LUTZ KILIAN
SIMONE MANGANELLI

The Central Banker as a Risk Manager: Estimating the Federal Reserve's Preferences under Greenspan

We derive a natural generalization of the Taylor rule that links changes in the interest rate to the balance of the risks implied by the dual objective of sustainable economic growth and price stability. This monetary policy rule reconciles economic models of expected utility maximization with the risk management approach to central banking. Within this framework, we formally test and reject the standard assumption of quadratic and symmetric preferences in inflation and output that underlies the derivation of the Taylor rule. Our results suggest that Fed policy decisions under Greenspan were better described in terms of the Fed weighing upside and downside risks to their objectives rather than simply responding to the conditional mean of inflation and of the output gap.

JEL codes: E31, E52, E58
Keywords: monetary policy, inflation, output, risk, preferences, policy rule, Greenspan.

IN MOTIVATING POLICY decisions, the Federal Reserve Board routinely makes reference to "the FOMC's [Federal Open Market Committee] consensus about the balance of risks to the attainment of its long-run goals of price stability and sustainable economic growth."[1] A good example of this language is the

1. Although this practice was only formalized in a press release dated January 19, 2000, it is important to note that this press release marked a change in language rather than substance. The Federal Reserve has consistently pursued the goal of price stability since the inflation crisis of the 1970s. After a long period in which the desired direction for inflation was always downward and the risk of deflation was so remote that it did not warrant explicit mention, by the year 2000 there was evidence that the risk of deflation was no longer negligible. The change of language by the Federal Reserve Board in 2000 was intended to remind the public that the Fed was aware of the existence of both upside and downside risks to inflation and intent on achieving a balance between those risks (see Bernanke 2003b).

The views expressed in this paper are our own and do not necessarily reflect the views of the European Central Bank (ECB) or its staff. We thank two anonymous referees, the editor, Chris House, and Barbara Rossi for helpful comments.

LUTZ KILIAN *is Professor of Economics, Department of Economics, University of Michigan and Centre of Economic Policy Research* (*E-mail:* lkilian@umich.edu). SIMONE MANGANELLI *is Principal Economist, Directorate General Research, European Central Bank* (*E-mail:* simone.manganelli@ecb.int).

Received March 17, 2008; and accepted in revised form March 18, 2008.

Journal of Money, Credit and Banking, Vol. 40, No. 6 (September 2008)

May 6, 2003, FOMC statement that "the Committee perceives that over the next few quarters the upside and downside risks to the attainment of sustainable growth are roughly equal. In contrast, over the same period, the probability of an unwelcome substantial fall in inflation, though minor, exceeds that of a pickup in inflation from its already low level." It therefore makes sense to view the FOMC's policy actions under Greenspan as the outcome of a risk management problem. The nature of this risk management problem is that the central banker changes the setting of the policy instrument such that the balance of these risks is minimized. This is indeed the way in which Greenspan has described the process of conducting monetary policy (see, e.g., Greenspan 2002, 2003, 2004, Meyer 2004). There is a striking discrepancy between this description of the conduct of the Federal Reserve's monetary policy and standard Taylor-type policy rules proposed in the academic literature. The objective of this paper is to reconcile these seemingly disparate accounts of the policy process, and to formalize and rationalize the approach to monetary policy referred to by the Federal Reserve.[2]

The paper makes three distinct contributions. Our first contribution is theoretical. We show that under certain conditions the risk management approach practiced by central bankers is equivalent to expected utility maximization as postulated by academic macroeconomists. This result is unexpected and important, since at first sight, what central bankers say they do seems completely at odds with what academic economists envision them doing.

Our second theoretical contribution is to show that the type of risk-based decision rule implicit in FOMC statements can be derived from this framework (and that this decision rule will coincide with the Taylor rule only under the restrictive assumption of quadratic and symmetric preferences). We show that such a policy rule is fully consistent with expected utility maximization. Generalizing monetary policy rules to the case of potentially asymmetric and non-quadratic preferences is a non-trivial task. Previous attempts in the literature to solve this problem have resulted in policy rules that were *ad hoc*, were not grounded in reality, and had no transparent economic interpretation.

Our third contribution is empirical. We show that the central bank's preference parameters can be estimated using our approach, and that the assumption of quadratic and symmetric preferences underlying the Taylor rule can be tested. Taylor rules have been among the leading empirical models used in macroeconomics for decades now. Using U.S. data for the Greenspan period, we provide highly statistically significant and robust evidence against the preference assumptions underlying the Taylor rule.

The Greenspan view that monetary policymaking is an exercise in risk management is shared by central bankers in many countries. While we focus on the Federal Reserve, the tools developed in this paper can be adapted easily to the risk management strategies pursued by other central banks, whether their objectives involve a range of

2. Our paper is an exercise in positive economics rather than normative economics. Rather than deriving optimal policy rules based on hypothetical model environments, our point is that central bankers have a well-documented preference for the risk management approach and that fact makes it worthwhile to study this approach. If we want to rationalize what central banks do, we need to show that their risk management approach can be understood as a different parameterization of the expected utility maximization problem.

inflation as in the case of the European Central Bank or a point target as in the case of the Bank of England.

The remainder of the paper is organized as follows. In Section 1, we propose to regard the central banker as a risk manager who aims at containing inflation and the deviation of output from potential within prespecified bounds. We formally state the objective of price stability and sustainable growth. We develop formal statistical measures that may be used to quantify the risks of failing to attain that objective. Risk measures inherently depend on the loss function of the user. The importance of allowing for flexible loss functions has recently been demonstrated in a variety of contexts by Elliott, Komunjer, and Timmermann (2005, 2008). We propose a simple, yet flexible class of loss functions that includes quadratic symmetric preferences as a special case but allows for more general asymmetric and nonquadratic preferences. This class of loss functions implies as special cases many of the measures of risk that have been discussed in the literature (see, e.g., Fishburn 1977, Holthausen 1981, Basak and Shapiro 2001, Kilian and Manganelli 2007). We show how to relate the risk measures that emerge under specific assumptions about the central banker's preferences to the language used by central bankers in describing the risks to price stability and real economic activity.

In Section 2, we present a model of central bank risk management. Our model dispenses with the assumption of quadratic and symmetric preferences that underlies the derivation of standard Taylor-type policy rules. Based on this model we derive a generalized policy rule that relates the change in the policy instrument to a suitably defined measure of the balance of risks, lending formal support to the language used by central bankers. We contrast our approach to the derivation of standard Taylor-type policy rules under the assumption of quadratic symmetric preferences over inflation and output. Our policy rule nests the standard case of quadratic and symmetric central bank preferences.

Both the risk measures of Section 1 and the policy rule of Section 2 depend on the parameters of the central banker's loss function. In Section 3, we show how these parameters under weak assumptions may be estimated from realizations for inflation and output gap data. Unlike previous attempts to estimate central bank preferences such as Ruge-Murcia (2003), our approach does not require the specification of a structural model of the macroeconomy. Moreover, rather than specifying an *ad hoc* asymmetric loss function as in Ruge-Murcia (2003) or taking higher-order expansions of Taylor rules as in Cukierman and Muscatelli (2003), the specific form of asymmetric and nonquadratic preferences we allow for is tightly linked to the risk language used by the Federal Reserve Board.

In Section 4, we present estimates of the risk aversion parameters of the Federal Reserve Board over the Greenspan period of 1987–2005 with respect to departures from the inflation and output objectives. We formally test for and reject the standard assumption of quadratic and symmetric preferences that underlies the derivation of the Taylor rule. Our analysis allows for alternative interpretations of the inflation objective in terms of a core Consumer Price Index (CPI) objective of 2.0% or a core personal consumption expenditures (PCE) inflation target of 1.5%. These choices are motivated by statements of current and former Fed policymakers. We also allow for alternative assumptions about the timing of the output gap measure in the policy rule. Finally, we compare the ability of the generalized policy rule to explain changes in the Fed Funds rate to that of the corresponding rule derived under quadratic symmetric preferences. We show that the differences not only are statistically significant but can be economically important. We conclude in Section 5.

1. RISK MEASURES

1.1 The Framework

What do central bankers mean when they refer to upside and downside risks to price stability and to deviations from potential output? Risks arise from the uncertainty surrounding future realizations of these variables. There is general consensus that "upside risks" are related to the event of the random variable of interest exceeding a certain threshold. Similarly, "downside risks" are commonly associated with realizations below a certain threshold. For example, Bernanke (2003a) defines price stability as avoidance of deflation as well as excessive inflation.[3] The central banker in this view is concerned about the possibility that inflation, denoted by π, may exceed some upper threshold $\bar{\pi}$, or fall below a lower threshold $\underline{\pi}$, where $\bar{\pi} \geq \underline{\pi}$. This allows for the possibility that $\bar{\pi} = \underline{\pi} = \pi^*$. In the former case, the central bank is said to have a zone target (or band target), in the latter case a point target (see, e.g., Orphanides and Wieland 2000, Mishkin and Westelius 2008). Similarly, the central bank's objective is to avoid deviations of output from potential, denoted by x, beyond an upper threshold $\bar{x}$ and a lower threshold $\underline{x}$, where typically $\underline{x} = \bar{x} = x^* = 0$. Thus, it is natural to frame the problem of managing the risks to output and price stability in terms of the objective of setting the policy instrument such that output and inflation remain within well-defined bounds set by the policymaker.

This approach requires that the central banker quantify the uncertainty associated with future output and inflation outcomes. From a purely statistical standpoint, knowledge of the joint probability distribution of the random variables of interest provides a complete and exhaustive description of their underlying uncertainty. Techniques to estimate and represent the probability distribution of inflation and of the deviation of output from potential are readily available, but little attention has been paid in the literature to the mapping from the probability distribution to the assessment of the risks inherent in the economic outlook.[4]

In this paper, our objective is to reduce the information contained in the probability distribution of inflation and output outcomes to indicators of risk that are easy to interpret and that at the same time effectively summarize the features of the predictive density that are most important to policymakers. We propose formal measures of output gap and inflation risks that are directly related to the decision problem of the central banker.

3. It is important to note that the notion of *price stability* referred to by central bankers such as Bernanke differs from the concept of *price level stability* sometimes used in the academic literature.

4. The "fan chart," popularized by the Bank of England, is one such example. The fan chart graphs the central 10% prediction interval as a dark band and graphs successively wider intervals in ever lighter shades of red. The selective shading of the intervals is intended to draw attention to the uncertainty of future inflation and output.

1.2 Formal Definition of the Risk Measures

There is a large and growing academic literature on risk measurement (see, e.g., Machina and Rothschild 1987 in *The New Palgrave Dictionary of Economics*). This literature postulates that 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. In the present context, the random variables of interest are the inflation rate and the output gap. The joint probability distribution function of future inflation and gap outcomes will be denoted by $F_{\pi,x}$. The corresponding marginal distributions will be denoted by F_π and F_x. The second requirement is that any measure of risk must be linked to the preferences of the economic agent, typically parameterized by a loss function. Except in special cases, this requirement rules out measures of risk based on the sampling distribution of the variables of interest alone such as the sample mean, sample variance, or sample skewness, range or interval forecasts, and more generally quantile-based risk measures such as value at risk or the corresponding tail conditional expectation.

Given the aforementioned objective of the central bank to keep inflation and output within prespecified bounds, a natural measure of risk in the context of our loss function is the probability-weighted average of the losses incurred when realizations fall outside their respective bounds of $[\underline{\pi},\ \bar{\pi}]$ and $[\underline{x},\ \bar{x}]$:

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

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

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

DEFINITION 1b [Risk of a negative gap (*NGR*) and risk of a positive gap (*PGR*)]:

$$NGR_\gamma \equiv -\int_{-\infty}^{\underline{x}} (\underline{x} - x)^\gamma \, dF_x(x), \quad \gamma \geq 0$$

$$PGR_\delta \equiv \int_{\bar{x}}^{\infty} (x - \bar{x})^\delta \, dF_x(x), \quad \delta \geq 0,$$

where the parameters α, β, γ, and δ are measures of risk aversion.[5]

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), Basak and Shapiro (2001), and Kilian and Manganelli (2007). As we will show in Section 1.6, many statistical measures of risk that have been proposed in other contexts (such as the probability of exceeding a threshold, the target semivariance, the target variance, and the tail conditional expectation or weighted expected shortfall) can be derived as special cases of Definition 1.

5. We adopt the convention of measuring downside risks as a negative number and upside risks as a positive number.

1.3 The General Risk Management Model

The quotes in the Introduction suggest that we may view the central banker as a risk manager. Risk management by definition involves a trade-off. In the mean-risk model of returns proposed by Fishburn (1977), for example, the trade-off is between achieving a higher expected return and increased downside risk of a return below a given threshold. In the central bank risk management model, in contrast, the trade-off is between the upside risk of inflation exceeding $\bar{\pi}$ and the downside risk of inflation falling below $\underline{\pi}$ and, analogously, between the risk of output exceeding $\bar{x}$ and falling below $\underline{x}$. There also is a trade-off between achieving the inflation and the output objective. To fix ideas, we will only focus on the inflation objective for now. In Section 1.5, we will return to the general model involving multiple objectives.

Let $DR_\alpha(F_\pi)$ denote the deflation risk computed under the distribution F_π and let $EIR_\beta(F_\pi)$ denote the corresponding risk of excessive inflation. Then the central banker's risk management problem can be described as follows:

DEFINITION 2 [Risk management problem]: *Let F_π and G_π denote two alternative probability distributions of inflation. Then F_π is weakly preferred to G_π if $|DR_\alpha(F_\pi)| \leq |DR_\alpha(G_\pi)|$ and $EIR_\beta(F_\pi) \leq EIR_\beta(G_\pi)$. If this condition does not hold, the central banker faces a risk management problem.*

This risk management problem cannot be solved in the absence of further information about preferences. Solving this trade-off requires a risk management model. The defining characteristic of the central bank risk management model that we propose is that preferences can be expressed as a function of DR_α and EIR_β.

DEFINITION 3 [Risk management model]: *We say that a central banker's preferences satisfy a risk management model if and only if there is a real valued function U in risks such that for all relevant distributions F_π and G_π, F_π is preferred to G_π if and only if $U(DR_\alpha(F_\pi), EIR_\beta(F_\pi)) > U(DR_\alpha(G_\pi), EIR_\beta(G_\pi))$.*

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

1.4 Congruence of the General Risk Management Model with the Expected Utility Model

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

Definition 4 [Congruence between risk management model and expected utility model]: *We say that the risk management model in Definition 3 is congruent with the expected utility model if*

$$U(DR_\alpha(F_\pi), EIR_\beta(F_\pi)) > U(DR_\alpha(G_\pi), EIR_\beta(G_\pi))$$
$$\Leftrightarrow \quad \int_{-\infty}^{\infty} L(\pi)\, dF_\pi(\pi) < \int_{-\infty}^{\infty} L(\pi)\, dG_\pi(\pi).$$

In other words, the expected utility model will be congruent with a risk management model only for suitably chosen loss functions. A natural choice that includes the leading examples of loss functions discussed in the academic literature is given by:

$$L = aI(\pi_t < \underline{\pi})(\underline{\pi} - \pi)^\alpha + (1-a)I(\pi > \bar{\pi})(\pi - \bar{\pi})^\beta,$$

where $0 \leq a \leq 1$, $\alpha \geq 0$ and $\beta \geq 0$. It follows immediately from integrating this loss function with respect to the distribution of π that $E[L] = -aDR_\alpha + (1-a)EIR_\beta$. In this example, the linear function $U(DR_\alpha, EIR_\beta) \equiv aDR_\alpha - (1-a)EIR_\beta$ equals $-E[L]$ and hence satisfies congruence. The significance of this result is that it clarifies that there need not be any contradiction between the way many central bankers have framed policy discussions and the way economists have traditionally modeled this problem. For a wide class of preference functions the use of a risk management strategy by the central banker is fully consistent with expected utility maximization.

1.5 Definition of the Central Banker's Loss Function

The congruence result of Section 1.4 readily generalizes to the case of multiple objectives. As the FOMC statement in the Introduction makes clear, the Federal Reserve's loss function can be stated in terms of inflation and a measure of the output gap. Let π_t denote the inflation rate and x_t the output gap at date t. Then the period loss function of the central banker can be written as $L_t = L(\pi_t, x_t)$. We focus on a specific parameterization of this loss function that offers two advantages. First, it provides a simple, yet flexible parameterization of preferences that nests as a special case the standard quadratic and symmetric central bank loss function commonly used in the macroeconomic literature (see, e.g., Blinder 1997, Clarida et al. 1999, Svensson 2002, Levin, Wieland, and Williams 2003). Second, as shown below, this class of loss functions ensures congruence with the risk management model and allows us to link the expected loss to commonly used statistical measures of risk.

Definition 5 [Loss function of the central banker]:

$$L_t = \left[aI(\pi_t < \underline{\pi})(\underline{\pi} - \pi_t)^\alpha + (1-a)I(\pi_t > \bar{\pi})(\pi_t - \bar{\pi})^\beta\right]$$
$$+ c\left[bI(x_t < \underline{x})(\underline{x} - x_t)^\gamma + (1-b)I(x_t > \bar{x})(x_t - \bar{x})^\delta\right],$$

where $0 \leq a, b \leq 1$ *and* $c \geq 0$.

The parameter c measures the relative weight of the inflation and output objective, and the parameters a and b measure the relative weight of the downside risk relative to the upside risk of inflation and the gap, respectively. Implicit in this definition is the standard assumption that the objectives of the central banker are additively separable, although the economic variables in question are jointly dependent. The loss function also postulates that the central banker is only concerned with realizations for inflation and the output gap that exceed some predetermined set of thresholds $(\underline{\pi}, \bar{\pi}, \underline{x}, \bar{x})$, where $\underline{\pi} \leq \bar{\pi}$ and $\underline{x} \leq \bar{x}$. Finally, the loss function states that the loss occurred in the event that a threshold is exceeded can be written as a power of the deviations from the threshold, where $\alpha \geq 0$, $\beta \geq 0$, $\gamma \geq 0$, and $\delta \geq 0$ are measures of the degree of risk aversion. The loss function allows for linear and quadratic loss as special cases, but the parameters are not restricted to integer values.

This loss function is general in that it allows the central banker to aim for inflation within a zone defined by $\underline{\pi}$ and $\bar{\pi}$ or aim for a specific inflation rate of $\bar{\pi} = \underline{\pi} = \pi^*$. Similarly, the output gap objective could be specified as a range or a specific value. A natural choice for the output gap thresholds is $\underline{x} = \bar{x} = x^* = 0$. The choice of $\underline{\pi}$ and $\bar{\pi}$ may or may not have been made explicit by the central bank. For example, the European Central Bank (π^* close to, but below 2%), the Bank of Canada ($\underline{\pi} = 1\%, \bar{\pi} = 3\%$), the Swedish Riksbank ($\underline{\pi} = 1\%, \bar{\pi} = 3\%$), the Reserve Bank of Australia ($\underline{\pi} = 2\%$, $\bar{\pi} = 3\%$), the Reserve Bank of New Zealand ($\underline{\pi} = 1\%$, $\bar{\pi} = 3\%$), and the Bank of England ($\pi^* = 2\%$) in stating their price stability objective explicitly make reference to such bounds on inflation, whereas the Bank of Japan and the Federal Reserve do not, although some Fed economists have discussed plausible ranges for the inflation objective.[6] In the empirical section, we will impose the assumption that the Federal Reserve aims for an inflation rate of $\bar{\pi} = \underline{\pi} = \pi^*$ consistent with the existing literature on policy rules, although our methodology would allow us to relax that assumption. We will set π^* in accordance with statements by policymakers.

Observing that the expected loss can be written as

$$E(L_t) = \left[a \int_{-\infty}^{\underline{\pi}} (\underline{\pi} - \pi_t)^{\alpha}\, dF_{\pi}(\pi_t) + (1-a) \int_{\bar{\pi}}^{\infty} (\pi_t - \bar{\pi})^{\beta}\, dF_{\pi}(\pi_t)\right]$$
$$+\, c\left[b \int_{-\infty}^{\underline{x}} (\underline{x} - x_t)^{\gamma}\, dF_x(x_t) + (1-b) \int_{\bar{x}}^{\infty} (x_t - \bar{x})^{\delta}\, dF_x(x_t)\right]$$

it follows immediately that congruence holds because

$$E(L_t) = [-aDR_{\alpha} + (1-a)EIR_{\beta}] + c[-bNGR_{\gamma} + (1-b)PGR_{\delta}].$$

6. For example, Bernanke (2003b) states that inflation between 1% and 2% per year is probably the *de facto* equivalent of price stability. See Section 4 for a more detailed discussion of the Federal Reserve's inflation objective and additional evidence.

This result illustrates that for the proposed loss function the risks as defined above are proportionate to the expected loss. Minimizing a suitable weighted average of these risks will also minimize the expected loss. This framework suggests that it makes sense for the central banker to be preoccupied with the upside and downside risks to inflation and to the gap, and that the language chosen by the Board may be given an economic interpretation that is consistent with standard economic models. We will make use of this feature in Section 2, when we derive a policy rule based on these measures of risk.

1.6 Interpretation of the Risk Measure in Special Cases

For appropriate choices of the risk aversion parameters, our general measure of risk reduces to measures of risk proposed in the literature in other contexts. For concreteness we focus on the upside and downside risks to inflation. First, consider the limiting case of $\alpha = \beta = 0$. In this case, we obtain:

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

Thus, the general measures of risk simply to reduce 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. This set of preferences would apply to a central banker concerned about avoiding at all cost the unpleasant task of having to report to parliament a failure to attain the inflation target (as would be required in some countries) or equivalently a central banker whose job contract stipulates that he will be fired in case inflation exceeds the target.

Although instructive, this limiting case is implausible in that most central bankers 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, Fed Governor Bernanke (2003b) notes that very low inflation and deflation pose qualitatively similar economic problems, 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 Fed preference function.

Second, a simple counterexample illustrates that few central bankers would be indifferent to whether inflation misses the threshold by a small amount or by a large amount. Consider a threshold of 2% and suppose that a central banker faces the

7. This interpretation of risk has been used in the press. For example, the *Financial Times* (2003) implies that in May 2003 the risk of deflation in the United Kingdom was low because the Bank of England's "forecasts imply only a one-in-twenty chance that inflation will be below 1.3% in 2 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.

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 policymakers 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 central bankers 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. 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 2003). 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

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

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.[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%.

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

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

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

a concept familiar from finance. Here the central banker 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, consistent with the objective of "low and stable inflation."

These three examples illustrate that the proposed risk measure includes as special cases many of the statistical measures of risk proposed in the academic literature such as the probability of exceeding a threshold, the tail conditional expectation or the target semivariance. There is of course no reason why the risk aversion parameters $\alpha,\ \beta,\ \gamma$, and δ should be restricted to integers of 0, 1, or 2, as in these examples. More generally, the proposed risk measure can be evaluated for any value of the risk aversion parameter by numerical methods. While these and other parameters of the loss function are presumably known to the central banker (or could be elicited from the central banker by appropriate questionnaires), they are in general unknown to the econometrician and to the general public. Without these parameter values, the measures of risk implied by the central banker's loss function cannot be computed. In Section 3, we will discuss methods for estimating these and other parameters of the central banker's loss function from observables. For now we will treat these parameters as known.

2. A MODEL OF CENTRAL BANK RISK MANAGEMENT

2.1 The Central Banker's Decision Problem

Suppose the following conditions are satisfied:

- (i) The objective of the central bank is to minimize $E_0 \sum_{t=0}^{\infty} \phi^t L_t$ by choosing the current interest rate, i_0, and the sequence of future interest rates, i_t, $t > 0$, where ϕ is a discount factor and the loss function L_t is given by Definition 5.
- (ii) The infinite horizon minimization problem in (i) can be decomposed into a series of separate one-period problems.
- (iii) The output gap and inflation processes can be broken up into a conditional mean component and a conditional innovation whose distribution does not depend on the interest rate.
- (iv) The central banker is risk averse (i.e., $\alpha > 1,\ \beta > 1,\ \gamma > 1,\ \delta > 1$).

Under these conditions, we can use the change-of-variable theorem and differentiation under the integral sign to obtain the following first-order condition at time t:

$$\frac{\partial E_t L_t(\pi_t, x_t)}{\partial i_t} = \frac{\partial E_t[\pi_t(i_t)]}{\partial i_t}\left[-a\alpha \int_{-\infty}^{\underline{\pi}} (\underline{\pi} - \pi_t)^{\alpha-1}\, dF_\pi(\pi_t) + (1-a)\beta \int_{\bar{\pi}}^{\infty} (\pi_t - \bar{\pi})^{\beta-1}\, dF_\pi(\pi_t)\right] + \frac{\partial E_t[x_t(i_t)]}{\partial i_t} c\left[-b\gamma \int_{-\infty}^{\underline{x}} (\underline{x} - x_t)^{\gamma-1}\, dF_x(x_t) + (1-b)\delta \int_{\bar{x}}^{\infty} (x_t - \bar{x})^{\delta-1}\, dF_x(x_t)\right] = 0. \quad (1)$$

Note that conditions (i) to (iv) are typically satisfied in standard settings used to derive Taylor rules. For example, in a New Keynesian, forward-looking sticky price model of the type recently discussed by Clarida et al. (1999), the economy is characterized by a Phillips Curve trade-off:

$$\pi_t = \lambda x_t + \xi E_t \pi_{t+1} + u_t \quad (2)$$

and an IS equation:

$$x_t = -\varphi(i_t - E_t \pi_{t+1}) + E_t x_{t+1} + g_t, \quad (3)$$

where i_t denotes the nominal interest rate, x_t denotes the output gap, and π_t denotes inflation. The terms u_t and g_t denote a cost shock and a demand shock, respectively, which for expository purposes are assumed to be zero mean and mutually and temporally independent. $E_t\pi_{t+1}$ and $E_t x_{t+1}$ denote the expected values of π_{t+1} and x_{t+1}, respectively, given information available in period t.

In this model, the central bank sets the nominal interest rate at the beginning of each time period. After this decision, realizations of the cost and demand shocks are drawn and agents make their optimal choices. Because neither the economic structure nor the objective function of the central bank contains lagged variables and since monetary policy is discretionary, the current interest rate setting does not affect expectations of future periods' inflation rates and output gaps. By iterating equations (2) and (3) forward and taking expectations, it can be shown that $E_t\pi_{t+1}$ and $E_t x_{t+1}$ only depend on future expected values, not on today's interest rate choices:

$$E_t \pi_{t+1} = \lambda E x_{t+1} + \xi E \pi_{t+2} + E_t u_{t+1}$$
$$E_t x_{t+1} = -\varphi(E_t i_{t+1} - E_t \pi_{t+2}) + E_t x_{t+2} + E_t g_{t+1}.$$

Thus, the central bank treats $E_t\pi_{t+1}$ and $E_t x_{t+1}$ in equations (2) and (3) as given, when setting the interest rate. This fact allows the central bank to decompose the infinite horizon minimization problem into a series of separate one-period problems (see, e.g., Cukierman and Muscatelli 2003). In each period *t,* the central bank chooses an interest rate that minimizes the expected loss:

$$\min_{i_t} E_t L_t(\pi_t, x_t) \quad t = 0, 1, 2, \ldots$$

where the date *t* expectation is taken with respect to the shocks u_t and g_t in equations (2) and (3). It may seem odd at first to treat $L_t(\pi_t, x_t)$ as a random variable in period *t*. The reason is that in this model the central banker does not know the realizations of π_t and x_t when setting the interest rate at date *t*.

In this New Keynesian framework, the output gap and inflation processes can be decomposed into a conditional mean component and a conditional innovation whose distribution does not depend on the interest rate:

$$\pi_t(i_t) = E_t\left[\pi_t(i_t)\right] + u_t + \lambda g_t$$
$$x_t(i_t) = E_t\left[x_t(i_t)\right] + g_t.$$

This feature follows directly from the reduced form representation of equations (2) and (3) but is consistent with other specifications of the New Keynesian model as well. The reduced form equations show that a change in the interest rate affects the first moment of the random variables of interest, but not the higher-order moments. As in Clarida et al. (1999), we do not allow the interest rate to affect higher moments of the distribution of the macroeconomic aggregates. This assumption is required to keep the analysis tractable.

2.2 *The Balance of Risks*

As discussed earlier, the FOMC stresses the need to balance the upside risks and the downside risks to price stability. This balance of risks can be thought of as a weighted average of the risks in Definitions 1a and 1b. Indeed, it is possible to recast the central banker's decision problem in terms of a suitably defined measure of the balance of risks, by rewriting the first-order condition in equation (1) as follows:

$$BR_t\left(i_t^*\right) \equiv \left[\alpha a DR_{t,\alpha-1}\left(i_t^*\right) + \beta(1-a)EIR_{t,\beta-1}\left(i_t^*\right)\right] + \tilde{c}\left[\gamma bNGR_{t,\gamma-1}\left(i_t^*\right) + \delta(1-b)PGR_{t,\delta-1}\left(i_t^*\right)\right] = 0, \tag{4}$$

where $\tilde{c} \equiv c[\partial E_t[\pi_t(i_t^*)]/\partial i_t]^{-1}[\partial E_t[x_t(i_t^*)]/\partial i_t]$ is a constant in the New Keynesian model. The optimal interest rate i_t^* is implicitly defined by this first-order condition.

Equation (4) provides a condition for optimal monetary policy expressed in terms of a weighted average of risk measures. It characterizes the need for interest rate changes in terms of the balance of risks. A positive balance would call for an increase in interest rates, whereas a negative balance would call for a reduction. 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. For example, a balance of risk of zero means that the central banker would not want to reduce one risk, say that of deflation, if doing so means disproportionately increasing another risk. It does not mean that the central banker would not prefer a reduction in both upside and downside risks, if given the choice.

2.3 Derivation of a Policy Rule in Terms of the Balance of Risks

Following the Taylor rule literature, we postulate that there will be some interest rate smoothing, so that the first-order condition in equation (4) is best viewed as implying a desired interest rate, i_t^*, to which the central bank is assumed to converge through a partial adjustment mechanism (see, e.g., Clarida et al. 2000):

$$i_t = \rho\, i_{t-1} + (1-\rho)\, i_t^*. \tag{5}$$

Evaluating the balance of risks at i_t (the current interest rate) and taking a first-order Taylor expansion around i_t^* (the optimal interest rate in the current period), we obtain

$$BR_t(i_t) \approx BR_t\left(i_t^*\right) + \left.\frac{\partial BR_t(i)}{\partial i}\right|_{i=i_t^*} \left(i_t - i_t^*\right). \tag{6}$$

Note that equation (5) implies that $(i_t - i_t^*) = \frac{-\rho}{1-\rho}(i_t - i_{t-1})$. Under our assumptions, after substituting for $(i_t - i_t^*)$ in equation (6), and observing that $BR_t(i_t^*) = 0$ by the first-order condition, we obtain the following relationship between the balance of risks and the change in interest rates:

$$\Delta i_t \approx kBR_t(i_t), \tag{7}$$

where $k \equiv -\frac{1-\rho}{\rho}[\frac{\partial BR_t(i)}{\partial i}|_{i=i_t^*}]^{-1}$.

Equation (7) is a nonlinear function of i_t. It implicitly defines the optimal interest rate setting under interest rate smoothing. Thus, in the risk management model of monetary policy, finding the optimal interest rate involves solving the policy rule in equation (7).

Unlike conventional policy rules, the structure of the rule in equation (7) closely mimics the way in which the Federal Reserve Board communicates its policy decisions. For example, the January 31, 2006, press release of the Federal Reserve directly links the change in the interest rate to the balance of risks (emphasis added):

> The Federal Open Market Committee decided today to *raise its target for the federal funds rate* by 25 basis points to 4–1/2 percent. [. . .] The Committee judges that some *further policy firming may be needed to keep the risks to the attainment of both sustainable economic growth and price stability roughly in balance*.

2.4 Special Case with Quadratic Symmetric Preferences

Our analysis so far has not taken a stand on whether preferences are quadratic or symmetric. We only imposed that the risk aversion parameters are greater than unity. If we were to add the assumption of quadratic preferences ($\alpha = \beta = \gamma = \delta = 2$) and symmetric preferences ($a = b = 0.5$) and the assumption that $\underline{\pi} = \bar{\pi} = \pi^*$ and $\underline{x} = \bar{x} = x^* = 0$, certainty equivalence would hold and the balance of risks would become a function of the deviations of the conditional means from their respective targets:

$$\begin{aligned} BR_t(i_t) &= [DR_{t,1}(i_t) + EIR_{t,1}(i_t)] + \tilde{c}[NGR_{t,1}(i_t) + PGR_{t,1}(i_t)] \\ &= -\int_{-\infty}^{\pi^*} \left(\pi^* - \pi_t\right) dF_\pi(\pi_t) + \int_{\pi^*}^{\infty} \left(\pi_t - \pi^*\right) dF_\pi(\pi_t) \\ &\quad + \tilde{c}\left[\int_{-\infty}^{0} (-x_t)\, dF_x(x_t) + \int_{0}^{\infty} x_t\, dF_x(x_t)\right] \\ &= \int_{\infty}^{\infty} \left(\pi_t - \pi^*\right) dF_\pi(\pi_t) + \tilde{c}\int_{-\infty}^{\infty} x_t\, dF_x(x_t) \\ &= \left(E_t[\pi_t] - \pi^*\right) + \tilde{c}E_t\,[x_t]\,. \end{aligned}$$

This implies a policy rule of the form:

$$\Delta i_t \approx k\left[E_t\left[\pi_t(i_t)\right] - \pi^*\right] + k\tilde{c}E_t\left[x_t(i_t)\right]. \tag{7'}$$

This result is intuitive, in that it suggests that the Federal Reserve increases interest rates when inflation is expected to be above the desired target and/or when the economy is expected to grow more than its potential. Note that the derivation of equations (7) and (7′) has not explicitly utilized equations (2) and (3). If we substitute the conditional expectations of (2) and (3), i.e., $E_t[\pi_t(i_t)] = \lambda E_t[x_t] + \xi E_t[\pi_{t+1}]$ and $E_t[x_t(i_t)] = -\varphi(i_t - E_t[\pi_{t+1}]) + E_t[x_{t+1}]$, into equation (7′), we obtain the standard Taylor rule with interest rate smoothing:

$$i_t = c_0 + c_1 i_{t-1} + c_2 E_t\left[\pi_{t+1}\right] + c_3 E_t\left[x_{t+1}\right], \tag{8}$$

where $c_0 \equiv -k\omega\pi^*$, $c_1 \equiv \omega$, $c_2 \equiv \omega k(\lambda\varphi + \xi + \tilde{c}\varphi)$, $c_3 \equiv \omega k(\lambda + \tilde{c})$, and $\omega \equiv [1 + k\varphi(\lambda + \tilde{c})]^{-1}$.[9]

9. Note that the derivation of the policy rules in equations (7) and (7′) at no point imposes that $\rho = 1$. In fact, by inspection, $\rho = 1$ would imply the counterfactual results that $k = 0$ and hence $\Delta i_t = 0\,\forall t$. Nor does the result in equation (7′) imply that the coefficient on the lagged interest rate in (8) is unity. That coefficient will be unity if and only if $\rho = 1$ and hence $k = 0$, as may be verified by exploiting the definition of k given below equation (5) and substituting into the definition of ω.

3. NONLINEAR GENERALIZED METHOD OF MOMENTS (GMM) ESTIMATION

In the empirical analysis, we impose that $a = b = 0.5$ for simplicity, which conforms to the standard assumption in the Taylor rule literature. Setting these weight parameters to 0.5 facilitates the comparison of the estimated risk aversion parameters with the standard quadratic symmetric loss function. In contrast, the relative weight of the inflation and output objectives remains unrestricted in the empirical analysis. Although our methodology allows for point as well as zone targets, in the empirical analysis of this paper we focus on point targets, as point targets seem more appropriate than zone targets for characterizing U.S. monetary policy. Specifically, we set $\underline{\pi} = \bar{\pi} = \pi^*$ and $\underline{x} = \bar{x} = x^* = 0$. This specification and our choice of π^* in the empirical application are motivated by the discussion of Fed objectives in Section 4.1 below. We estimate the remaining unknown parameters $\theta \equiv [\alpha, \beta, \gamma, \delta, \tilde{c}, k]'$ by iterated nonlinear GMM applied to equation (7). These parameters are all contained in the expression for $BR_t(i_t)$.

Estimation of the parameters in equation (7) thus requires that we first construct a measure of the balance of risks, conditional on a starting value for the parameter vector $\theta \equiv [\alpha, \beta, \gamma, \delta, \tilde{c}, k]'$. Note that we can equivalently rewrite relation (7) as a nonlinear regression model:

$$\Delta i_t = kE_t z_t + \tilde{\varepsilon}_t = kz_t + \varepsilon_t,$$

where $E_t z_t \equiv BR_t(i_t)$ and $\tilde{\varepsilon}_t$ is an exogenous control error and $\varepsilon_t \equiv \tilde{\varepsilon}_t + k(E_t z_t - z_t)$. This last equation makes it clear that we do not need to observe FOMC expectations to estimate the preference parameters. The parameters of interest can be estimated based on the *ex post* realizations of the balance of risks. The realizations of the individual risk measures may be easily computed from the unconditional distribution of inflation and output gap realizations over the 1987.III–2005.IV period. For expository purposes, consider the first element of the balance of risks: the realization of the upside risk of inflation, EIR_β is simply $I(\pi_t > \pi^*)(\pi_t - \pi^*)^\beta$. The realizations of the other three risk measures are computed analogously.

An important concern is that the GMM estimates of the parameters in equation (7) may be sensitive to the choice of moment conditions. Our candidate moment conditions involve a constant and $n \in \{4, 6, 8\}$ lags of suitably chosen powers $p \in \{0.5,\ 1,\ 1.5\}$ of the change in the Federal Funds rate, of the output gap, and of the deviation of the inflation rate from its target value. In total, this implies six alternative sets of moment conditions, each corresponding to a different model. Each model is estimated based on 25,000 independent random draws for the starting conditions and only the result that minimizes the GMM criterion is retained. In some cases, estimated models may violate the premise of risk aversion due to sampling error. We discard those model estimates. Finally, we select the preferred model among the set of these alternative model estimates based on a recently proposed information criterion (see Hall et al. 2007). This *relevant moment selection criterion* (RMSC) is based on the entropy of the limiting distribution of the GMM estimator. It is consistent under weak conditions, even when the parameters of interest are weakly identified for some combinations of the candidate moment conditions.

4. EMPIRICAL APPLICATION

In this section, we will apply the estimation methodology of Section 3 to data for the U.S. economy for the Greenspan period of 1987.III–2005.IV. Our objective is to estimate the parameters of the loss function that captures the FOMC's preferences during this period, most notably the degree of risk aversion to departures from the objectives of price stability and maximum sustainable employment.

4.1 Applicability of the Risk Management Model in the Greenspan Period

An implicit assumption in our empirical analysis is that the preferences of the Federal Reserve under Greenspan have remained constant since 1987. We carefully chose this period because the continuous leadership of one chairman suggests that Fed preferences over this period were as stable as they ever will be. This assumption is perhaps not without controversy but provides a reasonable working hypothesis in the absence of explicit statements by the Federal Reserve that would document a change in preferences.[10] If such a change were documented, on the other hand, it would be straightforward to incorporate this structural change into the analysis.

An important question is whether the risk management model of Section 2 is appropriate for the Greenspan period. It is important to be clear that we do not treat the Fed as an inflation targeter in the strict sense of the word. The defining characteristic of "inflation targeting" as practiced around the world is a hierarchical mandate that ensures that other objectives such as full employment can only be pursued after price stability has been achieved. In contrast, the Fed has a "dual" mandate, set by Congress, to address both inflation and full employment, which is reflected in our loss function (see Meyer 2004, p. 204).

We do, however, treat the Fed as though it were setting formal numerical targets for both inflation and for the output gap. Specifically, we impose a numerical target of zero for the output gap and an inflation target, π^*. It may be objected that the Fed has never announced such a target. Rather Greenspan has always insisted that such a numerical target would be undesirable. This does not mean that the Fed does not effectively have such a target.[11] In fact, Mankiw (2006) refers to Greenspan's policy as "covert inflation targeting."

10. For example, Erceg and Levin (2003) consider a model in which the Fed inflation target is time varying, but they do not formally test this proposition. Moreover, their analysis is conditional on the maintained assumption of quadratic and symmetric preferences.

11. There have been other instances in which the Fed has been deliberately obscure about its objectives. For example, in the years before Greenspan, the FOMC did not admit to having a target for the Federal Funds rate and did not publicly announce any changes in its target.

Greenspan's (2002) view is that "a specific numerical inflation target would represent an unhelpful and false precision," given the conceptual uncertainties and problems in measuring the general price level. In short, if policy decisions are based on faulty price indices, the Federal Reserve may inadvertently allow inflation to destabilize an economy. Greenspan, however, has also made it clear that the Federal Reserve nonetheless has been using traditional measures of price stability, and that "for the time being our conventional measures of the overall price level will remain useful" (see Greenspan 2002). Moreover, his opposition to numerical targets has not prevented Greenspan (2004, p. 35) from asserting that in mid-2003 "our goal of price stability was achieved by most analysts' definition." At that time CPI inflation at annual rates was running at about 2.1% and CPI inflation excluding food and energy at about 1.5%.

Several Fed economists including Governors Bernanke and Meyer have publicly favored the adoption of an explicit numerical objective for inflation. Most of the resistance from within the Fed to numerical inflation targets has focused on the danger that by announcing a numerical target for inflation, but not for employment, the Fed could over time overemphasize the inflation objective at the expense of the employment objective (see, e.g., Kohn 2003).[12] In this view, it is not the existence of a numerical target *per se* that is problematic but rather its public announcement. In fact, many observers feel that a possible policy shift toward announcing a numerical inflation target (while maintaining the dual objective) under the current chairman Ben Bernanke would be more of a change in the way the Fed communicates its objectives than a substantive change (see, e.g., Mankiw 2006, p. 183).

There is evidence that already many FOMC members have effectively used such numerical objectives. For example, former Fed Governor Laurence Meyer observes that at the July 1996 FOMC meeting, a majority agreed that a 2% inflation rate was consistent with its "price stability" mandate, although no target was ever discussed or announced given Greenspan's opposition (see Meyer 2004). It is less clear which inflation measure this definition of price stability referred to. According to Meyer, most members appeared to be talking about the core CPI measure. Some had in mind the core price index of PCE. Meyer's interpretation is that the consensus was for a 2% inflation rate for the core CPI, which would be consistent with a 1.5% rate for the core PCE index (see Meyer 2004, p. 237). These numbers also appear consistent with the earlier quote by Greenspan (2004). Thus, it makes sense to impose a numerical inflation objective of 1.5% core PCE inflation and 2% core CPI inflation in estimating the Federal Reserve's preference parameters in the Greenspan era.

Beyond the evidence that the Federal Reserve has an implicit inflation objective, there is more direct evidence that the Federal Reserve under Greenspan followed a risk management approach consistent with our econometric framework. Greenspan himself has repeatedly stressed that U.S. monetary policy under his chairmanship is

12. Note that we disagree with Kohn's (2003) assertion that even an implicit inflation target implies undue emphasis on the inflation objective. In the context of our loss function, the respective weights of inflation and the output gap are pinned down by the preferences of the central banker.

not well captured by a traditional Taylor rule. Such Taylor rules are predicated on a symmetric and quadratic loss function that implies that the conditional mean forecast is a sufficient statistic for policy analysis.

According to Meyer (2004), while Greenspan's vision for monetary policy—anchoring inflation expectations and responding aggressively to departures from full employment—is fully consistent with the traditional Taylor rule, his risk management approach emphasizes the importance of more flexible approaches. Specifically, Greenspan urged policymakers to expand their horizons to the full probability distribution of outcomes rather than limiting their focus to the outcome that seems most likely. Citing the example of the disinflation of 2002 and 2003, Meyer (2004) suggests that the FOMC at times cut rates by more than what the baseline forecast suggested, consistent with the view that it is asymmetric risks—rather than the most likely forecast outcome—that determine monetary policy. In Greenspan's own words, "The conduct of monetary policy in the United States at its core involves crucial elements of risk management, a process that requires an understanding of the many sources of risk . . . that policymakers face and the quantifying of those risks when possible" (see Greenspan 2003, 2004).[13]

The approach pursued in this paper formalizes this risk management paradigm and allows the estimation of the Fed's preferences parameters by nonlinear GMM. One advantage of this approach to estimating the Federal Reserve's preferences under Greenspan is that it allows for judgmental elements in the FOMC's forecasts of inflation and the output gap. A second important advantage of our approach is that it does not require any explicit forecasts, but only *ex post* realizations of inflation rates and output gaps. The key premise of our approach is that on average the Federal Reserve should be able to control inflation and output gap outcomes. This means that we do not rely on either econometric forecasts or Greenbook forecasts in our analysis. The latter forecasts are but one input into the FOMC's decision process and reflect the views of the Federal Reserve's staff as opposed to the FOMC (see Meyer 1998 for extensive discussion of this point). This fact makes them ill-suited for estimating the Federal Reserve's loss function parameters.

4.2 Data

The sample period is 1987.III–2005.IV. The estimation of equation (7) is based on seasonally adjusted data for the U.S. core CPI for all urban consumers from the Bureau of Labor Statistics or, alternatively, for the core PCE deflator from the National Income and Product Accounts. We also use monthly data for the Federal Funds rate from the Board of Governors and the Chicago Fed National Activity Index (CFNAI). The CFNAI is a principal components estimate of the output gap prepared by the

13. Note that the risk management paradigm outlined by Greenspan assumes that the probability distribution of outcomes is known (or can be consistently estimated as in our paper). It cannot be extended to deal with "Knightian uncertainty," i.e., situations in which the probability distribution of outcomes will forever remain unknown (also see Greenspan 2004, p. 36).

TABLE 1

GMM Estimates of Central Bank Loss Function Parameters Using Core CPI (with Newey–West Standard Errors)

	Model parameters							
	a	b	α	β	γ	δ	$\tilde{c}$	k
X_t ($n = 4$, $c = 1.5$)	0.5	0.5						
Estimate			7.01	3.20	1.94	3.99	21.08	0.05
SE			(2.89)	(0.53)	(0.16)	(1.29)	(15.50)	(0.04)
X_{t+1} ($n = 8$, $c = 0.5$)	0.5	0.5						
Estimate			8.22	1.10	1.39	3.31	24.56	0.02
SE			(2.42)	(1.16)	(0.07)	(0.54)	(24.22)	(0.02)

Notes: The moment conditions have been selected by the relevant moment selection criterion (RMSC) of Hall et al. (2007), allowing for powers of $c \in \{0.5, 1.5\}$ and for lags of $n \in \{4, 6, 8\}$. Each model has been estimated based on 25,000 random draws for the starting conditions. The results shown are for the estimate with the lowest GMM criterion value.

Federal Reserve Bank of Chicago.[14] We convert the monthly measures of the output gap to quarterly frequency. Our principal components approach is consistent with recent work on monetary policy reaction functions and measures of real economic activity (see, e.g., Evans, Liu, and Pham-Kanter 2002, Bernanke and Boivin 2003).

4.3 Estimating the Preferences of the Federal Reserve over the Greenspan Period

It is instructive to begin with a review of the point estimates of the four risk aversion parameters. How averse was the Fed under Greenspan to deflationary and to inflationary risks? How did it respond to risks of positive and negative output gaps? The answers depend to some extent on whether we think of policy as being framed in terms of core CPI inflation or core PCE inflation. Here we present results for both inflation concepts.

Table 1 shows the parameter estimates for the version of equation (7) based on the core CPI with $\pi^* = 2\%$. As there is some disagreement on whether the Fed responds to current or future output gaps, as evidenced by alternative specifications of Taylor rules in the literature, we present two sets of results. Starting with the equation involving the current output gap, x_t, Table 1 shows a point estimate for γ that is slightly below 2, whereas the estimates for α, β, and δ all exceed 2. These point estimates suggest that the Federal Reserve under Greenspan was highly risk averse with respect to deflationary developments, and somewhat less risk averse with respect to inflationary developments. On the output side, in contrast, there is more risk aversion toward positive gaps than negative gaps. Qualitatively similar results hold for x_{t+1}, except that the aversion to high inflation is much weaker.

Table 2 shows the corresponding results for the core PCE inflation measure favored by Greenspan with $\pi^* = 1.5\%$. The results using x_t in the policy rule are qualitatively

14. The CFNAI data are publicly available at: http://www.chicagofed.org/economic_research_and_data/files/data_series.xls.

TABLE 2

GMM Estimates of Central Bank Loss Function Parameters Using Core PCE (with Newey–West Standard Errors)

	Model parameters							
	a	b	α	β	γ	δ	$\tilde{c}$	k
X_t ($n = 6, c = 1.5$)	0.5	0.5						
Estimate			4.61	1.17	1.92	3.36	8.90	0.05
SE			(1.42)	(1.39)	(1.18)	(0.88)	(12.03)	(0.07)
X_{t+1} ($n = 8, c = 0.5$)	0.5	0.5						
Estimate			1.08	1.09	1.52	3.21	9.42	0.09
SE			(0.96)	(0.43)	(0.10)	(0.63)	(6.05)	(0.06)

Note: See Table 1.

similar to those in Table 1, except that there is less risk aversion against deflation and much less risk aversion against high inflation. The parameters for the output gap are almost identical to those in Table 1. In contrast, using x_{t+1} in the policy rule, the risk aversion toward deflation is reduced to close to 1, making the estimates of α and β nearly symmetric. On the other hand, the parameters for the output gap are almost identical to those in Table 1 with higher aversion to positive gaps than negative ones.

4.4 Testing the Assumption of Quadratic Symmetric Preferences

Given the sampling uncertainty in these estimates we proceed with formal statistical tests of various restrictions on the loss function parameters. The first row of Table 3 shows that we cannot reject at standard significance levels the hypothesis that the Fed has no output gap objective (H_0: $c = \tilde{c} = 0$). In other words, we cannot reject the hypothesis that the Federal Reserve was an inflation targeter in the narrow sense. Of course, this nonrejection may reflect low power of the test and does not necessarily mean that the null is true. Indeed, the strong emphasis the Fed puts on the output objective suggests otherwise. Since erroneously imposing $c = 0$ would result in inconsistent estimates, we proceed with the baseline specification.

The next seven rows consider various implications of the hypotheses of quadratic preferences for output and inflation. Table 3 shows that the joint null of quadratic symmetric preferences (H_0: $\alpha = \beta = \gamma = \delta = 2$) is rejected for all specifications at any reasonable significance level. This result is strong evidence in favor of the risk management paradigm. It is of some interest to examine whether this rejection is driven by the inflation objective or the output objective (or both). Table 3 shows that for two of the four models we reject the null of quadratic preferences in inflation (H_0: $\alpha = \beta = 2$) at the 5% level, with a third model marginally rejecting at the 10% level. The null of quadratic preferences in output (H_0: $\gamma = \delta = 2$) is rejected at the 1% level for the models involving x_{t+1} in the policy rule, but not for the models involving x_t. This leaves the possibility that at least some of the risk aversion parameters may

TABLE 3

SELECTED t- AND WALD-TEST STATISTICS FOR THE PREFERRED MODELS (WITH p-VALUES)

		CPI, X_t	CPI, X_{t+1}	PCE, X_t	PCE, X_{t+1}
Output weight	H_0: $c = \tilde{c} = 0$	1.85 (0.17)	1.03 (0.31)	0.55 (0.46)	2.43 (0.12)
Quadratic preferences	H_0: $\alpha = \beta = \gamma = \delta = 2$	**13.26** **(0.01)**	**117.99** **(0.00)**	**19.20** **(0.00)**	**35.73** **(0.00)**
	H_0: $\alpha = \beta = 2$	5.26 (0.72)	**7.33** **(0.03)**	**8.23** **(0.02)**	4.58 (0.10)
	H_0: $\gamma = \delta = 2$	2.39 (0.30)	**97.99** **(0.00)**	2.38 (0.30)	**25.43** **(0.00)**
	H_0: $\alpha = 2$	**2.99** **(0.08)**	**6.62** **(0.00)**	3.40 (0.06)	0.91 (0.34)
	H_0: $\beta = 2$	**5.06** **(0.02)**	0.59 (0.44)	0.36 (0.55)	**4.37** **(0.04)**
	H_0: $\gamma = 2$	0.13 (0.72)	**84.00** **(0.00)**	0.19 (0.66)	**23.42** **(0.00)**
	H_0: $\delta = 2$	2.39 (0.12)	**5.99** **(0.01)**	2.37 (0.12)	**3.74** **(0.05)**
Symmetric preferences	H_0: $\alpha = \beta$	2.13 (0.14)	**7.19** **(0.01)**	**7.82** **(0.005)**	0.00 (0.99)
	H_0: $\gamma = \delta$	2.37 (0.12)	**13.10** **(0.00)**	2.35 (0.12)	**6.90** **(0.01)**

NOTES: See Table 1. Rejections at the 5% significance level are in boldface.

equal 2. For α that hypothesis can be rejected at the 10% significance levels for three out of four models. At least two rejections each occur with the remaining risk aversion parameters. The last three rows of Table 3 consider the weaker assumption of symmetry without imposing quadratic loss. For two models each, we reject the hypothesis that H_0: $\alpha = \beta$ and H_0: $\gamma = \delta$ at any reasonable significance level.

These results suggest that there is strong evidence against standard loss functions used in the macroeconomic literature and against policy rules based on the assumption of quadratic symmetric preferences.

4.5 Comparing the Fit of the Policy Rule with and without Imposing Quadratic Symmetric Preferences

It is widely accepted that postwar policy behavior is well described by a standard Taylor-type rule. This raises the question of how much better the policy rule explains changes in the interest rate when the assumption of quadratic symmetric preferences is relaxed, and which aspect of the model is responsible for the improved fit. Below we show that the specification can make a difference for the implied interest rate setting. For illustrative purposes, the results shown are based on policy rules using core PCE inflation and x_{t+1}. The left panels of Figure 1 show the actual values of the changes in the Fed Funds rate as well as the fitted values implied by the policy rule with quadratic and symmetric preferences and those implied by the policy rule obtained with the freely estimated preference parameters.

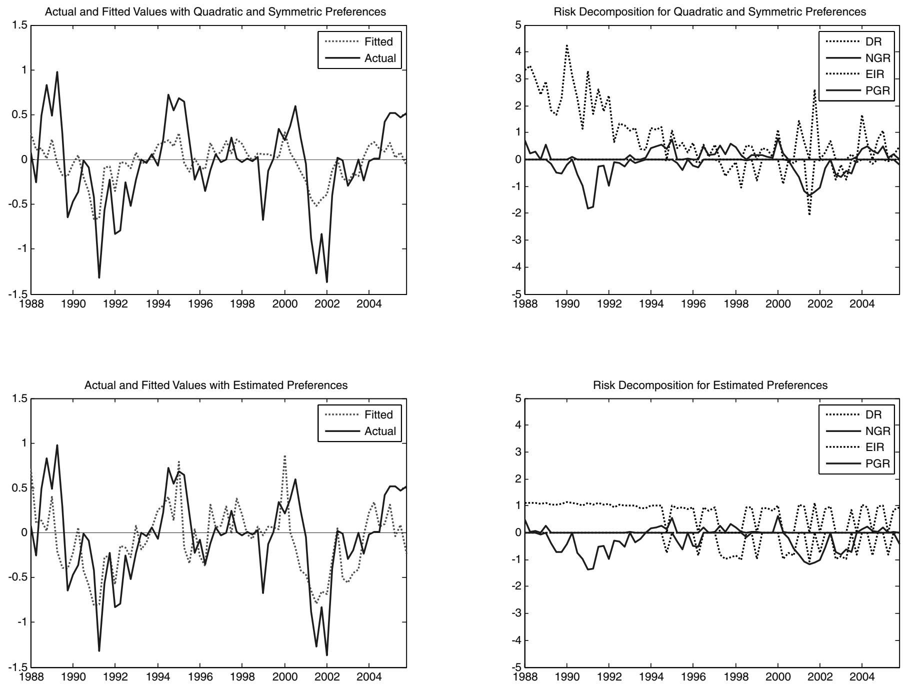

FIG. 1. Fitted Federal Funds Rate Changes Based on Core PCE Policy Rules and Actual Changes (Left Panel), Decomposition of the Risks Underlying the Policy Rules (Right Panel) (Federal Reserve Board 1987.IV–2005.III).

The right panels decompose the risks that are used in the construction of each policy rule.

Assessing the overall fit of the policy rules shown in the left panel of Figure 1 requires formal statistical analysis. Under a standard quadratic loss function, the fit of a regression model can be expressed in terms of the mean squared error (MSE). Note that we cannot simply compare the two models' MSEs because one model is nested in the other. That is why we report formal statistical tests in Table 3. A robust finding in Table 3 is that we reject the restricted model at conventional significance levels. In other words, there is not only a reduction in the MSE from relaxing the assumption of quadratic and symmetric preferences, but we are sure beyond a reasonable doubt that this reduction is not an artifact of random variation in the data. Hence, we conclude that Fed policy decisions under Greenspan are better described in terms of the Fed weighing upside and downside risks to their objectives rather than simply responding to the conditional mean of inflation and the output gap.

Statistical tests tell us whether the improvement in fit from relaxing the assumption of quadratic and symmetric preferences is precisely estimated. They do not tell us whether the improvement in fit is large enough to matter to economists. Whereas the differences in fit documented in the left panels of Figure 1 need not be large at all points in time, they can be large for some sample periods. For example, in 1995 the generalized policy rule has much better fit than the restricted rule. The improvement in fit is about 50 basis points. A move of the Fed funds rate by an extra 50 basis points seems economically significant by any reasonable standard. We conclude that not only does thinking of Fed policy decisions in terms of a risk management problem improve the fit of the policy rule on average but that doing so may distinctly improve our understanding of specific Fed policy decisions.

Figure 1 shows that there are three episodes for which the generalized policy rule has substantively different policy implications than the Taylor-type rule, resulting in a better fit: 1990–94, 1994–95, 2001–02. For example, Figure 1 shows that the generalized policy rule captures much better the sharp drop in the Fed Funds rate that occurred in 1990 and the subsequent smaller interest rate declines until 1994. These declines in the interest rate are associated with a rising risk of a negative output gap. The differential ability of the two policy rules to explain the decline in interest rates is driven not so much by differences in the magnitude of the inflation risks, although the risk of excessive inflation is much larger and volatile in the early 1990s under quadratic symmetric preferences, but by the higher weight attached to the risk of a negative output gap in the generalized policy rule. In contrast, the improved fit in the generalized policy rule in 1994–95 is driven by the higher weight attached to the risk of a positive output gap, which in turn is associated with a higher interest rate. Finally, the improved fit of the generalized policy rule in 2001–02 again primarily reflects a larger responsiveness to the risk of a negative output gap under that specification, which is associated with lower interest rates.

5. CONCLUSION

Motivated by policy statements of central bankers, we viewed the central banker as a risk manager who aims at containing inflation and the deviation of output from potential within prespecified bounds. We developed formal tools of risk management that may be used to quantify the risks of failing to attain that objective. This risk management framework helps us understand the way the Federal Reserve communicates its policy decisions. It also allows us to relax some of the assumptions implicit in the existing literature on monetary policy rules.

Risk measures inherently depend on the loss function of the user. We proposed a simple, yet flexible class of loss functions that nests the standard assumption of quadratic symmetric preferences, but is congruent with a risk management model. We showed how the parameters of this loss function under weak assumptions may be estimated from realizations for inflation and output gap data. We presented estimates of the FOMC's risk aversion parameters with respect to the inflation and output objectives during the Greenspan period. We formally tested for and rejected the standard assumption of quadratic and symmetric preference that underlies the derivation of the Taylor rule. Our results suggest that Fed policy decisions under Greenspan are better described in terms of the FOMC weighing upside and downside risks to their objectives rather than simply responding to the conditional mean of inflation and of the output gap. We derived a natural generalization of Taylor rules that links changes in the interest rate to the balance of the risks implied by the dual objective of sustainable economic growth and price stability. This generalized policy rule is much closer to the wording of policy decisions by the Federal Reserve than standard Taylor rules, it is consistent with expected utility maximization, and it has a transparent economic interpretation.

Our policy rule also has several advantages compared with previous efforts aimed at relaxing the linear specification of Taylor rules such as Ruge-Murcia (2003) and Cukierman and Muscatelli (2003). One advantage is that earlier nonlinear policy rules were *ad hoc*, whereas the rule proposed in this paper is directly motivated by the language used by policymakers. The second advantage is that the parameters of our policy rule have an economic interpretation, whereas the parameters in other nonlinear policy rules do not. While it is possible to view nonlinear policy rules based on higher-order expansions of the linear Taylor rule as an approximation to any nonlinear policy rule, including our rule, the parametric nature of the policy rule proposed in this paper not only facilitates estimation in practice, but it allows us to establish a tight link with the risk management paradigm of central banking.

Although our analysis in this paper has focused on the United States, the tools of risk management developed in this paper are quite general and could be adapted to reflect the policy objectives of other countries. The recognition that monetary policymaking involves a risk management problem has become increasingly common among central bankers since the 1990s. For example, the Bank of England in its *Inflation Report* makes explicit reference to the need to manage the risks to price stability and achieve a balance of risks. Similarly, the European Central Bank explicitly aims at setting the policy instrument to balance the risks to price stability. Issing (2002), in describing his experiences as a member of the Executive Board of the European Central Bank, stressed that "we analyse risks of deflation as well as inflation and will act accordingly to prevent both phenomena in case we detect risks in one or the other direction." An interesting avenue for future research will be a comprehensive study of central bank preferences. Even more important would be the development of operational tools that allow central bankers to formalize risk management strategies.

LITERATURE CITED

Basak, Suleyman, and Alexander Shapiro. (2001) "Value-at-Risk Based Management: Optimal Policies and Asset Prices." *Review of Financial Studies,* 14, 371–405.

Bernanke, Ben 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, Ben 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, July 23, 2003.

Bernanke, Ben S., and Jean Boivin. (2003) "Monetary Policy in a Data-Rich Environment." *Journal of Monetary Economics,* 50, 525–46.

Blinder, Alan. (1997) "What Central Bankers Could Learn from Academics—and Vice Versa." *Journal of Economic Perspectives,* 11, 3–19.

Board of Governors. (2000) FOMC, Press Release of January 19. http://www.federalreserve.gov/boarddocs/press/general/2000/20000119/default.htm.

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

Clarida, Richard, Jordi Galí, and Mark Gertler. (1999) "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature,* 37, 1661–707.

Clarida, Richard, Jordi Galí, and Mark Gertler. (2000) "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics,* 113, 147–80.

Cukierman, Alex, and Vito A. Muscatelli. (2003) "Do Central Banks Have Precautionary Demands for Expansions and for Price Stability? A New Keynesian Approach." Mimeo, Berglas School of Economics, Tel-Aviv University.

Elliott, Graham, Ivana Komunjer, and Allan Timmermann. (2005) "Estimation and Testing of Forecast Rationality under Flexible Loss." *Review of Economic Studies,* 72, 1107–25.

Elliott, Graham, Ivana Komunjer, and Allan Timmermann. (2008) "Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss." *Journal of the European Economic Association,* 6, 122–57.

Erceg, Christopher J., and Andrew T. Levin. (2003) "Imperfect Credibility and Inflation Persistences." *Journal of Monetary Economics,* 50, 915–44.

Evans, Charles, Chin T. Liu, and Genevieve Pham-Kanter. (2002) "The 2001 Recession and the Chicago Fed National Index: Identifying Business Cycle Turning Points." *Economic Perspectives,* Quarter 3, Federal Reserve Bank of Chicago, 26–43.

Financial Times. (2003) "Greenspan Plays Down the Threat of Deflation." May 22, p. 1.

Fishburn, Peter C. (1977) "Mean-Risk Analysis with Risk Associated with Below-Target Returns." *American Economic Review,* 67, 116–26.

Greenspan, Alan. (2002) "Transparency in Monetary Policy." *Federal Reserve Bank of St. Louis Review,* 84, 5–6.

Greenspan, Alan. (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. http://www.federalreserve.gov/boarddocs/speeches/2003/20030829/default.htm.

Greenspan, Alan. (2004) "Risk and Uncertainty in Monetary Policy." *American Economic Review,* 94, 33–40.

Hall, Alastair, Atsushi Inoue, Kalidas Jana, and Changmock Shin. (2007) "Information in Generalized Method of Moments Estimation and Entropy Based Moment Selection." *Journal of Econometrics,* 138, 488–512.

Holthausen, Duncan M. (1981) "A Risk-Return Model with Risk and Return Measured in Deviations from Target Return." *American Economic Review,* 71, 182–88.

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

Kilian, Lutz, and Simone Manganelli. (2007) "Quantifying the Risk of Deflation." *Journal of Money, Credit and Banking,* 39, 561–90.

Kohn, Donald L. (2003) Remarks by Governor Donald L. Kohn at the 28th Annual Policy Conference: Inflation Targeting: Prospects and Problems, Federal Reserve Bank of St. Louis, St. Louis, Missouri, October 17.

Kumar, Manmohan, Taimur Baig, J. Decressin, Chris Faulkner-MacDonagh, and Tarhan Feyzioğlu. (2003) "Deflation Determinants, Risks, and Policy Options." IMF Occasional Paper No. 221, June 2003, International Monetary Fund, Washington, DC.

Levin, Andrew T., Volker Wieland, and John C. Williams. (2003) "The Performance of Forecast-Based Mone- tary Policy Rules under Model Uncertainty." *American Economic Review,* 93, 622–45.

Machina, Mark J., and Michael Rothschild. (1987) "Risk." In *The New Palgrave Dictionary of Economics*, edited by J. Eatwell, M. Millgate, and P. Newman, pp. 203–205. London: MacMillan.

Mankiw, N. Gregory. (2006) "A Letter to Ben Bernanke." *American Economic Review, Papers and Proceedings,* 96, 182–4.

Meyer, Laurence H. (1998) "Come with Me to the FOMC." Remarks by Governor Laurence H. Meyer at Willamette University, Salem, OR, April 2.

Meyer, Laurence H. (2004) *A Term at the Fed: An Insider's View*. New York: Harper Business.

Mishkin, Frederic S., and Niklas J. Westelius. (2008) "Inflation Band Targeting and Optimal Inflation Contracts." *Journal of Money, Credit, and Banking*, 40, 557–82.

Orphanides, Athanasios, and Volker Wieland. (2000) "Inflation Zone Targeting." *European Economic Review,* 44, 1351–87.

Ruge-Murcia, Francisco. (2003) "Inflation Targeting under Asymmetric Preferences." *Journal of Money, Credit and Banking,* 35, 763–85.

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