Risk-Adjusted Optimal Policy for Scenario Analysis*

Isabel Cairó† Christopher Gust James Hebden

Edward Herbst Scott Konzem Giovanni Nicolò

May 2025

Abstract

Policy institutions often address uncertainty around a baseline outlook through a limited set of alternative scenarios, rather than specifying a complete predictive distribution. This paper examines optimal policy prescriptions within this setting with a focus on the effect of uncertainty about the outlook. We propose a framework where policymakers form beliefs about the likelihood of a baseline economic path and alternative scenarios, updating these beliefs over time using Bayes' rule. This "risk-adjusted optimal policy" allows policymakers to set policy rationally while learning which scenario is unfolding. We apply this framework to the U.S. economy at the end of 2024, considering risks of higher inflation and recession around a baseline projection.

JEL: E32, E47, D81, C63

Keywords: Optimal monetary policy, risk management, scenario analysis, Bayesian learning.

*We thank Ben Johannsen, David López-Salido and Bob Tetlow for valuable comments, and Jose Nino for excellent research assistance.
The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Board or the Federal Reserve System.

†The affiliation for all the authors is the Board of Governors of the Federal Reserve System.
E-mails: Cairó: isabel.cairo@frb.gov; Gust: christopher.j.gust@frb.gov; Hebden: james.s.hebden@frb.gov; Herbst: edward.p.herbst@frb.gov; Konzem: scott.r.konzem@frb.gov; Nicolò: giovanni.nicolo@frb.gov.

1 Introduction

Risk management is central to the conduct of monetary policy. In view of Greenspan (2004), it requires "*understanding as much as possible the many sources of risk and uncertainty that policymakers face, quantifying those risks when possible, and assessing the costs associated with each of the risks.*" In light of those risks, monetary policy devises a strategy aimed at "*maximizing the probabilities of achieving over time our goals of price stability and the maximum sustainable economic growth that we associate with it.*"

There are many frameworks for describing risks. In academic settings, it is common to characterize risk by specifying a complete predictive distribution or by focusing on worst-case outcomes under certain assumptions. In many policy institutions, however, it's common to use *scenario analysis* to capture alternative possible outcomes around a baseline (modal) forecast for both policy deliberations and communications. These scenarios can be seen as alternative conditional forecasts under different assumptions about the shocks likely to hit the economy, structural features in the economy, the nature of the expectations process, and more. Alternative scenarios have the benefits of providing a tangible description of relevant risks but their use poses challenges for optimal monetary policy analysis. Traditionally, these policies are derived using a macroeconometric model under the assumption that given forecast (baseline or alternative scenario) will materialize with certainty. Such policies serve as useful benchmarks for policymakers, but, except in special cases, omit meaningful risk-management considerations. The policy-rate prescriptions in each simulation do not account for the possibility alternative scenarios might materialize.

We propose an approach to embed risk-management considerations into the design of an optimal monetary policy strategy. The central feature of our approach is that policymakers form beliefs about the likelihood of the baseline path that the economy might follow, as well as a finite number of potential alternative scenarios capturing risks around that baseline path. As such, policymakers do not know which scenario is unfolding in real time and use all available information to set policy rationally, formulating beliefs about the alternative paths of economic outcomes. In particular, policymakers begin each period with prior beliefs that reflect all the available information up to that point and update their beliefs as new information arrives using Bayes' rule. Because policymakers set policy optimally while learning over time whether the contemplated risks are materializing or not, we call this policy a "risk-adjusted optimal policy."

Under risk-adjusted optimal policy, the optimal path for the policy rate depends on policymakers' beliefs about the likelihood of the various scenarios. And, because there are lags in the transmission of monetary policy to economic activity and inflation, the optimal setting of the current policy rate needs to be forward-looking and take into account the likelihood of future shocks that might push policymakers' target variables away from (or in line with) their macroeconomic stabilization objectives. As a result, under risk-adjusted optimal policy, policymakers have the incentive to "take out insurance" by setting the current policy rate to lean against potential risks that could emerge in the future and affect the achievement of their objectives. The extent of insurance (or risk-adjusted policy) hinges on several key elements. First, it depends on the scenarios themselves and how large of a departure they are from the baseline. Second, it depends on policymakers' priors (or initial beliefs) regarding the likelihood of the various scenarios. Third, it depends on the volatility of unanticipated shocks, which affects policymakers' ability to distinguish between likely paths along which the economy may evolve. If these unanticipated shocks are more volatile, it becomes more difficult for policymakers to learn over time which scenario is unfolding. We provide an intuition for these results using a canonical New-Keynesian (NK) model *à la* Galí (2008).

The proposed framework can offer policy-relevant insights when designing an optimal policy that embeds risk-management considerations. We focus on the period at the end of 2024 when the U.S. economy faced risks to the possibility that it continued its progress toward the goals of the Federal Open Market Committee (FOMC or the Committee) to achieve maximum employment and price stability—a possibility sometimes referred to as "soft landing." By then, inflation had made progress towards the Committee's objective, and economic activity had continued to expand at a solid pace. However, as noted in the minutes to the December 2024 FOMC meeting, participants highlighted upside risks to inflation and downside risks to the labor market. Against this backdrop, we use the median responses of the December 2024 Summary of Economic Projections (SEP) to construct a baseline projection for the U.S. economy. Moreover, in line with the discussion about risks to the economic outlook during that FOMC meeting, we consider two alternative scenarios as describing salient risks around that baseline outlook: the risk of higher and more persistent inflation (Inflation Surge) and the risk of a sharp decline in real activity (Recession). Using the linearized FRB/US model, we design the risk-adjusted optimal policies in the face of those risks.[1]

[1]The FRB/US model is a large-scale model of the U.S. economy that has been used at the Federal Reserve Board for decades and also in the existing literature to quantify the macroeconomic implications

We first examine the risk-adjusted optimal policy when the policymaker contemplates that only one of the risks can be realized instead of the baseline outlook. In this set of analysis, we consider cases in which the policymaker's initial prior is strongly in favor of either the baseline or the alternative scenario. This allows us to characterize the risk-adjusted optimal policy outcome in the case of miscalibrated beliefs in two benchmark cases. In the first setting, the policymaker initially views the alternative scenario as more likely, but instead the baseline outlook materializes. Borrowing from statistical decision theory, we call this a "Type I" error. The second setting describes a "Type II" error: the policymaker believes initially that the baseline outlook will materialize with high probability but instead the alternative scenario is realized. When policymakers view the Inflation Surge scenario as likely, policy initially tightens to be better positioned in case the risk of high inflation materializes. However, over time, this scenario does not come to pass, resource utilization and inflation come in lower than anticipated, and policy begins to ease quickly. On the other hand, if policymakers initially view the Inflation Surge as unlikely, but it in fact materializes, the decline in the policy rate occurs only slowly.

We next examine the case in which the policymaker's prior places positive probability on all three outcomes: the baseline, Inflation Surge, and Recession. By varying the initial weights assigned to each of the three possibilities, we generate a range of risk-adjusted optimal policy trajectories assuming the baseline projection in fact materializes. With a balanced prior on the three possibilities—giving the baseline outlook 50 percent probability and the scenarios 25 percent each—the optimal policy path follows closely the one associated with the optimal policy under perfect foresight of the baseline path. That is, there is only a small risk adjustment for optimal policy because the risks to inflation and economic activity—stemming from both an Inflation Surge and Recession—are roughly in balance. On the other hand, when a policymaker is relatively more worried about an inflation surge (or a recession)—while still assigning a 50 percent probability to the baseline outlook—policy will differ in a way resembling the analysis based on only one scenario at the time. Policymakers initially insure against the risk they judge more likely and then adjust their beliefs and policy as they observe that the baseline outlook is unfolding.

An important contribution of our methodology is its recursive representation that provides substantial advantages in terms of tractability to embed risk-management considerations into optimal policy strategies. In the spirit of Svensson and Williams (2005), risks

of alternative monetary policy strategies. In this paper, we use a linearized version of the FRB/US model described in Brayton and Reifschneider (2022) and available on the Federal Reserve Board's website.

are incorporated as policymakers judgement about likely future outcomes. A possible path of the economy is represented by a sequence of anticipated shocks that policymakers fully anticipate. However, differently from Svensson and Williams (2005), our framework ensures that policymakers take into account the contemplated risks when setting policy optimally. That is, they are aware that the anticipated shocks may have different effects on the economy depending on the path that the economy follows. Crucially, they make judgements about those likely paths and the associated likelihoods and update their beliefs over time in a Bayesian manner.

Incorporating risks as policymakers' judgments about likely future outcomes ensures that this approach is tractable and can be used in a wide range of macroeconomic models to study uncertainty as captured by a wide array of scenarios. Our approach is tractable because the optimal filtering problem of the policymaker (the best way to extract signals from noisy data) depends only on anticipated processes. As a result, their optimal filtering problem is independent of a policymaker's optimal decisions regarding the policy rate. Because of its tractability, the approach can be implemented also in the context of rich models that are commonly used by central banks for scenario analysis.

2 Risk-adjusted optimal policy: a simple example

Risks to the outlook. To build the intuition for our approach, we consider a simple example based on a canonical NK model *à la* Galí (2008), where

$$y_t = E_t(y_{t+1}) - \sigma^{-1}(i_t - E_t(\pi_{t+1})), \quad (1)$$

$$\pi_t = \beta E_t(\pi_{t+1}) + \kappa y_t + V_t, \quad (2)$$

represent a conventional IS equation and Phillips curve, respectively. The term V_t is described by

$$V_t = \xi_t + \theta_0(J)\varepsilon_{0,t} + \theta_1(J)\varepsilon_{1,t-1}, \quad (3)$$

and collects an unobserved *iid* disturbance $\xi_t \sim \mathcal{N}(0, \sigma_\xi^2)$ as well as two anticipated markup shocks: the shock $\varepsilon_{0,t}$ is announced at the beginning of date t and implemented in the same period, while the shock $\varepsilon_{1,t-1}$ has been announced at date $t-1$ and is implemented after one period, at time t. The shocks $\varepsilon_{0,t}$ and $\varepsilon_{1,t}$ are marginally distributed $\mathcal{N}(0,1)$ and are mutually independent with one another (and over time). The parameters $\theta_0(J)$ and $\theta_1(J)$ capture the effect of the corresponding shocks expected under alternative scenarios. The random variable J can take two values $j \in \{1, 2\}$, reflecting a policymakers' view that the economy may evolve according to either scenario. Although they cannot observe which scenario is unfolding, they can form and update their beliefs about the likelihood of each scenario as they receive incoming data.

Risk-adjusted optimal policy. The central bank takes as given the structure of the economy in (1) and (2) and commits to a forward-looking optimal policy that minimizes the quadratic loss function

$$E_0 \left\{ \frac{1}{2} \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda_y y_t^2) \right\}. \tag{4}$$

The coefficient λ_y captures the relative importance that a central bank assigns to movements in inflation and output gaps. It is well known that the constrained optimization problem leads to the following equation representing the dynamics of the (log) price level

$$\widehat{p}_t = a\widehat{p}_{t-1} + a\beta E_t(\widehat{p}_{t+1}) + aV_t, \tag{5}$$

where $a \equiv \lambda_y / [\lambda_y(1+\beta) + \kappa^2]$ and $\widehat{p}_t = \bar{p}_t - \bar{p}_{-1}$ is the inflation rate relative to the (log) price level prevailing one period before the central bank chooses the optimal plan.[2]

The central bank has beliefs over the likelihood of each scenario. While we discuss below how those beliefs are updated as incoming data is received, it is important to recognize how those beliefs affect the design of the optimal policy. In particular, policymakers have beliefs $p_t = \{p_{1t}, p_{2t}\}$ over the two possible scenarios. So, the expected effect that a shock announced today will have on tomorrow's macroeconomic outcomes is $E_t(V_{t+1}) = \widetilde{\theta}_{1,t}\varepsilon_{1,t}$, where $\widetilde{\theta}_{1,t} \equiv \sum_{j=1}^{2} p_{jt}\theta_1(j)$ is the average of the effect of the shock under the two scenarios weighted by the policymakers' beliefs. To solve the optimality condition in (5), we postulate that the solution takes the form

$$\widehat{p}_t = \delta\widehat{p}_{t-1} + r_0 V_t + r_1 \widetilde{\theta}_{1,t}\varepsilon_{1,t}, \tag{6}$$

where V_t is defined in (3). The method of undetermined coefficients ensures that the postulated solution holds so long as

$$\delta = \frac{1 - \sqrt{1 - 4a^2\beta}}{2a\beta}, \qquad r_0 = \frac{a}{1 - a\beta\delta}, \qquad r_1 = \frac{a\beta}{1 - a\beta\delta} r_0.$$

[2]Appendix A provides detailed derivations of the risk-adjusted optimal policy in this NK model.

The risk-adjusted optimal policy to which the central bank commits is

$$i_t = \left(1 - \frac{\kappa}{\lambda_y}\sigma\right)\left[(\delta - 1)\widehat{p}_t + r_0\theta_1(J)\varepsilon_{1,t}\right], \tag{7}$$

where $\widehat{p}_t$ is defined in (6). Therefore, the optimal policy is to stabilize the price level while accounting for the risk of anticipated markup shocks.

2.1 Forward-looking and data-dependent optimal policy

An important advantage of the proposed framework is that policymakers' beliefs about the risks to the outlook evolve over time and are updated as incoming data is received. Specifically, at the beginning of date t, policymakers has prior beliefs $p_{t-1} = \{p_{1t-1}, p_{2t-1}\}$ regarding the likelihood of the two scenarios. Later in the period, they observe V_t and update their beliefs about scenario j according to Bayes rule,

$$p_{jt} = \frac{\varphi\left(V_t|\mu_t(j), \sigma_\xi^2\right) p_{jt-1}}{\sum_{j=1}^{2}\varphi\left(V_t|\mu_t(j), \sigma_\xi^2\right) p_{jt-1}}, \tag{8}$$

where $\varphi(\cdot)$ is the normal pdf for the variables V_t whose conditional mean is $\mu_t(j) = \theta_0(j)$ and the announced shock is normalized to 1. To provide some intuition, we assume that the first scenario actually materializes. So, the updated probability for this scenario is

$$p_{1t} = \frac{1}{1 + \frac{p_{2t-1}}{p_{1t-1}}(1/exp\left\{\frac{[\theta_0(1)-\theta_0(2)]^2}{2\sigma_\xi^2}\right\})}, \tag{9}$$

and $p_{2t} = 1 - p_{1t}$. Equation (9) is useful particularly to identify the features that are key for the updating of the beliefs.

Lemma 1 *Suppose scenario 1 materializes and* $0 < \sigma_\xi < \infty$. *If* $\theta_0(1) = \theta_0(2)$, *then* $p_{jt} = p_{jt-1}$. *Instead, if* $|\theta_0(1) - \theta_0(2)| \to \infty$, *then* $p_{1t} \to 1$.

This lemma clarifies that policymakers update their beliefs whenever they expect different impacts of the anticipated shocks across scenarios. If this is not the case, the updated probabilities are unchanged relative to those held at the beginning of the period. Conversely, for a well-defined volatility of the unobserved shocks, the more distinct the effect of the anticipated shock between the two scenarios, the faster the policymakers uncover that the first scenario is unfolding.

Lemma 2 *Suppose scenario 1 materializes and* $|\theta_0(1) - \theta_0(2)| > 0$. *If* $\sigma_\xi \to \infty$, *then* $p_{jt} \to p_{jt-1}$. *Instead, if* $\sigma_\xi \to 0$, *then* $p_{1t} \to 1$.

The lemma establishes that although the anticipated effect of the shock are different under the two scenarios, the speed at which policymakers update their beliefs depends on whether the signal taken from the incoming data is obfuscated by the unobserved shock ξ_t. If the noise is substantial, policymakers hardly update their beliefs. By contrast, if the signal from the received data is not polluted by the unanticipated shock, then they quickly realize that the economy is evolving according to the first scenario.

The lemmas transparently show an important feature of monetary policymaking in our framework. The risk-adjusted optimal policy in (7) is both data dependent and forward-looking. The information received over time helps policymakers to update their beliefs over the risks to the outlook and the likely effects of the anticipated shocks on the economy. Crucially, as the beliefs evolve, policymakers adjust the optimal policy accordingly.

2.2 Beliefs and preemptive optimal policy

The proposed risk-adjusted optimal policy embeds risk-management considerations that are relevant for policymaking: policymakers can "take out insurance" in the face of salient risks to the economic outlook. To provide this intuition, we consider an example in which the economy described by the stylized NK model can follow two alternative paths. Under the baseline, no shock hits the economy, that remains at its steady state. Under the alternative scenario, a one-percent markup shock impacts the economy in period 1, and both policymakers and the private sector perfectly anticipate it.[3] This scenario illustrates how the economy is expected to evolve when an inflation surge occurs and policy responds optimally. Given these two alternative paths for the economy, we are interested in designing the risk-adjusted optimal policy for some initial beliefs about their likelihood. In this example, the risk of an inflation surge does not materialize in period 1, but policymakers and the private sector are highly concerned about that possibility—that is, in period 0, they assign a likelihood of 90 percent to the inflation-surge scenario and of 10 percent to the baseline. To further simplify the intuition, noise shocks are basically absent in period 1. Therefore, when policymakers and the private sector observe that the inflationary shock does not materialize, they are certain that the baseline is unfolding.

Figure 1 shows the optimal policies and associated macroeconomic outcomes in three cases. When the baseline scenario occurs (solid blue line), the economy simply remains at

[3]Appendix A details the parameterization used for this example that is in line with Galí (2008).

Figure 1: Optimal Policy and Risk of Inflation Surge

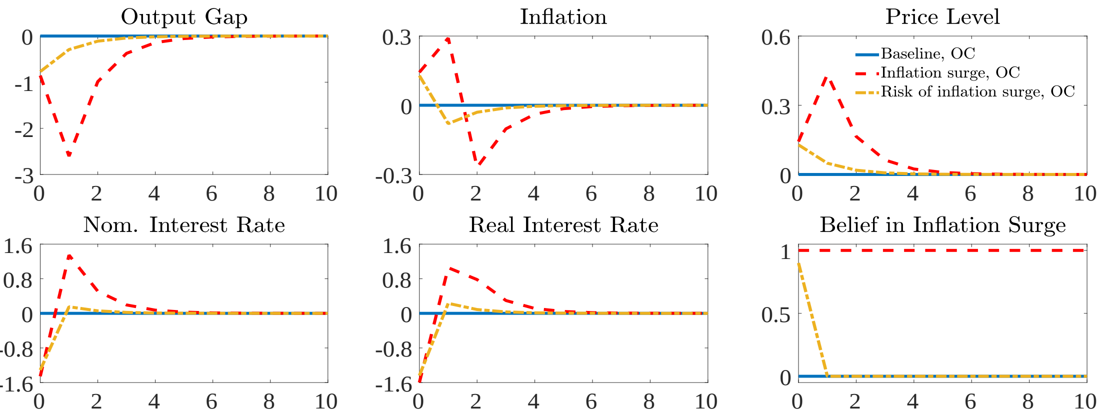

Notes: The figure shows the optimal policies and associated macroeconomic outcomes in three cases: under the baseline scenario (solid blue line), under an inflation-surge scenario in which a one-percent markup shock is anticipated to hit the economy in period 1 (dashed red line), and under the risk of an inflation surge (dash-dotted yellow line). In the latter case, the risk does not materialize.

its steady state as no shock affects the economy. Under the inflation-surge scenario (dashed red line), the policymaker and the private sector fully anticipate in period 0 that the markup shock will affect the economy in period 1. As a result, the optimal policy is accommodative in period 0 to avoid a large drop in economic activity and restrictive thereafter to return the price level to its steady-state level. Under the risk of an inflation surge (dash-dotted yellow line), policymakers have an incentive to take out insurance in period 0 and lean against the possibility of weaker economic activity resulting from the concern—also shared by the private sector—that a markup shock may occur in the following period. However, in period 1, policymakers and the private sector know with certainty that the baseline scenario is unfolding, as shown in the bottom-right panel of figure 1. Crucially, in that period, the risk-adjusted optimal policy does not immediately coincide with the optimal policy under the baseline because it needs to make up for the preemptive accommodation provided in period 0. By being slightly restrictive over the subsequent few periods, the risk-adjusted optimal policy returns the price level to target.

3 Risk-adjusted optimal policy in a general framework

3.1 Judgment and beliefs about risks to the outlook

We study optimal monetary policy under uncertainty, incorporating central-bank judgment about likely future outcomes. Our approach is similar to Svensson and Williams (2005) who incorporate central bank judgment to consider optimal policies around a baseline projection. In Svensson and Williams (2005), the central bank does not take into uncertainty about economic outcomes. In contrast, under our approach, a central bank takes into account uncertainty and makes judgments about several likely paths that the economy may follow. In addition, the central bank makes judgments about the likelihood of these alternative projected paths and updates these beliefs over time in a Bayesian manner.

We add central bank judgment to a model whose private sector equilibrium conditions are given by:

$$A_0 Y_t = A_1 Y_{t-1} + A_2 E[Y_{t+1}|\Omega_t] + A_3 X_t + A_4 E[X_{t+1}|\Omega_t] + A_5 V_t \tag{10}$$

where Y_t is a $N_y \times 1$ vector of endogenous variables, $X_t = (i_t, i_{t-1})'$ where i_t is the policy rate, and ω_t is a central bank's information set. The $n_v \times 1$ vector of exogenous shocks, V_t, is assumed to evolve according to:

$$V_t = \xi_t(J) + \sum_{k=0}^{K} \theta_k(J)\varepsilon_{k,t-k}, \tag{11}$$

Here, $\xi_t(J)$ is $n_v \times 1$ vector of innovations that occurs at time t. These innovations are unobserved by the central bank. Their distribution depends on J, a discrete random variable that takes on N different values, corresponding to the N scenarios that a central bank judges as likely. A central bank does not observe J but knows that it takes on values from $j \in \{1, 2, ..., N\}$. Conditional on $J = j$, the distribution of $\xi_t(j)$ is given by:

$$\xi_t(j) \sim \mathcal{N}(\mu(j), \Sigma_\xi)$$

where Σ_ξ is a diagonal matrix. The $n_v \times 1$ random vector, $\varepsilon_{k,t-k}$, in equation (11) is composed of mean-zero random variables realized at the beginning of period $t-k$ that affect the economy at date t, that is after k period. At date t, a central bank observes $\varepsilon_{k,t}$ and uses it to form judgments about the projected paths of the economy, taking into account that at date $t+k$ the news $\varepsilon_{k,t}$ have an impact effect of $\theta_k(J)$, where $\theta_k(J)$ is a $n_v \times n_v$ diagonal matrix of parameters. Because $\theta_k(J)$ depends on J, a central bank does not observe the impact effect of the news shocks that it receives; it can only judge their likely impact.

A few remarks about these parameters are important. First, the variable J seeks to capture the potential risks that a central bank and the private sector could consider. However, if a risk j^* is not deemed salient for the design of an optimal policy, the anticipated shocks under this scenario are assumed to have no effect, that is $\theta_k(j^*) = 0$. Second, policymakers and the private sector know that there exist N likely paths that the economy could follow but do not observe the scenario J that materializes. Over time, they will update the beliefs about the likelihood of each scenario as incoming data becomes available. Third, the shocks $\varepsilon_{k,t-k}$ do not depend on the scenario and whenever they take a value different from their zero mean, risk-management considerations are incorporated into the optimal policy strategy. In our framework, we normalize the shocks to one and the effect of the anticipated shocks under each scenario is captured by scaling the associated parameters accordingly.

3.1.1 Beliefs updating

It is convenient to collect the news shocks that a central bank observes at date t into the vector, $\varepsilon^t = (\varepsilon'_{0,t}, ..., \varepsilon'_{K,t})'$. A central bank's information set, Ω_t, includes ε^t and all of its lags. It also includes V_t and the model's endogenous variables and all of their lags. Because a central bank does not observe $\xi_t(J)$, or $\theta_k(J)$ for $k = 0, 1, ..., K$, it does not know which scenario is unfolding in real time. Accordingly, a central bank's must make a judgment about the likelihood of the various scenarios. As the economy evolves over time, a policymaker receives incoming data and updates its beliefs about the likely scenarios using Bayes rule. Formally, at the beginning of date t, a policymaker has prior belief, $p_{t-1} = (p_{1t-1}, ..., p_{Nt-1})$, regarding the likelihood of the N scenarios where $p_{jt-1} = \text{Prob}(J = j|\Omega_{t-1})$. Later in the period after observing V_t, a central bank updates their beliefs about scenario j according to:

$$p_{jt} = \frac{\varphi\left(V_t|\mu_t(j), \Sigma_\xi\right) p_{jt-1}}{\sum_{j=1}^{N} \varphi\left(V_t|\mu_t(j), \Sigma_\xi\right) p_{jt-1}}, \tag{12}$$

where $\varphi(\cdot)$ is the normal pdf for a vector of random variables and $\mu_t(j)$ is given by:

$$E[V_t|\Omega_t, J = j] \equiv \mu_t(j) = \mu(j) + \sum_{k=0}^{K} \theta_k(j)\varepsilon_{k,t-k}. \tag{13}$$

Equation (12) governs how a central bank updates its priors about the scenarios in response to incoming data.

3.1.2 Policymakers' views about risks in a recursive form

A central bank uses these probabilities to form expectations about future shocks. To keep the notation compact, let us define the weighted average of the effect of a news shock implemented k periods after its announcement by $\tilde{\theta}_{k,t} \equiv \sum_{j=1}^{N} p_{jt}\theta_k(j)$. As shown in Appendix B, the central bank's expectation of the shocks at any future period $t+s$ with $1 \leq s \leq K$ is then given by:

$$E[V_{t+s}|\Omega_t] = \tilde{\theta}_{s,t}\,\varepsilon_{s,t} + E[V_{t+s}|\Omega_{t-1}] + \sum_{k=s+1}^{K} \Delta\tilde{\theta}_{k,t}\,\varepsilon_{k,t+s-k}, \tag{14}$$

where $\Delta\tilde{\theta}_{k,t} = \left(\tilde{\theta}_{k,t} - \tilde{\theta}_{k,t-1}\right)$ captures the evolution of the central bank's beliefs about the impact of the news shocks. In a similar vein, the shocks V_t in (11) can be written as

$$V_t = \xi_t(J) + \left[\theta_0(J)\varepsilon_{0,t} + E[V_t|\Omega_{t-1}] + \sum_{k=1}^{K}\left(\theta_k(J) - \tilde{\theta}_{k,t-1}\right)\varepsilon_{k,t-k}\right]. \tag{15}$$

It is convenient to write the expected future process for V_t in companion form, defining the vectors:

$$S_t = (V_t', E_tV_{t+1}', ..., E_tV_{t+K}'; \varepsilon_{1,t}', \ldots, \varepsilon_{K,t-K+1}'; \varepsilon_{2,t}', \ldots, \varepsilon_{K,t-K+2}'; \ldots; \varepsilon_{K,t}')',$$

where for convenience we have defined $E_t \equiv E[\cdot|\Omega_t]$. Note that the shocks in S_t are grouped depending on the time when the shocks are implemented. For example, the vector $(\varepsilon_{1,t}', \ldots, \varepsilon_{K,t-K+1}')'$ collects the shocks implemented at time $t+1$ that have been announced at time t or earlier. So, this vector collects Kn_v shocks. The subsequent vector $(\varepsilon_{2,t}', \ldots, \varepsilon_{K,t-K+2}')'$ represents shocks implemented at time $t+2$ announced at time t or earlier. Note however that this vector collects only $(K-1)n_v$ rather than Kn_v as the previous vector. Following the same logic, the last vector $\varepsilon_{K,t}'$ collects the only n_v shocks that can be implemented at $t+K$. So, in total, there are $n_E \equiv [K(K+1)/2]n_v$ shocks in the vector S_t. In addition, the vector S_t collects V_t and its expectations, totaling $n_V \equiv (K+1)n_v$

variables. With these vectors defined, we can express S_t and $E[S_{t+1}|\Omega_t]$ as:

$$S_t = G_\varepsilon(p_t)\varepsilon^t + G_S(p_t, p_{t-1}; J)S_{t-1} + G_\xi \xi_t \tag{16}$$

$$E[S_{t+1}|\Omega_t] = \widetilde{G}_S S_t + G_\xi G_\mu(p_t), \tag{17}$$

where $G_\mu(p_t) \equiv \sum_{j=1}^N p_{jt}\mu(j)$ and details about the matrices in (16) and (17) are provided in Appendix B. Note that the policymaker does not observe ξ_t or the matrix $G_S(p_t, p_{t-1}; J)$, which depends on J. Instead, the policymaker observes $G_\varepsilon(p_t)\varepsilon^t$ and the sum, $G_S(p_t, p_{t-1}; J)S_{t-1} + G_\xi \xi_t$, given their imperfect information.

3.2 Risk-adjusted optimal policy under commitment

We start by deriving the solution for optimal policy under commitment in linear rational expectations models. We modify the optimal monetary solution in Dennis (2007) to incorporate interest rate smoothing and show how the solution for optimal policy under commitment changes as a result. The details of the derivations are in Appendix C.

Given that the companion form for V_t is expressed in (16) in terms of S_t, we rewrite the linear rational expectations model in (10) that describes the structure of the economy as

$$A_0 Y_t = A_1 Y_{t-1} + A_2 E_t Y_{t+1} + A_3 X_t + A_4 E_t X_{t+1} + \widetilde{A}_5 S_t. \tag{18}$$

We then represent optimal policy under commitment as a policymaker's problem at date 0 to minimize the quadratic loss function:

$$\min_{\{Y_t, i_t\}_{t=0}^\infty} E_0 \sum_{t=0}^\infty \beta^t \left\{ \frac{1}{2} \left[Y_t' W Y_t + X_t' Q X_t \right] \right\} \tag{19}$$

subject to the structure of the economy in (18). The first order conditions with respect to Y_t are given by:

$$W Y_t + A_0' \lambda_t - \beta^{-1} A_2' \lambda_{t-1} - \beta A_1' E_t \lambda_{t+1} = 0, \tag{20}$$

where λ_t denotes the vector of lagrange multipliers on the model's N_y private-sector equilibrium conditions. We solve the model using the "timeless perspective" and impose $\lambda_{-1} = 0$. As a result, these conditions are consistent with optimality at $t \geq 0$. With interest rate smoothing in the loss function, the matrix, Q, satisfies:

$$Q = \rho_i \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

where ρ_i determines the weight on the change in the policy rate in the loss function. The first order condition with respect to the policy rate is given by

$$\rho_i(i_t - i_{t-1}) - \beta\rho_i(E_t i_{t+1} - i_t) - A'_{3,\cdot 1}\lambda_t - \beta A'_{3,\cdot 2}E_t\lambda_{t+1} = 0. \tag{21}$$

where $A_{3,\cdot i}$ denotes the ith column of the matrix A_3. Here we have assumed that $A_{4,\cdot 2} = 0$, which holds without loss of generality, since $A_{4,\cdot 2}$ is redundant with $A_{3,\cdot 1}$.

With these first order conditions in hand, we can write the model's equilibrium conditions as:

$$CZ_t = BZ_{t-1} + FE_tZ_{t+1} + DS_t \tag{22}$$

where $Z_t = (Y_t, i_t, \lambda_t)'$ and the C, B, and F matrices function of the structural parameters of the model. These equilibrium conditions form the basis for solving the model when optimal policy embeds risk-management considerations.

Finally, equations (22), (16) and (17) are the equilibrium conditions defining the model. Using the method of undetermined coefficients, the solution takes the form:

$$Z_t = TZ_{t-1} + RS_t + R_0(p_t), \tag{23}$$

$$S_t = G_\varepsilon(p_t)\varepsilon^t + G_S(p_t, p_{t-1}; J)S_{t-1} + G_\xi\xi_t, \tag{24}$$

where the matrices T, R, $R_0(p_t)$ are defined in Appendix C.

4 Optimal policy and risks to a soft landing

By the end of 2024, inflation had made progress towards the FOMC's inflation objective, and economic activity had continued to expand at a solid pace, possibly suggesting that the U.S. economy may experience a soft landing. However, this possibility did not come without risks to the economic outlook. For instance, the minutes of the December 2024 FOMC meeting provide a characterization of the risks that the U.S. economy faced in that period. In particular, the risks to the attainment of the Committee's dual-mandate objectives of maximum employment and price stability were judged to be roughly in balance. Inflation risks were tilted to the upside. Inflation had remained somewhat elevated, and some readings on inflation came in stronger than expected. Conversely, the risks to the labor market were tilted to the downside. Labor market conditions had generally eased since earlier in the year and could weakened unexpectedly.

Against this backdrop, we apply our approach to design optimal policy strategies that embed risk-management considerations around the possibility of a soft landing for the U.S. economy.[4] To this end, three key elements are necessary to apply our methodology. First, a baseline outlook that we choose to be consistent with the median responses to the December 2024 SEP. Second, alternative scenarios capturing the risk of either an inflation surge or a recession. Finally, given the tractability of our approach, we implement our methodology to the linearized FRB/US model.

One of the benefits of using the FRB/US model, as described in Brayton and Reifschneider (2022), is the availability of modeling expectations of agents in different sectors of the economy in two ways: model consistent (or "rational") expectations and VAR-based expectations. Under VAR-based expectations, the expectation formation of households and firms is backward-looking and does not take into account the uncertainty regarding alternative future trajectories of the economy. Importantly, this assumption helps simplify the analysis on a number of dimensions. First, policymakers' ability to influence private-sector expectations through future commitments is not possible, and the distinction between risk-adjusted optimal policy under commitment and discretion is eliminated. Second, the use of VAR-based expectations helps isolate the role of risk-management considerations for optimal policy, because households and firms do not take into account the uncertain future trajectories of the economy in their own spending and pricing decisions. While it is straightforward to extend the methodology to the case in which the private sector has rational expectations and takes into account uncertainty, like the policymakers, we leave to future work the study of optimal policy under uncertainty when the policymaker has the ability to influence private-sector expectations formation and when the risk-management considerations of policymakers interact with those of households and firms.[5]

[4]All optimal control simulations shown in this section posit that policymakers choose a path for the federal funds rate to minimize a discounted equally-weighted sum of squared inflation gaps (measured as the difference between four-quarter headline PCE price inflation and the Committee's 2 percent objective), squared unemployment gaps (measured as the difference between the unemployment rate and the median December 2024 SEP longer-run level of the unemployment rate), and squared changes in the federal funds rate.

[5]In this regard, it is possible to consider versions of the FRB/US model with model-consistent expectations. In these versions of the model, financial market participants and/or wage and price setters could be interpreted as having model consistent expectations but incomplete knowledge of the distribution of shocks like the policymakers. They would update their beliefs in a Bayesian manner as well and the setting of optimal policy under uncertainty would take into account the updating of their beliefs and the effect it has on the economy.

4.1 Inflation surge

In our first example, policymakers in the model judge the economy will likely evolve following two potential paths. One is the baseline outlook path that is consistent with the median responses to the December 2024 SEP under optimal policy. The alternative path seeks to capture an upside risk to an inflation, consistent with, for example, the description in the minutes of the December 2024 FOMC meeting. In particular, the "Inflation Surge" scenario provides a quantitative (but illustrative) judgment about what an inflation surge, if it were to occur, would likely entail. In the FRB/US model, this judgment is implemented via an anticipated path of adverse shocks to the main price equation, that act as a cost push shock, and generate an increase in inflation and in the unemployment rate. At the beginning of the simulation, policymakers view this path as the most likely sequence of shocks to occur in the future conditional on there being an inflation surge. In addition, policymakers initially assign some probability to shocks emanating from the distribution underlying the Inflation Surge scenario. In this example, we separately consider the cases in which either the baseline or the inflation surge materialize. In each case, policymakers update their beliefs about the outlook using Bayes' rule and in line with the data they observe.

Baseline is realized. We start by focusing on the case in which the baseline outlook materializes—that is, the inflationary shocks do not occur. However, in the first quarter of 2025, policymakers do not anticipate that path, as their initial beliefs are set so that they view the alternative scenario as very likely. In the simulation, for illustrative purposes, we assume that policymakers assign a 90 percent probability to shocks that emanate from the distribution underlying the Inflation Surge scenario. When policymakers observe that no shock has been realized in the first quarter of 2025, they downweight the likelihood of the distribution of shocks underlying the alternative scenario; however, given that the noise shocks (ξ_t) obfuscate their learning process, they cannot rule it out completely. The calibration of the volatility of these shocks determines the ability of the policymakers to distinguish between the two distributions of shocks underlying the scenarios. The values of the diagonal matrix, Σ_ξ, are chosen to ensure gradual learning in this exercise given the initial probability.

Figure 2 shows the optimal policy-rate prescriptions, beliefs, and macroeconomic outcomes (inflation and the unemployment rate) for three sets of simulations. The solid blue lines represent the optimal policy and its associated outcomes under the December 2024

Figure 2: Optimal Policy under Risk of Inflation Surge

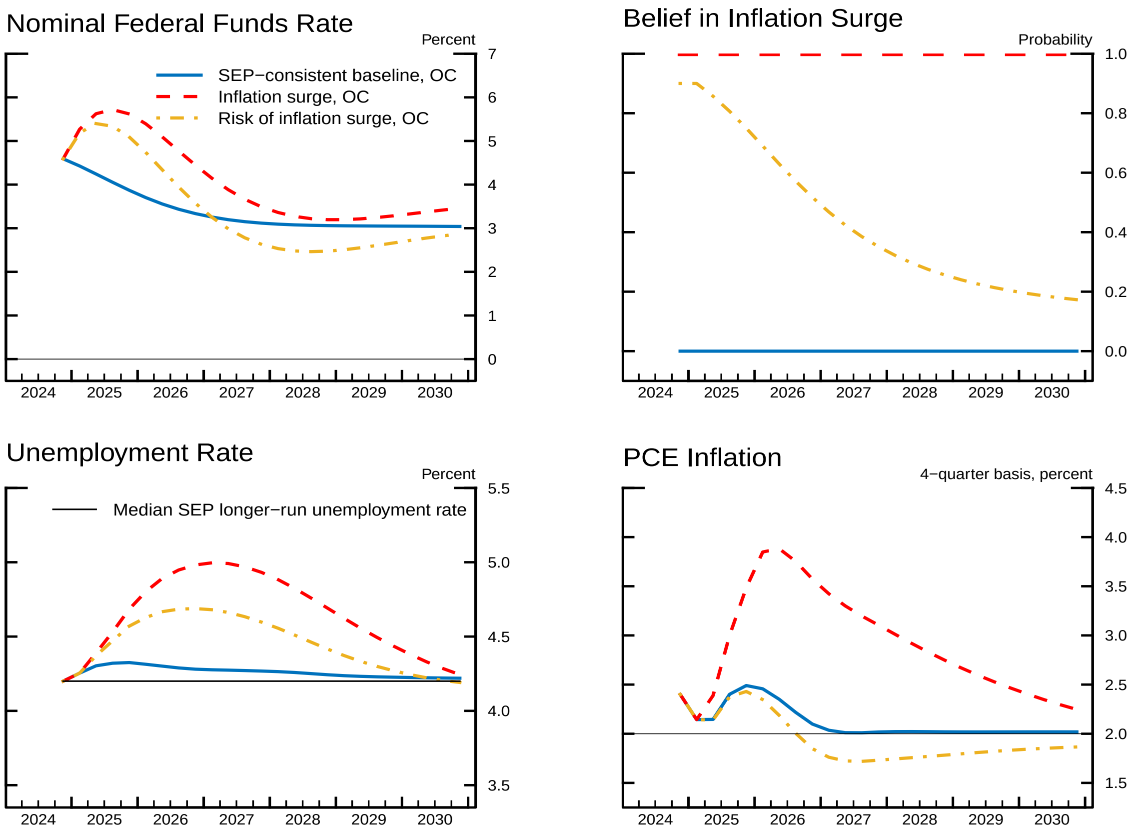

Note: The optimal control simulations are conducted with equal weights in the loss function on squared inflation gaps, squared unemployment rate gaps, and squared changes in the federal funds rate.

SEP-consistent baseline assuming perfect foresight. As shown in the top-right panel, the beliefs that the inflationary scenario is unfolding are always zero because the policymakers correctly anticipate that the inflationary shocks will not occur. Overall, optimal policy under the SEP-consistent baseline prescribes a gradual decrease in the federal funds rate as the benign outlook does not project large deviations of inflation from 2 percent nor of the unemployment rate from its longer-run value.

The dashed red lines in Figure 2 illustrate the optimal policy and its associated outcomes when policymakers have perfect foresight and know that the inflationary scenario will in fact occur (i.e., policymakers' belief in the inflation-surge scenario are equal to 1). In this case, the policy rate rises to contain the rapid and persistent increase in inflation, and as result, the unemployment rate nearly reaches 5 percent. However, as inflation returns to target and the unemployment rate to its longer-run level, the policy rate declines gradually.

Figure 2 also illustrates the risk-adjusted optimal policy and outcomes (dash-dotted yellow lines) when policymakers initially believe with high probability that the inflation surge will occur but learn over time that the scenario's shocks do not occur. With a high probability attached to the possibility of inflationary shocks, the federal funds rate under risk-adjusted optimal control initially increases. However, as the expected adverse shocks do not occur, data on resource utilization and inflation come in lower than anticipated, and policymakers assign over time a smaller likelihood to the possibility that future shocks will be inflationary. As a result, the pace at which policy eases accelerates. With unemployment higher and inflation lower than if policy had not accounted for the risk of an inflation surge, the policy rate is subsequently repositioned to a level below what it would have been under the baseline projection. Unemployment and inflation then gradually return to the baseline optimal control paths.

Of note, the results in Figure 2 clearly show that policy-rate prescriptions and macroeconomic outcomes of the risk-adjusted optimal policy are not simply a weighted-average of the perfect foresight scenarios, where the weights reflect policymakers' beliefs shown in the top-right panel. For example, as the policymakers realize that their concern of an inflation surge is not materializing, the policy rate path under the risk-adjusted optimal policy falls below that path under both the baseline or alternative scenarios.

Inflation Surge scenario is realized. In Figure 2, policymakers' initial strong beliefs about the likelihood of the alternative scenario were inconsistent with the realization of the disturbances (which were themselves consistent with the baseline projection.) Borrowing from statistical decision theory, we characterize this as a "Type I" error; that is, the error made rejecting a null hypothesis when it is in fact true. In this analogy, the baseline projection is taken to be the null hypothesis and the projection underlying an alternative scenario as the alternative one. In this framework, a "Type II" error is one made when failing to reject a null hypothesis that is in fact false.

Figure 3 displays a set of simulations exploring this possibility in the context of the Inflation Surge scenario. The blue and red lines remain identical to Figure 2. The risk-adjusted optimal policy and related outcomes shown in the dash-dotted yellow lines differ from Figure 3 in two keys ways. First, the inflation surge embedded in the alternative scenario actually occurs. Second, policymakers initially assign a probability of 10 percent to the alternative scenario, implying that their beliefs heavily favor the baseline projection. In this case, policymakers misplace their initial confidence that the baseline projection will materialize and do not respond as aggressively as if they had perfect foresight to the inflation surge. However, as they slowly learn over time that future inflationary shocks are

Figure 3: Optimal Policy under Unexpected Inflation Surge

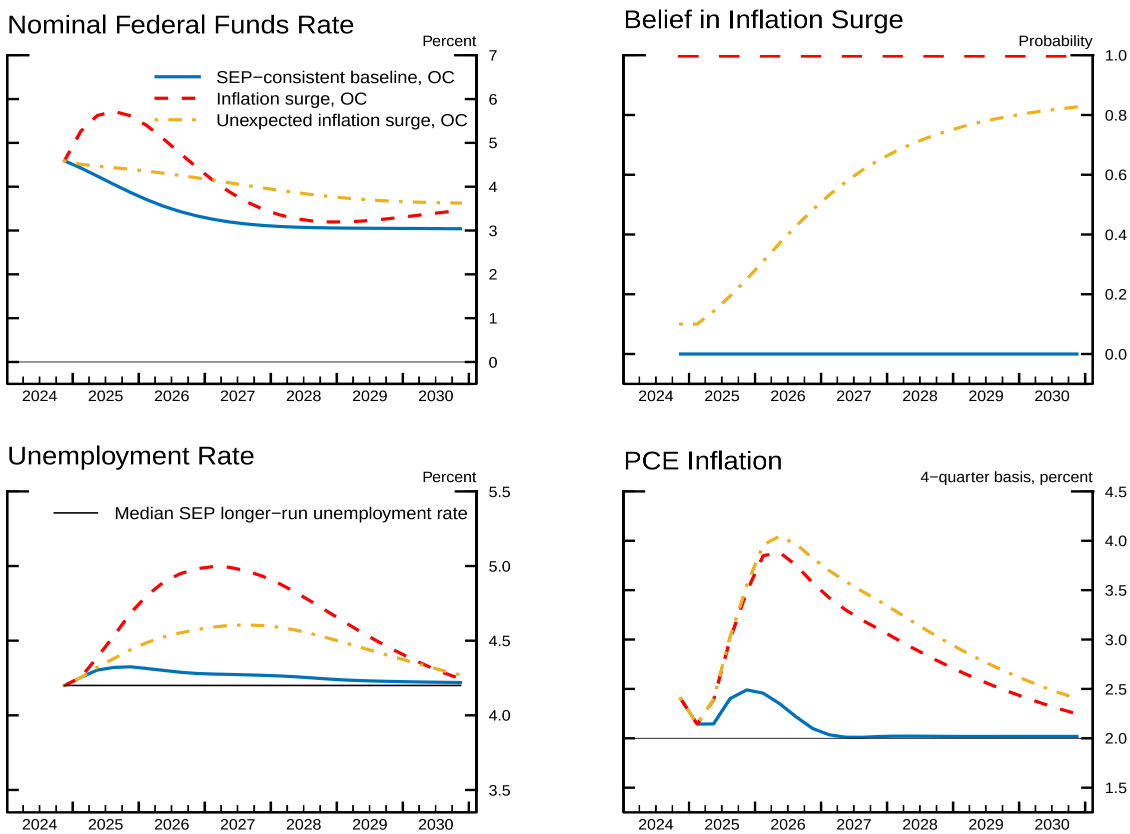

Note: The optimal control simulations are conducted with equal weights in the loss function on squared inflation gaps, squared unemployment rate gaps, and squared changes in the federal funds rate.

likely, they then decrease the policy rate more gradually than prescribed by optimal policy under the inflationary scenario because inflation remains persistently higher.

Like in statistical decision theory, in setting risk-adjusted appropriate monetary policy, both Type I and Type II errors are relevant. Carefully assessing risks to the economic outlook and updating beliefs about their likelihood in response to incoming data are key features of risk-adjusted optimal policy.

4.2 Recession

Our second example illustrates the downside risk to the labor market described in the minutes of the December 2024 FOMC meeting. As in the previous example, the economy is expected to follow two possible paths. The baseline outlook is consistent with the December 2024 SEP and therefore unchanged relative to the previous example. However, the

Figure 4: Optimal Policy under Risk of Recession

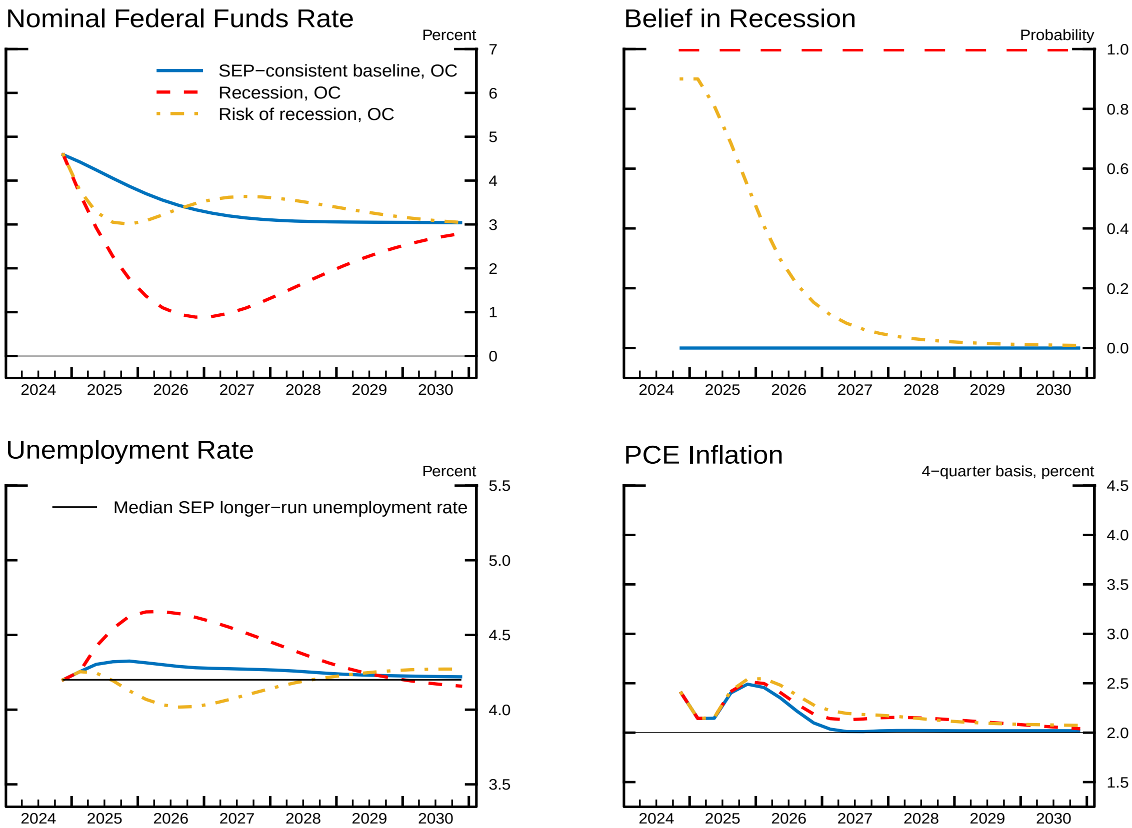

Note: The optimal control simulations are conducted with equal weights in the loss function on squared inflation gaps, squared unemployment rate gaps, and squared changes in the federal funds rate.

alternative path is a recession driven by anticipated adverse shocks to households' consumption. Policymakers view the "Recession" scenario as the most likely evolution of the economy if the recession unfolds. Facing the risk of a recession, policymakers initially assess the likelihood of the two paths for the economy and update their beliefs over time as the economy evolves. We consider the case in which the recession does not occur, but policymakers initially see that possibility as very likely, assigning a probability of 90 percent to the recession scenario.

Figure 4 shows the optimal policy, beliefs, and macroeconomic outcomes for three cases. The baseline (solid blue line) is the same as in figure 2. When the recession occurs—and policymakers know it with certainty—the federal funds rate under the Recession scenario (dashed red line) declines rapidly in response to weak resource utilization. The swift policy response avoids a much deeper recession and boosts the private sector's inflation expectations. As a result, inflation is slightly above target notwithstanding high unemployment.

Figure 5: Beliefs about the Risks

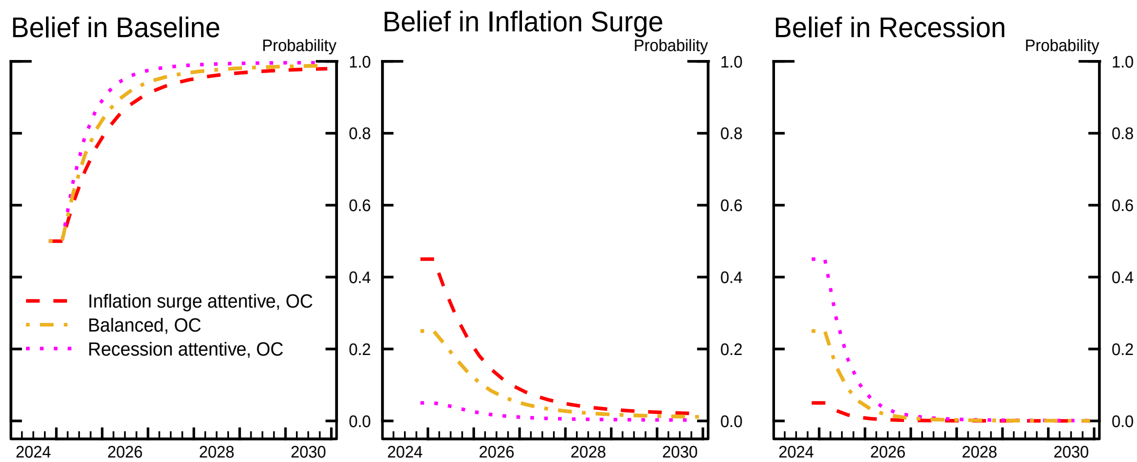

When policymakers consider the elevated risk of a recession (dash-dotted yellow line), the optimal policy initially calls for policy-rate cuts, implying an overshooting of resource utilization and inflation. However, as policymakers learn that the recessionary shocks are not materializing, they revise their beliefs accordingly and reverse their policy actions by hiking the policy rate to bring the unemployment rate and inflation back to target.

4.3 Risks to the soft landing

This section answers a key question for policymaking: What are the implications for optimal monetary policy of embedding considerations about risks to the so-called "soft landing"? As described in the minutes to the December 2024 FOMC minutes, the U.S. economy faced upside risks to inflation and downside risks to the labor market, that were judged to be roughly in balance. To capture these risks, we use the Inflation Surge scenario discussed in section 4.1 and the Recession scenario described in section 4.2.

Differently from the previous examples, the central feature of this exercise is that the policymakers in the model now account for two-sided risks—that is, the risk of both an inflation surge and a recession—when choosing the optimal policy-rate path. We consider three cases that differ in terms of the initial beliefs assigned to each of the risks. As Figure 5 shows, all cases share the SEP-consistent baseline as the modal outlook with an initial 50 percent probability assign to it. The "inflation surge attentive" case characterizes a situation where policymakers assign more probability to the inflation-surge risk than to the recession risk. Conversely, the "recession attentive" case characterizes a situation where policymakers assign more probability to the recession risk than to the inflation-surge

Figure 6: Optimal Policy under Two-Sided Risks

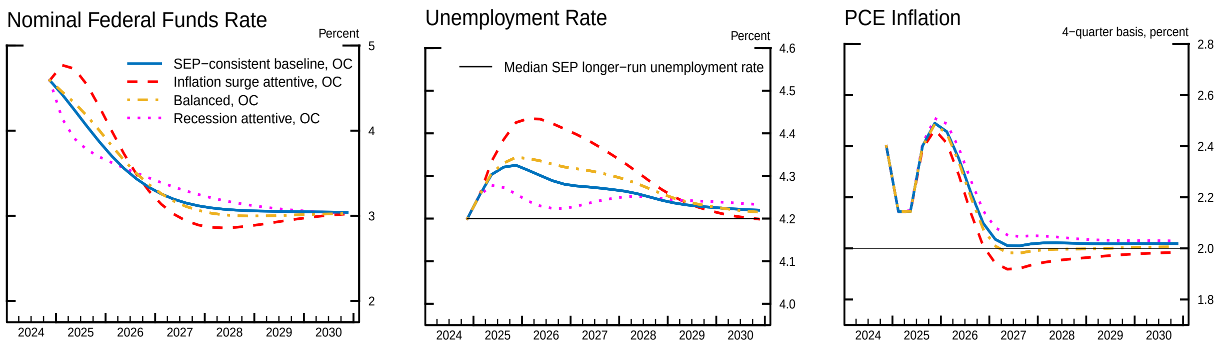

Note: The optimal control simulations are conducted with equal weights in the loss function on squared inflation gaps, squared unemployment rate gaps, and squared changes in the federal funds rate.

risk. Finally, the “balanced" case is consistent with policymakers assigning equal initial probabilities of 25 percent to each alternative scenario.

Figure 6 shows that by varying the initial weights on each of the three paths for the economy (baseline, Inflation Surge, and Recession), we generate a range of risk-adjusted optimal policy trajectories assuming the baseline projection in fact materializes. With a balanced prior on the three possible outcomes, the optimal policy rate path, depicted by the yellow dash-dotted lines, follows closely the one associated with perfect foresight of the baseline outlook (solid blue lines). That is, there is only a small risk-adjustment for optimal policy when weighing equally both the possibility of an Inflation Surge and Recession. As a result, macroeconomic outcomes are similar as those achieved under the baseline outlook. When policymakers are relatively more worried about an inflation surge (red dashed lines), policy will evolve more in line with the analysis based on the risk of high inflation only—shown in Figure 2. In particular, the risk-adjusted optimal control path for the federal funds rate initially increases. As the expected adverse shocks are not realized, the tighter policy stance results in lower inflation and higher unemployment rate. Over time, policymakers assign a greater likelihood that the baseline outlook materializes and ease policy accordingly. On the other hand, when policymakers are relatively more worried about the recession risk (purple dotted lines), policy will evolve more in line with the analysis based on the risk of a recession only—shown in Figure 4. As a result, the risk-adjusted optimal control path for the federal funds rate initially declines. As the expected adverse shocks are not realized, the looser policy stance results in higher inflation and lower unemployment rate. As policymakers assign over time a greater likelihood that the baseline outlook materializes, they keep the policy rate higher for longer to gradually return the economy to its longer-run equilibrium.

5 Conclusion

This paper presents an approach to embed risk-management considerations into the optimal design of monetary policy strategies when characterizing uncertainty around a baseline outlook using a finite set of alternative scenarios. The approach highlights the way that the assessment of the balance of risks can shape the construction of optimal policy. The methodology is tractable and can be applied to both academic style models and larger workhorse models used by central banks.

While this paper considers that uncertainty only arises from anticipated disturbances, our methodology can be extended in different dimensions. In particular, those anticipated shocks could be characterized by asymmetric distributions or could be non-additive. In addition, the methodology could be adjusted to allow the analysis of broader sources of uncertainty such as those coming from uncertainty about structural parameters of the model economy or uncertainty coming from different models. For instance, Svensson and Williams (2008) discuss how a methodology similar to ours can be applied to model many different types of uncertainties. As a result, we view this paper as an important first step forward in incorporating risks as captured by alternative scenarios into the optimal design of monetary policy strategies.

References

Brayton, Flint and David L. Reifschneider, "LINVER: The Linear Version of FRB/US," Finance and Economics Discussion Series 2022-053, Board of Governors of the Federal Reserve System (U.S.) August 2022.

Dennis, Richard, "Optimal Policy in Rational Expectations Models: New Solution Algorithms," *Macroeconomic Dynamics*, 2007, *11*, 31–55.

Galí, Jordi, *Monetary Policy, Inflation and the Business Cycle*, New Jersey: Princeton University Press, 2008.

Greenspan, Alan, "Risk and Uncertainty in Monetary Policy," *American Economic Review*, May 2004, *94* (2), 33–40.

Svensson, Lars E. O. and Noah Williams, "Monetary Policy with Model Uncertainty: Distribution Forecast Targeting," 2005. NBER Working Paper No. 11733.

A Risk-adjusted optimal policy in a simple NK model

The policymaker seeks to minimize expected losses

$$E_0 \left\{ \frac{1}{2} \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda_y y_t^2) \right\}, \tag{25}$$

subject to the Phillips curve and the IS curve

$$\pi_t = \beta E_t(\pi_{t+1}) + \kappa y_t + V_t \tag{26}$$

$$y_t = E_t(y_{t+1}) - \sigma^{-1}(i_t - E_t(\pi_{t+1})), \tag{27}$$

where $V_t = \xi_t + \theta_0(J)\varepsilon_{0,t} + \theta_1(J)\varepsilon_{1,t-1}$. The FOCs resulting from the minimization problem subject to the Phillips curve are

$$\lambda_y y_t - \kappa\gamma_t = 0, \tag{28}$$

$$\pi_t + \gamma_t - \gamma_{t-1} = 0, \tag{29}$$

where γ_t is the Lagrangian multiplier associated with the Phillips curve and $\gamma_{-1} = 0$. Combining the FOCs, the tradeoff between inflation and changes in the output gap is

$$y_t = y_{t-1} - \frac{\kappa}{\lambda_y}\pi_t. \tag{30}$$

Defining $\widehat{p}_t \equiv \bar{p}_t - \bar{p}_{-1}$ as the inflation rate relative to the (log) price level prevailing one period before the central bank chooses the optimal plan, then $\pi_t = \widehat{p}_t - \widehat{p}_{t-1}$ and from equation (30), it becomes evident the tradeoff between output gap and price level

$$y_t = -\frac{\kappa}{\lambda_y}\widehat{p}_t. \tag{31}$$

Plugging (31) in the Phillips curve (26), the following equation represents the dynamics of the price level

$$\widehat{p}_t = a\widehat{p}_{t-1} + a\beta E_t(\widehat{p}_{t+1}) + aV_t, \tag{32}$$

where $a \equiv \lambda_y/[\lambda_y(1+\beta)+\kappa^2]$. To solve (32), we need to guess a solution and use the method of undetermined coefficients. Crucially, in our framework, policymakers have beliefs $p_t = \{p_{1t}, p_{2t}\}$ over the two possible scenarios. So, an educated guess for the solution also accounts for the expectations of future anticipated shocks $E_t(V_{t+1}) = \widetilde{\theta}_{1,t}\varepsilon_{1,t}$, where $\widetilde{\theta}_{1,t} \equiv \sum_{j=1}^{2} p_{jt}\theta_1(j)$ is the average of the effect of the shock under the two scenarios weighted by the policymakers' beliefs. So, we guess that the solution takes the form

$$\widehat{p}_t = \delta\widehat{p}_{t-1} + r_0 V_t + r_1\widetilde{\theta}_{1,t}\varepsilon_{1,t}. \tag{33}$$

The method of undetermined coefficients ensures that the postulated solution holds if

$$\delta = \frac{1-\sqrt{1-4a^2\beta}}{2a\beta}, \qquad r_0 = \frac{a}{1-a\beta\delta}, \qquad r_1 = \frac{a\beta}{1-a\beta\delta}r_0.$$

To derive the Taylor-type rule that implements the risk-adjusted optimal policy, the IS curve in (27) can be written as

$$i_t = \sigma\left[E_t(y_{t+1}) - y_t\right] + E_t(\pi_{t+1}). \tag{34}$$

From (31), then $E_t(y_{t+1}) - y_t = -\kappa/\lambda_y E_t(\pi_{t+1})$, and

$$i_t = \left(1 - \frac{\kappa}{\lambda_y}\sigma\right) E_t(\pi_{t+1}). \tag{35}$$

Recalling the solution for $\widehat{p}_t$ in (33), the risk-adjusted optimal policy to which the central bank commits is

$$i_t = \left(1 - \frac{\kappa}{\lambda_y}\sigma\right)\left[(\delta-1)\widehat{p}_t + r_0\theta_1(J)\varepsilon_{1,t}\right]. \tag{36}$$

For the simulations shown in figure 1, we use a canonical parameterization of the simple NK model that is in line with Galí (2008). We set $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.17$ (consistent with a fraction $\theta = 2/3$ of firms keeping their price unchanged in each period), the elasticity of substitution among goods ϵ to 6 and the weight on output-gap deviations λ_y in the policymakers' loss function to κ/ϵ.

B Judgment about alternative scenarios: Derivations

Recall from Subsection 3.1 that the central bank's expectation of the shocks at any future period $t+s$ with $1 \leq s \leq K$ is given by:

$$\begin{aligned} E[V_{t+s}|\Omega_t] &= \sum_{k=s}^{K} \tilde{\theta}_{k,t}\, \varepsilon_{k,t+s-k} \\ &= \tilde{\theta}_{s,t}\, \varepsilon_{s,t} + \sum_{k=s+1}^{K} \tilde{\theta}_{k,t}\, \varepsilon_{k,t+s-k} \\ &= \tilde{\theta}_{s,t}\, \varepsilon_{s,t} + E[V_{t+s}|\Omega_{t-1}] + \sum_{k=s+1}^{K} \Delta\tilde{\theta}_{k,t}\, \varepsilon_{k,t+s-k}, \end{aligned} \tag{37}$$

where $\tilde{\theta}_{k,t} \equiv \sum_{j=1}^{N} p_{jt}\theta_k(j)$ is the weighted average of the effect of a news shock implemented k periods after its announcement and $\Delta\tilde{\theta}_{k,t} = \left(\tilde{\theta}_{k,t} - \tilde{\theta}_{k,t-1}\right)$ captures the evolution of the central bank's beliefs about the impact of the news shocks. In a similar vein, the shocks V_t can be written as

$$V_t = \xi_t(J) + \sum_{k=0}^{K} \theta_k(J)\varepsilon_{k,t-k} \tag{38}$$

$$\tag{39}$$

$$= \xi_t(J) + \left[\theta_0(J)\varepsilon_{0,t} + E[V_t|\Omega_{t-1}] + \sum_{k=1}^{K} \left(\theta_k(J) - \tilde{\theta}_{k,t-1}\right) \varepsilon_{k,t-k}\right]. \tag{40}$$

Defining the vector

$$S_t = (V_t', E_t V_{t+1}', ..., E_t V_{t+K}'; \varepsilon_{1,t}', \ldots, \varepsilon_{K,t-K+1}'; \varepsilon_{2,t}', \ldots, \varepsilon_{K,t-K+2}'; \ldots; \varepsilon_{K,t}')',$$

it is convenient to write the expected future process for V_t in companion form

$$S_t = G_\varepsilon(p_t)\varepsilon^t + G_S(p_t, p_{t-1}; J)S_{t-1} + G_\xi \xi_t \tag{41}$$

$$E[S_{t+1}|\Omega_t] = \widetilde{G}_S S_t + G_\xi G_\mu(p_t), \tag{42}$$

where $G_\mu(p_t) \equiv \sum_{j=1}^{N} p_{jt}\mu(j)$. The matrix $G_\varepsilon(p_t)$ takes the following form

$$\underset{(n_V+n_E)\times n_V}{G_\varepsilon} = \begin{bmatrix} \underset{n_V\times n_V}{G_{V\varepsilon}} \\ \underset{n_E\times n_V}{G_{\varepsilon\varepsilon}} \end{bmatrix}, \qquad \underset{n_V\times n_V}{G_{V\varepsilon}} = \begin{bmatrix} \theta_0(J) & \mathbf{0}_{n_v} & \dots & \mathbf{0}_{n_v} \\ \mathbf{0}_{n_v} & \tilde{\theta}_{1,t} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0}_{n_v} \\ \mathbf{0}_{n_v} & \dots & \mathbf{0}_{n_v} & \tilde{\theta}_{K,t} \end{bmatrix},$$

$$\underset{n_E\times n_V}{G_{\varepsilon\varepsilon}} = \begin{bmatrix} \underset{Kn_v\times n_v}{0} & G_{\varepsilon\varepsilon}^{1} & \underset{Kn_v\times n_v}{0} & \dots & \underset{Kn_v\times n_v}{0} \\ \underset{(K-1)n_v\times n_v}{0} & \underset{(K-1)n_v\times n_v}{0} & G_{\varepsilon\varepsilon}^{2} & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \underset{2n_v\times n_v}{0} \\ \mathbf{0}_{n_v} & \mathbf{0}_{n_v} & \dots & \mathbf{0}_{n_v} & G_{\varepsilon\varepsilon}^{K} \end{bmatrix}, \qquad G_{\varepsilon\varepsilon}^{k} = \begin{bmatrix} \mathcal{I}_{n_v} \\ \underset{(K-k)n_v\times n_v}{0} \end{bmatrix}, \qquad k = \{1,\dots,K\}.$$

In addition, the matrix $G_S(p_t, p_{t-1}; J)$ that establishes the relationship between S_t and its lag S_{t-1} is

$$\underset{(n_V+n_E)\times(n_V+n_E)}{G_S} = \begin{bmatrix} \underset{n_V\times n_V}{G_{11}} & \underset{n_V\times n_E}{G_{12}} \\ \underset{n_E\times n_V}{0} & \underset{n_E\times n_E}{G_{22}} \end{bmatrix}, \qquad \underset{n_V\times n_V}{G_{11}} = \begin{bmatrix} \underset{Kn_v\times n_v}{0} & \mathcal{I}_{Kn_v} \\ \mathbf{0}_{n_v} & \underset{n_v\times Kn_v}{0} \end{bmatrix},$$

$$\underset{n_V\times n_E}{G_{12}} = \left[\begin{array}{ccc|ccc|c|c} (\theta_1(J)-\tilde{\theta}_{1,t-1}) & \dots & (\theta_K(J)-\tilde{\theta}_{K,t-1}) & & \underset{n_v\times(K-1)n_v}{0} & & \dots & \mathbf{0}_{n_v} \\ & \underset{n_v\times Kn_v}{0} & & \Delta\tilde{\theta}_{2,t} & \dots & \Delta\tilde{\theta}_{K,t} & \ddots & \vdots \\ & \vdots & & & \underset{n_v\times(K-1)n_v}{0} & & \ddots & \mathbf{0}_{n_v} \\ & \vdots & & & \vdots & & \ddots & \Delta\tilde{\theta}_{K,t} \\ & \underset{n_v\times Kn_v}{0} & & & \underset{n_v\times(K-1)n_v}{0} & & \dots & \mathbf{0}_{n_v} \end{array}\right],$$

$$\underset{n_E\times n_E}{G_{22}} = \begin{bmatrix} \mathbf{0}_{Kn_v} & G_{22}^{1} & \underset{Kn_v\times(K-2)n_v}{0} & \dots & \underset{Kn_v\times n_v}{0} \\ \vdots & \mathbf{0}_{(K-1)n_v} & G_{22}^{2} & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \underset{3n_v\times n_v}{0} \\ \vdots & \vdots & \dots & \mathbf{0}_{2n_v} & G_{22}^{K-1} \\ \underset{n_v\times Kn_v}{0} & \underset{n_v\times(K-1)n_v}{0} & \dots & \underset{n_v\times 2n_v}{0} & \mathbf{0}_{n_v} \end{bmatrix},$$

$$G_{22}^{k} = \begin{bmatrix} \underset{n_v\times(K-k)n_v}{0} \\ \mathcal{I}_{(K-k)n_v} \end{bmatrix}, \qquad k = \{1,\dots,K-1\}.$$

Finally, the matrices in (42) are $G_\xi = \left(\mathcal{I}_{nv}, \underset{n_v\times(Kn_v+n_E)}{0'}\right)'$ and $\underset{(n_V+n_E)\times(n_V+n_E)}{\tilde{G}_S} = \begin{bmatrix} \underset{n_V\times n_V}{G_{11}} & \underset{n_V\times n_E}{0} \\ \underset{n_E\times n_V}{0} & \underset{n_E\times n_E}{G_{22}} \end{bmatrix}$.

C Risk-adjusted optimal policy under commitment

We modify the optimal monetary solution in Dennis (2007) to incorporate interest rate smoothing and show how the solution for optimal policy under commitment changes as a result. We represent optimal policy under commitment as a policymaker's problem at date 0 as:

$$\mathcal{L} = \min_{\{Y_t,i_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{1}{2}\left[Y_t'WY_t + X_t'QX_t\right] + \lambda_t'\left[A_0Y_t - A_1Y_{t-1} - \right.\right.$$
$$\left.\left. A_2Y_{t+1} - A_3X_t - A_4X_{t+1} - \nu_t\right] \right\} \tag{43}$$

where Y_t is a $N_y \times 1$ vector of endogenous variables, $X_t = (i_t, i_{t-1})'$ where i_t is the policy rate. The vector, λ_t, denotes the vector of lagrange multipliers on the model's N_y private-sector equilibrium conditions:

$$A_0Y_t = A_1Y_{t-1} + A_2E_tY_{t+1} + A_3X_t + A_4E_tX_{t+1} + \widetilde{A}_5S_t \tag{44}$$

where

$$V_t = \xi_t + \sum_{k=0}^{K} \theta_k(J)\varepsilon_{k,t-k} \tag{45}$$

is described in detail in Section 3.1. The rational expectations error ν_t satisfies:

$$\begin{aligned} \nu_t &= A_5V_t - A_2\eta_{t+1} - A_4\mu_{t+1} \\ \eta_{t+1} &= Y_{t+1} - E_tY_{t+1} \\ \mu_{t+1} &= X_{t+1} - E_tX_{t+1} \end{aligned} \tag{46}$$

The first order conditions with respect to Y_t are given by:

$$WY_t + A_0'\lambda_t - \beta^{-1}A_2'\lambda_{t-1} - \beta A_1'E_t\lambda_{t+1} = 0 \tag{47}$$

We solve the model using the "timeless perspective" and impose $\lambda_{-1} = 0$. As a result, these conditions are consistent with optimality at $t \geq 0$. With interest rate smoothing in the loss function, the matrix, Q, satisfies:

$$Q = \rho_i \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

where ρ_i determines the weight on the change in the policy rate in the loss function. The first order condition with respect to the policy rate is given by:

$$\rho_i(i_t - i_{t-1}) - \beta\rho_i(E_t i_{t+1} - i_t) - A'_{3,\cdot 1}\lambda_t - \beta A'_{3,\cdot 2}E_t\lambda_{t+1} = 0 \tag{48}$$

where $A_{3,\cdot i}$ denotes the ith column of the matrix A_3. Here we have assumed that $A_{4,\cdot 2} = 0$, which holds without loss of generality, since $A_{4,\cdot 2}$ is redundant with $A_{3,\cdot 1}$.

With these first order conditions in hand, we can write the model's equilibrium conditions as:

$$CZ_t = BZ_{t-1} + FE_t Z_{t+1} + DS_t \tag{49}$$

where $Z_t = (Y_t, i_t, \lambda_t)'$ and the C, B, and F matrices are given by:

$$C = \begin{bmatrix} A_0 & -A_{3,\cdot 1} & 0 \\ 0 & \rho_i(1+\beta) & -A'_{3,\cdot 1} \\ W & 0 & A'_0 \end{bmatrix}, \qquad B = \begin{bmatrix} A_1 & A_{3,\cdot 2} & 0 \\ 0 & \rho_i & 0 \\ 0 & 0 & \beta^{-1}A'_2 \end{bmatrix}$$

$$F = \begin{bmatrix} A_2 & A_{4,\cdot 1} & 0 \\ 0 & \beta\rho_i & \beta A'_{3,\cdot 2} \\ 0 & 0 & \beta A'_1 \end{bmatrix}, \qquad D = \begin{bmatrix} \widetilde{A}_5 \\ 0 \\ 0 \end{bmatrix}$$

These equilibrium conditions form the basis for solving the model when optimal policy embeds risk-management considerations.

Equations (49), (41) and (42) are the equilibrium conditions used to solve the model. Using the method of undetermined coefficients, the solution takes the form:

$$Z_t = TZ_{t-1} + RS_t + R_0(p_t) \tag{50}$$

$$S_t = G_\varepsilon(p_t)\varepsilon^t + G_S(p_t, p_{t-1}; J)S_{t-1} + G_\xi\xi_t \tag{51}$$

The matrices T, R and $R_0(p_t)$ can be computed as:

$$T = (C - FT)^{-1}B$$
$$R = (C - FT)^{-1}\left[FR\widetilde{G}_S + D\right]$$
$$R_0(p_t) = (C - FT)^{-1}F\left[RG_\xi G_\mu(p_t) + R_0(p_t)\right].$$

Equivalently,

$$FT^2 - CT + B = 0 \tag{52}$$
$$vec(R) = \left[(\mathcal{I}_{n_V+n_E} \otimes (C - FT)) - (\widetilde{G}'_S \otimes F)\right]^{-1} vec(D) \tag{53}$$
$$R_0(p_t) = (C - FT - F)^{-1}\left[FRG_\xi G_\mu(p_t)\right]. \tag{54}$$