ASEAN +3 MACROECONOMIC RESEARCH OFFICE

Working Paper (WP/26-02)

Monetary Policy under High Uncertainty:
The Option Value of Waiting

Haobin Wang, Allen Ng

April 2026

[This page is intentionally left blank]

**Monetary Policy under High Uncertainty:
The Option Value of Waiting**

Prepared by Haobin Wang and Allen Ng[1 2]

March 2026

Abstract

We embed the option value of waiting into a standard monetary policy framework to analyze optimal monetary policy under disturbance uncertainty. Under high uncertainty, especially when shocks are assessed to be transitory or carry high trade-offs, there is an option value of waiting. We derive a threshold rule that formalizes this trade-off, determining when the benefits of acting outweigh the option value of waiting for more information. If shocks are assessed to be persistent, there is a case for front-loading policy actions instead of waiting. The risk of hitting the effective lower bound reduces the option value of waiting. State-contingent communication dominates rigid commitments under high uncertainty.

JEL classification: E52, E58, D81

Keywords: Monetary policy; uncertainty; option value; forward guidance; effective lower bound

---

[1] Authors' e-mail: haobin.wang@amro-asia.org, allen.ng@amro-asia.org; This paper was authorized for distribution by Dong He, Chief Economist.

[2] The authors would like to thank Yuhong Wu and Jun Ee Yohnsen Ang for their excellent research assistance. We thank Aruhan Shi, Xiangguo (Jerry) Huang, Catherine Kho, Seung Hyun (Luke) Hong, Runchana Pongsaparn, Sekar Utami Setiastuti, Yijia Chen for useful comments. All errors remain the authors'.

Contents

I. Introduction

Recent years have seen an extraordinary confluence of shocks that have reshaped the global macroeconomic landscape and challenged conventional approaches to monetary policy. From the COVID-19 pandemic to the cascading effects of supply chain disruptions, energy price spikes, and geopolitical tensions, central banks have found themselves navigating an environment marked by historically high levels of uncertainty.

Of course, uncertainty is not new to central banking – policymakers have always had to make decisions under incomplete information. What distinguishes the current environment is the sheer confluence of shocks, differing in nature, origin, and persistence, that have hit the global economy in rapid succession. This has left central banks increasingly needing to grapple with the key questions: whether to act, when to act, and how forcefully to respond. A useful starting point for thinking about monetary policy in this environment is to distinguish between three broad categories of uncertainty, each of which has been vividly illustrated in recent years[3].

First, structural uncertainty concerns the fundamental workings of the economy – the strength of policy transmission, the slope of the Phillips curve, or the level of the neutral rate of interest. The post-pandemic inflation surge, for instance, has reignited fierce debates over the stability of the Phillips curve and sparked a resurgence of interest in estimating the neutral rate of interest (r^*) even in advanced economies like the United States and Japan, where it was long thought to be better understood.

Second, radical uncertainty, in the tradition of Knight (1921), describes situations where we cannot even specify the relevant states of the world or assign meaningful probabilities to them. These are the true "unknown unknowns": the COVID-19 pandemic, major geopolitical conflicts, and the increasing frequency of extreme weather events are all shocks whose macroeconomic consequences are difficult to model using standard probabilistic methods.

The third category is disturbance uncertainty: ambiguity regarding the specific shocks hitting the economy – their nature, size, and persistence. This was the central challenge of the 2021–2023 inflation episode, when policymakers struggled to disentangle the effects of a confluence of shocks, including supply-chain disruptions, shifts in consumer demand from services to goods, and large-scale fiscal stimulus.

This paper focuses on disturbance uncertainty. While structural uncertainty requires continuous research and radical uncertainty calls for contingency planning, it is disturbance uncertainty that policymakers must confront in real time to make high-frequency decisions. The diagnostic questions at its core – is this a demand or supply shock? Is it transitory or persistent? – directly determine whether, when, and how forcefully to respond.

To this end, we develop a tractable framework for analyzing optimal monetary policy under diagnostic ambiguity. By embedding real-options logic, first applied to investment by Dixit and Pindyck (1994) and Bernanke (1983), into a stylized IS-Phillips Curve framework, we derive analytical conditions under which waiting becomes the optimal strategy. This allows

[3] The following three-part taxonomy is our own synthesis for the purpose of classifying policy challenges, drawing on the broader uncertainty literature. See AMRO (2026), Chapter 1, Special Feature: "Decision-Making in the Dark – ASEAN+3 Monetary Policy in a High-Uncertainty Environment" for a more in-depth discussion on the three types of uncertainty.

us to formally explore the trade-off between acting immediately and preserving policy space until the nature of the economic disturbance becomes clearer.

II. Related Literature

This paper draws on two main strands of the economics literature: the theory of real options and the value of waiting, and the emerging literature on diagnostic ambiguity in monetary policy.

The case for policy caution under uncertainty has a long and distinguished history in economics. The seminal contribution is from Brainard (1967), who showed that uncertainty about the effectiveness of policy instruments — that is, multiplicative uncertainty over the transmission mechanism — should lead policymakers to attenuate their responses. A related but distinct insight emerges in the literature on irreversible investment under uncertainty, where Bernanke (1983) demonstrated that when investment decisions are irreversible, waiting to gather more information has value. This insight was formalized and extended by Dixit and Pindyck (1994), who developed a comprehensive framework for analyzing investment under uncertainty using real-options methods. [4]

While this literature was developed primarily in the context of firm investment, its logic applies directly to monetary policy. A central bank's decision to change interest rates is not literally irreversible, but it is costly to reverse: policy pivots can undermine credibility, confuse market expectations, and generate unnecessary volatility. A separate strand of literature has focused on the practical challenges central banks face when diagnosing the shocks hitting the economy in real time. The 2021–2023 inflation episode brought this challenge into sharp relief, motivating a growing body of work on monetary policy under diagnostic uncertainty (Amaral and others 2025; Bauer and others 2025).

Two key dimensions of this uncertainty are the type of shock (demand vs. supply) and its persistence (transitory vs. permanent). The optimal policy response depends critically on both. Clarida, Galí, and Gertler (1999) provide the canonical treatment, showing that central banks should accommodate supply shocks but offset demand shocks. More recently, Shapiro (2024) develops a methodology for decomposing inflation into its demand- and supply-driven components, highlighting the empirical difficulty of this task in real time.

Likewise, the persistence of shocks is a crucial, yet uncertain, input into policy. Brandão-Marques, Meeks, and Nguyen (2024) argue that when inflation persistence is uncertain, a robust approach that guards against the worst-case scenario (high persistence) is warranted. They find that underestimating persistence is a particularly costly error, suggesting a bias toward more front-loaded policy responses.

This paper bridges these two strands of literatures by embedding the option value of waiting into a stylized monetary policy model where the central bank faces uncertainty about the type and persistence of shocks. Existing literature has also explored the option value in central bank decisions. Lei and Tseng (2019) show that fixed adjustment costs generate an option value of waiting, while Hasui (2021) demonstrates that reversal aversion produces a similar effect. Both studies draw on option pricing theory from finance. Our approach differs

[4] As will be illustrated in the next section, our paper addresses a different source of uncertainty: additive disturbance uncertainty about the shocks themselves, combined with partial irreversibility of policy actions. This leads to a qualitatively different prescription; not a graduated response, but a binary act-or-wait threshold.

in that we embed the option value within a standard IS-Phillips Curve framework, and the option value arises directly from uncertainty about the nature and persistence of demand and supply shocks, rather than from assumed adjustment costs or reversal aversion. Our analysis also goes beyond the option value itself, exploring how policy space constraints (the effective lower bound) and communication strategies should adapt under high uncertainty.

III. Model Framework

The model comprises two periods, t=1,2, and is built around static representations of the IS and Phillips curves:

$$\begin{aligned} y_t &= -\frac{1}{\varphi}(r_t - r^n) + \varepsilon_t^d \quad \text{(IS Curve)} \\ \pi_t &= \kappa y_t + \varepsilon_t^s \quad \text{(Phillips Curve)} \end{aligned}$$

where y_t is the output gap, π_t is inflation, r_t is the policy rate, and r^n is the natural rate of interest (normalized to zero). For simplicity, we normalize the steady-state inflation target to zero, so that π_t represents both the inflation rate and the deviation from its target (the inflation gap). Following standard convention, a positive supply shock ($\varepsilon^s > 0$) represents a cost-push shock that raises inflation directly, while a positive demand shock ($\varepsilon^d > 0$) raises both output and inflation. These are static, reduced-form representations of the IS and Phillips curves. This is a deliberate modeling choice; it strips the framework down to its essential elements—the central bank's real-time diagnostic problem—which would be obscured by the additional complexity of forward-looking expectations.

The shock structure follows:

- In period 1, the shocks $(\varepsilon_1^d, \varepsilon_1^s)$ are realized and perfectly observed.
- In period 2, the shocks $(\tilde{\varepsilon}_2^d, \tilde{\varepsilon}_2^s)$ are unknown at $t = 1$. The central bank holds prior beliefs about their distribution.

The central bank can either:

- **Act Now**: Set $r_1 \neq 0$ based on observed period-1 shocks. We assume this commits the central bank to $r_2 = r_1$, reflecting a desire for policy continuity and the cost of adjusting policy in the short term. This assumption can be relaxed later to allow for policy adjustments that carry a reputational cost.
- **Wait**: Set$r_1 = 0$, and choose r_2 optimally after observing period-2 shocks. Under this strategy, the central bank retains the ability to set r_2 optimally conditional on the realized shocks in period 2. From the perspective of period 1, the central bank evaluates the expected loss by integrating over all possible shock realizations, in each case assuming it will respond optimally. The linear-quadratic structure of the model ensures this is analytically tractable: the optimal r_2 is a linear function of the realized shocks, and the expected minimized loss reduces to a function of the shock variances..

The policymaker minimizes the expected quadratic loss over two periods:

$$L = \sum_{t=1}^{2} \left(\frac{1}{2} y_t^2 + \frac{\lambda}{2} \pi_t^2 \right)$$

where $\lambda > 0$ reflects the relative aversion to inflation versus output deviations.

Solving the optimization problem yields an optimal policy rate under immediate action:

$$r^* = \frac{\varphi}{1 + \lambda\kappa^2}\left(\varepsilon_1^d + p \cdot \mu_d + \lambda\kappa\varepsilon_1^s + \lambda\kappa(1 - p) \cdot \mu_s\right)$$

where μ_d and μ_s are the expected future demand and supply shocks. Note that the coefficient on the supply shock is positive: a cost-push shock that raises inflation calls for a higher interest rate to partially offset the inflationary pressure.

Proposition 1 (Optimal Waiting Threshold): There exists a threshold condition for acting versus waiting. The central bank should act if and only if the squared benefit of acting, weighted by the structure of the economy, exceeds the expected variance of future shocks.

$$((1 + \lambda\kappa^2)\varepsilon_1^d + \lambda\kappa\varepsilon_1^s)^2 > 2(1 + \lambda\kappa^2)((1 + \lambda\kappa^2)\sigma_d^2 + \lambda\sigma_s^2)$$

The result follows from comparing the expected loss from acting immediately (L^{act}) versus the expected loss from waiting (L^{wait}). Acting is optimal if $L^{\text{act}} < L^{\text{wait}}$. The left-hand side of Equation represents the squared magnitude of the optimal policy response to current shocks, capturing the benefit of acting. The right-hand side represents the option value of waiting, which is proportional to the expected variance of future shocks. The full derivation is provided in Appendix A.

IV. Extensions

A. Shock Persistence

A key aspect of the uncertainty surrounding recent economic shocks has been their persistence. To analyze this rigorously, we extend the model to include persistence in both demand and supply shocks. The shocks now evolve according to an AR(1) process:

$$\varepsilon_{2,i} = \rho\varepsilon_{1,i} + \eta_{2,i}, \quad i \in \{d, s\}$$

For simplicity, we assume both shocks share the same persistence parameter, $\rho \in [0,1]$. The innovations $\eta_{2,d}$ and $\eta_{2,s}$ are i.i.d. with variances $\sigma_{\eta,d}^2$ and $\sigma_{\eta,s}^2$.

Proposition 2 (Shock Persistence): Higher shock persistence (ρ) increases the benefit of acting.

With persistent shocks, the central bank's optimal policy under commitment must account for the expected future path of shocks. The revised threshold condition, derived in Appendix A.2, is:

$$\Gamma(\rho) \cdot ((1 + \lambda\kappa^2)\varepsilon_1^d + \lambda\kappa\varepsilon_1^s)^2 > (1 + \lambda\kappa^2)\left((1 + \lambda\kappa^2)\sigma_{\eta,d}^2 + \lambda\sigma_{\eta,s}^2\right)$$

where $\Gamma(\rho) = 1 - \frac{(1-\rho)^2}{2}$. The right-hand side represents the option value of waiting, which now depends on the variance of the future shock innovations. The left-hand side is the benefit of acting, which is the baseline benefit from Proposition 1 scaled by the persistence factor $\Gamma(\rho)$.

Since $\Gamma'(\rho) = (1 - \rho) > 0$ for $\rho < 1$, the benefit of acting is strictly increasing in the persistence of the shocks. A more persistent shock (higher ρ) increases the left-hand side of the inequality, making the condition for acting easier to satisfy and thus lowering the threshold. The implication is clear: persistent shocks warrant earlier action, while transitory shocks favor waiting.

Figure 1 illustrates how the optimal policy response varies across four canonical shock scenarios, distinguished by shock type (demand vs. supply) and persistence (high vs. low).

Figure 1. Model-Implied Optimal Responses to Different Types of Shocks

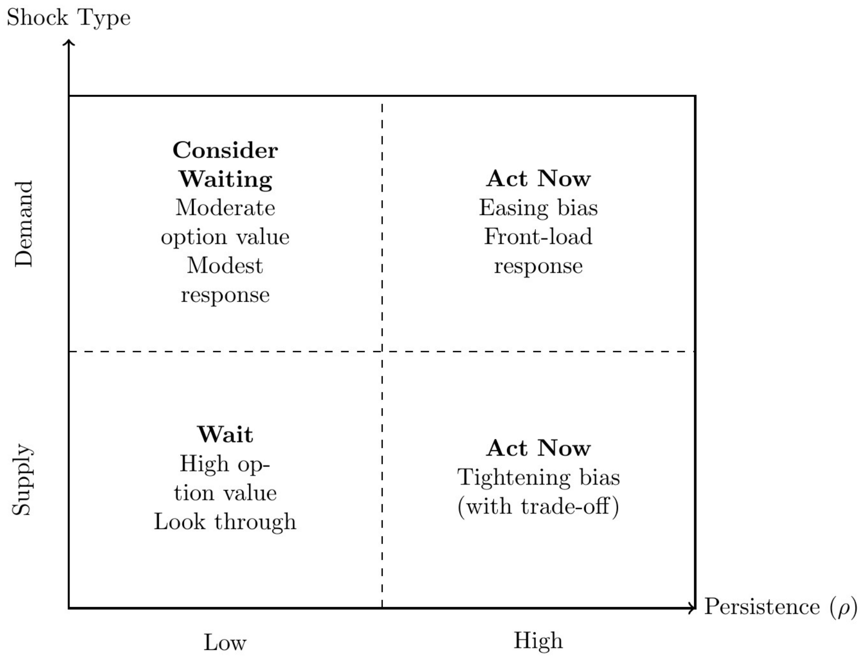

Source: AMRO staff.

B. The Effective Lower Bound

In recent years, central banks in some economies have faced the challenge of needing to ease monetary policy when policy space is already limited. This situation is further complicated when high uncertainty calls for preserving policy space for the future. We therefore analyze how the risk of exhausting policy space at the effective lower bound (ELB) affects the optimal policy response. When interest rates approach zero, the central bank's ability to respond to future negative shocks becomes constrained. This section examines how the effective lower bound (ELB) affects the option value of waiting.

In period 2, the nominal interest rate is constrained:

$$r_2 \geq \bar{r} \geq 0$$

Proposition 3 (Effective Lower Bound): The risk of hitting the effective lower bound (ELB) reduces the option value of waiting.

Without the ELB, the optimal policy fully closes the output gap: $r_2^* = \varepsilon_2/\varphi$. With the ELB, the optimal policy becomes:

$$r_2^{\text{opt}} = \max\left\{\bar{r}, \frac{\varepsilon_2}{\varphi}\right\}$$

Since $\varepsilon_2 \sim N(\rho\varepsilon_1, \sigma_\eta^2)$, the probability that the ELB binds is:

$$P_{\text{ELB}} = \Phi\left(\frac{\varphi\bar{r} - \rho\varepsilon_1}{\sigma_\eta}\right)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function.

The expected loss from waiting now includes an additional term reflecting the cost of being constrained at the ELB:

$$\mathbb{E}_1[L_2^{\text{wait}}] = P_{\text{ELB}} \cdot \lambda \cdot \mathbb{E}[(\varepsilon_2 - \varphi\bar{r})^2 \mid \varepsilon_2 < \varphi\bar{r}]$$

Without the ELB, optimal policy in period 2 closes the gap completely (zero loss). With the ELB, there is positive expected loss when the constraint binds.

The modified threshold condition becomes:

$$\underbrace{c \cdot \text{Var}_1(x_2, \pi_2)}_{\text{Option value of waiting}} \leq \underbrace{\left[1 - \frac{(1-\rho)^2}{2}\right](x_1^2 + \lambda\pi_1^2)}_{\text{Benefit of acting now}} - \underbrace{P_{\text{ELB}} \cdot \lambda \cdot \mathbb{E}[(\varepsilon_2 - \varphi\bar{r})^2 \mid \text{ELB binds}]}_{\text{ELB adjustment}}$$

The ELB adjustment term is always positive, reflecting the additional expected loss when the ELB binds. Thus, the condition for waiting becomes harder to satisfy – the presence of ELB risk strengthens the incentive to act early. This finding is consistent with Eggertsson and Woodford (2003), who show that a credible commitment to a history-dependent policy can largely mitigate the distortions created by the effective lower bound. Our result can be seen as a specific application of this principle: by easing more aggressively today, the central bank signals its commitment to offset future negative shocks, thereby anchoring expectations and stimulating current activity.

C. Forward Guidance

Forward guidance became a widely embraced central bank communication strategy to anchor expectations, particularly after the Global Financial Crisis. However, in an environment where the outlook is highly uncertain and volatile, rigid, time-based forward guidance can become ineffective or even counterproductive. This section explores how communication strategies should adapt under high uncertainty.

Suppose at $t = 1$, the central bank provides forward guidance by communicating an intended policy stance $\hat{r}_2$ for period 2. We assume the central bank incurs a reputational loss when it deviates from prior guidance:

$$L = \lambda[(x_1)^2 + \lambda(\pi_1)^2] + \lambda[(x_2)^2 + \lambda(\pi_2)^2] + \theta(r_2 - \hat{r}_2)^2$$

where $\theta > 0$ captures the weight on reputation.

Proposition 4 (State-Contingent Communication): Under high uncertainty, state-contingent communication dominates rigid, time-based commitments.

When the central bank acts in period 1 and sets r_1, it implicitly signals $\hat{r}_2 = r_1$. In period 2, the optimal rate balances stabilization against reputation:

$$r_2^* = \frac{\gamma}{\gamma + \theta} \cdot \frac{\varepsilon_2}{\varphi} + \frac{\theta}{\gamma + \theta} \cdot r_1$$

where $\gamma = \lambda\varphi^2(1 + \lambda\kappa^2)$. The optimal r_2 is a weighted average of the stabilization-optimal rate and the forward-guided rate.

The threshold condition for acting now becomes:

$$\mathbb{E}_1[\theta(r_2 - r_1)^2] \leq \text{Gain from waiting (in terms of better stabilization).}$$

If the central bank waits and provides no strong guidance in $t = 1$, it retains full flexibility in $t = 2$ without reputational cost:

$$L^{\text{wait}} = \lambda[(x_1)^2 + \lambda(\pi_1)^2] + \mathbb{E}_1[\lambda(x_2^2 + \lambda\pi_2^2)].$$

The forward guidance penalty effectively reduces the value of waiting, since delayed action avoids reputational cost but forgoes the benefits of commitment. This condition shows that:

- Higher θ increases the reputational cost of deviating from prior guidance, thereby making waiting less attractive.
- When the policymaker expects large deviations between r_1 and the optimal r_2, the reputational cost is more likely to bind.
- The presence of reputational cost may lead to either premature tightening or premature easing, depending on whether the initial signal over- or under-estimates future needs.

This extension helps rationalize why central banks may under-react to evolving data in order to maintain consistency with past forward guidance, even when full flexibility would be optimal ex post. The presence of reputational concerns interacts with uncertainty and persistence to alter the timing and magnitude of optimal policy moves. A more robust approach under uncertainty is to adopt contingent forward guidance, whereby the central bank pre-commits to conditional policy paths:

$$r_2 = \begin{cases} r_2^D, & \text{if a demand shock materializes} \\ r_2^S, & \text{if a supply shock materializes} \end{cases}$$

These paths are still commitments, but they allow for flexibility and avoid credibility loss as long as the central bank follows its pre-announced rule. The expected credibility loss under this scheme becomes:

$$\text{Credibility Loss}^{\text{conditional}} = \sum_{j \in \{D,S\}} \pi_j \left[L(r_2^j) - L(r_2^{*,j})\right],$$

where j indexes the type of shock, π_j is the probability of shock j, and $r_2^{*,j}$ is the ex-post optimal rate under shock j.

A rigid commitment to a single future policy path incurs a large expected credibility loss when the future is uncertain. A state-contingent commitment, which specifies different policy paths for different future states of the world, allows the central bank to tailor its response to the realized shock, thus minimizing the deviation from the ex-post optimal policy. The expected credibility loss is therefore significantly lower:

$$\text{Credibility Loss}^{\text{conditional}} \ll \text{Credibility Loss}^{\text{rigid}}.$$

This demonstrates the superiority of communication strategies that preserve flexibility.

1. This setup highlights the benefit of data-dependent and state-contingent communication strategies, allowing central banks to maintain credibility while preserving policy flexibility under high uncertainty. This finding is consistent with the arguments in Woodford (2012), who shows that threshold-based forward guidance that links future policy to observable economic conditions (such as unemployment or inflation thresholds) is more credible and effective than rigid, time-based commitments, particularly at the ELB. The effectiveness of forward guidance, particularly at the ELB, depends crucially on its credibility (Woodford, 2013).

V. Conclusion

This paper has explored the nuances of monetary policy under disturbance uncertainty. By embedding the logic of real options into a stylized IS-Phillips Curve framework, we formalize the trade-offs inherent in making policy in real time, when the nature, size, and persistence of economic shocks are not fully known.

Our analysis yields a nuanced view of policy caution. When faced with ambiguous and potentially transitory shocks, a strategy of waiting is not passive but optimal, as it preserves flexibility and reduces the risk of costly policy errors. However, this cautious approach has clear limits. If a shock is diagnosed as persistent or particularly large, the cost of inaction rises sharply, shifting the calculus toward a more front-loaded and decisive policy response to prevent the economy from drifting further from its objectives.

The decision is further complicated by two important factors in modern central banking. First, the presence of the effective lower bound (ELB) asymmetrically constrains future policy space. The risk of being unable to respond to future negative shocks strengthens the case for pre-emptive action, particularly when downside risks dominate. Second, communication strategies become critical. Our framework highlights the superiority of state-contingent forward guidance, which allows central banks to communicate their policy reaction function under different scenarios. This approach enhances credibility and anchors expectations without sacrificing the flexibility needed to respond to evolving conditions – a marked advantage over rigid, time-based commitments. We apply the propositions derived from the model to the context of Thailand to illustrate their practical application (see Appendix B).

Ultimately, in an environment of recurring and overlapping shocks, the ability to accurately diagnose the nature of economic disturbances is paramount. A disciplined, data-dependent, and diagnostic-driven approach to monetary policy, supported by a clear and credible communication strategy, offers the most reliable guide for navigating the uncertainties ahead.

VI. References

Amaral, E., Avalos, F., and Shim, I. 2025. Monetary policy decision-making and communication under high uncertainty. BIS Papers No 163, Bank for International Settlements.

ASEAN+3 Macroeconomic Research Office (AMRO). Forthcoming. *ASEAN+3 Regional Economic Outlook (AREO) 2026*. ASEAN+3 Macroeconomic Research Office, Singapore.

Bauer, M., Berge, T., Fiori, G., Loria, F., and Zhong, M. 2025. Accounting for Uncertainty and Risks in Monetary Policy. Finance and Economics Discussion Series 2025-073, Board of Governors of the Federal Reserve System.

Bernanke, B. S. 1983. Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics*, 98(1):85–106.

Brainard, W. C. 1967. Uncertainty and the effectiveness of policy. American Economic Review, 57(2):411–425.

Brandão-Marques, L., Meeks, R., and Nguyen, V. 2024. Monetary Policy with Uncertain Inflation Persistence. IMF Working Paper 2024/047, International Monetary Fund.

Clarida, R., Galí, J., and Gertler, M. 1999. The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature*, 37(4):1661–1707.

Dixit, A. K. and Pindyck, R. S. 1994. *Investment under uncertainty*. Princeton University Press.

Eggertsson, G. B. and Woodford, M. 2003. The effective lower bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 2003(1):139–233.

Hasui, K. 2021. Optimal irreversible monetary policy. *European Economic Review*, 134:103701.

Knight, F. H. 1921. *Risk, uncertainty and profit*. Houghton Mifflin.

Lei, X. and Tseng, M. C. 2019. Wait-and-see monetary policy. *Macroeconomic Dynamics*, 23(5):1793–1814.

Shapiro, A. H. 2024. Decomposing Supply- and Demand-Driven Inflation. *Journal of Money, Credit and Banking*, 56(S1):233–263.

Woodford, M. 2012. Methods of policy accommodation at the interest-rate lower bound. In *The Changing Policy Landscape: 2012 Jackson Hole Symposium*, pages 185–288. Federal Reserve Bank of Kansas City.

Woodford, M. 2013. Forward guidance by inflation-targeting central banks. In *Two Decades of Inflation Targeting: Main Lessons and Remaining Challenges*, pages 81–120. Sveriges Riksbank.

VII. Appendices

Appendix A: Technical Appendix

This appendix provides detailed derivations for the propositions in the main text.

A.1 Detailed Derivation of the Baseline Optimal Waiting Threshold (Proposition 1)

We begin with the full two-shock model. The central bank minimizes the two-period loss function:

$$L = \frac{1}{2}\sum_{t=1}^{2}(y_t^2 + \lambda \pi_t^2)$$

where $y_t = -\frac{1}{\varphi}r_t + \varepsilon_t^d$ and $\pi_t = \kappa y_t + \varepsilon_t^s$. We assume $r^n = 0$ and that period-2 shocks are zero-mean, independent, and have variances σ_d^2 and σ_s^2.

Case 1: Act Now. The central bank sets r_1 and commits to $r_2 = r_1$. The total expected loss is $L^{\text{act}} = L_1 + \mathbb{E}_1[L_2]$.

$$\begin{aligned} L_1 &= \frac{1}{2}\left[\left(-\frac{r_1}{\varphi} + \varepsilon_1^d\right)^2 + \lambda\left(\kappa(-\frac{r_1}{\varphi} + \varepsilon_1^d) + \varepsilon_1^s\right)^2\right] \\ \mathbb{E}_1[L_2] &= \frac{1}{2}\mathbb{E}_1\left[\left(-\frac{r_1}{\varphi} + \varepsilon_2^d\right)^2 + \lambda\left(\kappa(-\frac{r_1}{\varphi} + \varepsilon_2^d) + \varepsilon_2^s\right)^2\right] \\ &= \frac{1}{2}\left[\left(-\frac{r_1}{\varphi}\right)^2 + \sigma_d^2 + \lambda\left(\kappa^2(-\frac{r_1}{\varphi})^2 + \kappa^2\sigma_d^2 + \sigma_s^2\right)\right] \end{aligned}$$

Combining terms, the total loss is a quadratic function of r_1. Minimizing with respect to r_1 yields the optimal rate under commitment:

$$r_1^* = \frac{\varphi(1+\lambda\kappa^2)\varepsilon_1^d + \lambda\varphi\kappa\varepsilon_1^s}{2(1+\lambda\kappa^2)}$$

Case 2: Wait. The central bank sets $r_1 = 0$ and chooses r_2 optimally in period 2. The loss is:

$$\begin{aligned} L^{\text{wait}} &= L_1(r_1 = 0) + \mathbb{E}_1[\min_{r_2} L_2(r_2)] \\ L_1(r_1 = 0) &= \frac{1}{2}\left[(\varepsilon_1^d)^2 + \lambda(\kappa\varepsilon_1^d + \varepsilon_1^s)^2\right] \end{aligned}$$

In period 2, the central bank observes the realized shocks and sets r_2 optimally. For any given realization of $(\varepsilon_2, \varepsilon_2)$, the optimal rate takes the same linear form as in the "Act Now" case, evaluated at the realized shocks. From the perspective of period 1, the ex-ante optimal plan is $r_2 *= 0$ (since $E_1[\varepsilon_2] = E_1[\varepsilon_2] = 0$), but the expected loss accounts for the fact that the central bank will respond optimally to whatever shocks materialize. The expected second-period loss is determined by the residual variance of future shocks that optimal policy cannot fully eliminate:

$$\mathbb{E}_1[L_2^{\text{wait}}] = \frac{1}{2}\mathbb{E}_1\left[(\varepsilon_2^d)^2 + \lambda(\kappa\varepsilon_2^d + \varepsilon_2^s)^2\right] = \frac{1}{2}((1+\lambda\kappa^2)\sigma_d^2 + \lambda\sigma_s^2)$$

Threshold Condition. The central bank acts if the net benefit is positive, $L^{\text{wait}} - L^{\text{act}^*} > 0$. This can be expressed as the immediate *Benefit of Acting* being greater than the *Option Value of Waiting*.

$$\underbrace{L_1(r_1 = 0) - L_1(r_1^*)}_{\text{Benefit of Acting}} > \underbrace{\mathbb{E}_1[L_2(r_1^*)] - \mathbb{E}_1[L_2^{\text{wait}}]}_{\text{Option Value of Waiting}}$$

The option value of waiting is the expected future loss from being locked into r_1^* versus being able to set r_2 optimally. This simplifies to:

$$\text{Option Value} = \frac{1}{2}(1 + \lambda\kappa^2)\left(-\frac{r_1^*}{\varphi}\right)^2 = \frac{1}{8}\frac{((1 + \lambda\kappa^2)\varepsilon_1^d + \lambda\kappa\varepsilon_1^s)^2}{1 + \lambda\kappa^2}$$

After substituting and significant algebraic simplification, the condition $L^{\text{wait}} > L^{\text{act}^*}$ reduces to:

$$((1 + \lambda\kappa^2)\varepsilon_1^d + \lambda\kappa\varepsilon_1^s)^2 > 2(1 + \lambda\kappa^2)((1 + \lambda\kappa^2)\sigma_d^2 + \lambda\sigma_s^2)$$

This is the explicit form of the threshold condition in Proposition 1.

A.2 Threshold Condition under Shock Persistence (Proposition 2)

We introduce persistence with a common parameter ρ for both shocks:

$$\varepsilon_{2,i} = \rho\varepsilon_{1,i} + \eta_{2,i}, \quad i \in \{d, s\}$$

Case 1: Act Now. The central bank sets r_1 and commits to $r_2 = r_1$. The total loss $L^{\text{act}} = L_1 + \mathbb{E}_1[L_2]$ is minimized. The optimal rate r_1^* is:

$$r_1^* = \frac{\varphi(1 + \rho)}{2(1 + \lambda\kappa^2)}\left[(1 + \lambda\kappa^2)\varepsilon_1^d + \lambda\kappa\varepsilon_1^s\right]$$

This is the baseline optimal rate scaled by $(1 + \rho)$.

Case 2: Wait. The central bank sets $r_1 = 0$. The expected loss in period 2, $\mathbb{E}_1[L_2^{\text{wait}}]$, now includes terms for the persistent shocks $\rho\varepsilon_1^d$ and $\rho\varepsilon_1^s$ as well as the innovation variances $\sigma_{\eta,d}^2$ and $\sigma_{\eta,s}^2$.

Threshold Condition. Comparing the minimized loss from acting, L^{act^*}, with the loss from waiting, L^{wait}, yields the threshold condition. Let $X_1 = (1 + \lambda\kappa^2)\varepsilon_1^d + \lambda\kappa\varepsilon_1^s$. The condition $L^{\text{wait}} > L^{\text{act}^*}$ simplifies to:

$$\left(1 - \frac{(1 - \rho)^2}{2}\right)X_1^2 > (1 + \lambda\kappa^2)\big((1 + \lambda\kappa^2)\sigma_{\eta,d}^2 + \lambda\sigma_{\eta,s}^2\big)$$

This is the explicit form of the threshold condition in Equation. When $\rho = 0$, the condition becomes $\frac{1}{2}X_1^2 > (1 + \lambda\kappa^2)((1 + \lambda\kappa^2)\sigma_{\eta,d}^2 + \lambda\sigma_{\eta,s}^2)$, which is consistent with the baseline derivation where the future shock variance is now interpreted as the innovation variance.

A.3 Effective Lower Bound Extension (Proposition 3)

With the ELB constraint $r_2 \geq \bar{r}$, the optimal policy in period 2 becomes $r_2^{\text{opt}} = \max\{\bar{r}, r_2^*\}$, where r_2^* is the unconstrained optimal rate. The probability of the ELB binding is $P_{\text{ELB}} = \Phi((\varphi\bar{r} - \mathbb{E}_1[\varepsilon_2])/\sigma_\eta)$, where Φ is the standard normal cumulative distribution function.

The expected loss from waiting now includes an additional term reflecting the cost of being constrained:

$$\mathbb{E}_1[L_2^{\text{wait}}] = \mathbb{E}_1[L_2^{\text{unconstrained}}] + P_{\text{ELB}} \cdot \mathbb{E}[L_2(r_2 = \bar{r}) - L_2(r_2 = r_2^*) \mid \text{ELB binds}]$$

This additional positive term reduces the net option value of waiting, making early action more attractive when ELB risk is elevated.

A.4 Forward Guidance with Reputational Costs (Proposition 4)

When the central bank deviates from its prior guidance ($\hat{r}_2 = r_1$), it incurs a reputational cost $\theta(r_2 - r_1)^2$. The period-2 optimization problem becomes:

$$\min_{r_2} \left[\frac{1}{2}(y_2^2 + \lambda\pi_2^2) + \frac{\theta}{2}(r_2 - r_1)^2\right]$$

Let $\gamma = (1 + \lambda\kappa^2)/\varphi^2$. The first-order condition yields the new optimal rate:

$$r_2^* = \frac{\gamma\varphi((1 + \lambda\kappa^2)\varepsilon_2^d + \lambda\kappa\varepsilon_2^s) + \theta r_1}{\gamma(1 + \lambda\kappa^2) + \theta}$$

The optimal r_2 is a weighted average of the unconstrained optimal rate and the forward-guided rate r_1. Higher reputational cost θ pulls r_2 closer to r_1, reducing responsiveness to new information and effectively lowering the option value of waiting.

Appendix B: Application: A Scenario-Based Approach for Thailand

The model results suggest that when faced with high uncertainty, a key challenge is to diagnose the underlying shock's type and nature. In practice, an economy can be faced with a multitude of demand and supply shocks, and analyzing all of them may be ineffective. One approach is to adopt a scenario-based approach, which is to focus on a specific risk scenario most relevant to policymaking.

Take Thailand during 2023–2025 as an example. While, like other countries, it faced various demand and supply shocks, one of the most pressing risk scenarios during this period was a persistent weakness in private domestic demand. Household debt remained high at about 90 percent of GDP, limiting consumption amid sluggish income growth and fragile confidence. Private consumption growth slowed markedly, from around 7 percent in 2023 to approximately 2 percent by 2025. Credit growth remained negative across most sectors for a prolonged period.

Weak income growth and confidence, tightened credit conditions, and ongoing political uncertainty could create a negative feedback loop that perpetuates domestic demand weakness. Persistent credit constraints may erode borrower confidence and credit quality, reinforcing the downturn. If inflation expectations fall further, rising real debt burdens could add to the drag on demand, amplifying the adverse feedback loop.

The model results in the earlier sections show that when downside risks are persistent, front-loaded monetary easing becomes more optimal – delayed action can lead to deeper and more protracted output losses. However, if the demand weakness is deemed transitory and uncertainty remains high, it may be more prudent to preserve monetary space until clearer signals emerge.

So how should we decide whether Thailand is facing a transitory or persistent demand shock? Evaluating this risk requires a systematic approach to monitoring demand-side conditions using timely and targeted indicators. The Bank of Thailand's Private Consumption and Private Investment Indices, disaggregated by durables vs. non-durables and machinery vs. construction, can serve as core indicators for monitoring. External signals such as export orders from PMI surveys and the regional electronics cycle can offer complementary insights into demand momentum.

To extract the underlying demand shock, we can use a simple process of the form:

$$x_t = \mu_t + \varepsilon_t, \quad \varepsilon_t = \phi \varepsilon_{t-1} + u_t$$

where x_t denotes the log-level of private consumption or investment, μ_t captures the trend, and ε_t is the autocorrelated demand shock with persistence ϕ. The innovation term u_t reflects sudden changes in demand. This framework allows us to detect persistent deviations from trend, which are more relevant for policy calibration.

Recent data for Thailand suggests that the slowdown is relatively persistent. Our estimates show that after rebounding to above its pre-COVID trend, private consumption has recently fallen below the trendline (Figure B1.1). The estimated shocks have turned negative, with a persistence coefficient around 0.7–0.8 (Figure B1.2). The weakness is more pronounced in durable goods consumption, where negative shocks have emerged with high persistence (above 0.8) (Figure B1.3). Durables are particularly sensitive to credit availability and forward-looking household sentiment, making them a leading indicator of consumption trends.

Figure B1.1. Private Consumption Index

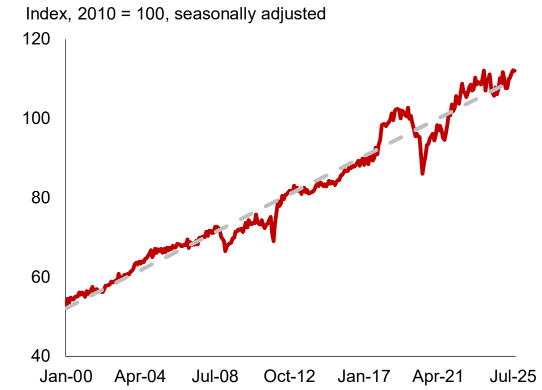

Source: BOT; AMRO staff calculation

Figure B1.2. Autocorrelated Shocks for Private Consumption Index

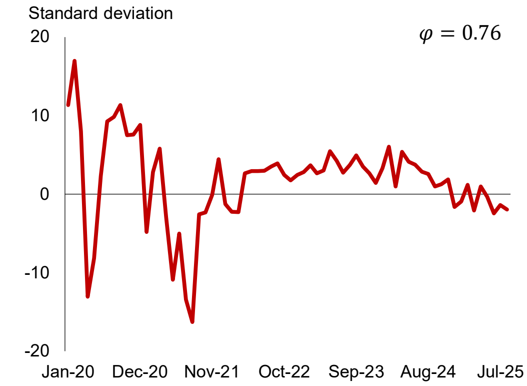

Source: BOT; AMRO staff calculation

Figure B1.3. Autocorrelated Shocks for Durables Index

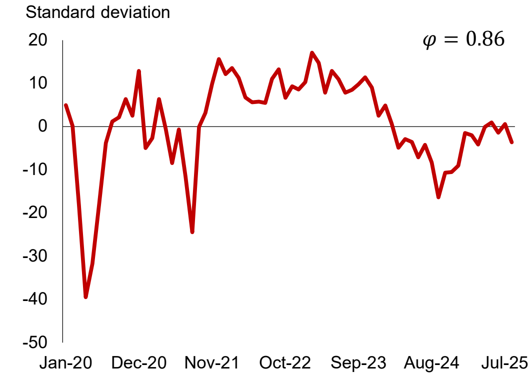

Source: BOT; AMRO staff calculation
Note: Durables index is the composite index comprising of sales of commercial cars, sales of Passenger cars, and sales of motorcycles.

Figure B1.4. Credit-to-GDP Gaps

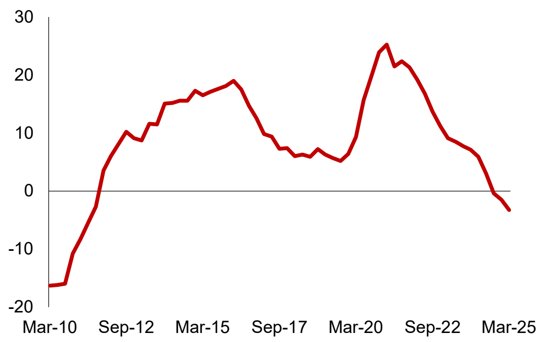

Source: BIS
Note: The credit-to-GDP gap is defined as the difference between the credit-to-GDP ratio and its long-run trend. The trend is derived using a one-sided Hodrick-Prescott filter. The credit-to-GDP ratio uses total credit to the private non-financial sector.

These consumption patterns are mirrored in credit dynamics. Thailand's credit-to-GDP gap has steadily narrowed and recently turned negative, suggesting that credit availability is falling behind its historical trend, despite still-elevated household debt levels (Figure B1.4). Tighter credit in an already leveraged household sector risks reinforcing demand weakness through reduced borrowing and spending.

These signals warrant close monitoring. Our framework provides a simple guide for policy response based on shock type and persistence. If demand shocks are both negative and persistent – indicated by sustained declines and high persistence estimates – our framework suggests a strong case for proactive easing to preempt a prolonged slowdown. This is because when downside risks are persistent, delayed action can lead to deeper and more protracted output losses. This insight aligns with existing literature on effective lower bound constraints (Eggertsson and Woodford 2003), which suggests frontloading policy accommodation when the risk of hitting the bound is elevated. Conversely, if the shocks appear transitory or mild, a wait-and-see approach would be more appropriate, preserving policy space for more severe scenarios. This tailored strategy equips policymakers with a simple and transparent framework to guide monetary policy decisions amid high uncertainty and evolving conditions.

The evolution of the Bank of Thailand's (BOT) policy response during this period offers an instructive real-world example of this dynamic. Official data show that year-on-year private consumption growth peaked at 7.9 percent in the third quarter of 2023 before beginning a steady decline – to 6.6 percent in the first quarter of 2024 and 3.3 percent by the third quarter of 2024. Throughout this initial period of decline, the Monetary Policy Committee (MPC) held the policy rate unchanged at 2.50 percent, assessing that private consumption remained robust. A central question at the time was whether the weakness was a temporary moderation or a more persistent trend.[5] It was not until August 2024, after nearly a year of decelerating consumption data, that the MPC's communication shifted, noting an "expected slowdown in private consumption going forward."[6] This change in assessment preceded the first policy rate cut of the cycle in October 2024. A series of subsequent rate cuts followed throughout 2025 – bringing the policy rate from 2.25 percent to 1.25 percent by December 2025 – as the MPC increasingly acknowledged the persistent nature of the economic headwinds, noting that the economy had "slowed more than expected" and that domestic demand had continued to weaken.[7]

[5] At a press briefing following the February 2024 MPC meeting, Deputy Governor Piti Disyatat remarked: "The question is whether this is a temporary thing or something more persistent that we need to be concerned about."

[6] MPC Press Release, Meeting 4/2024 (August 21, 2024): The Committee assessed that there would be an "expected slowdown in private consumption going forward after a period of strong growth."

[7] MPC Press Release, Meeting 1/2025 (February 26, 2025): The Committee noted that "the Thai economy slowed more than expected." By August 2025, the MPC described private consumption as "expected to be subdued due to weakening consumer confidence and income trajectory" (MPC Meeting 4/2025, August 13, 2025).

AMRO

ASEAN +3 MACROECONOMIC RESEARCH OFFICE

Address: 10 Shenton Way, #15-08
MAS Building, Singapore 079117
Website: www.amro-asia.org
Tel: +65 6323 9844
Email: enquiry@amro-asia.org
LinkedIn | **Twitter** | **Facebook**