



# **Implicit Quantile Preferences of the Fed and the Taylor Rule**

**Gabriel Montes-Rojas (IIEP-UBA-CONICET)**

**Fernando Toledo (UNLP)**

**Nicolás Bertholet (IIEP-UBA)**

**Kevin Corfield (UBA)**

**DOCUMENTO DE TRABAJO N° 384**

**Enero de 2026**

Los documentos de trabajo de la RedNIE se difunden con el propósito de generar comentarios y debate, no habiendo estado sujetos a revisión de pares. Las opiniones expresadas en este trabajo son de los autores y no necesariamente representan las opiniones de la RedNIE o su Comisión Directiva.

The RedNIE working papers are disseminated for the purpose of generating comments and debate, and have not been subjected to peer review. The opinions expressed in this paper are exclusively those of the authors and do not necessarily represent the opinions of the RedNIE or its Board of Directors.

**Citar como:**

Montes-Rojas, Gabriel, Fernando Toledo, Nicolás Bertholet y Kevin Corfield (2026). Implicit Quantile Preferences of the Fed and the Taylor Rule. Documento de trabajo RedNIE N°384.

# Implicit quantile preferences of the Fed and the Taylor rule

Gabriel Montes-Rojas\*    Fernando Toledo†    Nicolás Bertholet‡    Kevin Corfield§

January 13, 2026

## Abstract

We study optimal monetary policy when a central bank maximizes a quantile utility objective rather than expected utility. In our framework, the central bank's risk attitude is indexed by the quantile index level, providing a transparent mapping between hawkish/dovish stances and attention to adverse macroeconomic realizations. We formulate the infinite-horizon problem using a Bellman equation with the quantile operator. Implementing a Euler-equation approach, we get Taylor-rule-type reaction functions. Using an indirect inference approach, we derive an implicit quantile index of central bank risk aversion. An empirical implementation for the US is outlined based on reduced-form laws of motion with conditional heteroskedasticity, enabling estimation of the new monetary policy rule and its dependence on the Fed risk attitudes. The results reveal that the Fed has mostly a dovish-type behavior but with some periods of hawkish attitudes.

Word count: 7022

**Keywords:** Taylor rule; inflation; output gap; quantile preferences; dynamic programming; recursive model.

**JEL:** C22; C61; E52; E58.

---

\*Instituto Interdisciplinario de Economía Polática, Universidad de Buenos Aires and CONICET, Ciudad Autónoma de Buenos Aires, Argentina. [gabriel.montes@economics.uba.ar](mailto:gabriel.montes@economics.uba.ar)

†Universidad Nacional de La Plata, La Plata, Argentina. [fernando.toledo@econo.unlp.edu.ar](mailto:fernando.toledo@econo.unlp.edu.ar)

‡Instituto Interdisciplinario de Economía Política, Universidad de Buenos Aires, Ciudad Autónoma de Buenos Aires, Argentina. [nicolasbertholet2008@economics.uba.ar](mailto:nicolasbertholet2008@economics.uba.ar)

§Universidad de Buenos Aires, Ciudad Autónoma de Buenos Aires, Argentina. [kevincorfield@economics.uba.ar](mailto:kevincorfield@economics.uba.ar)

# 1 Introduction

The debate between hawkish and dovish monetary policy authorities has received increased attention (see [Tobback, Nardelli, and Martens \(2017\)](#), [Hack, Istrefi, and Meier \(2025\)](#)). For example, the London Investment Service Company In Touch Capital Markets has introduced a relative scale to explain the position of Fed's regarding their relative aversion to inflationary pressures: the lowest bound on the scale indicates very hawkish members and the upper bound specifies very dovish ones. In the empirical literature, CBs are often classified as hawkish or dovish based on projected Taylor rules, particularly the interest rate's long-term response to inflation (see [Castro \(2011\)](#); [Wilson \(2020\)](#); [Malmendier, Nagel, and Yan \(2021\)](#); [González-Astudillo and Tanvir \(2023\)](#)).

The responses of monetary authorities to the post-pandemic period and the Ukraine war are paramount to understand coordinated monetary policy actions in terms of interest rate hikes to recent global inflation synchronization ([Ha, Kose, Ohnsorge, and Yilmazkuday \(2024\)](#)). Although this empirical regularity, the intensity of restrictive monetary policy measures differs among countries. This stylized fact reminds us how important it is for CBs to find an optimal policy monetary rule that reflects the distinct degree of risk aversion to inflation and economic fluctuations. For example, the Fed's approach was shaped by empirical evidence that gathered in the years before the pandemic, and the results were impacted by a decrease in inflation persistence, a flattening of the Phillips curve's slope, and inaccurate assessments of real-time output or unemployment ([Sargent and Williams \(2025\)](#)).

Taylor rules have become a cornerstone of modern monetary policy science ([Woodford \(2003\)](#), [Taylor \(1993\)](#)). Optimal monetary policy involves setting short-term nominal interest rates to stabilize the economy by managing inflation and output gaps, often guided by rules that recommend increasing rates when inflation is high or output exceeds its potential and lowering them otherwise ([Clarida, Galí, and Gertler \(1999\)](#)).

Although Taylor rules provide a simple and robust framework, true optimal policy depends on specific economic models, policy objectives, and whether policymakers prioritize inflation, output, or a combination of both, leading to ongoing debates and refinements of such rules. In that regard, during the last Jackson Hole Symposium (August 2025, 21-23), [Nakamura, Riblier, and Steinsson \(2025\)](#) presented a study that remarks the descriptive nature of Taylor rules instead of their prescriptive one. These authors have also observed deviations from the Taylor principle after exploring the recent Fed's behavior and also pointed out the coexistence of early and late policy interest rate hikers.

In the case of New Keynesian macroeconomic models, the discussion about the theoretical validity of optimal Taylor rules usually includes these and other relevant issues (see [Boehm and House \(2014\)](#)). The main point we would like to state here is the discontent that some contributions express in terms of the theoretical form of optimal Taylor rules (see [Cochrane \(2007\)](#), [Benhabib, Schmitt-Grohé, and Uribe \(2001\)](#)).

Theoretical contributions have usually approached the different preferences toward the traditional monetary policy trade-off (inflation versus output fluctuations) by deriving an optimally monetary policy Taylor rule minimizing a quadratic loss intertemporal function ([Woodford \(2003\)](#)). However, problems arise in applying these functions when the underlying economic model is nonlinear ([Benigno and Eggertsson \(2023\)](#)), shocks are non-normal ([Hofmann, Manea, and Mojon \(2024\)](#)), or the quadratic assumption is too simplistic ([al Nowaihi and Stracca \(2002\)](#)). Furthermore, although linear-quadratic models offer analytical tractability for deriving optimal rules, they may fail to capture key asymmetric preferences that drive

actual monetary policy behavior in the real world ([Svensson \(2003\)](#))<sup>1</sup>. So, traditional methods often look at the average policy reaction function of a CB.

After the 2008 Global Financial Crisis, many authors argued that CBs should pay more attention to tail risks (especially downside risks) rather than just average outcomes (see [Demirguc-Kunt, Detragiache, and Merrouche \(2013\)](#)). For instance, regardless of the US economy's condition, [Barci \(2025\)](#) notes that monetary policy can increase downside risk; nevertheless, this ability is significantly diminished during economic expansions. Such disparity, if not appropriately taken into consideration, could cause monetary authorities to be too cautious when it comes to tightening during booms.

In the present paper, our theoretical contribution is to introduce quantile utility (QU) preferences into the intertemporal minimization of the loss function of CBs.

In essence, a QU agent makes decisions based on maximizing the quantile of random utility (see [de Castro and Galvao \(2022\)](#)), or equivalently minimizing the quantile of a random loss function. We study how CBs minimize their intertemporal loss function using dynamic programming through the Bellman equation solution as in [de Castro and Galvao \(2019\)](#). In contrast to the usual framework of minimizing a quadratic intertemporal loss function using expected utility to attain the optimal policy Taylor rule, we incorporate QU preferences to study how CBs respond to undesirable macroeconomic outcomes. An important feature of QU is that it allows for complete separation of the risk attitude and the intertemporal substitution. For QU, risk aversion is unambiguously measured by the quantile index.

Quantile preferences in dynamic models represent a generalization of the standard Bellman framework that can justify more precautionary and asymmetric monetary policy potentially providing microfoundations for why CBs might want to deviate from simple Taylor type rules in the face of tail risks. This framework helps explain why CBs often appear more dovish than simple Taylor rules would suggest and they may be intrinsically more concerned with downside risks (unemployment spikes, financial instability) than with symmetric inflation and output gaps.

Our research adds new arguments to the theoretical discussion about the accurate form of the optimal Taylor rule. We depart from the conventional dynamic optimization problem faced by CBs to propose a new closed analytical form for their reaction function. The main advantage of the new Taylor rule is that CBs are concerned not only with average inflation-output trade-off but also in analyzing scenarios of inflation and output gaps.

Moreover, by comparing the observed policy actions (i.e. interest rate) with the entire myriad of available actions for all quantiles, we can infer the type of the CB at each point in time. We define this as the implicit quantile preference index, which is a dynamic index of CB risk aversion. In turn, we interpret this index as a dovish/hawkish scale.

We use US quarterly data from 1954-Q4 to 2025-Q2. This framework suggests that the Fed often appear more dovish than simple Taylor rules would suggest and they may be intrinsically less concerned with downside risks (unemployment spikes, financial instability). However, this is not a general description, and there are specific periods where the Fed is characterized with hawkish behavior. Our analysis delivers an index of the Fed's risk aversion attitudes across time that is based on the implicit quantile preference.

Our argument is that QU may offer a proper framework to incorporate these preferences. In particular, to model policy actions to accommodate specific concerns on certain potential

---

<sup>1</sup>[El-Shagi \(2025\)](#) has recently shown that the Fed prioritizes business cycle stabilization over containing inflation.

outcomes, measured by locations in the distribution of utility (or loss), which correspond to the quantile index. It should be noted that a competing alternative is to estimate risk through a specific functional form of the utility function. Here, however, we consider a canonical simple approach towards the functional form of the utility function and explore the implied risk behavior from quantiles.

Although the Fed does not have an explicit policy of quantile preferences, we argue that it may behave as though it does because of three factors: i) a contemporary, inclusive interpretation of its maximum employment objective<sup>2</sup>); ii) a potent asymmetric risk management approach that places a high priority on avoiding deep recessions<sup>3</sup>); and iii) an awareness that preventing catastrophic tail risks is what financial stability is all about<sup>4</sup>.

Since the Fed's role is to stabilize the business cycle and prevent booms and busts, it is naturally risk adverse<sup>5</sup>. However, the Fed became exceedingly risk apprehensive about two left-tail risks following the 2008 financial crisis and the COVID-19 shock: i) the deflationary tail risk; and ii) the zero lower bound (ZLB) tail risk. The Fed is compelled to embrace a low quantile preference due to its severe risk aversion toward these particular catastrophic possibilities. Even though it results in a slightly less ideal outcome in the median scenario (a normal economy), its policy shift to average inflation targeting and its tolerance for a hot labor market can be interpreted as a strategy to maximize the outcome in the worst economic scenarios (those involving deflation and ZLB).

In addition, the COVID-19 pandemic in 2020 and the Global Financial Crisis in 2008 served as sobering reminders that tail risks do occur. Normal and well-behaved distribution models could not work. This has forced the Fed to embrace frameworks that specifically take fat tails and extraordinary events into consideration. Because it does not imply symmetry and can reveal information about the full distribution of possible outcomes rather than just the center, QU preferences are an ideal tool for this.

This study relates to three branches of the literature. First, we add to the literature on optimal monetary policy the idea that CBs adjust their monetary policy actions to more complex optimal rules than traditional ones. We show how our theoretical framework relates to the New Keynesian macroeconomic models by getting a new Taylor rule that allows different risk aversion attitudes towards inflation and output combinations. Second, we contribute to the quantile preferences literature (see [de Castro and Galvao \(2022\)](#)) with an innovative

---

<sup>2</sup>A robust job market is especially advantageous for low and moderate income areas, and those with lower levels of education, as Chairs Janet Yellen and Jerome Powell have made clear. These groups are frequently the first to be let go during a downturn and the last to be hired during a recovery. Through the implementation of a high pressure economy, the Fed can directly enhance the results for these susceptible groups, which are the lower quantiles of the market distribution.

<sup>3</sup>A severe, protracted recession or deflationary spiral may ensue if the economy falters while interest rates are already zero. The most vulnerable are severely harmed by this disastrous left tail incident. The Fed has a dependable tool: it can raise interest rates if the economy overheats and inflation increases. The quantile preference is imposed by this asymmetry. The Fed will inevitably be biased toward policies that aggressively guard against that left tail risk, even if it means temporarily tolerating a higher risk of inflation, because the costs of a deep recession, such as high unemployment and widespread bankruptcies, are so devastating, particularly for the lower end of the income distribution.

<sup>4</sup>The ultimate example of a left-tail occurrence is a financial crisis. They start in certain industries (subprime mortgages in 2008, for example), but they ruin the whole economy, with the poorest people bearing the brunt of it. A type of quantile preference, the Fed's post-crisis emphasis on regulation and systemic risk monitoring is an effort to stop the financial meltdown's excessive negative tail.

<sup>5</sup>The Fed's preference to avoid recessions during the Volcker-Greenspan era is compatible with asymmetric interest rate responses, which can instead be linked to expansions and recessions (see [Wolters \(2012\)](#)).

application: the formal derivation of a closed form for a new Taylor rule. We use dynamic programming methods and the Bellman equation to optimize the CB intertemporal loss function and obtain the new optimal reaction function of policymakers (see [de Castro and Galvao \(2019\)](#); [Hills, Nakata, and Sunakawa \(2020\)](#)). Third, we present an indirect inference approach to estimate the Fed’s risk attitude across time, thus contributing to the analysis of parameter and/or model uncertainty (see [Cogley, Colacito, Hansen, and Sargent \(2008\)](#)).

It should be noted that our approach is different from empirical papers that estimate heterogeneous responses in a Taylor rule regression model as in [Chevapatrakul, Kim, and Mizen \(2009\)](#), [Chevapatrakul and Paez-Farrell \(2014\)](#), [Chen and Kashiwagi \(2017\)](#) and [Christou, Naraidoo, Gupta, and Kim \(2018\)](#), among others. In those papers, the key goal is to evaluate heterogeneous responses of the interest rate to inflation and output gap (and others) using a quantile regression framework. In our paper, quantiles relate to a structural preference parameter of the CB and not to the conditional quantiles of the conditional distribution of the interest rate.

The paper proceeds as follows. Section 2 summarizes the quantile utility framework. Section 3 applies dynamic programming to the intertemporal QU maximization problem to derive Taylor rules. Section 4 describes the empirical implementation strategy. Section 5 presents the estimation results. Section 6 concludes.

## 2 Quantile preferences and risk attitude

### 2.1 Quantile preferences for univariate outcomes

An expected utility (EU) maximizer with utility function  $u : \mathbb{R} \rightarrow \mathbb{R}$  prefers lottery  $X$  to  $Y$  if  $E[u(X)] \geq E[u(Y)]$ . This refers to a case when a decision maker (DM) that is faced with uncertain outcomes chooses the action that maximizes the expected average outcome. Quantile utility<sup>6</sup> (QU) is based on a framework where optimal decisions and allocations correspond to maximizing a specific quantile of the distribution of outcomes or returns. For a given a random (univariate) variables,  $Y$ , let  $F(y) = F_Y(y) = \Pr(Y \leq y)$  denote the conditional cumulative distribution function (c.d.f.) of  $Y$ . If the function  $y \mapsto F_Y(y)$  is strictly increasing and continuous in its support, its inverse is the quantile of  $Y$ , that is,  $Q_\tau[Y] = F_Y^{-1}(\tau)$ , for  $\tau \in (0, 1)$ . If  $y \mapsto F_Y(y)$  is not invertible, we can still define the quantile as one of its generalized inverses. Following the standard practice, we define the quantile as the left-continuous version of the generalized inverse:

$$Q_\tau[Y] \equiv \inf\{y \in \mathbb{R} : \Pr[Y \leq y] \geq \tau\}.$$

Quantile preferences are defined by simply substituting the expectation by the quantile operator, that is,

$$X \succeq Y \iff Q_\tau[u(X)] \geq Q_\tau[u(Y)]. \quad (1)$$

---

<sup>6</sup>Quantile preferences were first introduced by [Manski \(1988\)](#). [Rostek \(2010\)](#) and [Chambers \(2009\)](#) provide axioms for the static case, and [de Castro and Galvao \(2022\)](#) formally axiomatize both the static and dynamic quantile preferences. [Giovannetti \(2013\)](#) studies a two-period economy for an intertemporal consumption model under quantile utility maximization. [de Castro and Galvao \(2019\)](#) establish the properties of a general dynamically consistent quantile preferences model. We refer to this type of preference modeling as quantile utility (QU hereafter).

The intuition is that, in the presence of uncertainty, a QU maximizer makes decisions based on maximizing a given  $\tau$  quantile of the distribution of potential outcomes.

For univariate random variables (i.e. monetary outcomes), quantiles enjoy the following property: for any continuous and increasing function  $f : \mathbb{R} \rightarrow \mathbb{R}$ ,  $f(Q_\tau[X]) = Q_\tau[f(X)]$ . If  $u : \mathbb{R} \rightarrow \mathbb{R}$  is strictly increasing and continuous, as usual, then we can take its inverse and apply to (1), to obtain:

$$X \succeq Y \iff u^{-1}(Q_\tau[u(X)]) \geq u^{-1}(Q_\tau[u(Y)]) \iff Q_\tau[X] \geq Q_\tau[Y].$$

In other words, for the univariate case, the utility functional form is not necessary to model DM behavior.

As noted by [Manski \(1988\)](#), under QU risk attitudes can be indexed by  $\tau$  itself. Intuitively, we can map risk aversion into the  $\tau$  scale, such that a  $\tau$ -DM is more risk averse than  $\tau'$ -DM if  $\tau \leq \tau'$ . In sum, a  $\tau$ -DM evaluates lotteries and actions based on choosing the ones with the highest  $\tau$  quantile, and then, the lower  $\tau$  is, the more the DM is concerned with low values or losses.

A canonical competing alternative is to measure risk by a risk aversion parameter inside the utility function. Optimal decisions of individuals with constant absolute risk aversion and constant relative risk aversion preferences may be very different, although, in practice, it may be difficult to differentiate between both attitudes towards risk from real data. A robust approach within the EU paradigm is stochastic dominance. This theory allows one to rank risky alternatives without relying on specific forms of the individuals' utility function. However, the equivalence between EU maximization and stochastic dominance is only satisfied, under risk aversion, for well-behaved (increasing and concave) utility functions. Moreover, as shown in the next subsection, stochastic dominance in multivariate settings may not be satisfied.

## 2.2 Quantile preferences for multivariate outcomes

For a multivariate random outcome variable, say  $m$  dimensional vectors of the form  $Y = (Y_1, Y_2, \dots, Y_m)$  with domain  $\mathcal{Y} \subseteq \mathbb{R}^m$ ,  $\inf\{y \in \mathcal{Y} : \tau \leq F(y)\}$  is (in general) not unique.

Take the bivariate case  $Y = (Y_1, Y_2)$  with domain  $\mathcal{Y} \subseteq \mathbb{R}^2$ . Quantiles are themselves then defined on regions, contours and depths ([Hallin and Konen, 2024](#)). A quantile on the bivariate domain is any pair  $(q_1, q_2) \in \mathcal{Y}$  such that  $P(Y_1 \leq q_1, Y_2 \leq q_2) = \tau$ ,  $\tau \in (0, 1)$ . Figure 1a plots the contour plot for the probability density function of a bivariate distribution and adds two points that correspond to the same quantile  $\tau$ . Figure 1b plots the same contour plot but considers the curves corresponding to two different  $\tau$ 's.

Note that the preceding analysis of QU cannot be applied to the multivariate domain unless additional considerations are taken. Consider two random variables  $X = (X_1, X_2)$  and  $Y = (Y_1, Y_2)$  on the bivariate domain. Figure 2 plots two different cases for the same quantile  $\tau$  with (a) and without (b) crossing. As such, there is no natural ordering that can be used in terms of the distribution function or its inverse, the quantiles.

A consequence of this is that the QU model cannot evaluate random utility based on the multivariate distribution of the arguments determining the utility. For our purposes, a QU-maximizer CB that has preferences over inflation and output gap cannot resort to the joint distribution of these variables to evaluate policies. On the contrary, it does require the utility function and the relative valuation of each component.

Following [Hallin, Paindaveine, and Šiman \(2010\)](#) multivariate models can be decomposed into a series of univariate models in terms of quantile analysis. Quantiles are analyzed in terms

Figure 1: (a) Two points representing quantile  $\tau$ .  $P(Y_1 \leq q_1, Y_2 \leq q_2) = \tau$ . (b) Two contour lines for  $\tau < \tau'$

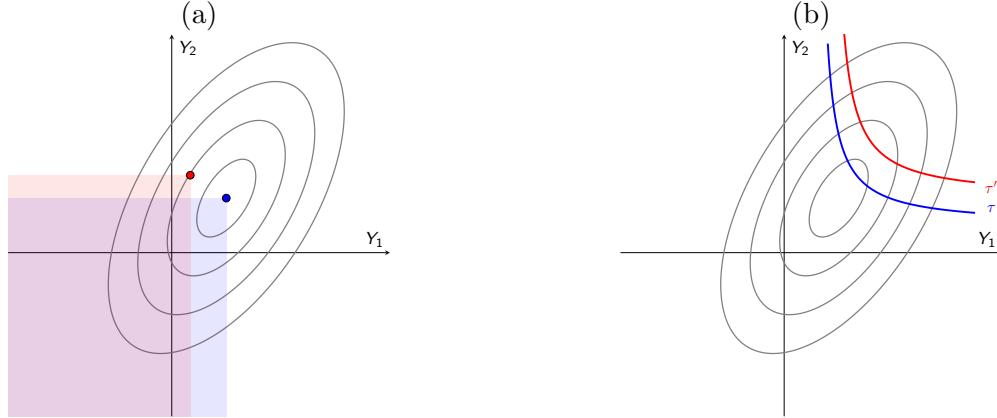
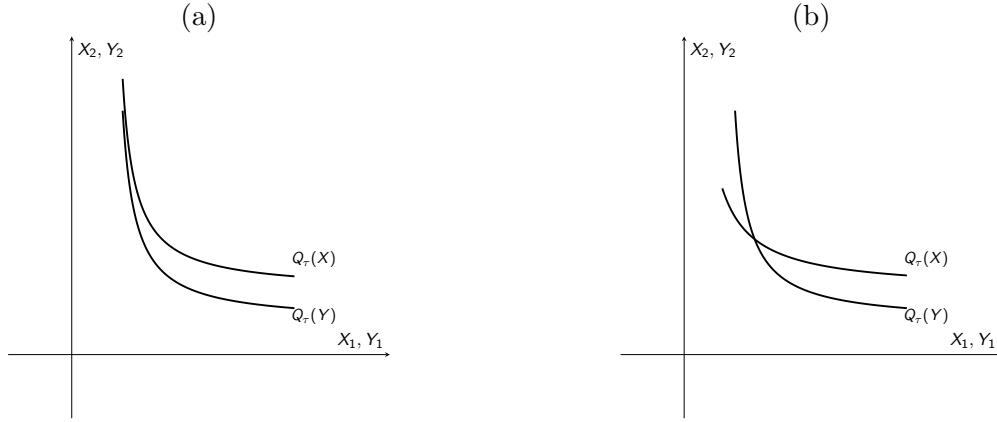


Figure 2: (a)  $Q_\tau(X) >> Q_\tau(Y)$ . (b)  $Q_\tau(X) >< Q_\tau(Y)$

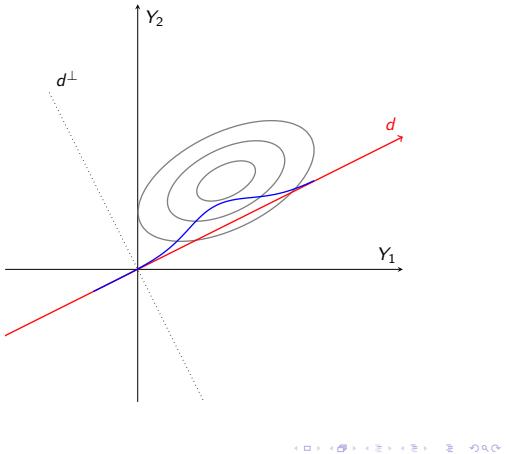


of a *magnitude* and a *direction*. We define  $T = (\tau_1, \tau_2, \dots, \tau_m) \in (0, 1)^m$  as a set of quantile indices. The vector  $T$  can be factorized as  $T \equiv \tau v$ , where  $\tau = \|T\| \in (0, 1)$  represents the *magnitude*, and  $d \in \mathbb{R}^{m-1} \equiv \{d \in \mathbb{R}^m : \|d\| = 1\}$  represents the *direction* expressed as a unit vector in the Euclidean framework.

In this model,  $\tau$  is a scalar quantile index that specifies the position along the distribution, while  $v$  is a unit vector that determines the direction in the  $m$ -dimensional space. This vector can be interpreted as an  $(m - 1)$ -dimensional directional component that captures how quantile changes unfold across variables. This decomposition allows for an intuitive and geometric interpretation of multivariate quantiles in terms of distance and orientation within the variable space.

Vector directional quantile proposes to study univariate variables of the form  $d \cdot y$ , where  $\cdot$  represent element-by-element vector multiplication. Figure 3 plots this idea for the bivariate case (in the figure  $d^\perp$  is an orthonormal basis of the subspace orthogonal to  $d$ ). Once a

Figure 3: Vector directional quantile



direction is fixed, the problem becomes one of univariate quantiles, and QU analysis can be applied.

In the QU setup, the direction  $d$  can be interpreted as a linearization of utility function over the multivariate domain. The direction reduces the dimensionality of preferences into a linear univariate model. Note, however, that is only valid in a local sense. Non-local comparisons require the use of the utility function to fulfill this role for all cases.

## 2.3 Dynamic models

Many applications of intertemporal maximization use the standard recursive EU. These models have been workhorses in several economic fields. EU is simple and amenable to theoretical modeling. The assumption of maximization of average utility, the average being a simple measure of centrality, has intuitive appeal as a behavioral postulate. Nevertheless, the usual EU framework has been subjected to a number of criticisms, including in its dynamic version. For example, it has been well documented in the literature that it is not possible to separate the intertemporal substitution from the risk attitude parameters when using standard dynamic models based on the EU (see, e.g., [Hall, 1988](#)). The framework proposed by [Kreps and Porteus \(1978\)](#) to study temporal resolution of uncertainty was one of the first efforts to go beyond EU in the dynamic setting. An expanding literature considers alternative recursive models.<sup>7</sup>

de Castro and Galvao (2019) developed a new alternative to the EU recursive model based on QU. In their model, the economic agent chooses the alternative that leads to the highest  $\tau$ -quantile of the stream of future utilities for a fixed  $\tau \in (0, 1)$ . The dynamic quantile preferences for intertemporal decisions are represented by an additively separable quantile model with standard discounting. The associated recursive equation is characterized by the sum of the current period utility function and the discounted value of the certainty equivalent, which is obtained from a quantile operator. This intertemporal model is tractable and simple to interpret, since the value function and Euler equation are transparent, and easy to calculate

<sup>7</sup>We refer the reader to Epstein and Zin (1989, 1991), Weil (1990), Grant, Kajii, and Polak (2000), Epstein and Schneider (2003), Hansen and Sargent (2004), Maccheroni, Marinacci, and Rustichini (2006), Klibanoff, Marinacci, and Mukerji (2009), Marinacci and Montrucchio (2010), Strzalecki (2013), Bommier, Kochov, and Le Grand (2017), Sarver (2018), and Dejarnette, Dillenberger, Gottlieb, and Ortoleva (2020) among others.

(analytically or numerically). This framework allows for the separation of the risk attitude from the intertemporal substitution, which is not possible with EU, while maintaining important features of the standard model, such as dynamic consistency and monotonicity.

### 3 Preferences of a CB and Taylor rule for QU maximizer

#### 3.1 General set-up

Consider now a CB that has  $\tau$ -QU preferences based on  $(\pi_t - \pi^*)$  where  $\pi_t$  is inflation and  $\pi^*$  is the target inflation rate, and output gap  $y_t = (\bar{y}_t - y_t^*)$  where  $\bar{y}_t$  is output and  $y_t^*$  is a measure of output long-run trend and potential output. Moreover, we assume that the CB has a preference for smoothing policy variables over time (i.e. the interest rate).

The CB decisions can be represented along a utility function  $u(y, \pi, i, z)$  that typically is a trade-off between inflation and output gaps, and it may depend on interest rate and current shocks. In general, we could assume that the CB values more inflation and output closer to the target values.

We assume a quadratic utility function of the form

$$u(\pi_t, y_t, i_t, z_t) = -\frac{(\pi_t - \pi^*)^2}{2} - \frac{\lambda(y_t)^2}{2} - \frac{\delta(i_t - i_{t-1})^2}{2}. \quad (2)$$

In this model, the CB has preferences for state variables close to the target values and for avoiding fluctuations in the interest rate. This utility function is in fact a loss function multiplied by  $-1$ , and it can be derived from micro-foundations as in [Woodford \(2003\)](#).  $\lambda$  and  $\delta$  are structural parameters that correspond to the degree of substitution of the inflation gap, the output gap and the variations in interest rate along indifference curves.

Let  $x \in \mathcal{X}$  denote the particular state and the state space,  $i \in \mathcal{I}$  be the action and the set of possible actions the CB may take, and  $z \in \mathcal{Z}$ , the range of the shocks. For our purposes,  $x_t = (y_t, \pi_t)$ ,  $i_t$  is the interest rate and  $z_t$  represents random components that affect  $x_t$ . Moreover, we consider that  $z_t = (z_{\pi,t}, z_{y,t})$  a bivariate random vector. Although we do not explicitly consider it to reduce notation, the state variables may include lags of the variables. For our particular case, we use the lag of  $i$  inside the utility function.

The next period state,  $x_{t+1}$ , is defined by a law of motion function  $\phi : \mathcal{X} \times \mathcal{I} \times \mathcal{Z} \rightarrow \mathcal{X}$  that satisfies  $x_{t+1} = \phi(x_t, i_t, z_{t+1})$ . Given the current state  $x_t$  and current shock  $z_t$ ,  $\Gamma(x_t, z_t)$  denotes the set of possible choices  $i_t$ , that is, the feasibility constraint set.

#### 3.2 Infinite horizon and recursive maximization problem

In the proposed CB model, the uncertainty with respect to the future realizations of  $z_t$  is given by a quantile applied to potential values of the utility function. In line with QU theory, the quantile index  $\tau$  represents CB attitudes towards risk. We refer to a  $\tau'$ -CB to be more risk averse than a  $\tau$ -CB one if  $\tau' < \tau$ . That is, the  $\tau'$ -CB is more concerned with worse outcomes scenarios (i.e. high inflation, low output) than a  $\tau$ -CB.

In the QU framework, optimal decisions are taken to maximize the  $\tau$  quantile of intertemporal utility in an infinite horizon problem. This framework does not allow for the same solution strategies as in the EU case, because we cannot apply the law of iterated expectations. However, under certain conditions described in [de Castro and Galvao \(2019\)](#), these

dynamic intertemporal choices can be represented by the maximization of a value function  $v: \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$  that satisfies the recursive Bellman equation:

$$v(x, z) = \sup_{i \in \Gamma(x, z)} \left\{ u(x, i, z) + \beta Q_\tau[v(\phi(x, i, z'), z') | z] \right\}, \quad (3)$$

where  $z'$  indicates the next period shock.

Note that this is similar to the usual dynamic programming problem, in which the expectation operator  $E[\cdot]$  is in place of  $Q_\tau[\cdot]$ . [de Castro and Galvao \(2019\)](#) and [de Castro, Galvao, and Nunes \(2025\)](#) endorse the construction of this type of recursive models from dated preferences. Those authors prove uniqueness of the solution to problem (3), under a set of regularity conditions similar to those in dynamic programming set-up and some specific restrictions for the use of quantiles. The solution is a policy function  $i_\tau^*: \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{I}$ , that associates to each  $(x_t, z_t)$  the optimal solution  $i_\tau^* = i^*(x_t, z_t)$ .

### 3.3 Law of motion

Now consider a location-scale law of motion  $\phi(\cdot)$  for inflation and output gap, using autoregressive processes of order 1.

$$\pi_{t+1} = \phi_\pi(x_t, i_t, z_{\pi, t+1}) = \alpha_{\pi 0} + \alpha_{\pi\pi}\pi_t + \alpha_{\pi y}y_t + \alpha_{\pi i}i_t + h_\pi(\pi_t, y_t, i_t)z_{\pi, t+1},$$

$$y_{t+1} = \phi_y(x_t, i_t, z_{y, t+1}) = \alpha_{y 0} + \alpha_{y\pi}\pi_t + \alpha_{yy}y_t + \alpha_{yi}i_t + h_y(\pi_t, y_t, i_t)z_{y, t+1},$$

where the  $\alpha$  coefficients capture location mean effects, i.e. the persistence of inflation and output gap and how sensitive inflation and output gap are to changes in the interest rate  $i_t$ , and  $z_{\pi, t+1}$  and  $z_{y, t+1}$  are random shocks, possibly correlated to each other but assumed to be independent of the state and interest rate variable, i.e.  $(z_{\pi, t+1}, z_{y, t+1}) \perp (\pi_t, y_t, i_t) | z_t$ . They have zero conditional mean  $E[z_{a, t+1} | z_{\pi, t}, z_{y, t}] = 0$ ,  $a = \pi, y$  and unit variance  $E[z_{a, t+1}^2 | z_{\pi, t}, z_{y, t}] = 1$  (which anyway cannot be identified separately from  $h$  skedastic function). This is a reduced form that may be the result of intertemporal IS curve and a New Keynesian Phillips curve as in [Woodford \(2003\)](#) and [Galí \(2015\)](#).

Functions  $h_\pi(\cdot)$  and  $h_y(\cdot)$  are skedastic strictly positive functions that control the conditional heteroskedasticity of the state variables, affecting the scale. In turn, they determine the structure of heterogeneity in the law of motion and whether the quantiles are not parallel to each other. We can refer to location shift only models to those where the  $h$  functions are constant, and to location-scale shift models where the  $h$  functions depend on  $(\pi, y, i)$ .

Several parameterizations can be applied, see for instance [Romano and Wolf \(2017\)](#). Different specifications used in the heteroskedasticity literature to model location-scale shift effects are

$$h_a(\pi_t, y_t, i_t) = (\gamma_{a0} + \gamma_{a\pi}\pi_t + \gamma_{ay}y_t + \gamma_{ai}i_t)^{1/2}, \quad a = \pi, y.$$

or

$$h_a(\pi_t, y_t, i_t) = \exp \left[ (\gamma_{a0} + \gamma_{a\pi}\pi_t + \gamma_{ay}y_t + \gamma_{ai}i_t)^{1/2} \right], \quad a = \pi, y.$$

The  $\gamma$  coefficients control whether the random shocks affect the scale impact of these shocks on inflation and output gap. Thus models with  $\gamma = 0$  have only location shifts, while  $\gamma \neq 0$  characterize location-scale shift ones.

In the context of quantile regression specifications, this representation allows for a random-coefficient model indexed by a quantile index, i.e.,

$$Q_{\tau_\pi}(\pi_{t+1}|x_t, i_t) = \alpha_{\pi 0}(\tau_\pi) + \alpha_{\pi\pi}(\tau_\pi)\pi_t + \alpha_{\pi y}(\tau_\pi)y_t + \alpha_{\pi i}(\tau_\pi)i_t, \quad \tau_\pi \in (0, 1),$$

$$Q_{\tau_y}(y_{t+1}|x_t, i_t) = \alpha_{y 0}(\tau_y) + \alpha_{y\pi}(\tau_y)\pi_t + \alpha_{y y}(\tau_y)y_t + \alpha_{y i}(\tau_y)i_t, \quad \tau_y \in (0, 1),$$

where  $\alpha_{ab}(\tau_a) = \alpha_{ab} + \frac{\partial h_a(x_t, i_t)}{\partial b} Q_{\tau_a}(z_a|x_t, i_t)$ ,  $a = \pi, y$  and  $b = 0, \pi, y, i$ . Here  $\tau_\pi$  and  $\tau_y$  reflects different conditional responses of inflation and output gap to current state variables and policy choices. Both indexes are not necessarily independent and they need to be considered in a multivariate quantile model as in [Montes-Rojas \(2017, 2019, 2022\)](#).

### 3.4 Euler equations

Under certain conditions, the Taylor rule can be derived analytically from this function by implementing the Euler equation as in [de Castro, Galvao, and Nunes \(2025\)](#) Theorem 3.18. Consider an application of the theorem to get the Euler equation as:

$$\begin{aligned} & \frac{\partial u(x_t, i_t, z_t)}{\partial i} + \\ & \beta Q_\pi \left[ \frac{\frac{\partial u(x_{t+1}, i_{t+1}, z_{t+1})}{\partial \pi} \frac{\partial \phi_\pi(x_t, i_t, z_{t+1})}{\partial i}}{\partial \pi} + \right. \\ & \left. \frac{\frac{\partial u(x_{t+1}, i_{t+1}, z_{t+1})}{\partial y} \frac{\partial \phi_y(x_t, i_t, z_{t+1})}{\partial i}}{\partial y} \Big| z \right] = 0. \end{aligned} \quad (4)$$

For this derivation to be applied, it requires that differentiability and the quantile operator can be interchanged, and that the shocks have an increasing monotonic effect. In particular, we need the following univariate component

$$\frac{\partial u(x_{t+1}, i_{t+1}, z_{t+1})}{\partial \pi} \frac{\partial \phi_\pi(x_t, i_t, z_{t+1})}{\partial i} + \frac{\partial u(x_{t+1}, i_{t+1}, z_{t+1})}{\partial y} \frac{\partial \phi_y(x_t, i_t, z_{t+1})}{\partial i}$$

to be monotonically increasing on  $z$ . Note, however, that  $z$  is bivariate, and therefore the monotonicity requirement has to be evaluated at particular vector directions. For the case in eq. (2),

$$\begin{aligned} \frac{\partial u(x_{t+1}, i_{t+1}, z_{t+1})}{\partial \pi} &= -(\pi_{t+1} - \pi^*), \\ \frac{\partial u(x_{t+1}, i_{t+1}, z_{t+1})}{\partial y} &= -\lambda y_{t+1}, \\ \frac{\partial u(x_t, i_t, z_t)}{\partial i} &= -\delta(i_t - i_{t-1}). \end{aligned}$$

Then the monotonicity requirement is that the random variable  $q(x, i) := -\phi'_{\pi, i}(x, i)z_\pi - \lambda\phi'_{\pi, i}(x, i)z_y$  with  $\phi'_{\pi, i}(x, i) = \frac{\partial \phi_\pi(x, i)}{\partial i}$  and  $\phi'_{y, i} = \frac{\partial \phi_y(x, i)}{\partial i}$ , has a well defined quantile function.

We derive here Euler equations solutions. For simplicity we assume that  $\alpha_{\pi 0} = \alpha_{y 0} = 0$  (but this is not assumed in the empirical application).

### Location shift only

Suppose first that  $\gamma_{ab} = 0$   $a, b = \pi, y, i$ , that is, the quantiles of the random shocks are not affected by the state variables nor by the control variable.

Then

$$-\delta(i_t - i_{t-1}) + \beta Q_\tau [-(\alpha_{\pi\pi}\pi_t + \alpha_{\pi y}y_t + \alpha_{\pi i}i_t + z_{\pi,t+1} - \pi^*)\alpha_{\pi i} \\ - \lambda(\alpha_{y\pi}\pi_t + \alpha_{yy}y_t + \alpha_{yi}i_t + z_{y,t+1})\alpha_{yi} | z_t] = 0. \quad (5)$$

Note that by the requirements on the validity of the Euler implementation for QU,  $Q_\tau(-z_{\pi,t+1}\alpha_{\pi i} - \lambda z_{y,t+1}\alpha_{yi} | z_t)$  needs to be monotonically increasing in both components  $(z_{\pi,t+1}, z_{y,t+1})$ .

Thus we obtain the following Taylor rule for the  $\tau$ -QU problem

$$i_\tau^*(\pi_t, y_t, i_{t-1}) = \\ (\delta + \beta(\alpha_{\pi i}^2 + \lambda\alpha_{yi}^2))^{-1} \times \{\delta i_{t-1} \\ - \beta [(\alpha_{\pi\pi}\alpha_{\pi i} + \alpha_{y\pi}\alpha_{yi})\pi_t + \lambda(\alpha_{\pi y}\alpha_{\pi i} + \alpha_{yy}\alpha_{yi})y_t - \alpha_{\pi i}\pi^*] \\ + \beta Q_\tau(-z_{\pi,t+1}\alpha_{\pi i} - \lambda z_{y,t+1}\alpha_{yi} | z_t)\}.$$

This is similar to the typical Taylor rule derivation as in [Giannoni and Woodford \(2003\)](#). In the standard model, since the expectation of the random shocks is zero, the second term becomes zero.

For QU, however, the quantiles need to be computed on a case-by-case basis. For the location shift case, the quantile index  $\tau$  determines the quantiles of the two shocks in a particular direction given by  $d_\pi z_{\pi,t+1} + d_y z_{y,t+1}$  with  $d_\pi = \frac{-\alpha_{\pi i}}{\sqrt{\alpha_{\pi i}^2 + \lambda^2\alpha_{yi}^2}}$  and  $d_y = \frac{-\lambda\alpha_{yi}}{\sqrt{\alpha_{\pi i}^2 + \lambda^2\alpha_{yi}^2}}$ .

### Location-scale on state variables only

Suppose now that  $\gamma_{ai} = 0$   $a = \pi, y$ , that is, the interest rate exerts no scale effect on the random shocks, which may be affected by the state variables.

Then

$$-\delta(i_t - i_{t-1}) + \beta Q_\tau [-(\alpha_{\pi\pi}\pi_t + \alpha_{\pi y}y_t + \alpha_{\pi i}i_t + h_\pi(\pi_t, y_t)z_{\pi,t+1} - \pi^*)\alpha_{\pi i} \\ - \lambda(\alpha_{y\pi}\pi_t + \alpha_{yy}y_t + \alpha_{yi}i_t + h_y(\pi_t, y_t)z_{y,t+1})\alpha_{yi} | z_t] = 0, \quad (6)$$

$$i_\tau^*(\pi_t, y_t, i_{t-1}) = \\ (\delta + \beta(\alpha_{\pi i}^2 + \lambda\alpha_{yi}^2))^{-1} \times \{\delta i_{t-1} \\ - \beta [(\alpha_{\pi\pi}\alpha_{\pi i} + \alpha_{y\pi}\alpha_{yi})\pi_t + \lambda(\alpha_{\pi y}\alpha_{\pi i} + \alpha_{yy}\alpha_{yi})y_t - \alpha_{\pi i}\pi^*] \\ + \beta Q_\tau(-h_\pi(\pi_t, y_t)z_{\pi,t+1}\alpha_{\pi i} - \lambda h_y(\pi_t, y_t)z_{y,t+1}\alpha_{yi} | z_t)\}.$$

For the location-scale shift case, the quantile index  $\tau$  determines the quantiles of the two shocks in a particular direction given by  $d_\pi z_{\pi,t+1} + d_y z_{y,t+1}$  with

$$d_\pi = -h_\pi(\pi_t, y_t)\alpha_{\pi i} (\alpha_{\pi i}^2 h_\pi(\pi_t, y_t)^2 + \lambda^2 h_y(\pi_t, y_t)^2 \alpha_{yi}^2)^{-1/2}$$

and

$$d_y = -\lambda h_y(\pi_t, y_t) \alpha_{yi} (\alpha_{\pi i}^2 h_\pi(\pi_t, y_t)^2 + \lambda^2 h_y(\pi_t, y_t)^2 \alpha_{yi}^2)^{-1/2}.$$

For this case, the direction is state dependent, that is, the QU analysis is  $(\pi_t, y_t)$ -specific.

### Location-scale on state and control variables

Finally, for the general case when there are no restrictions on  $\gamma_{ab}$ ,  $a, b = \pi, y, i$  we have

$$\begin{aligned} \beta Q_\tau & \left[ -(\alpha_{\pi\pi}\pi_t + \alpha_{\pi y}y_t + \alpha_{yi}i_t + h_\pi(\pi_t, y_t, i_t)z_{\pi,t+1} - \pi^*)(\alpha_{\pi i} + \frac{\partial h_\pi(\pi_t, y_t, i_t)}{\partial i} z_{\pi,t+1}) \right. \\ & \left. - \lambda(\alpha_{yy}\pi_t + \alpha_{yy}y_t + \alpha_{yi}i_t + h_y(\pi_t, y_t, i_t)z_{y,t+1})(\alpha_{yi} + \frac{\partial h_y(\pi_t, y_t, i_t)}{\partial i} z_{y,t+1}) \mid z_t \right] \\ & = 0. \end{aligned} \quad (7)$$

Here, there is no analytical solution because the quantiles will depend on both  $z_a$  and  $z_a^2$ ,  $a, b = \pi, y$ . (Note that for the expected utility case, the expectation of the square is just replaced by its variance, and thus we could still derive a Taylor rule type model).

## 4 Empirical implementation

### 4.1 Algorithm for empirical implementation

Consider time-series data  $\{\pi_t, y_t, i_t\}_{t=0}^T$  and set a target value  $\pi^*$  and parameters  $(\beta, \lambda, \delta)$ . Note that we are implicitly defining that the target value for output gap is 0. Define  $\mathcal{T}$  as a discrete grid on the interval  $(0, 1)$ .

1. Estimate law of motion reduced form VAR(1) models for  $(\pi_t, y_t)$  using  $i_t$  as an exogenous variable to get the  $\alpha$  coefficients.
2. Estimate the skedastic functions  $h_\pi$  and  $h_y$  by running reduced form VAR(1) models of squared OLS residuals  $\{\hat{u}_{\pi,t+1}^2, \hat{u}_{y,t+1}^2\}$  on  $\{\pi_t, y_t, i_t\}$  to get  $\gamma$  coefficients.
3. Compute  $\{\hat{z}_{\pi,t}, \hat{z}_{y,t}\}_{t=1}^T$  stochastic shocks estimates, i.e.  $\hat{z}_{a,t} = \hat{u}_{a,t}/\hat{h}_{a,t}$ ,  $a = \pi, y$ . Then compute the empirical quantiles,  $\hat{Q}_\tau$ .
4. Solve for  $i_\tau^*(\pi_t, y_t, i_{t-1})$  for  $\tau \in \mathcal{T}$ .

Consider now the evaluation of the underlying preferences of the CB. Here we follow an indirect inference procedure. For each time period  $t$ , we can evaluate the optimal response for all quantile indexes and then infer the  $\tau$  that produces the closest value of the observed policy variable. In other words

$$\hat{\tau}_t = \arg \min_{\tau \in \mathcal{T}} |i_t - i_\tau^*(\pi_t, y_t, i_{t-1})|.$$

This procedure delivers an index of the implicit risk aversion quantile preferences of the CB.

## 4.2 Data sources and model calibration

### 4.2.1 Data

The empirical estimation involved in step 1 of algorithm is based on three macroeconomic variables constructed from raw data obtained from the Federal Reserve Economic Data of St. Louis (FRED St. Louis) database. The original dataset comprised Real GDP (GDPC1), Potential GDP (GDPPOT), the Effective Federal Funds Rate (FEDFUNDS), and the Personal Consumption Expenditures Chain-type Price Index (PCECTPI). For consistency across series, monthly observations (FEDFUNDS) were converted to quarterly frequency using arithmetic averages. From these sources, we derived the following variables

1. Output gap ( $y_t$ )

$$y_t = \left( \frac{GDPC1_t}{GDPPOT_t} - 1 \right) * 100 \quad (8)$$

2. Inflation ( $\pi_t$ )

$$\pi_t = 100 \times \Delta \ln(PCECTPI_t) \quad (9)$$

3. Interest rate ( $i_t$ ), proxied by the Effective Federal Funds rate.

Table 1 reports summary statistics of the variables used in this paper.

Table 1: Descriptive Statistics

Statistic	$i$	$y$	$\pi$
Mean	4.62	-0.27	0.78
Minimum	0.06	-9.02	-1.61
1st Quartile	1.94	-1.58	0.41
Median	4.33	-0.20	0.66
3rd Quartile	6.24	1.25	1.01
Maximum	17.78	5.68	2.96

The three constructed series – the output gap ( $y_t$ ), inflation ( $\pi_t$ ), and the nominal interest rate ( $i_t$ ) – are employed in step 1 of the empirical algorithm. Specifically, they serve as the input variables for estimating the reduced-form law of motion VAR(1) models, from which the coefficients  $\alpha$  are obtained. They are then used in step 2 to compute the skedastic functions and the  $z$  shocks components in step 3.

These estimates provide the foundation for the subsequent steps of the empirical implementation. The final dataset used for estimation spans from the fourth quarter of 1954 to the second quarter of 2025, yielding a total of 283 quarterly observations.

The second VAR(1) model estimation incorporates two dummy variables to control for the extraordinary shocks associated with major global crises. The first dummy captures the exceptional effects of the global financial crisis, covering the period from the fourth quarter of 2007 to the fourth quarter of 2009. The second dummy accounts for the economic disruptions linked to the COVID-19 pandemic, spanning from the first quarter of 2020 to the first quarter of 2021. Including these variables ensures that the estimated relationships among the models endogenous variables are not biased by these exceptional and exogenous events.

### 4.2.2 Calibration

We calibrate the remaining parameters of the model in step 4 of the algorithm, following the literature on the Taylor rule for the US. The calibration strategy consists of setting some non-target structural parameters based on empirical evidence.

In particular, following [Dennis \(2004\)](#) the policy discount factor is set to  $\beta = 0.99$ <sup>8</sup>, the relative weight on output gap stabilization is set to  $\lambda = 1$ <sup>9</sup>, and the interest rate smoothing is set to  $\delta = 0.1$ <sup>10</sup>. Finally, we assume that estimates that the average implied inflation target of the Fed is around 2 percent of annual inflation (0.496 percent in quarterly log differences). Among others, [Andrade, Gali, Le Bihan, and Matheron \(2019\)](#) and [Bianchi \(2019\)](#) calibrate their models using an inflation target consistent with the Fed's 2 percent objective.<sup>11</sup>

The baseline calibration is summarized in Table 2.

Table 2: Quarterly calibrated parameters

Parameter	Value	Description	Source
$\beta$	0.99	Policy discount factor	Dennis (2004)
$\lambda$	1	Relative weight on output gap stabilization	Dennis (2004)
$\delta$	0.1	Interest rate smoothing parameter	Sack and Weiland (2000)
$\pi^*$	0.496	Quarterly inflation target	Andrade et al. (2019); Bianchi (2019)

## 5 Results

### 5.1 Baseline model

Tables 3 and 4 report the reduced form VAR(1) models and the skedastic functions, respectively. In both cases we use two different models. A baseline model without COVID and Global Financial Crisis (GFC) dummies and another with those dummies included.

<sup>8</sup>This represents the CB's time preference, indicating how much it values future welfare compared to current welfare. A higher discount factor 0.99 means the CB is more patient and cares more about long-term stability.

<sup>9</sup>The quarterly calibration of lambda for a CB loss function with output equal to 1 involves using the loss function's sensitivity to output deviations to determine the weight (lambda) on output in the loss function, relative to inflation. The value of lambda is adjusted to prioritize output stability and it is a matter of a CB deciding how to weigh output versus inflation.

<sup>10</sup>Interest-rate smoothing is the tendency for CBs, including the Fed, to adjust interest rates in small steps over time. A value of delta equal to  $\delta = 0.1$  in a policy rule represents the weight on a smoothing term in a simplified model, suggesting that about 10 percent of the adjustment in the desired interest rate is reflected in the policy rate within a quarter, indicating a very gradual policy response. The value of 0.1 is a hypothetical calibration that would imply a very fast adjustment compared to the historical norm (historically estimated to be closer to 0.8 in the US), though still gradual. This behavior can stem from optimal policy choices to reduce volatility and manage expectations, or from practical considerations like market reaction and uncertainty (see [Sack and Weiland \(2000\)](#)). For our purposes, it helps in evaluating heterogeneity across quantiles in a relatively short period of time.

<sup>11</sup>Under former Chair Ben Bernanke, the Fed officially adopted a 2 percent inflation target in January 2012. This move brought the Fed in line with many other central banks and was based on a strategy of price stability, aiming for 2 percent inflation as the longer-run goal for achieving both maximum employment and price stability. Although the 2 percent target was made public and official in 2012, the Fed had been operating with a similar goal behind the scenes since 1996.

The VAR model reveals that lagged  $i$  has a positive effect on inflation but a negative effect on output gap. Note that these estimated parameters do not imply a structural relationship among the variables, but they are only a reduced-form result.

In a reduced-form VAR model, the sign of the output gap's coefficient in the inflation equation is typically positive, consistent with the New Keynesian Phillips curve. However, this relationship can be obscured, unstable, or even appear with the wrong sign (i.e. the flat Phillips curve puzzle) due to the nature of reduced-form estimation and the influence of other shocks. While the underlying structural relationship is positive, the sign we get from a simple reduced-form VAR estimation is not a reliable estimate of the Phillips curve slope (see [Mavroeidis, Plagborg-Møller, and Stock \(2014\)](#)).

A similar argument applies for understanding the sign of the interest rate on inflation. In that regard, the price puzzle is the empirical finding that interest rate hikes can be followed by rising inflation (see [Sims \(1992\)](#)). It is primarily a statistical illusion caused by the econometrician's model failing to account for the fact that the CB is raising rates in anticipation of future inflation. When models are properly specified to include the CB's information, the puzzle usually vanishes<sup>12</sup>, and the standard theoretical relationship holds (see [Stock and Watson \(2001\)](#); [Bernanke and Mihov \(1998\)](#); and [Leeper, Sims, and Zha \(1996\)](#)).

In addition, our empirical findings indicate that inflation and output gap have high autoregressive coefficients (close to 0.7 for inflation, 0.9 for output gap). The inclusion of the COVID and the GFC dummies do not change the sign of the estimated coefficients.

For the skedastic function, a preliminary analysis (not reported; available upon request) reveals that lagged  $i$  is not statistically significant in the skedastic functions, and therefore we impose that  $\gamma_{ai} = 0$   $a = \pi, y$  for the computation of the optimal Taylor rule. In turn, this determines that we follow the location-shift model with control variables only in the skedastic function, and that analytical derivations can be used.

---

<sup>12</sup>A typical way to eliminate the price puzzle would be to create a monetary policy shock series free of it using the high-frequency identification method or the proxy SVAR methodology of [Gertler and Karadi \(2015\)](#). However, the Taylor rule aims to capture the Fed's endogenous response. Using the shock series, which is the exogenous component, would eliminate our object of study.

Table 3: VAR(1) Results Full Sample

	Baseline		Baseline + dummies	
	Inflation ( $\pi$ )	Output gap ( $y$ )	Inflation ( $\pi$ )	Output gap ( $y$ )
$i_{t-1}$	0.024*** (0.008)	-0.029 (0.022)	0.024*** (0.008)	-0.037* (0.022)
$\pi_{t-1}$	0.719*** (0.045)	-0.130 (0.125)	0.718*** (0.045)	-0.149 (0.123)
$y_{t-1}$	0.006 (0.010)	0.904*** (0.027)	0.007 (0.010)	0.891*** (0.027)
Constant	0.114*** (0.037)	0.215** (0.105)	0.110*** (0.039)	0.314*** (0.107)
COVID			0.131 (0.164)	-0.922** (0.451)
GFC			-0.025 (0.123)	-0.979*** (0.337)
Observations	282	282	282	282
R <sup>2</sup>	0.668	0.804	0.669	0.812
Adjusted R <sup>2</sup>	0.664	0.802	0.663	0.808

Note: \*p&lt;0.1; \*\*p&lt;0.05; \*\*\*p&lt;0.01

Table 4: Skedastic Models - Full Sample

	Baseline		Baseline + dummies	
	$\hat{u}_\pi^2$	$\hat{u}_y^2$	$\hat{u}_\pi^2$	$\hat{u}_y^2$
$\pi_{t-1}$	0.045 (0.040)	-0.311 (0.463)	0.078** (0.038)	0.101 (0.337)
$y_{t-1}$	-0.010 (0.011)	-0.217* (0.127)	-0.001 (0.010)	-0.111 (0.093)
COVID			0.355** (0.176)	19.274*** (1.568)
GFC			0.775*** (0.132)	0.305 (1.175)
Constant	0.087** (0.041)	1.153** (0.464)	0.033 (0.039)	0.467 (0.345)
Observations	281	281	281	281
R <sup>2</sup>	0.007	0.013	0.126	0.368
Adjusted R <sup>2</sup>	-0.0004	0.006	0.113	0.359

Note: \*p&lt;0.1; \*\*p&lt;0.05; \*\*\*p&lt;0.01

Figure 4a plots the observed Effective Federal Funds Rate (EFFR) together with the optimal interest rate response for a QU maximizer CB (i.e. the Fed), for different representative quantile indexes  $\tau \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$ . The graph represents the wide variety of optimal conditional reactions that may arise for any given QU preference.

Figure 4b plots the implied  $\tau$  that represent the closest match to the observed interest rate using a discrete grid search  $\tau \in \{0.01, 0.02, \dots, 0.98, 0.99\}$ . Overall, the results indicate that most of the time the Fed has an implied behavior that is consistent with high values of  $\tau$ . However, in some periods, the implied  $\tau$  is drastically reduced.

Figures 5a and 5b presents the same exercise for the model with GFC and COVID dummies. Note that the results are very similar to the baseline model, thus suggesting that the periods highlighted by the dummies are not driving the main results.

Figure 4: Baseline results: Taylor interest-rate rule and Fed risk aversion.

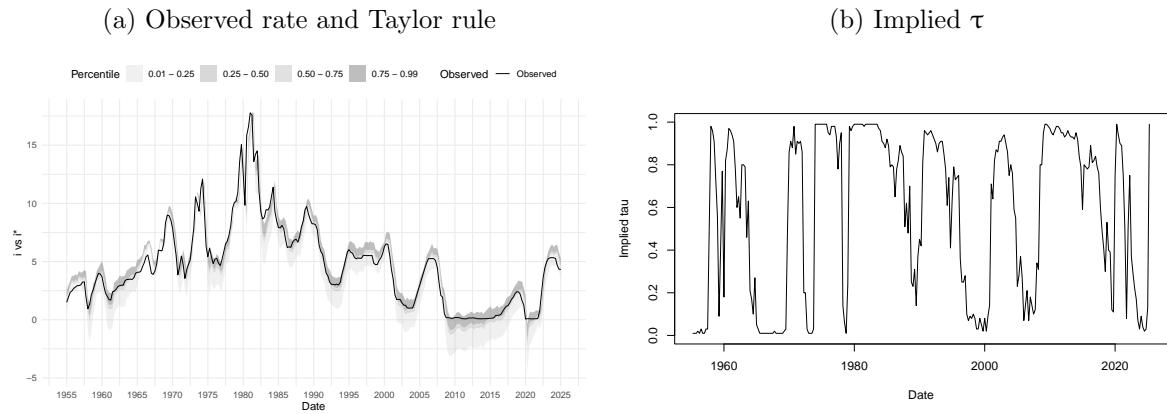
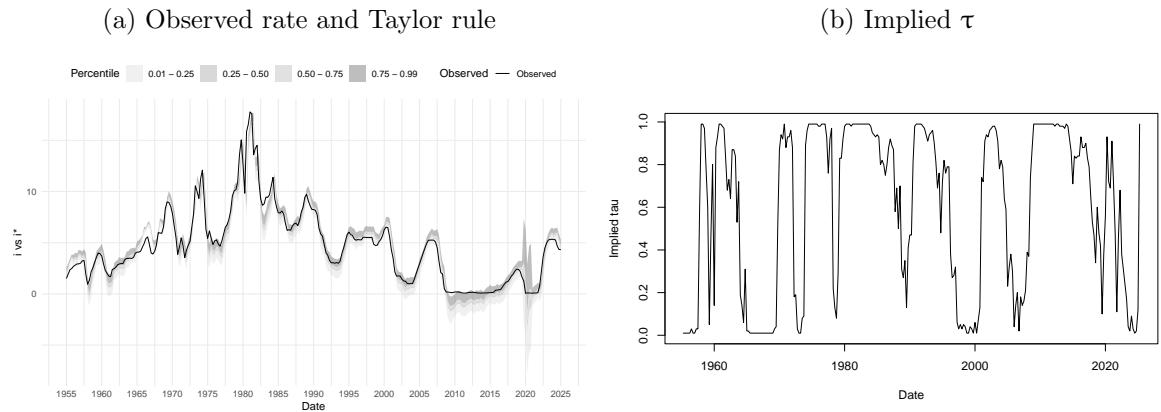


Figure 5: Results with dummies: Taylor interest-rate rule and Fed risk aversion.



In economic terms, a higher  $\tau$  means a relatively lower risk aversion by the Fed's authorities. In turn, a lower risk aversion relates to an implied QU preference of the Fed that gives more weight to good macroeconomic outcomes. These findings are consistent with the empirical evidence that shows a significant reduction in inflation and output volatility during the Great Moderation period (see [Stock and Watson \(2012\)](#), [Bernanke \(2004\)](#)). By contrast, relatively lower values of  $\tau$  means higher risk aversion from monetary policy authorities and more weight

on potential QU losses.

Here are a few notable instances where the Fed raised the EFFR dramatically. Generally speaking, these are associated with lower  $\tau$  values and shifts in the conduct of US policymakers with respect to their increased risk aversion to unfavorable macroeconomic outcomes. More specifically, as more hawkish positions are adopted.

The Great Inflation Wars, which took place in the late 1960s and early 1980s: the Fed's battle against consistently high inflation, which resulted in the most dramatic rate hikes in its history, characterized this era. In order to fight inflation brought on by the Vietnam War and the 1973 OPEC oil embargo, the Fed started hiking interest rates. From about 3.5 per cent in 1972 to a peak of almost 13 percent in July 1974, the EFFR increased. The inflation rate rose into double digits between 1977 and 1980. The Fed grew more active under Fed Chairs Paul Volcker and G. William Miller. Beginning at around 7 percent in early 1977, the rate skyrocketed to a startling 20 percent in April 1980. Inflation remained high following a temporary easing. Paul Volcker is renowned for creating a deep recession in order to permanently curb inflation. The Fed pushed the rate up from about 10 percent in the middle of 1980 to a second high of 19 percent in June 1981. The Fed's most well-known and forceful tightening cycle to date is this one.

The Preserving Trustworthiness (1983–1984): now that inflation was declining, the Fed had to demonstrate that it was determined to keep it that way while the economy improved. To avoid a recurrence of inflation, the Fed increased the funds rate from about 8.5 percent to about 11.5 percent.

The 1987–1989 Soft-Landing Attempt: as inflation pressures started to rise once more, the Fed tightened policy under new Chair Alan Greenspan. The rate increased to almost 9.75 percent from about 6.5 percent.

The 1994–1995 Preemptive Strike: this is an excellent illustration of a proactive approach to combating inflation. The Fed raised rates before inflation had occurred because it was worried about future inflation and the economy was rapidly improving. The EFFR increased from about 3 percent to roughly 6 percent over a series of sharp increases. A gentle landing is a common term used to describe this successful technique.

The 1999–2000 Tech Bubble Era: the Fed tightened monetary policy as a result of the rising economy and concerns about asset bubbles. The funding rate increased from approximately 4.75 percent to roughly 6.5 percent.

The Measured Pace (2004–2006): the Fed started a lengthy, predictable cycle of 0.25 percent raises to normalize rates after lowering them to historic lows of 1 percent after the dot-com bust. The funds rate rose from 1 percent to 5.25 percent at the same measured pace of 17 straight rises.

The Liftoff Following the Financial Crisis (2015–2018): the Fed started a very modest and cautious tightening cycle after seven years at zero following the 2008 crisis. The rate increased from about 0.25 percent to a peak of over 2.5 percent in 2018 through a series of modest, well-publicized increases.

The Battle Against Post-Pandemic Inflation (2022–2023): the Fed started its most severe tightening cycle since the 1980s in reaction to the biggest inflation in 40 years, which was caused by supply chain problems, the war in Ukraine, and pandemic stimulation. The Fed increased the EFFR from approximately 0.25 percent to a target range of 5.25 to 5.50 percent in less than a year, and it is still at this level as of mid-2024.

These facts are revealing in terms of showing that optimal interest rate response by the

Fed relates to a specific value of  $\tau$  and, accordingly, to a certain degree aversion to undesirable macroeconomic scenarios in terms of inflation and economic activity deviations from their target values.

Additionally, the implied  $\tau$  series also report a consistent behavior with the US monetary policy history observed in the estimation period (see Figure 4b). We observe how higher values of  $\tau$  and low degrees of risk aversion by the Fed are the norm, with exceptional lower values of  $\tau$  and higher degrees of risk aversion related to macroeconomic critical and known events (for a review of the history of monetary policy rules in the US, see [Taylor \(1999\)](#)).

These empirical findings allow us to remark some novel evidence. The dovish or hawkish Fed's behavior varies through time according to the economic, institutional and political context prevailing in the US (see [Eijffinger and Masciandaro \(2018\)](#)). Although we observe that during the great part of estimation period (1959-2025), the Fed's monetary policy authorities display a conduct that approximates more to dovish patterns, we also notice that in critical circumstances, the Fed shows a hawkish stance.

Our theoretical and empirical contributions differ from the existent ones in the following terms. We introduce a more flexible theoretical framework that maps the dovish/hawkish stances of monetary policy regarding not only the variations of interest rates. Indeed, we consider a more complex analytical framework which allows to define an undesirable macroeconomic scenario in terms of inflation and output deviations from optimal targets jointly.

## 5.2 Robustness and sensitivity analysis

To assess the robustness of our empirical findings, we conducted a battery of sensitivity checks focusing on two main aspects: (i) the relative weight assigned to output stabilization in the CB loss function, parameterized by  $\lambda$ , and (ii) potential structural breaks in monetary policy behavior associated with the Volcker regime shift. The results appear in the Tables 5 and 6 and Figure 6.

First, we explored the sensitivity of the estimated Taylor-type quantile rule to alternative values of  $\lambda \in \{0.5, 1, 2\}$ . This exercise allows us to gauge whether the implied policy stance and the inferred degree of risk aversion are contingent upon the assumed trade-off between inflation and output stabilization. The results remained qualitatively stable across specifications. Lower values of  $\lambda$  (0.5) led to slightly more aggressive responses to inflation deviations, reflecting a relatively more hawkish stance, while higher values (2) induced smoother interest rate paths, consistent with greater tolerance toward output fluctuations. Nevertheless, the implied quantile-based preferences ( $\tau_t$ ) maintained the same cyclical pattern, confirming that the estimated policy reaction functions are not overly sensitive to moderate changes in the structural weighting scheme.

Second, to account for potential regime shifts in the conduct of US monetary policy, we re-estimated the baseline VAR and conditional heteroskedasticity models using a truncated sample starting in the fourth quarter of 1979. This subsample captures the onset of the Volcker disinflation episode, a well-documented structural change in the Fed's reaction to inflationary pressures. The results reported show that the estimated coefficients remain broadly consistent with those of the full sample, with minor quantitative adjustments reflecting a stronger disinflationary response of the Fed in the post-Volcker era. The inferred quantile preferences confirm a temporary decline in  $\tau_t$  during the early 1980s, signaling a shift towards higher risk aversion and more hawkish policy attitudes, followed by a gradual return to higher  $\tau_t$  values consistent with a more dovish stance in subsequent decades.

Table 5: VAR(1) Results - Post 1979

	Baseline		Baseline + dummies	
	Inflation ( $\pi$ )	Output gap (y)	Inflation ( $\pi$ )	Output gap (y)
$i_{t-1}$	0.023*** (0.008)	-0.017 (0.023)	0.023*** (0.009)	-0.026 (0.023)
$\pi_{t-1}$	0.615*** (0.061)	-0.093 (0.165)	0.614*** (0.061)	-0.128 (0.162)
$y_{t-1}$	-0.016 (0.014)	0.879*** (0.037)	-0.015 (0.014)	0.863*** (0.037)
Constant	0.144*** (0.048)	0.054 (0.130)	0.145*** (0.050)	0.177 (0.133)
COVID			0.074 (0.171)	-0.872* (0.452)
GFC			-0.051 (0.128)	-0.914*** (0.338)
Observations	182	182	182	182
R <sup>2</sup>	0.545	0.760	0.546	0.773
Adjusted R <sup>2</sup>	0.537	0.756	0.533	0.766

Note:

\*p&lt;0.1; \*\*p&lt;0.05; \*\*\*p&lt;0.01

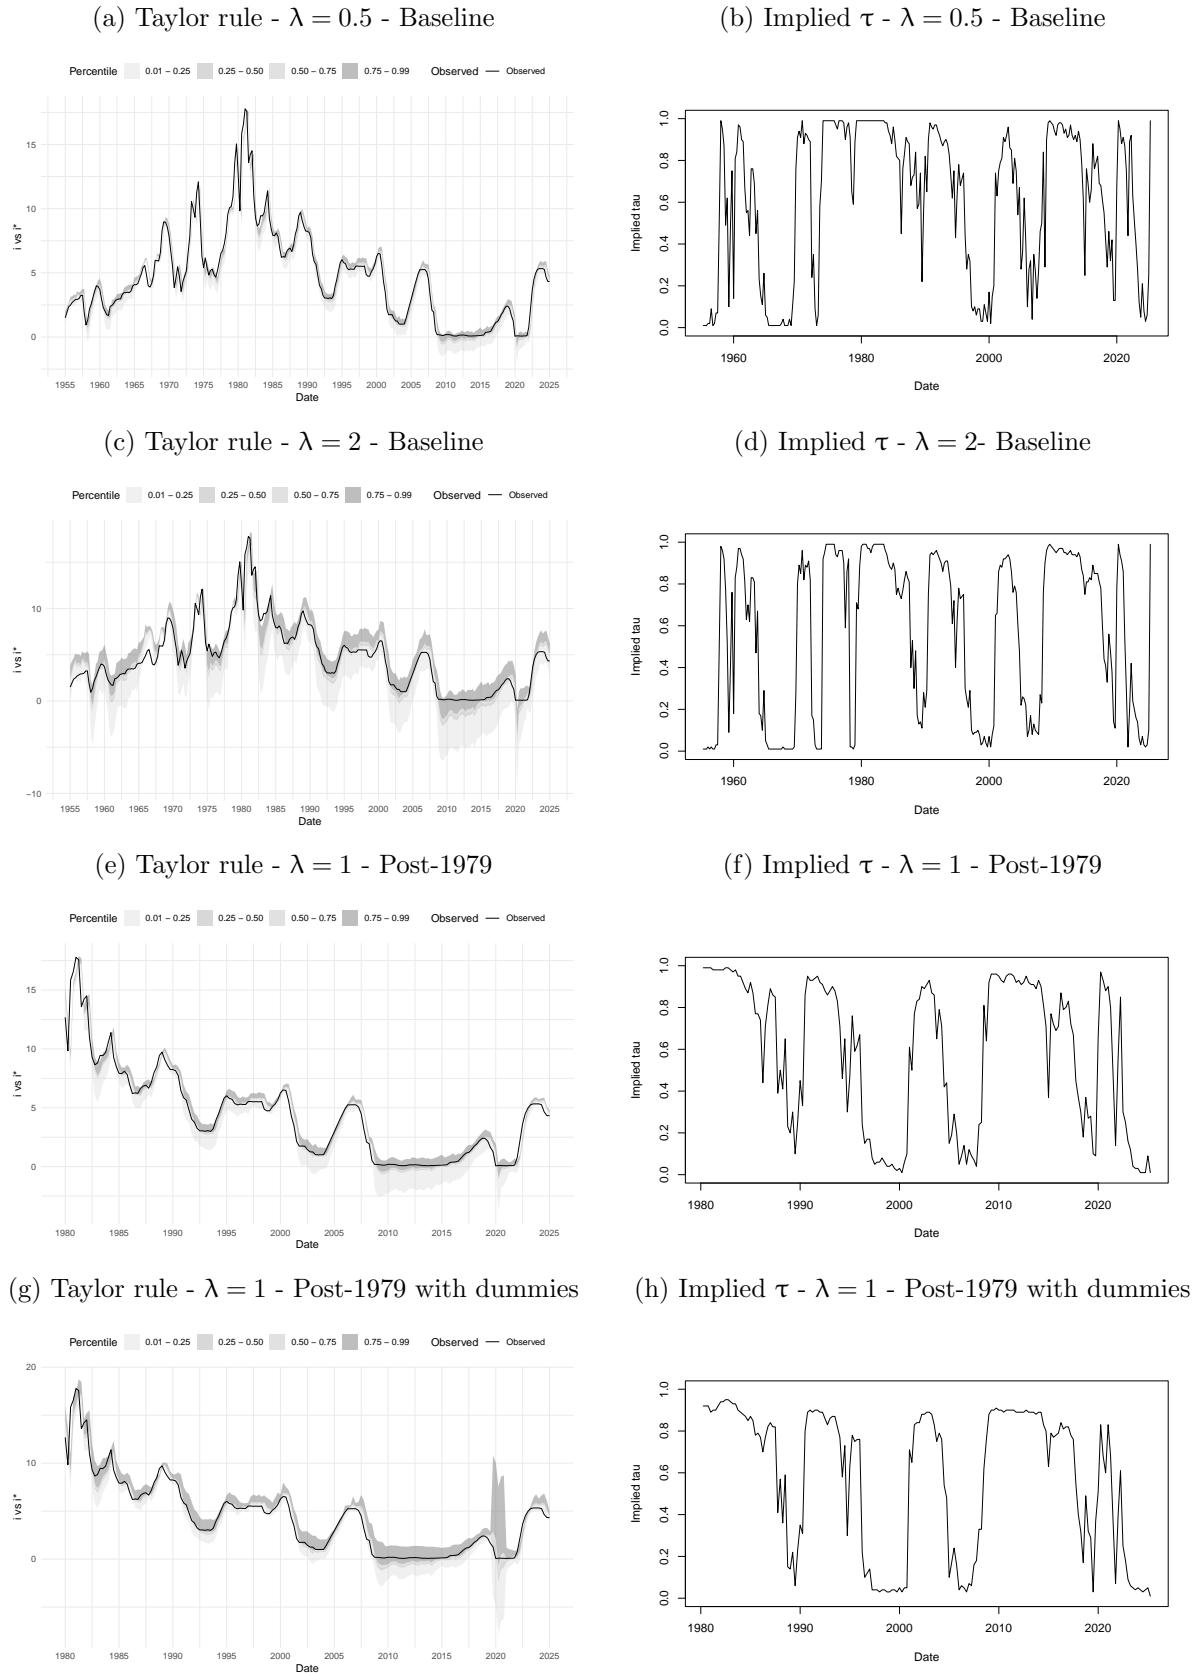
Table 6: Skedastic Models - Post 1979

	Baseline		Baseline + dummies	
	$\hat{u}_\pi^2$	$\hat{u}_y^2$	$\hat{u}_\pi^2$	$\hat{u}_y^2$
$\pi_{t-1}$	0.059 (0.065)	-0.514 (0.781)	0.104* (0.059)	0.131 (0.560)
$y_{t-1}$	-0.004 (0.017)	-0.358* (0.209)	0.006 (0.016)	-0.216 (0.150)
COVID			0.294 (0.195)	18.803*** (1.836)
GFC			0.708*** (0.146)	0.483 (1.378)
Constant	0.082 (0.058)	1.049 (0.707)	0.015 (0.055)	0.119 (0.518)
Observations	181	181	181	181
R <sup>2</sup>	0.005	0.020	0.128	0.396
Adjusted R <sup>2</sup>	-0.006	0.009	0.108	0.382

Note:

\*p&lt;0.1; \*\*p&lt;0.05; \*\*\*p&lt;0.01

Figure 6: Robustness and sensitivity analysis



Overall, both exercises confirm the internal consistency and empirical robustness of our quantile-based Taylor rule. The results suggest that the main findings, the predominance of dovish-type behavior with episodic hawkish responses are not artifacts of parameter calibration or sample selection, but rather reflect persistent structural features of the US monetary policy.

## 6 Conclusions and discussion

The study of Taylor rules through quantile methods highlights that a one-size-fits-all linear rule may be inadequate, with more nuanced, quantile-aware approaches needed to understand and formulate policy, especially in diverse economic environments. A CB may not react linearly to economic variables; its responses can vary significantly at different quantiles of the distribution.

The QU framework allows to study how independent variables (like policy tools) impact the dependent variable (like interest rates) differently across the entire distribution of outcomes, revealing heterogeneity in policy reactions. QU models explore how agents make decisions under uncertainty by focusing on outcomes at different parts of the probability distribution, while Taylor rules describe how the CBs set interest rates based on inflation and economic output.

In this paper, we study optimal monetary policy when a CB maximizes a QU objective rather than expected utility operator. In our framework, the CB's risk attitude is indexed by the quantile level  $\tau$ , providing a transparent mapping between hawkish/dovish stances and attention to adverse macroeconomic realizations. We formulate the infinite-horizon problem using a Bellman equation with the quantile operator. Implementing an Euler-equation approach, we derive Taylor-rule-type reaction functions. The Taylor rule is recovered as a special case when quantiles replace the expectation operator. An empirical implementation is outlined based on reduced-form VAR(1) laws of motion with conditional heteroskedasticity, enabling estimation of the new rule and its dependence on risk attitudes. It is important to note that our analytical methodology makes the assumption that the policymaker's statements are entirely credible. Determining optimality criteria for establishing a Taylor rule inside the quintile preference framework in the setting of monetary policy pronouncements lacking credibility is a tenable extension of this study.

If the CB is relatively more risk averse, it would react more aggressively to downside risks than to equivalent upside risks. If the CB is relatively less risk averse, policies are guided by the possibility of good economic results. This could justify the risk management approach often discussed at CBs.

Quantile preferences naturally incorporate concerns about tail risks that expected utility might underweight. Moreover, risk attitudes might not be constant across time. Our model allows us to identify risk aversion behavior through indirect inference, by mapping the observed policy variables with the corresponding value in the optimal CB behavior.

Our empirical results for the US show that the Fed has mostly a dovish attitude over long periods of time (higher  $\tau$ ), but with hawkish attitudes in specific periods (lower  $\tau$ ). These periods coincide with regime changing events, like the oil crisis in the middle of the 1970s, the Volcker new approach towards fighting inflation at the end of the 1970s, the global financial crash of 2008, and the post COVID pandemic. As such, the implied risk aversion estimates reveal important changes in the Fed preferences.

The policy implications are that focusing solely on the mean in economic modeling, as many traditional Taylor rules applications do, can be misleading, as non-linear relationships

and risk preferences (represented by QU) can significantly alter optimal monetary policy and under specific conditions.

## References

AL NOWAIHI, A., AND L. STRACCA (2002): “Non-Standard Central Bank Loss Functions, Skewed Risks, and Certainty Equivalence,” ECB Working Paper 129, European Central Bank.

ANDRADE, P., J. GALI, H. LE BIHAN, AND J. MATHERON (2019): “The Optimal Inflation Target and the Natural Rate of Interest,” *Brookings Papers on Economic Activity*, 50(2), 173–255.

BARCI, G. (2025): “The effects of monetary policy on macroeconomic downside risk: state-dependence matters,” *Journal of Economic Dynamics & Control*, Journal Pre-Proof, <https://doi.org/10.1016/j.jedc.2025.105201>.

BENHABIB, J., S. SCHMITT-GROHÉ, AND M. URIBE (2001): “The Perils of Taylor Rules,” *Journal of Economic Theory*, 96, 40–69.

BENIGNO, P., AND G. E. EGGERTSSON (2023): “It is baaack: The surge in inflation in the 2020s and the return of the non-linear Phillips curve,” NBER Working Paper 31197.

BERNANKE, B. S. (2004): “The Great Moderation,” Remarks at the meetings of the Eastern Economic Association.

BERNANKE, B. S., AND I. MIHOV (1998): “Measuring Monetary Policy,” *Quarterly Journal of Economics*, 113(3), 869–902.

BIANCHI, F. (2019): “Hitting the Elusive Inflation Target,” NBER Working Paper 26279, National Bureau of Economic Research.

BOEHM, C., AND C. L. HOUSE (2014): “Optimal Taylor rules in New Keynesian models,” NBER Working Paper 20237.

BOMMIER, A., A. KOCHOV, AND F. LE GRAND (2017): “On Monotone Recursive Preferences,” *Econometrica*, 85, 1433–1466.

CASTRO, V. (2011): “Can central banks’ monetary policy be described by a linear (augmented) Taylor Rule or by a nonlinear rule?,” *Journal of Financial Stability*, 7(4), 228–246.

CHAMBERS, C. P. (2009): “An Axiomatization of Quantiles on the Domain of Distribution Functions,” *Mathematical Finance*, 19, 335–342.

CHEN, J.-E., AND M. KASHIWAGI (2017): “The Japanese Taylor rule estimated using censored quantile regressions,” *Empirical Economics*, 52, 357–371.

CHEVAPATRAKUL, T., T.-H. KIM, AND P. MIZEN (2009): “The Taylor principle and monetary policy approaching a zero bound on nominal rates: Quantile regression results for the United States and Japan,” *Journal of Money, Credit and Banking*, 41(8), 1705–1723.

CHEVAPATRAKUL, T., AND J. PAEZ-FARRELL (2014): “Monetary Policy Reaction Functions in Small Open Economies: A Quantile Regression Approach,” *The Manchester School*, 82(2), 237–256.

CHRISTOU, C., R. NARAIODOO, R. GUPTA, AND W. J. KIM (2018): “Monetary policy reaction functions of the TICKs: A quantile regression approach,” *Emerging Markets Finance and Trade*, 54(15), 3552–3565.

CLARIDA, R., J. GALÍ, AND M. GERTLER (1999): “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37, 1661–1707.

COCHRANE, J. H. (2007): “Determinacy and Identification with Taylor Rules,” NBER Working Paper 13409.

COGLEY, T., R. COLACITO, L. P. HANSEN, AND T. J. SARGENT (2008): “Robustness and U.S. Monetary Policy Experimentation,” *Journal of Money, Credit and Banking*, 40(8), 1599–1623.

DE CASTRO, L., AND A. F. GALVAO (2019): “Dynamic Quantile Models of Rational Behavior,” *Econometrica*, 87, 1893–1939.

——— (2022): “Static and Dynamic Quantile Preferences,” *Economic Theory*, 73, 747–779.

DE CASTRO, L., A. F. GALVAO, AND D. NUNES (2025): “Dynamic economics with quantile preferences,” *Theoretical Economics*, 20, 353–425.

DEJARNETTE, P., D. DILLENBERGER, D. GOTTLIEB, AND P. ORTOLEVA (2020): “Time Lotteries and Stochastic Impatience,” *Econometrica*, 88, 619–656.

DEMIRGUC-KUNT, A., E. DETRAGIACHE, AND O. MERROUCHE (2013): “Bank capital: Lessons from the financial crisis,” *Journal of Money, Credit and Banking*, 45(6), 1147–1164.

DENNIS, R. (2004): “The Policy Preferences of the U.S. Federal Reserve,” Working Paper 2001-08, Federal Reserve Bank of San Francisco.

EIJFFINGER, S., AND D. MASCIANDARO (2018): *Hawks and Doves: Deeds and Words Economics and Politics of Monetary Policymaking*. CEPR Press.

EL-SHAGI, M. (2025): “Does the Fed adhere to its mandate? Estimating the Federal Reserve’s objective function,” *Economic Analysis and Policy*, 87, 1782–1796.

EPSTEIN, L. G., AND M. SCHNEIDER (2003): “Recursive Multiple-Priors,” *Journal of Economic Theory*, 113, 1–31.

EPSTEIN, L. G., AND S. E. ZIN (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 57, 937–969.

——— (1991): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Investigation,” *Journal of Political Economy*, 99(2), 263–286.

GALÍ, J. (2015): *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, Princeton, NJ.

GERTLER, M., AND P. KARADI (2015): “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, 7(1), 44–76.

GIANNONI, M. P., AND M. WOODFORD (2003): “Optimal Interest-Rate Rules: I. General Theory,” NBER Working Paper 9419.

GIOVANNETTI, B. C. (2013): “Asset Pricing under Quantile Utility Maximization,” *Review of Financial Economics*, 22, 169–179.

GONZÁLEZ-ASTUDILLO, M., AND R. TANVIR (2023): “Hawkish or Dovish Fed? Estimating a Time-Varying Reaction Function of the Federal Open Market Committees Median Participant,” Federal Open Market Committees Median Participant, Finance and Economics Discussion Series 2023-070, Board of Governors of the Federal Reserve System.

GRANT, S., A. KAJII, AND B. POLAK (2000): “Temporal Resolution of Uncertainty and Recursive Non-Expected Utility Models,” *Econometrica*, 68(2), 425–434.

HA, J., M. A. KOSE, F. OHNSORGE, AND H. YILMAZKUDAY (2024): “What explains global inflation,” *IMF Economic Review*.

HACK, L., K. ISTREFI, AND M. MEIER (2025): “The Systematic Origins of Monetary Policy Shocks,” Working Paper.

HALL, R. E. (1988): “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, 96, 339–357.

HALLIN, M., AND D. KONEN (2024): “Multivariate Quantiles: Geometric and Measure-Transportation-Based Contours,” arXiv:2401.02499.

HALLIN, M., D. PAINDAVEINE, AND M. ŠIMAN (2010): “Multivariate quantiles and multiple-output regression quantiles: From  $L_1$  optimization to halfspace depth,” *The Annals of Statistics*, 38(2), 635–669.

HANSEN, L. P., AND T. J. SARGENT (2004): “Misspecification in Recursive Macroeconomic Theory,” manuscript.

HILLS, T., T. NAKATA, AND T. SUNAKAWA (2020): “A Promised Value Approach to Optimal Monetary Policy,” *Oxford Bulletin of Economics and Statistics*, 83(1), 173–198.

HOFMANN, B., C. MANEA, AND B. MOJON (2024): “Targeted Taylor rules: Monetary policy responses to demand- and supply-driven inflation,” Discussion paper, Bank for International Settlements.

KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2009): “Recursive Smooth Ambiguity Preferences,” *Journal of Economic Theory*, 144(3), 930–976.

KREPS, D. M., AND E. L. PORTEUS (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46, 185–200.

LEEPER, E. M., C. A. SIMS, AND T. ZHA (1996): “What Does Monetary Policy Do?,” *Brookings Papers on Economic Activity*, 1996(2), 1–78.

MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): “Dynamic Variational Preferences,” *Journal of Economic Theory*, 128, 4–44.

MALMENDIER, U., S. NAGEL, AND Z. YAN (2021): “The making of hawks and doves,” *Journal of Monetary Economics*, 117, 19–42.

MANSKI, C. F. (1988): “Ordinal Utility Models of Decision Making under Uncertainty,” *Theory and Decision*, 25, 79–104.

MARINACCI, M., AND L. MONTRUCCHIO (2010): “Unique Solutions for Stochastic Recursive Utilities,” *Journal of Economic Theory*, 145, 1776–1804.

MAVROEIDIS, S., M. PLAGBORG-MØLLER, AND J. H. STOCK (2014): “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve,” *Journal of Economic Literature*, 52(1), 124–188.

MONTES-ROJAS, G. (2017): “Reduced Form Vector Directional Quantiles,” *Journal of Multivariate Analysis*, 158, 20–30.

——— (2019): “Quantile Impulse Response Functions,” *Journal of Time Series Analysis*, 40, 739–752.

——— (2022): “Estimating Impulse-Response Functions for Macroeconomic Models using Directional Quantiles,” *Journal of Time Series Econometrics*, 14(2), 199–225.

NAKAMURA, E., V. RIBLIER, AND J. STEINSSON (2025): “Beyond the Taylor Rule,” NBER Working Paper 34200.

ROMANO, J. P., AND M. WOLF (2017): “Resurrecting Weighted Least Squares,” *Journal of Econometrics*, 197, 1–19.

ROSTEK, M. (2010): “Quantile Maximization in Decision Theory,” *Review of Economic Studies*, 77, 339–371.

SACK, B., AND V. WEILAND (2000): “Interest-rate smoothing and optimal monetary policy: a review of recent empirical evidence,” *Journal of Economics and Business*, 52(1–2), 205–228.

SARGENT, T. S., AND N. WILLIAMS (2025): “Conquest Lost and Regained: American Inflation in the 2020s,” Mimeo graph.

SARVER, T. (2018): “Dynamic Mixture-Averse Preferences,” *Econometrica*, 86, 1347–1382.

SIMS, C. A. (1992): “Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy,” *European Economic Review*, 36(5), 975–1000.

STOCK, J. H., AND M. W. WATSON (2001): “Vector Autoregressions,” *Journal of Economic Perspectives*, 15(4), 101–115.

——— (2012): “Disentangling the Channels of the 2007–09 Recession,” *Brookings Papers on Economic Activity*, Spring 2012, 81–156.

STRZALECKI, T. (2013): “Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion,” *Econometrica*, 81, 1039–1074.

SVENSSON, L. E. (2003): “What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules,” *Journal of Economic Literature*, 41, 426–477.

TAYLOR, J. B. (1993): “Discretion versus policy rules in practice,” *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

——— (1999): *Monetary Policy Rules*. University of Chicago Press.

TOBACK, E., S. NARDELLI, AND D. MARTENS (2017): “Between hawks and doves: Measuring central bank communication,” ECB Working Paper Series 2085, European Central Bank.

WEIL, P. (1990): “Nonexpected Utility in Macroeconomics,” *Quarterly Journal of Economics*, 105, 29–42.

WILSON, L. (2020): “A dove to hawk ranking of the Martin to Yellen federal reserves,” *Southwestern Economic Review*, 47, 101–109.

WOLTERS, M. H. (2012): “Estimating monetary policy reaction functions using quantile regressions,” *Journal of Macroeconomics*, 34(2), 342–361.

WOODFORD, M. (2003): *Interest and Prices*. Princeton University Press, Princeton, NJ.