 1997). Similarly, the Swedish Riksbank uses a horizon one or two years ahead (see Heikensten 1999). The Bank of England publishes forecasts for a horizon of 10 quarters. In many cases, one may be interested in even longer horizons. We do not pursue this possibility in this paper because it is widely recognized that it is difficult to generate credible forecasts for horizons exceeding two years. This is not a problem of the risk methodology, but of the available data and of the forecast model.

3. Computing Measures of Inflation risk in practice

The measures of risk proposed in section 2 are fully operational and can be computed easily in practice, as long as these forecasts can be expressed in the form of a distribution $F(\cdot)$. Here we will focus on their implementation in the context of a simple econometric forecasting model. We are interested in forecasting inflation at horizons of one year and two years. Given that forecast models must generally be considered misspecified, it is natural to fit the conditional mean model directly at the horizon of interest. This approach also dispenses with the need for a multivariate model, and it avoids the well-known problems of time aggregation of volatility dynamics.

Let π_τ, $\tau = 1,2,...,\mathrm{T}$, denote nonoverlapping observations for the inflation rate measured at the horizon of interest. Then:

$$\text{(2)} \qquad \begin{aligned} &\pi_{\tau+1} = \mu_{\tau+1} + u_{\tau+1}, \quad u_{\tau+1} = \varepsilon_{\tau+1}\sqrt{h_{\tau+1}}, \qquad \varepsilon_{\tau+1} \mid \Omega_\tau \sim i.i.d.(0,1) \\ &h_{\tau+1} = au_\tau^2 + bh_\tau + c \end{aligned}$$

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 $h_{\tau+1}$, 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 include other variables that may affect future inflation, such as the percent change of oil prices or monetary aggregates. In that case, the model for $\mu_{\tau+1}$ would take the form:

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

where the additional predictors, $x_{k,\tau}$, $k = 1,...,s,$ 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 (2), we proceed as follows. In general, all proposed risk measures may be computed by bootstrap simulation. Consider $EIR_{\beta,\tau+1}(\overline{\pi}) = \int_{\overline{\pi}}^{\infty} (\pi_{\tau+1} - \overline{\pi})^{\beta} dF(\pi_{\tau+1})$, for example. Denote the estimated standardized residuals of model (2) 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, conditional on the parameter estimates of model (2), based on $r = 1,...,R$ bootstrap replications of $\{\pi_{\tau+1}^{r}\}$.[10] Then we make use of the fact that

$$EIR_{\beta,\tau+1}(\overline{\pi}) = \Pr(\pi_{\tau+1}^{r} > \overline{\pi}) E[(\pi_{\tau+1}^{r} - \overline{\pi})^{\beta} \mid \pi_{\tau+1}^{r} > \overline{\pi}] ,$$

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

where the probability $\Pr(\pi^r_{\tau+1} > \bar{\pi})$ can be computed as the fraction of observations $(\pi^r_{\tau+1} - \hat{\mu}_{\tau+1}) / \sqrt{\hat{h}_{\tau+1}}$ that exceed $(\bar{\pi} - \hat{\mu}_{\tau+1}) / \sqrt{\hat{h}_{\tau+1}}$, and where the expectation $E[(\pi^r_{\tau+1} - \bar{\pi})^\beta \mid \pi^r_{\tau+1} > \bar{\pi}]$ can be estimated by the sample moment $\frac{1}{\#(\pi^r_{\tau+1} > \bar{\pi})} \sum_{\pi^r_{\tau+1} > \bar{\pi}} (\pi^r_{\tau+1} - \bar{\pi})^\beta$ (or alternatively by a smoothed estimator).

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 $EIR_{1,\tau+1}(\bar{\pi}) =$ $\Pr(\pi_{\tau+1} > \bar{\pi}) E[(\pi_{\tau+1} - \bar{\pi}) \mid \pi_{\tau+1} > \bar{\pi}]$ note that

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

where the last expectation can be computed by taking averages of the standardized residuals that satisfy the inequality condition.[11] The probability $\Pr(\pi_{\tau+1} > \bar{\pi})$ can be computed, as discussed before. Similarly, for $EIR_{2,\tau+1}(\bar{\pi}) =$ $\Pr(\pi_{\tau+1} > \bar{\pi}) E[(\pi_{\tau+1} - \bar{\pi})^2 \mid \pi_{\tau+1} > \bar{\pi}]$ we have

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

$$2(\mu_{\tau+1} - \bar{\pi}) \sqrt{h_{\tau+1}} E[\varepsilon_{\tau+1} \mid \varepsilon_{\tau+1} > (\bar{\pi} - \mu_{\tau+1}) / \sqrt{h_{\tau+1}}],$$

which may be computed analogously.

4. EMPIRICAL ANALYSIS

As noted in the introduction, there is considerable debate about the magnitude of the risk of deflation for the three largest OECD economies. Whereas some observers have downplayed the risks of deflation for the U.S. and German economy, others have warned of serious risks. For example, Issing (2002a) made the case that there are no apparent risks of deflation in the Euro area or for that matter in Germany. In contrast, the IMF concluded that the risk for Germany is high (see Kumar et al. 2003). This is also the majority view in the financial press (see, e.g., The Economist 2003c , the Financial Times 2003a, Business Week 2003a). There has been no formal analysis of these risks, however, for any of the major OECD countries, and there is reason to question the relevance of the empirical

[11] A similar technique has been used by Manganelli and Engle (2001) to compute the expected shortfall of the distribution of portfolio returns.

evidence presented in this regard. As we have shown, neither the actual inflation rate today nor the conditional mean forecast of inflation is an appropriate measure of risk. Rather risks are related to the tails of the distribution of inflation forecasts. Our methodology is designed to provide answers to precisely these types of questions.

We will use data for Germany, Japan and the United States to estimate these risks for horizons of one year and of two years, as they existed in September of 2002, when the issue of deflation risks was first raised. We will assess these risks from the point of view of a private sector company or investor. 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 2002a, Business Week 2003a). Second, there are no data for Euro area wide inflation that extend far enough back in time to allow the construction of risk forecasts for the horizons of interest here.[12]

Our starting point is a forecast model of inflation. Risk forecasts of course are only as good as the underlying forecast model. In practice, we use statistical model selection criteria to help us select the best possible forecast model among a number of time series forecast models involving lags of inflation and lags of other variables such as money growth rates and percent changes in oil prices.

Forecast Model Specification

Our data source for the U.S. CPI inflation 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.[13] Our estimation sample starts in January 1960 and ends in September 2002. We consider three alternative

---

[12] For the Euro area the short time span of inflation data since 1999 makes it difficult 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.

[13] 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.

specifications of the conditional mean in model (2): 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):

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

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

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

The use of oil prices and of monetary aggregates as potential additional predictors is a natural choice.[14] Although theory does not restrict the set of conditioning information, we did not investigate the use of additional predictors, given the large number of parameters involved.

All models are estimated by least-squares. Table 1 presents the model diagnostics by country. We first use the Schwarz Information Criterion (SIC) to select the optimal number of lags for each model. For simplicity we restrict the number of lags to be the same across predictors. The upper bound is three lags. 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).[15] 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.[16]

---

[14] 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.

[15] See Inoue and Kilian (2003) for a discussion of alternative criteria for forecast model selection.

[16] 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

A Snapshot of the Risks to Price Stability

The risk forecasts are shown in Tables 2, 3 and 4. 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. As discussed earlier, we cannot assess risks without reference to a preference function. Here we explore three alternative preference settings for illustrative purposes. The baseline is a set of quadratic symmetric preferences, as postulated by many macroeconomists. Later we will consider asymmetric preferences incorporating either a strong aversion to deflation or a strong aversion to excessive inflation. Throughout, we presume for expository purposes that the decision-maker is concerned about inflation exceeding $\underline{\pi} = 1\%$ and $\overline{\pi} = 3\%$, and we set $a = 0.5$.

Table 2 shows the projected risks of deflation (DR_1), the risks of excessive inflation (EIR_1) and their balance ($BR_{1,1}$) by country and horizon. This balance measure would be appropriate for $\alpha = \beta = 2$. We also include the conditional mean forecast for comparison. Recall that DR_1 and EIR_1 are statistical measures of the expected excess inflation (in percent) relative to the thresholds of 1% and 3%. The balance of risks is simply a weighted average of these risks defined as $BR_{1,1} = DR_1 + EIR_1$.

The last panel of Table 2 shows 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. It may be tempting to interpret the mere existence of deflation risks as a reason for concern, but this interpretation ignores the simultaneous presence of risks of excessive inflation (see first panel of Table 2). One-sided attention to one of these risks at the expense of the other clearly will give a misleading impression of the overall risks to price stability. It therefore makes sense to put these numbers into perspective by focusing on the

---

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.

degree to which the balance of risks is tilted in favour of the inflationary or the deflationary region. The second panel of Table 2 shows the balance of these risks ($BR_{1,1}$). 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.

How robust are these results? Suppose for expository purposes the existence of a strong aversion to deflation in the form of preference parameters of $\alpha = 3$ and $\beta = 2$ ("deflation paranoia"). Table 3 shows that in that case there is some evidence at the one-year horizon that the balance is tilted in favour of deflation not just for Japan, but also for Germany. In contrast, the United States on balance remain in the inflationary region even for this strong form of deflation aversion. For the two-year horizon there is some evidence of the balance being shifted toward deflation for both Germany and the United States, but the magnitudes of the imbalances are negligible for Germany and small for the United States (not only compared to Japan, but also compared to its own past). Thus, one would be hard pressed to make a case for serious risks of deflation in Germany or in the United States at the horizons that matter to policy makers, whereas for Japan there is overwhelming evidence of deflation risks.

Conversely, we may take the point of view of an observer with strong aversion to excessive inflation in the form of preference parameters of $\alpha = 2$ and $\beta = 3$ ("inflation paranoia"). It is only with such extreme inflation aversion that we might hope to explain away the earlier evidence of disproportionate deflation risks for Japan. As Table 4 shows, in that case Japan indeed is on balance well inside the inflationary region, but interestingly Germany is not, although only by a small margin. Intuitively, this happens because German inflation rate predictions are so close to the target range that raising them to higher powers lowers the disutility, rather than compounding it (see Figure 1). This result is a consequence of the assumption of risk aversion. For the United States, as expected, the results are qualitatively the same as in Table 2, except that the risks are much more tilted toward excessive inflation. Thus, on balance none of three countries would appear to suffer from important deflation risks under this form of "inflation paranoia".

This example highlights the importance of preference parameters in assessing the risks to price stability. It shows that unless one appeals to extreme forms of "deflation paranoia" the risks of deflation appear to be rather negligible for the United States and Germany, especially when we extend the horizon to two years.

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 2 into historical perspective. In taking this historical point of view, we implicitly read the historical evidence through the lens of today's preferences. For expository purposes we impose the baseline symmetric quadratic preferences in this historical analysis. Figure 3 shows the evolution of the risks of deflation and of excessive inflation at the one-year horizon since the 1960s. Table 5 presents the same evidence as historical averages by decade together with the average balance of risks.

As discussed earlier, these risks may be interpreted as statistical measures of the expected excess inflation relative to the upper threshold of 3% and of the expected excess deflation relative to the lower threshold of 1%. Alternatively, the same risk measures also have an interpretation as indices of the expected disutility associated with departures from the lower and the upper threshold. Thus, the time series plots allow us to address, at least in part, the recent claim in The Economist (2002b, 2003b) that the risk of falling prices in the United States and in Germany is greater than at any time since the 1930s.

Figure 3 and Table 5 show that deflation risks for the United States today are not unlike those in the 1960s. Then as now the balance of risks is well inside the inflationary region on average. For Germany, deflation risks were much higher in the late 1980s than they are today, but interestingly these risks went largely unnoticed. Although the balance of risks for Germany recently has been closer to zero than in any previous decade, on average it remains slightly above zero. Only for Japan is there evidence that deflation risks have reached unprecedented levels by post-war historical standards. Since the 1990s the balance has been tilted toward deflation. Moreover, the risks of deflation in Japan are highly persistent.

Relationship Between the Balance of Risks and the Conditional Mean Forecast

It is common to judge the overall risks to price stability by the deviation of conditional mean forecasts or other baseline forecasts of inflation from the threshold of 1% below which deflation is said to occur (see, e.g., Issing 2002a, Kumar et al. 2003). It is immaterial from our point of view whether such forecasts are generated from an expectations-augmented Phillips curve or some other forecast model. Such approaches are inconsistent with the utility-based approach to risk measurement as well as inconsistent with stochastic dominance criteria. In addition, measures of the central tendency of forecasts may be poor indicators of risks. We will illustrate this point below using Japanese data.

Figure 4 shows estimates of a conditional mean measure (CMM) of the risks to price stability where the definition of

$$CMM = \max(\hat{\mu}_{\tau+1} - 3, 0) + \min(\hat{\mu}_{\tau+1} - 1, 0)$$

is motivated by the use of the conditional mean as an indicator of risk by practitioners. Figure 4 plots actual inflation, the conditional mean measure, and the balance of risks under quadratic symmetric preferences. As before, the balance of risks is computed under the assumption that $\underline{\pi} = 1\%$, $\overline{\pi} = 3\%$ and $a = 0.5$ based on the model selected in Table 1. Figure 4 illustrates that there is no simple relationship between the conditional mean measure of risks and the balance of risks even under quadratic symmetric preferences. Nevertheless, the qualitative implications are broadly similar in this specific example. This conclusion changes if we allow for asymmetry in the user's preferences.

Figure 5 contrasts the conditional mean measure with measures of the balance of risk that would be obtained if the economic agent had a strong aversion to deflation ($\alpha = 3$ and $\beta = 2$) or to excessive inflation ($\alpha = 2$ and $\beta = 3$). Figure 5 suggests two main conclusions. First, the deviation of the conditional mean from the threshold will typically misrepresent the degree to which risks are tilted, when the agent has asymmetric preferences. It may even obscure the existence of deflation risks in the data. For example, during 1984-1989 the $BR_{2,1}$ measure is mostly negative, suggesting the presence of deflation risks, whereas *CMM* is mostly zero, suggesting the absence of deflation risks. Second, properly computed measures of the balance of risks may convey a very different picture of the change in the balance of risks over time. For example, $BR_{2,1}$ shows a much more dramatic drop in the balance of risks in the 1990s than *CMM*, suggesting a much more severe threat to price stability. This example illustrates that preferences and (in particular the degree of risk aversion) matter in assessing the risk of deflation. Even in the simplest cases, the risks to price stability cannot be quantified without explicit reference to the preference parameters of the economic agent.

5. CONCLUSION

The ability to measure and forecast the risks to price stability is important to consumers, investors and businesses alike. We proposed formal and quantitative measures of the risk that future inflation will be excessively high or low relative to the economic agent's preferences. Unlike alternative measures of deflation risk, our measures are designed to be consistent with the expected utility framework and make explicit the dependence of risk measures on the private sector's preferences for inflation. We showed that our measures of risk have intuitive interpretations, and we discussed their relationship to existing measures of risk. We also provided economic content for the widely used notion of a balance of upside and downside risks to price stability.

We illustrated our methodology by estimating the risks of deflation for the United States, Germany and Japan for horizons of up to two years. We found that, as of September 2002, there was no evidence of substantial deflation risks for the United States and for Germany, contrary to a common perception at the time. In contrast, there was evidence of substantial deflation risks in Japan. The latter conclusion is robust to allowing for some asymmetry in inflation preferences, but depended on the absence of extreme forms of inflation aversion in the preference function of the economic agent.

We also put the estimates of deflation risk into historical perspective. Our analysis refuted the common claim that the risk of deflation in the United States and Germany is greater today than at any time since the 1930s (see, e.g., The Economist 2002b, 2003b). We found that only for Japan there is evidence of deflation risks that are unusually high by historical standards. Finally, we contrasted the use of conditional mean forecasts and of our measure of the balance of risks as indicators of the risks to price stability under alternative preference settings.

The methodology proposed in this paper is quite general and can be adapted to other private sector forecasting problems. An interesting question for future research is how the risk methodology employed in this paper could be adapted for the purposes of policy decisions at central banks. This problem requires nontrivial modifications of the approach described here and is being addressed in ongoing research.

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. (2003a), "Constrained Discretion and Monetary Policy", Remarks by Governor Ben S. Bernanke, Federal Reserve System, before the Money Market Marketeers o New York University, New York, February 3, 2003.

Bernanke, B.S. (2003b), "An Unwelcome Fall in Inflation", Remarks by Governor Ben S. Bernanke, Federal Reserve System, before the Economics Roundtable, University of California, San Diego, La Jolla, CA, July 23, 2003.

Board of Governors (2003a), FOMC Statement, Press Release of October 28. http://www.federalreserve.gov/boarddocs/press/monetary/2003.

Board of Governors (2003b), FOMC Statement, Press Release of December 9. http://www.federalreserve.gov/boarddocs/press/monetary/2003.

Board of Governors (2003c), FOMC Statement, Press Release of May 6. http://www.federalreserve.gov/boarddocs/press/monetary/2003.

Blinder, A. (1997), "What Central Bankers Could Learn from Academics – and vice versa," *Journal of Economic Perspectives*, 11, 3-19.

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.

Business Week (2003a), "Deflation's Deepening Shadows", June 2.

Business Week (2003b), "Why Greenspan Had to Cut Rates", July 7.

Cao, R., M. Febrero-Bande, W. González-Manteiga, J.M. Prada-Sánchez 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.

The Economist (2003a), "Downhill Dollar" February 1.

The Economist (2003b), "The Joy of Inflation," May 17.

The Economist (2003c), "Letting Off Steam," July 12.

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.

Financial Times (2003a), "Global Slowdown Raises Spectre of Germany and U.S. Following Japan into Deflation", May 8.

Financial Times (2003b), "Greenspan Plays Down the Threat of Deflation" May 22.

Granger, C.W. and P. Newbold (1986), *Forecasting Economic Time Series.* 2nd ed., Academic Press, San Diego.

Greenspan, A. (2003), "Monetary Policy and Uncertainty," Remarks by Chairman Alan Greenspan at a Symposium Sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, WY, August 29, 2003.

Heikensten, L. (1999), "The Riksbank's Inflation Target – Clarification and Evaluation," *Sveriges Riksbank Quarterly Review*, 1, 5-17.

Holthausen, D.M. (1981), "A Risk-Return Model with Risk and Return Measured in 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. (2000), Why Price Stability?, in *Why Price Stability?* (ed. by Herrero, Gaspar, Hoogduin, Morgan, Winkler), *Conference Proceedings of the First ECB Central Banking Conference*.

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

Issing, O. (2002b), "Monetary Policy and the Role of the Price Stability Definition" Panel Panel speech at the ECB Watchers Conference in Milan on June 10, 2002.

Kilian, L., and S. Manganelli (2003), "The Central Banker as a Risk Manager: Quantifying and Forecasting Inflation Risks", revised version of ECB Working Paper Nr. 226, April 2003, European Central Bank, Frankfurt/M, Germany.

Kumar, M.S., T. Baig, J. Decressin, C. Faulkner,-MacDonagh, and T. Feyzioğlu (2003), "Deflation Determinants, Risks, and Policy Options", IMF Occasional Paper No 221, June 2003, International Monetary Fund, Washington, DC.

Krugman, P. (2003), "Fear of a Quagmire?" *New York Times*, May 24.

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.

Manganelli, S. and R. Engle (2001), Value at Risk Models in Finance, ECB Working Paper No. 75.

Reifschneider, D.L., D.J. Stockton and D.W. Wilcox (1997), "Econometric Models and the Monetary Policy Process", *Carnegie-Rochester Conference Series on Public Policy*, 47, 1-37.

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	*Inflation*		*Inflation & Oil Prices*		*Inflation & Money*	
	Horizon	*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
	SIC	**1.37**	**1.91**	1.47	2.03	1.44	1.97
	$LB(15): u_t$	0.47	0.74	7.49	0.40	6.09	0.38
	(p-value)	(0.99)	(0.98)	(0.19)	(1.00)	(0.30)	(1.00)
	$LB(15): u_t^2$	2.58	4.61	4.06	4.26	4.03	5.55
	(p-value)	(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
	SIC	**0.58**	0.94	0.66	1.26	0.61	**0.35**
	$LB(15): u_t$	3.51	1.00	3.89	9.36	3.42	0.77
	(p-value)	(0.62)	(0.96)	(0.57)	(0.10)	(0.64)	(0.98)
	$LB(15): u_t^2$	6.78	6.58	6.51	1.41	5.77	1.68
	(p-value)	(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
	SIC	2.22	2.40	2.32	2.56	**1.97**	**1.42**
	$LB(15): u_t$	3.14	0.42	3.21	0.46	2.36	2.18
	(p-value)	(0.68)	(0.99)	(0.67)	(0.99)	(0.80)	(0.82)
	$LB(15): u_t^2$	13.84	0.65	14.04	0.66	2.41	0.69
	(p-value)	(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 – Risks to Price Stability as of September 2002

$\alpha = 2$ and $\beta = 2$ (Symmetric Preferences)

		Horizon of 1 Year	Horizon of 2 Years
Inflation risk	*United States*	1.03	0.82
forecast	*Germany*	0.07	0.00
EIR_1	*Japan*	0.10	0.10
Balance of Risks	*United States*	1.00	0.78
forecast	*Germany*	-0.22	-0.18
$BR_{1,1}$	*Japan*	-2.09	-2.09
Deflation risk	*United States*	-0.02	-0.04
forecast	*Germany*	-0.29	-0.18
DR_1	*Japan*	-2.19	-2.19
Conditional	*United States*	3.68	3.06
mean forecast	*Germany*	1.45	1.20
$\hat{\mu}_{\tau+1}$	*Japan*	-0.92	-0.99

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

Table 3 – Risks to Price Stability as of September 2002

$\alpha = 3$ and $\beta = 2$ ("Deflation Paranoia")

		Horizon of 1 Year	Horizon of 2 Years
Inflation risk	*United States*	1.03	0.82
forecast	*Germany*	0.07	0.00
EIR_1	*Japan*	0.10	0.10
Balance of Risks	*United States*	0.68	-0.51
forecast	*Germany*	-0.72	-0.07
$BR_{2,1}$	*Japan*	-23.58	-14.07
Deflation risk	*United States*	-0.37	-0.89
forecast	*Germany*	-0.53	-0.05
DR_2	*Japan*	15.79	-9.45

NOTE: See Table 2.

Table 4 – Risks to Price Stability as of September 2002

$\alpha = 2$ **and** $\beta = 3$ **("Inflation Paranoia")**

		Horizon of 1 Year	Horizon of 2 Years
Inflation risk	*United States*	5.85	15.75
forecast	*Germany*	0.16	0.00
EIR_2	*Japan*	4.30	2.31
Balance of Risks	*United States*	8.75	23.58
forecast	*Germany*	-0.05	-0.18
$BR_{1,1}$	*Japan*	4.26	1.27
Deflation risk	*United States*	-0.02	-0.04
forecast	*Germany*	-0.29	-0.18
DR_1	*Japan*	-2.19	-2.19

NOTE: See Table 2.

Table 5 – Average Historical Year-on-Year Risks by Decade and Country

Symmetric preferences: $\alpha = 2$ **and** $\beta = 2$

	1960s	**1970s**	**1980s**	**1990s**	**2000s**
Inflation Risk EIR_1					
United States	0.75	3.00	2.91	1.02	0.84
Germany	0.40	1.42	0.80	0.55	0.21
Japan	3.92	4.83	0.79	0.34	0.11
Deflation Risk DR_1					
United States	-0.13	0.00	-0.04	-0.06	-0.04
Germany	-0.06	-0.02	-0.15	-0.07	-0.18
Japan	0.00	0.00	-0.21	-1.27	-2.01
Balance of Risks $BR_{1,1}$					
United States	0.62	3.00	2.87	0.95	0.80
Germany	0.34	1.40	0.65	0.47	0.03
Japan	3.92	4.83	0.58	-0.92	-1.90

NOTE: See Table 2.

Figure 1 – Alternative Central Bank Preferences

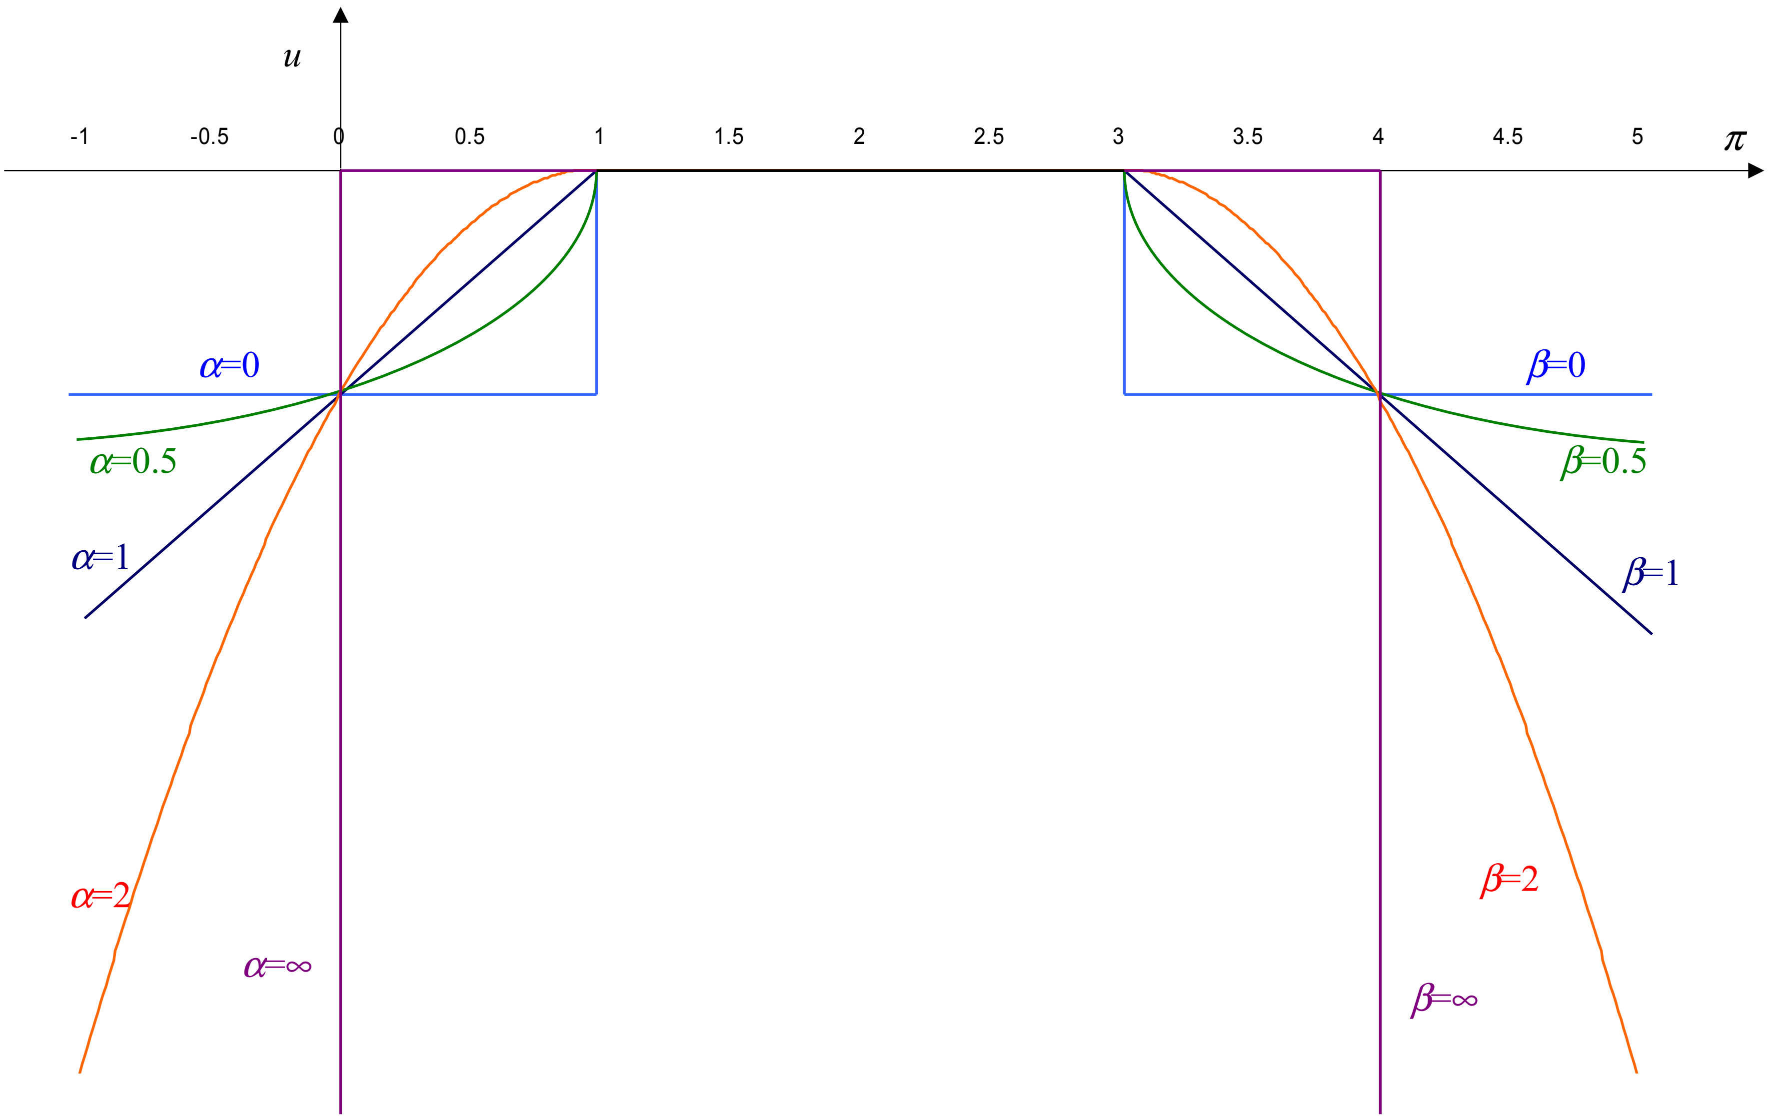

NOTE: For expository purposes we set $\underline{\pi} = 1\%$ and $\bar{\pi} = 3\%$ and impose symmetry.

Figure 2 – Areas of the Distribution of Inflation that Determine the Risks to Price Stability

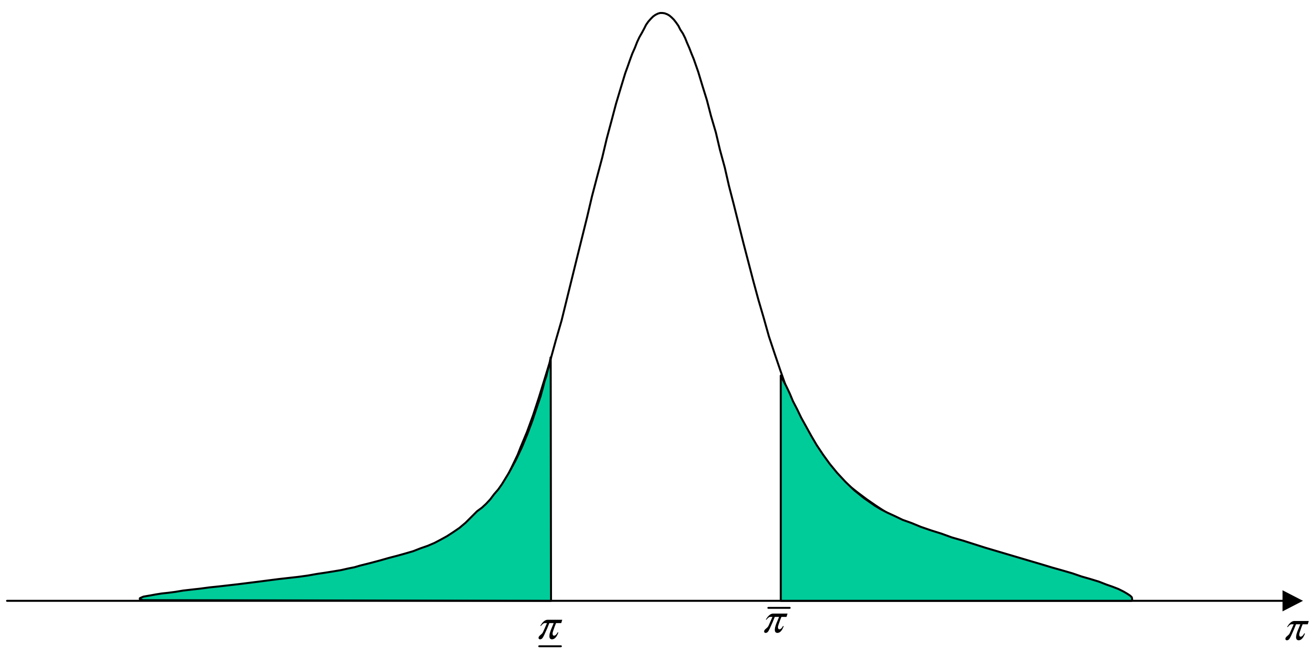

Figure 3 – Historical Evolution of Year-on-Year Risks by Country

Symmetric preferences: $\alpha = 2$ and $\beta = 2$

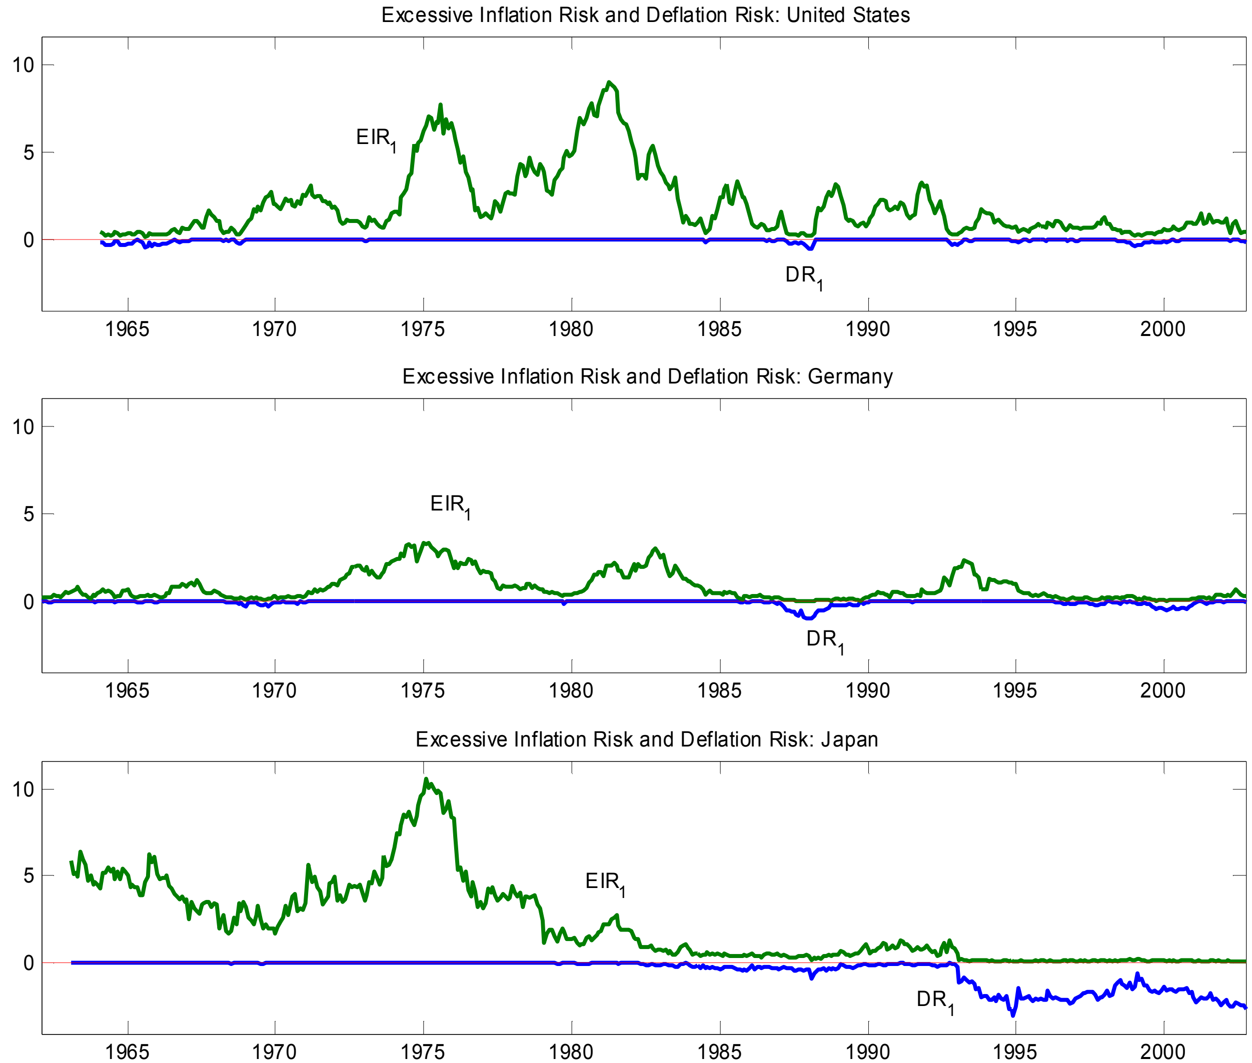

Figure 4 – Balance of Risks Under Quadratic Symmetric Preferences and Conditional Mean Forecasts One Year Ahead Plotted Against Actual CPI Inflation in Japan

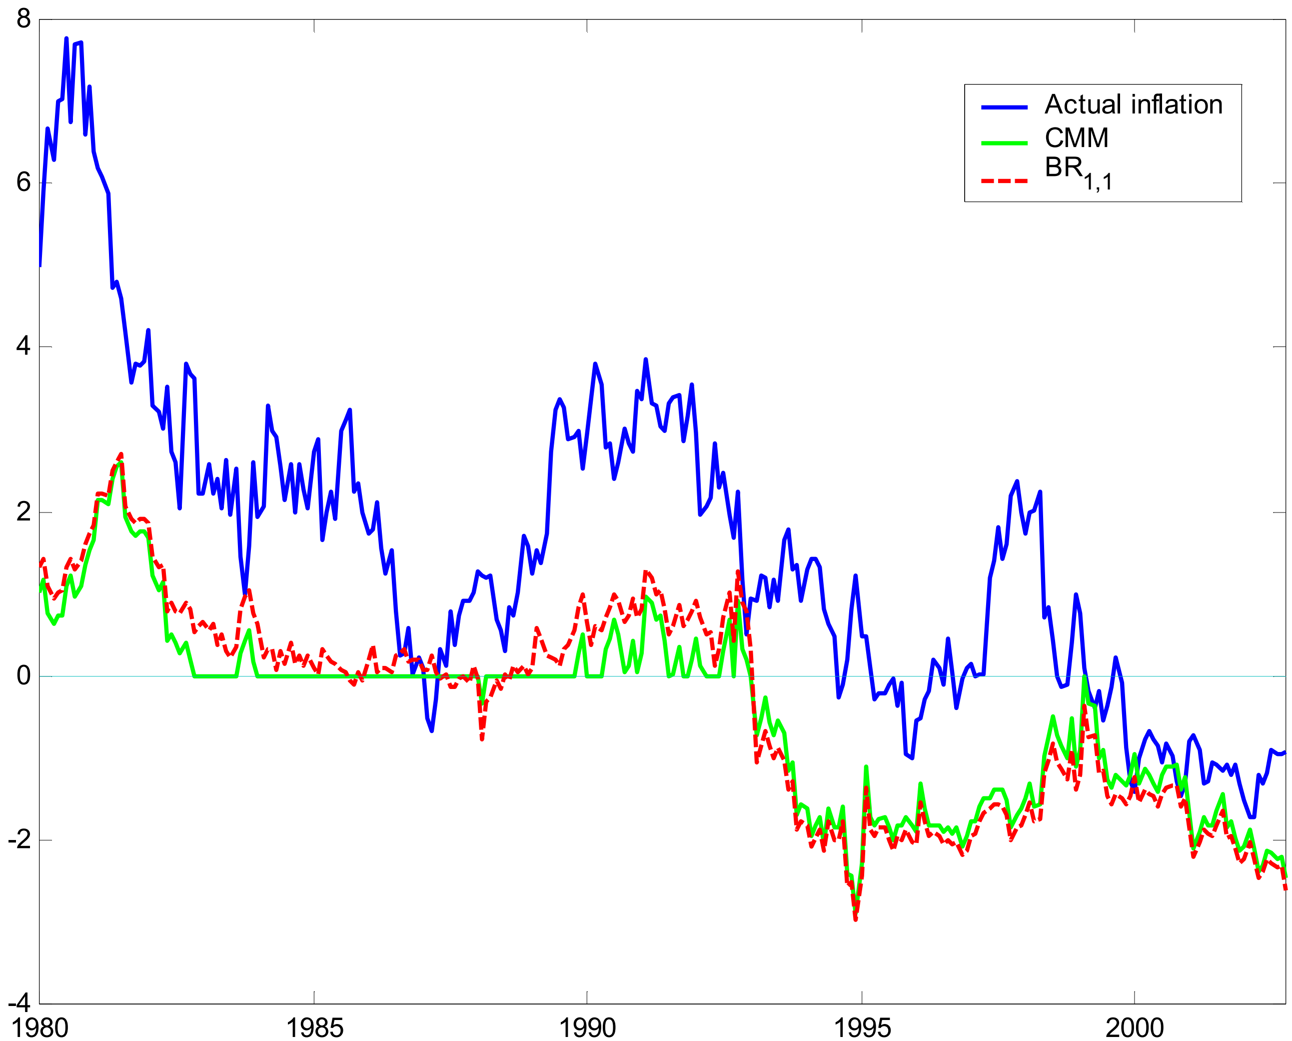

Figure 5 – Balance of Risks Under Asymmetric Preferences and Conditional Mean Forecasts One Year Ahead Plotted Against Actual CPI Inflation in Japan