LUTZ KILIAN

SIMONE MANGANELLI

Quantifying the Risk of Deflation

We propose formal and quantitative measures of the risk that future inflation will be excessively high or low relative to the range preferred by a private sector agent. Unlike alternative measures of risk, our measures are designed to 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 2 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.

As of late 2002, there was much public debate about whether the risks to price stability were tilted toward deflation in 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" (The Economist 2002b, 2003a, 2003b, 2003c).[1]

In response to these concerns in the financial press, central bankers such as Bernanke (2002, 2003a, 2003b) 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 were 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 signaled that "the probability of an unwelcome fall in inflation has diminished over recent months" (see Board of Governors 2003b).

The risk of deflation was also on the minds of consumers, investors, and businesses. The mere mention of the word "deflation" appeared 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 was concern that consumers, fearful of deflation, might curtail their spending, further aggravating countries' economic stagnation (see, e.g., Krugman 2003).

Despite the apparent importance of the risks of deflation for the private sector, 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. In this paper we propose formal measures of the risks to price stability that can be used by private sector agents to assess the risks to price stability in general and the risk of deflation in particular. The measures of risk that we propose are consistent with the standard treatment of risk in the economic literature, as summarized, for example, by Machina and Rothschild (1987) in *The New Palgrave Dictionary of Economics*. Also see Fishburn (1977), Holthausen (1981), and Basak and Shapiro (2001) for related approaches in other contexts.

Our formal measure of the risk of deflation is consistent with a loss-function approach to risk measurement. Our premise is that every private sector agent can express

1. Similar concerns were expressed in *Business Week*, the *Wall Street Journal*, and the *Financial Times*, among others (see Issing 2002a for details). Also see *The Economist* (2002a, 2002c) and the article by Wilson (2002) in the Goldman-Sachs *Global Economics Weekly*. 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.

his or her preferences for inflation in the form of a loss function subject to weak restrictions. Given a suitable parameterization of the loss function, 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.

It is important to stress that the risk of deflation is fundamentally different from the danger of a "deflationary spiral," which may or may not occur following a prolonged period of deflation, as the experience of Japan demonstrates (see Bernanke 2003b). 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. However, this leaves the central bank with the option of expanding the money supply by other means (see Bernanke 2003b). 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.

The remainder of the paper is organized as follows. In Section 1, we outline a loss-function-based approach to measuring the risks to price stability. In Section 2 we show how measures of risk may be computed by private sector agents in practice. In Section 3, 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 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 degree of risk aversion on the risk estimates. We conclude in Section 4.

1. THE LOSS-FUNCTION 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 New Palgrave Dictionary of Economics*). 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. This distribution function may be estimated by the empirical distribution of inflation forecasts.

Second, following Machina and Rothschild (1987), 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.

The term price stability here does not refer to *price level* stability (as in models of price level targeting in the academic literature). Rather "price stability" in this paper will be equated with periods of low inflation, in line with the definition of price stability used by policy makers (see Bernanke 2003b, and Issing 2000, 2002b, among others).[2] More formally, we define *low inflation* as inflation in the range $[\underline{\pi}, \bar{\pi}]$, where $\underline{\pi} \leq \bar{\pi}$. We define *excessive inflation* as inflation in excess of $\bar{\pi}$ and *deflation* as inflation below $\underline{\pi}$.[3] Our analysis allows for the possibility that $\underline{\pi} = \bar{\pi} \equiv \pi^*$. Users of inflation forecasts tend to associate risk with the failure of inflation to remain inside the range $[\underline{\pi}, \bar{\pi}]$. There is an upside risk of inflation exceeding $\bar{\pi}$ as well as a downside risk that inflation falls below $\underline{\pi}$.

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

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

where π denotes the inflation rate over the horizon of interest, $0 \leq a \leq 1$ represents the weight given to the losses associated with deflation and the parameters $\alpha \geq 0$, and $\beta \geq 0$ measures the degree of risk aversion of the economic agent. This loss 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 loss functions of interest in practice. It also provides the basis for an intuitive and economically appealing

---

2. Price stability in this sense clearly will not be the preferred outcome for every private sector agent under all circumstances. For example, unexpected inflation will favor borrowers at the expense of lenders. Nevertheless, there is a sense that departures from price stability ultimately are harmful to the private sector, a general attitude that is exemplified by the "cult of the conservative banker." Thus, measures of the risks to price stability are likely to be a useful summary statistic for many private sector agents.

3. 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."
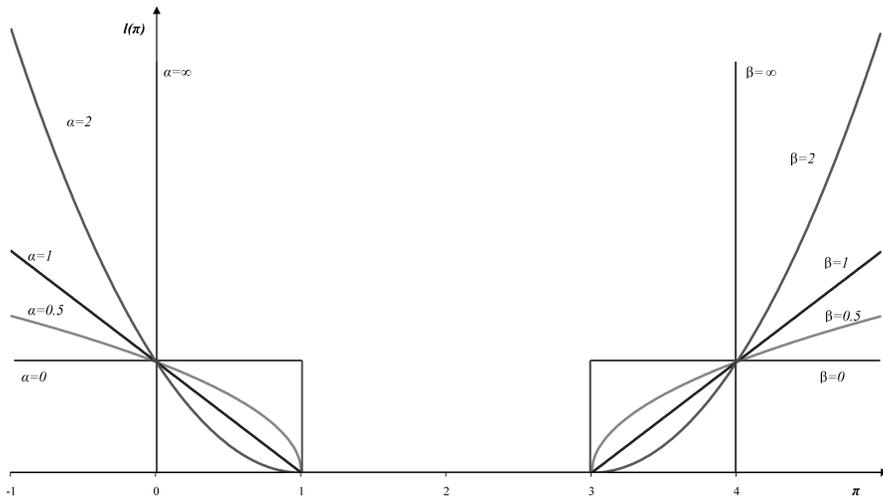
FIG. 1. Alternative Inflation Preferences.

NOTE: For expository purposes we set $\underline{\pi} = 1\%$ and $\bar{\pi} = 3\%$ and impose $a = 0.5$.

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 dissatisfaction that the agent experiences, as the inflation rate varies. For expository purposes we set $\underline{\pi} = 1\%$, $\bar{\pi} = 3\%$, and $a = 0.5$. Figure 1 shows that the loss is zero when inflation remains in the range consistent with price stability. As inflation exceeds the thresholds (whether it is too high or too low), the agent experiences dissatisfaction, represented by the upward-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.

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
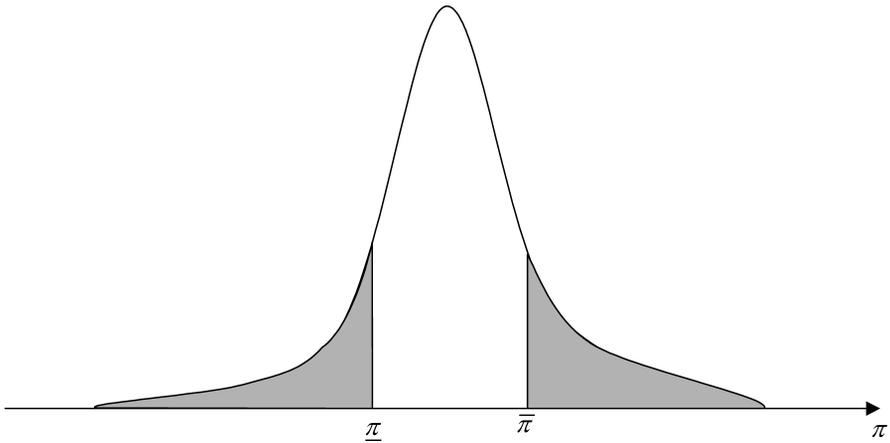
FIG. 2. Areas of the Distribution of Inflation that Determine the Risks to Price Stability.

forecasters. Rather the analyst 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 ($\bar{\pi}$) 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.

1.1 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 $\bar{\pi}$:

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

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

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

We adopt the convention of defining the risk of deflation as a negative number and the risk of excessive inflation as a positive number. 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 $\underline{\pi}$ and $\bar{\pi}$, and a fixed parameter governing risk aversion,

denoted by α and β. The parameters α and β may be different.[4] 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.

- 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.
- 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.
- The proposed measures of risk are directly linked to the expected loss of the agent, as will be shown below.
- In addition, 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 minimizes the expected loss of the economic agent.
- 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

---

4. A value of $\alpha = 1$ implies risk neutrality in the relevant range. 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.

unknown von Neumann–Morgenstern utility functions for inflation. 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).

1.2 On the Interpretation of the Proposed Risk Measures

Under suitable assumptions, there is a close link between the risk measures of Definition 1 and the expected loss of the economic agent. Let DR_α denote the deflation risk computed under the distribution F and let EIR_β denote the corresponding risk of excessive inflation as in Definition 1. Then $E[l] = -aDR_\alpha + (1-a)EIR_\beta$ if the agent's preferences can be expressed as in (1). This result implies that DR_α and EIR_β are proportionate to the expected loss 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 in which 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.

1.3 Interpretation of Risk in Some Interesting Special Cases

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

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

Thus, the general measures of risk simply reduce to the probabilities of exceeding the target range at either end.[5] 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.

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 for the private sector,

---

5. 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 UK 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 "the magnitude of the associated costs can be expected to increase sharply as deflationary pressures intensify." This statement implies that we would expect $\alpha > 1$ in the private sector 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. Consider a threshold of 2% and suppose that an agent faces the choice between two situations: (i) 2.001% inflation with probability 100%, and (ii) 10% inflation with probability 20% and inflation below 2% with probability 80%. If we go by the probabilities of missing the threshold, situation (i) is clearly worse. In practice, most agents would prefer (i) over (ii), 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. Greenspan's language is actually much closer to what our general risk measure would imply for $\alpha = \beta = 1$. In that case, we obtain

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

By construction DR_1 is a measure of expected deflation, and EIR_1 is 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.[6] This language closely mimics that used by Greenspan. In

6. Interestingly, our measure of deflation risk, for $\alpha = 1$, is formally equivalent to the measure of

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 $(4-2\%) \, 1/2 = 1\%$.[7]

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

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

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 loss when $\underline{\pi} = \bar{\pi} \equiv \pi^*$:

$$l(\pi) = 0.5(\pi - \pi^*)^2.$$

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

1.4 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 zero-mean conditional innovation term (u_{t+1}), whose distribution does not depend on μ_{t+1}:

$$\pi_{t+1} = \mu_{t+1} + u_{t+1} \quad u_{t+1} \sim F_{t+1}^u. \tag{2}$$

Further suppose that, for given F_{t+1}^u, the agent could choose the level of expected inflation μ_{t+1} that minimizes expected loss. Then under the assumptions of the Proposition below, the loss-minimizing 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: *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 loss is minimized for the level of expected inflation μ_{t+1}^* that solves the equation*

(downside) risk recently proposed by Basak and Shapiro (2001) in the context of value-at-risk applications under the name of *weighted expected shortfall*.

7. It may be tempting to conclude that the expected inflation conditional on missing the target would be 3% (= 2% + 1%) therefore. That conclusion would not be correct. The expected inflation rate conditional on exceeding the 2% target is 4%.

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

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

This result follows immediately from the first-order condition of the expected loss minimization problem. The proposition 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 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 $BR_{1,1} = DR_1 + EIR_1$.

One implication of this proposition is that quantitative comparisons of the balance of risks can be made over time and across space. Under the expected utility paradigm, the expected utility form is preserved under increasing linear transformation; i.e., for $b > 0$ we can write $l^*(\pi) = bE[l(\pi)] + c \Rightarrow BR_{\alpha-1,\beta-1}^*(\pi) = bBR_{\alpha-1,\beta-1}(\pi)$. Thus the ratio of the balances is invariant to positive linear transformations, provided the baseline balance differs from zero. This fact allows us to make intertemporal and spatial comparisons of the balance of risks for a given set of preferences.[8]

Note that the balance of risks is not an indicator of the overall extent of risk (which would be measured by the expected loss), 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 3. 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% 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^*:

$$BR_{1,1} = DR_1(\pi^*) + EIR_1(\pi^*) = -\int_{-\infty}^{\pi^*} (\pi^* - \pi)\, dF(\pi)$$

$$+ \int_{\pi^*}^{\infty} (\pi - \pi^*)\, dF(\pi) = \int_{-\infty}^{\infty} \pi\, dF(\pi) - \pi^* = E(\pi) - \pi^*.$$

8. When the balance of risks changes sign, the absolute values of the balances may be compared instead.

1.5 Implementation Issues

In this paper, we take as given that private sector agents have exogenously given preferences for price stability as shown in expression (1). We do not address the deeper question of where those preferences come from, but simply observe that preferences for inflation are derived preferences and in general will depend on the country-specific economic structure and historical experience.[9]

An important question in practice is how the agent's loss function $l(\pi)$ should be parameterized. Even though the focus of this paper is the private sector, it seems important that measures of the risks to price stability can be justified both from a private sector and from a central bank point of view. The overriding objective of all modern central banks is price stability. Most central banks in recent years have been quite specific about their definition of price stability. At one extreme, there is the European Central Bank, which by law is pursuing a given price stability objective. At the other extreme are central banks such as the Federal Reserve Board that are committed to pursuing a price stability target as reflected in the moderate inflation outcomes since the late 1980s, but that have not stated an explicit inflation target range.[10]

The preference for low inflation articulated by central banks reflects broad public support that was built in the 1970s during a period of excessively high inflation from the point of view of the public. This public support in some cases led to the inflation objective being enshrined by law, in many other countries it is reflected in official central bank policies, and in the United States it is at least implicit in official statements, actions, and speeches of Board officials. For example, Fed Governor Ben Bernanke (2003b) notes that "since the inflation crisis of the 1970s, the Federal Reserve has consistently pursued the goal of price stability in the United States." He further suggests that inflation "in the general range of 1 to 2 percent per year . . . is probably the *de facto* equivalent of price stability."[11] Given the broad public support for the objective of price stability since the early 1980s, exemplified by the cult of the conservative central banker, we may view the central bank's definition of price stability as representative for the private sector. This view is also supported by the close

9. Our approach is similar to that used in the empirical finance literature. Clearly, asset returns are not part of the agent's direct utility function, rather consumption is. Yet, many empirical researchers in finance have found it useful to specify the loss (or utility, to use the term preferred in the finance literature) of the agent directly in terms of the returns or wealth (see, e.g., Campbell and Thompson 2005, Brandt, Santa-Clara, and Valkanov 2005). Our analysis is in the same spirit as this literature in that we postulate that the utility that the economic agent derives from price stability (or equivalently the loss the agent incurs from its violation) can be expressed in terms of inflation directly, rather than in terms of the agent's consumption. Our approach also mirrors the loss function approach taken in the forecasting literature, except that our concern in this paper is how to evaluate a given forecasting distribution for inflation, rather than assessing the optimality of the forecasting model for inflation.

10. For example, Fed Governor Ben Bernanke on July 23, 2003, in a speech on the risk of deflation acknowledged that "although the Federal Reserve does not have an explicit target range for measured inflation, FOMC behavior and rhetoric have suggested . . . an implicit preferred range of inflation" (see www.federalreserve.gov/boarddocs/speeches/2003/20030723/default.htm).

11. Similarly, the Bank of Japan's stated mission is "to pursue price stability, in other words to maintain an economic environment in which there is neither inflation nor deflation," although no numerical targets have been made public (see http://www.boj.or.jp/en/about/about_f.htm).

attention the private sector pays to central bank statements related to the achievement of its price stability objective, as illustrated by the examples in footnote 1 such as the Goldman-Sachs *Global Economics Weekly* article by Wilson (2002) or the July 2003 *Business Week* article.

There is surprising agreement among the definitions of price stability across countries. Most central banks prefer a range of 1–3% (Bank of Sweden, Bank of New Zealand, Bank of Canada); the Bank of England is similarly committed to "2% on average"; some central banks such as the European Central Bank prefer a somewhat lower value of below 2%, but close to 2% in the medium run (see Issing 2002b, Duisenberg 2003); others prefer a somewhat higher range of 2–3% (Reserve Bank of Australia). This evidence suggests that there is near agreement on what reasonable upper and lower bounds on inflation are, both among central bankers and among the general public. In the empirical application in the next section we will postulate $\underline{\pi} = 0.01$ and $\bar{\pi} = 0.03$ for expository purposes, consistent with the definition of price stability most commonly used by central banks. It is understood that our approach could be adapted to any other definition of price stability, if necessary.

Even if there is a consensus on what price stability means, there still is the question of what specific form the loss function should take, conditional on $\underline{\pi}$ and $\bar{\pi}$. The specific loss function we focus on has three advantages. One advantage, as noted earlier, is that it implies as special cases measures of risk that are actually discussed by central bankers and in the financial press. This is indirect evidence that this loss function provides a reasonable representation of public views. A second reason is that this loss function allows us to rationalize the popular notion that upside and downside risks ought to be balanced. A third reason is that the measures of risk that are implied by our loss function under weak conditions will produce preference rankings of inflation distributions that are consistent with stochastic dominance. These features set our risk measures apart from alternative proposals. Finally, we note that the specific loss function proposed in this paper may be viewed as a flexible approximation to a wider class of loss functions in the neighborhood of the price stability range.

Given $\underline{\pi}$ and $\bar{\pi}$, the particular loss function underlying our risk measures requires the user to specify the parameters a, α, and β, and the forecast horizon of interest. In this paper, we choose specific values for a, α, and β to illustrate the computation of risk measures in practice and their dependence on the parameter values. These parameter values are chosen to be economically plausible, but it is understood that the substantive results may be affected by changes in the parameterization. Specifically, the parameter a that measures the relative weight of inflation in excess of $\bar{\pi}$ and inflation below $\underline{\pi}$ is set to 0.5, reflecting the fact that most central banks assign equal importance to departures from price stability in either direction (see, e.g., Issing 2002a). Our choice for a is also consistent with the quadratic loss functions typically specified in the macroeconomic literature (see, e.g., Svensson 1997).

The parameters α and β measure the degree of risk aversion of the economic agent. There is no direct evidence on central bankers' degree of risk aversion, although sometimes these parameters can be bounded based on policy statements, as discussed earlier. There is no direct evidence on the degree of risk aversion of private sector

agents either. 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 follow the finance literature in assuming that in practice the economic agent knows his or her degree of risk aversion.

Rather than restricting ourselves to one set of risk aversion parameters, we focus on illustrating how changes in risk aversion may affect the perception of risk. A virtue of our framework is that it is consistent with a wide range of assumptions about risk aversion. For the empirical application in the next section, we follow the academic literature in postulating $\alpha = \beta = 2$ as our benchmark case (see, e.g., Blinder 1997, Svensson 1997, 2002). We also consider larger values for α and β, and we allow for asymmetry in α and β. As we show, in many cases it is not necessary to know the precise degree of risk aversion because the qualitative results are robust for a wide range of these parameters.

The last parameter to be pinned down is the forecast horizon h. 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. In other words, we are interested in forecasting the cumulative change in inflation between now and h months from now, denoted by $\pi_{t+h}^{h} \equiv \sum_{i=1}^{h} \pi_{t+i}^{1}$, where π_{t}^{1} is the inflation rate recorded over a 1-month period and observed at month t. The choice of forecast horizon h thus affects the definition of π used in the risk assessment. In practice, most users of inflation forecasts are likely to be interested in risk assessments at medium-term rather than short-term horizons. This is especially true in the context of assessing the risk of deflation.

In the empirical application, we focus on horizons of 1 and 2 years. The upper bound of 2 years is roughly 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 of 1 or 2 years ahead (see Heikensten 1999). The Bank of England publishes forecasts for a horizon of 10 quarters. In some 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 2 years. This is not a problem of the risk methodology, but of the available data and of the forecasting model.

2. COMPUTING MEASURES OF INFLATION RISK IN PRACTICE

The measures of risk proposed in Section 1 are fully operational and can be computed easily in practice, as long as the forecasts can be expressed in the form of a distribution F. Here we will focus on their implementation in the context of a simple econometric forecasting model. Given that forecasting models must generally be considered mis-specified, it is natural to fit the conditional mean model directly at

the horizon of interest. This approach also avoids the well-known problems of time aggregation of volatility dynamics, and it dispenses with the need for a multivariate model.[12]

Let π_τ, $\tau = 1, 2, \ldots, T$, denote non-overlapping observations for the inflation rate measured at the horizon of interest, where we have suppressed the superscript h for notational convenience. Then

$$\pi_{\tau+1} = \mu_{\tau+1} + u_{\tau+1}, \quad u_{\tau+1} = \varepsilon_{\tau+1}\sqrt{h_{\tau+1}}, \quad \varepsilon_{\tau+1} \mid \Omega_\tau \sim \text{i.i.d.}(0, 1)$$
$$h_{\tau+1} = \theta_0 + \theta_1 u_\tau^2 + \theta_2 h_\tau, \tag{3}$$

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 $\theta_1 = 0$ and $\theta_2 = 0$, this model reduces to the homoskedastic case with $h_{\tau+1} = \theta_0 \forall \tau$.

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

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

12. In generating forecasts from models fit at higher frequency than the target horizon, it becomes necessary to recursively iterate the model forward. Since explanatory variables are known only for the current period, their future values have to be forecast, which necessitates the inclusion of an additional forecast equation for each additional variable to be used in predicting inflation. By fitting the model directly at the horizon of interest, the problem of forecasting reduces to one-step-ahead prediction, and no iteration is required.

where the additional predictors, $x_{k,\tau}, k = 1, \ldots, s$, 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 (see Inoue and Kilian 2006).

The conditional mean model may be estimated consistently by ordinary least squares. 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 (3), 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 (3) by $\hat{\varepsilon}_{\tau+1} \equiv (\pi_{\tau+1} - \hat{\mu}_{\tau+1})/(\hat{h}_{\tau+1})^{1/2}$. Given the empirical distribution of $\hat{\varepsilon}_{\tau+1}$, we first generate the bootstrap predictive density, conditional on the parameter estimates of model (3), based on $r = 1, \ldots, R$ bootstrap replications of $\{\pi_{\tau+1}^r\}$.[13] Then we make use of the fact that

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

where the probability $\Pr(\pi_{\tau+1}^r > \bar{\pi})$ can be computed as the fraction of observations $(\pi_{\tau+1}^r - \hat{\mu}_{\tau+1})/(\hat{h}_{\tau+1})^{1/2}$ that exceed $(\bar{\pi} - \hat{\mu}_{\tau+1})/(\hat{h}_{\tau+1})^{1/2}$, and where the expectation $E[(\pi_{\tau+1}^r - \bar{\pi})^{\beta} \mid \pi_{\tau+1}^r > \bar{\pi}]$ can be estimated by the sample moment $\frac{1}{\#(\pi_{\tau+1}^r > \bar{\pi})} \sum_{\pi_{\tau+1}^r > \bar{\pi}} (\pi_{\tau+1}^r - \bar{\pi})^{\beta}$ (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

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

where the last expectation can be computed by taking averages of the standardized residuals that satisfy the inequality condition.[14] 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

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

which may be computed analogously.

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

14. A similar technique has been used by Manganelli and Engle (2004) to compute the expected shortfall of the distribution of portfolio returns.

3. EMPIRICAL ANALYSIS

As noted in the introduction, there has been 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 economies, 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 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 1 and 2 years, as they existed in September 2002, when the issue of deflation risks was first raised. We will assess these risks from the point of view of a private sector agent. 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.[15]

Our starting point is a forecasting model of inflation. Risk forecasts of course are only as good as the underlying forecasting model. In practice, we use statistical model selection criteria to help us select the best possible forecasting model among a number of time series forecasting models involving lags of inflation and lags of other variables such as money growth rates and percent changes in oil prices. This set of predictors and class of forecasting models were chosen for illustrative purposes. Clearly, our risk analysis tools could be applied to any forecast distribution for inflation, whether the forecasts are judgmental or based on econometric forecast models, whether the forecasting model is structural or expressed in reduced form, and whether the forecast distribution is based on one model or on several probability-weighted models.

---

15. 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.

3.1 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.[16] Our estimation sample starts in January 1960 and ends in September 2002. We consider three alternative specifications of the conditional mean in model (3): 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.[17] 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 forecasting 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, suggesting that our AR model does a good job in describing the dynamics of the mean equation in all cases. We also test for

16. 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.

17. The monetary aggregates are M2 for 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.

TABLE 1

FORECAST MODEL DIAGNOSTICS BY COUNTRY

Predictors horizon	Inflation		Inflation and oil prices		Inflation and money	
	1 year	2 years	1 year	2 years	1 year	2 years
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.

conditional heteroskedasticity in the model residuals. The LB test shows no evidence of conditional heteroskedasticity.[18]

3.2 A Snapshot of the Risks to Price Stability

Table 2 provides risk estimates computed using information available in September 2002, when concern over deflation risks first arose. Clearly, these results should be viewed with some caution, as the effective sample size is small, especially for the

18. An important concern in modeling 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 modeling the inflation process that underlies the risk estimates.

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

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

2-year horizon. Nevertheless, it is of interest to obtain at least a preliminary and tentative assessment of these risks. As discussed earlier, in this section we define price stability as inflation inside the band defined by $\underline{\pi} = 1\%$ and $\bar{\pi} = 3\%$, and we attach equal weight to departures from this objective in either direction by setting $a = 0.5$. The baseline is a set of quadratic symmetric preferences. Later we will consider asymmetric preferences incorporating either a strong aversion to deflation or a strong aversion to excessive inflation.

An important question is how to interpret these magnitudes. As shown in Section 1, in many cases risk measures have an intuitive interpretation, for example, as probabilities of exceeding a target, when the risk aversion parameter is zero, or as tail conditional expectations or semi-target variances when the risk aversion parameter is 1 or 2, respectively. Even in cases in which the risk measure reduces to an object as simple as a probability, however, it may not be clear whether a given probability is small or large. The problem of evaluating risk estimates for inflation is analogous to that of evaluating risks for alternative portfolios at a point in time or of evaluating time series of risk estimates for the same portfolio. One way of evaluating risk estimates is to compare the risk at a point in time for one country to that of another. The key assumption in comparing risks across countries is that the evaluator applies the same loss function across countries much like an investor would compare the risks on two alternative portfolios. Another possibility is to compare today's risk in a given country to historical risk levels in the same country (again using the same loss function), much like an investor would compare changes in the risk of a given portfolio over time. In other words, we can think of risk measures as indices, the level of which is arbitrary, but whose change over time or space conveys information. We will illustrate both approaches below.

An obvious limitation of our approach of comparing risks across countries conditional on a given preference function is that agents in different countries may have

different preferences for inflation. This seems the more likely, the more different the historical experiences and economic structures of the countries in question are. A private sector agent (or for that matter a central banker) in Argentina, for example, may have a different loss function and hence a different perception of the risks to price stability in Argentina than an agent (or central banker) in the United States would when confronted with the same inflation forecast distribution F. Given such differences in preferences, one cannot compare the resulting risk assessments across countries; one can only judge the inflation and deflation risks relative to each country's own past. In contrast, in this paper we evaluate the inflation forecast data for different countries based on the same preference function. Implicit in our approach below is the presumption that the countries in question are similar enough for international risk comparisons based on the same preference function to be meaningful. Our approach can also be motivated from the perspective of a global investor who is interested in comparing inflation and deflation risks across countries. This perspective is consistent with the approach taken by investor newsletters such as the Goldman-Sachs *Global Economics Weekly* (see Wilson 2002).

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 third panel of Table 2 shows that for all countries there is some evidence of deflation risk, although the magnitudes differ greatly. To put these numbers in perspective, it is useful to focus on the relative magnitudes across countries. For example, we may be concerned that the U.S. deflation risk is approaching Japanese levels. Table 2 rejects that notion. The U.S. estimate of –0.02 at the 1-year horizon is likely to be negligible, not just compared with Japan's estimate of -2.19, but even compared with Germany's estimate of -0.29. The risk of deflation in the United States is only 0.9% of that in Japan and only 6.9% of that in Germany. The practical significance of the estimate of –0.29 for Germany also seems small compared with Japan, where the risk of deflation is about eight times higher.

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 the 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 degree to which the balance of risks is tilted in favor 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 1-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 2-year horizon indicate a slight improvement of

TABLE 3

RISKS TO PRICE STABILITY AS OF SEPTEMBER 2002: $\alpha = 3$ AND $\beta = 2$ ("STRONG DEFLATION AVERSION")

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

NOTE: See Table 2.

the balance of risks, strengthening the conclusion that deflation risks are not a serious concern in September 2002.

How would this assessment change if conducted by an agent with asymmetric preferences? Suppose for expository purposes the existence of a strong aversion to deflation in the form of preference parameters of $\alpha = 3$ and $\beta = 2$ ("strong deflation aversion"). Table 3 shows that in that case there is some evidence at the 1-year horizon that the balance is tilted in favor of deflation not just for Japan, but also for Germany. In contrast, the United States on balance remains in the inflationary region even for this strong form of deflation aversion. For the 2-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 their own past, as we will show below). 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 $\beta = 3$ ("strong inflation aversion"). 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 loss, 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 the three countries would appear to suffer from important deflation risks under this form of strong inflation aversion.

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

TABLE 4

RISKS TO PRICE STABILITY AS OF SEPTEMBER 2002: $\alpha = 2$ AND $\beta = 3$ ("STRONG INFLATION AVERSION")

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

NOTE: See Table 2.

aversion" the risks of deflation appear to be negligible for the United States and Germany, especially when we extend the horizon to 2 years.

3.3 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 1-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 risk estimates 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 loss associated with departures from the lower and the upper threshold. Thus, the time series plots in Figure 3 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. With the exception of Japan there is no evidence in Figure 3 to support *The Economist*'s conjecture. There is some evidence that deflation risks were unusually high in 1987, following the stock market crash, for example, but no evidence of high risks of deflation in 2002. For the United States, the 2002 risk estimates look very much like they did in the early 1960s. For Germany, the risk of deflation if anything has declined. In the year 2000 it was only half as high as in 1987. At no point in time have German or U.S. risks of deflation in Figure 3 been even close to the level of Japan's deflation risk since 1993.

Since time series estimates are noisy, we also computed historical averages. Table 5 shows that deflation risks for the United States today are not unlike those in the 1960s.
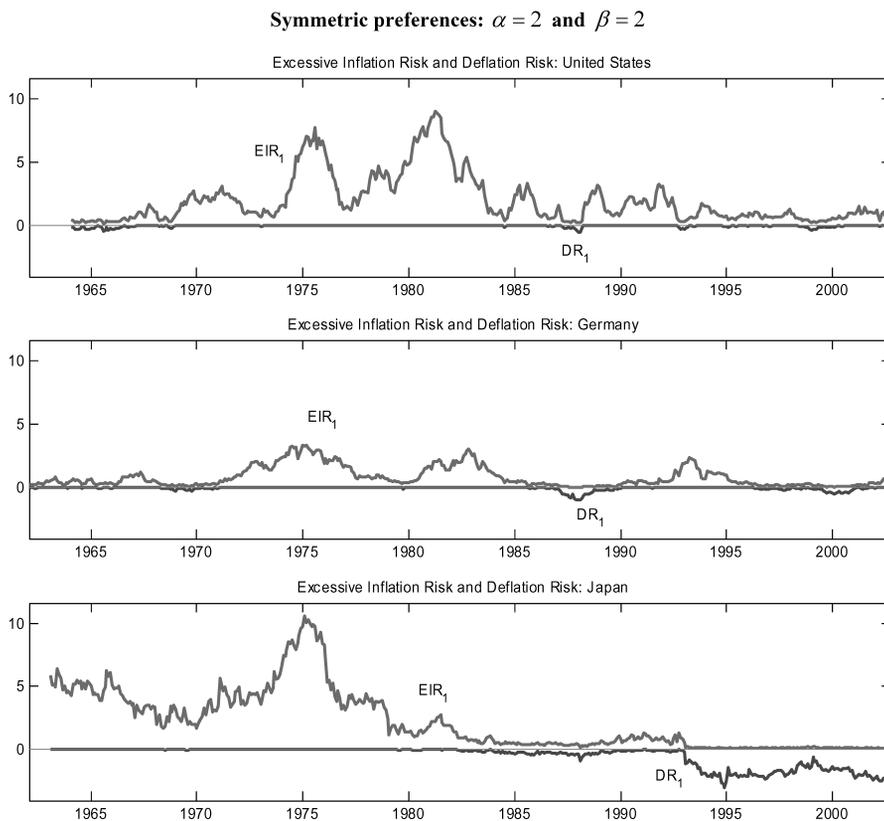
Symmetric preferences: $\alpha = 2$ **and** $\beta = 2$

Excessive Inflation Risk and Deflation Risk: United States

Excessive Inflation Risk and Deflation Risk: Germany

Excessive Inflation Risk and Deflation Risk: Japan

Fɪɢ. 3. Historical Evolution of Year-on-Year Risks by Country: Symmetric Preferences: $\alpha = 2$ and $\beta = 2$.

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.

3.4 The 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).

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.

It is immaterial from our point of view whether such forecasts are generated from an expectations-augmented Phillips curve or some other forecasting model. Such approaches are inconsistent with the loss-function-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 the balance 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 CMM, and the balance of risks under quadratic symmetric preferences. As before, the balance of risks is computed under the assumption that $\underline{\pi} = 1\%$, $\bar{\pi} = 3\%$, and $a = 0.5$ based on the model selected in Table 1. Figure 4 illustrates that there is no simple relationship between the CMM 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 CMM 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 1984–89 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,
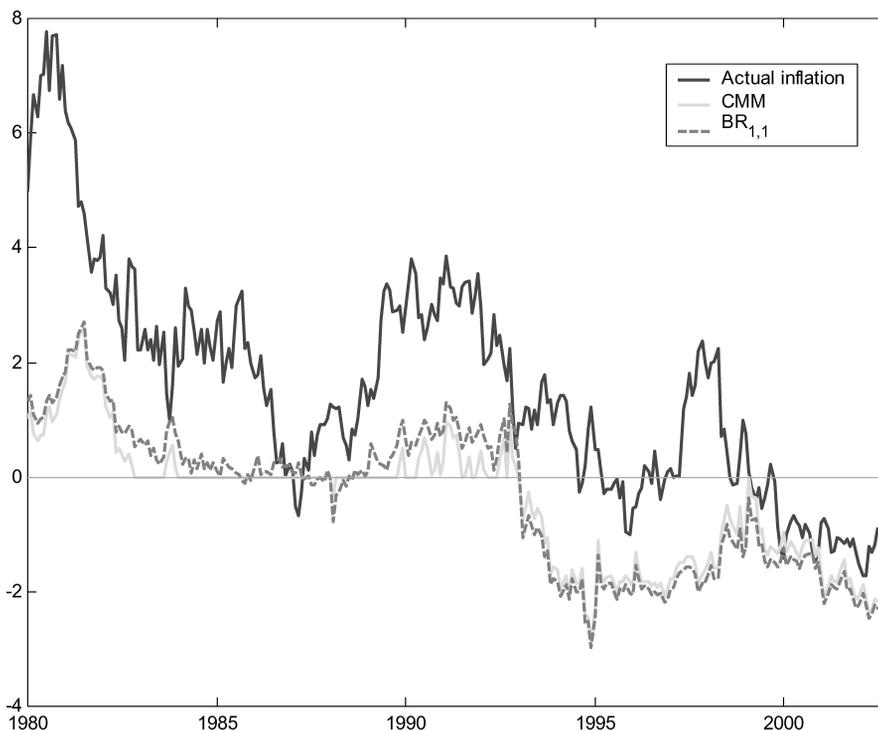
Fɪɢ. 4. Balance of Risks under Quadratic Symmetric Preferences and Conditional Mean Forecasts 1 Year Ahead Plotted against Actual CPI Inflation in Japan.

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.

4. CONCLUSION

The ability to measure and forecast the risks to price stability is important to consumers, investors, and businesses, as the extensive discussion of these risks in the financial press illustrates. 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. The proposed measures of risk are consistent with the standard treatment of risk in the economic literature, as summarized, for example, by Machina
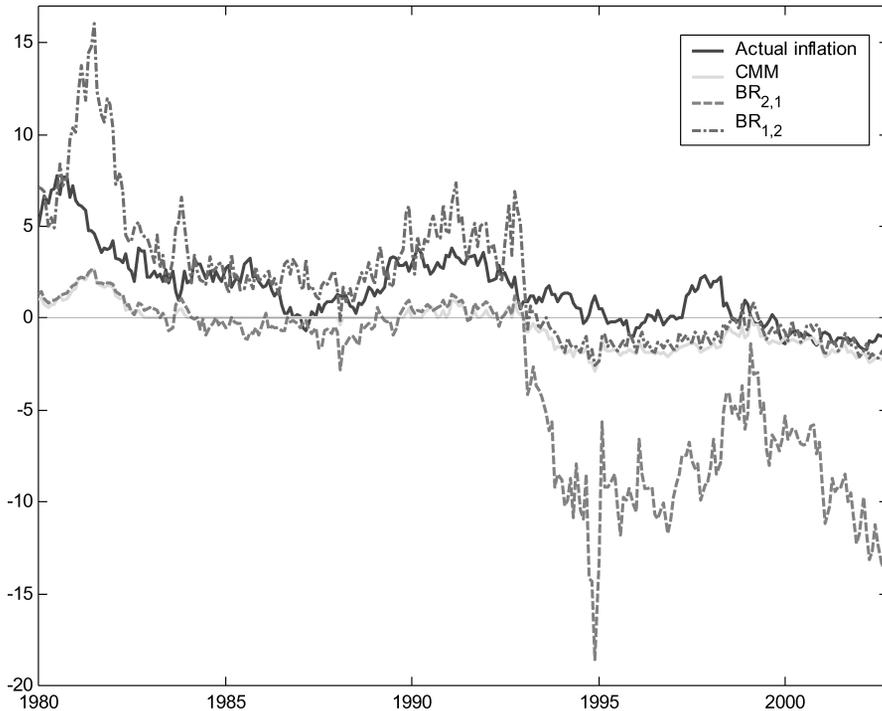
Fig. 5. Balance of Risks under Asymmetric Preferences and Conditional Mean Forecasts 1 Year Ahead Plotted against Actual CPI Inflation in Japan.

and Rothschild (1987). Unlike alternative measures of deflation risk, our measures are designed to 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 2 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 depends 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 such an analysis of the risks to real GDP growth. 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.

LITERATURE CITED

Basak, S., and A. Shapiro. (2001) "Value-at-Risk Based Management: Optimal Policies and Asset Prices." *Review of Financial Studies,* 14, 405–7.

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 of 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, 307–28.

Bollerslev, T., and J.M. Wooldridge. (1992) "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances." *Econometric Reviews,* 11, 143–72.

Brandt, M.W., P. Santa-Clara, and R. Valkanov. (2005) "Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns." Mimeo, Fuqua School of Business, Duke University.

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

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

Campbell, J., and S. Thompson. (2005) "Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average?" Mimeo, Department of Economics, Harvard University.

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, 961–78.

Drost, F.C., and T.E. Nijman. (1993) "Temporal Aggregation of GARCH Processes." *Econometrica,* 61, 909–27.

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, 116–26.

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. San Diego, CA: Academic Press.

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, 182–88.

Inoue, A., and L. Kilian. (2006) "On the Selection of Forecasting Models." *Journal of Econometrics,* 130, 273–303.

Issing, O. (2000) "Why Price Stability?" In *Why Price Stability?* edited by A.G. Herrero, V. Gaspar, L. Hoogduin, J. Morgan, and B. Winkler, pp. 179–202. Frankfurt a.M., Germany: European Central Bank.

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 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." European Central Bank Working Paper No. 226.

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

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

Machina, M.J., and M. Rothschild. (1987), "Risk." In *The New Palgrave Dictionary of Economics*, edited by J. Eatwell, M. Millgate, and P. Newman, pp. 203–5. London, UK: MacMillan.

Manganelli, S., and R.F. Engle. (2004) "A Comparison of Value at Risk Models in Finance." In *Risk Measures in the 21st* Century, edited by G. Szegö. Chichester, UK: John Wiley.

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 *NBER Macroeconomics Annual 2001*, edited by B. Bernanke and K. Rogoff, pp. 137–83. Cambridge, MA: MIT Press.

Stock, J.H. (2002) "Comment on Cogley and Sargent: Evolving Post-World War II U.S. Inflation Dynamics." In *NBER Macroeconomics Annual 2001*, edited by B. Bernanke and K. Rogoff, pp. 137–83. Cambridge, MA: MIT Press.

Svensson, L.E.O. (1997) "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." *European Economic Review,* 41, 1111–46.

Svensson, L.E.O. (2002) "Inflation Targeting: Should It Be Modeled as an Instrument Rule or a Targeting Rule?" *European Economic Review,* 46, 771–80.

Wilson, D. (2002) "Deflation—An Inflated Concern?" *Global Economics Weekly*, 02(40), Goldman-Sachs.