LEARNING, UNCERTAINTY AND MONETARY POLICY

PABLO GARCIA

ABSTRACT. I present a New Keynesian model in which the central bank's anti-inflationary preferences change over time. Agents do not observe the current monetary regime, but rationally learn about it using Bayes theorem. The model reproduces the contractionary effects of monetary policy uncertainty shocks recently documented in the empirical literature. In addition, the model shows that learning reduces the effects of monetary policy on the economy by softening the link between fundamentals and equilibrium prices and allocations.

JEL Codes: C11, D83, E52.

Keywords: Monetary policy, Uncertainty, Bayesian learning.

E-mail: pablo.garciasanchez@bcl.lu. Banque Centrale du Luxembourg, Département Économie et Recherche, 2 boulevard Royal, L-2983 Luxembourg. For useful comments, I thank Paolo Guarda, Paolo Fegatelli, Patrick Fève, Alban Moura, Julien Pascal, and Olivier Pierrard. This paper should not be reported as representing the views of the Banque centrale du Luxembourg or the Eurosystem. The views expressed are those of the author and may not be shared by other research staff or policymakers at BCL or the Eurosystem.

1. Introduction

The uncertainty that surrounds public perceptions of central bank policy actions matters for business cycles. Using time-series econometrics, Creal and Wu (2017) and Husted et al. (2019) document the contractionary effects of increases in monetary policy uncertainty (MPU) in the United States, while Azqueta-Gavaldon et al. (2020) does so for the euro area. Why does MPU dampen economic activity?

Learning dynamics are key. In the context of an otherwise standard New Keynesian model, I consider an environment in which households and firms have limited information about the central bank reaction function.[1] As in Bianchi and Melosi (2018), the latter alternates between periods of active inflation stabilisation, when the Taylor principle is satisfied, and periods of passive inflation stabilisation, when the Taylor principle is violated.[2] Each quarter, the central bank sets the nominal interest rate consistent with past inflation, its desired anti-inflationary attitude, and discretionary monetary policy. While private agents know the structure of the economy and observe all endogenous variables, they do not know the current desired anti-inflationary attitude, for it is obscured by discretionary monetary policy. Instead, agents make an inference about the current monetary regime by combining what they see happening in the economy with their own past beliefs using the Hamilton filter [Hamilton (1989)]. Lastly, they include their state likelihood estimates into their expectations, and hence, into their decision-making process.

This belief structure accounts for the contractionary effects of increases in MPU. Spikes in uncertainty about the current monetary regime obscure the likely path of future interest rates, and thus, of the marginal product of capital. This unpredictability tempers agents' choices, thereby making it rational to defer consumption and investment until uncertainty is resolved.

Yet the sensitivity of consumption to MPU is lower than that of investment, for households attempt to smooth consumption. Moreover, the lower the intertemporal elasticity of substitution is, the lower the effects of MPU on consumption become. This makes sense: in the limit, when this intertemporal elasticity of substitution approaches zero, households are infinitely unwilling to substitute consumption over time.

[1] Recent survey data confirms that private agents' knowledge about monetary policy is imperfect; see e.g., Carvalho and Nechio (2014), van der Cruijsen et al. (2015) and Candia et al. (2021).

[2] According to Bernanke and Mishkin (1997), many central banks follow a constrained discretion strategy that provides them with some flexibility to de-emphasise inflation stabilisation to pursue secondary objectives.

I also explore how the unknown monetary policy regime affects the transmission of monetary policy shocks. As in Eusepi and Preston (2011), beliefs are a function of historical data, introducing an additional state variable. The model dynamics therefore evolve over time in response to central bank actions and private sector inferences. Relative to a Rational Expectations analysis of the model, learning makes the economy less sensitive to changes in the central bank's inflation response. Households are aware that they do not know the current monetary regime. As a result, they base their actions on the choices they would make under either monetary regime. By softening the link between fundamentals and equilibrium prices and allocations, learning reduces the effects of monetary policy on the economy.

As mentioned earlier, my paper relates to the growing empirical literature documenting the contractionary effects of MPU shocks; see e.g., Creal and Wu (2017), Husted et al. (2019), and Azqueta-Gavaldon et al. (2020). Specifically, my model reproduces these effects.

In addition, my work contributes to the literature on policy uncertainty in general equilibrium; see e.g., Davig (2004), Bi et al. (2013), Richter and Throckmorton (2015) and Davig and Foerster (2019) for fiscal policy uncertainty[3]; and Leeper and Zha (2003), Schorfheide (2005), Eusepi and Preston (2010) and Cogley et al. (2015) for monetary policy uncertainty. The closest paper is Bianchi and Melosi (2018), which also develops a general equilibrium model where the central bank deviates from active inflation stabilisation and households face uncertainty about the nature of these deviations. I differ from Bianchi and Melosi (2018) in at least three ways: (i) households in my model are uncertain about the current monetary policy regime, while agents in their model do not know whether the central bank is engaging in a short or a long lasting deviation from active inflation stabilisation. (ii) I focus on the effect of MPU shocks, while Bianchi and Melosi (2018) study the welfare consequences of limited information. (iii) I use a global solution method, capturing the endogenous non-linearities linked to the learning process.

Also, my paper relates to the stochastic volatility literature [e.g., Bloom (2009), Bloom et al. (2018) and Bachmann and Bayer (2013)]. This body of work mostly focuses on how uncertainty about productivity influences the business cycle, and has one main disadvantage: it is silent on the exact source of uncertainty. In my model, however, the origin of uncertainty is explicit: the monetary policy regime is unobserved. This clarity is essential to obtain structural interpretations.

[3]This literature on fiscal policy uncertainty finds that expectations about future policy affect agents' choices, thus generating deviations from the full information equilibrium. This insight is in line with my results concerning monetary policy uncertainty.

Besides, my model connects with the least squares learning literature (see Evans and Honkapohja (2009) and references therein). In this class of models, agents do not know the structure of the economy. Instead, they use historical data to infer the unknown components. In my work, agents do have full information about the structure of the economy, but are uncertain about the monetary regime. Hence, they filter a noisy signal using the correct model, and then rationally form expectations.

Lastly, my finding that spikes in MPU dampen investment recalls the seminal works by Bernanke (1983) and Abel and Eberly (1994). These authors stress the importance of delaying investment until economic uncertainty is resolved. My work, however, has a different emphasis. To be more precise, I present a general equilibrium model, where uncertainty stems from the monetary policy stance. Furthermore, I assess how this MPU impacts the transmission of monetary policy shocks.

The remainder of the paper is organised as follows. Section 2 introduces the model, and the different information sets. Section 3 presents the calibration strategy, the solution method and the accuracy of beliefs. Section 4 shows the results. Section 5 concludes.[4]

2. THE MODEL

I adopt a New Keynesian model with Rotemberg price-setting frictions to study the consequences of recurring unobserved changes in the central bank's reaction function. The first subsection describes the standard part of the model. Though this description can be found in many sources [e.g. Christiano et al. (2010) and Gavin et al. (2015)], I include it to make my presentation self-contained. The second subsection explains the use of Bayesian learning to endogenies monetary policy uncertainty.

2.1. **Standard part of the model.**

2.1.1. *Households.* A representative household chooses $\{c_t, h_t, b_t, i_t, k_t\}_{t=1}^{\infty}$ to maximise expected lifetime utility given by:

$$\mathrm{E}_t \sum_{t=1}^{\infty} \beta^{t-1} \left[\frac{c_t^{1-\omega}}{1-\omega} - \chi \frac{h_t^{1+\eta}}{1+\eta} \right],$$

where $\beta \in [0,1[$ is the discount factor, $1/\omega$ is the intertemporal elasticity of substitution, $\chi > 0$ is a scale parameter, $1/\eta$ is the Frisch elasticity of labor supply, c_t is real consumption, h_t is labor hours, b_t is the quantity of nominal bonds purchased by the household, i_t is real investment, k_t is

[4] A key issue to get straight from the outset: my model is too stylised for a full quantitative analysis, so I shall focus instead on its qualitative implications and insights.

the capital stock and E_t is the expectation operator conditional on information available at time t. These choices are constrained by:

$$c_t + i_t + \Gamma(i_t/k_{t-1})k_{t-1} + b_{t+1} = w_t h_t + r_t^k k_{t-1} + \Pi_t + \frac{r_{t-1}b_t}{1+\pi_t},$$

$$k_t = (1-\delta)k_{t-1} + i_t. \tag{1}$$

Here π_t is the rate of inflation from $t-1$ to t, w_t is the real wage, r_t^k is the capital rental rate, Π_t denotes dividends received from ownership of retail firms, r_t denotes the one-period gross nominal rate of interest on a bond purchased in period t, and δ is the depreciation rate of capital. The model features adjustment costs: $\Gamma(i_t/k_{t-1})$ is a positive, increasing and convex function measuring the cost of adjusting the capital stock. Specifically, I assume that $\Gamma(i_t/k_{t-1}) = \gamma(i_t/k_{t-1} - \delta)^2/2$, where γ governs how sensitive investment is to the price of capital.

Solving the household's problem yields standard optimality conditions:

$$w_t = \chi c_t h_t^{\eta}, \tag{2}$$

$$1 = \beta \mathrm{E}_t \left[\left(\frac{c_t}{c_{t+1}} \right)^{\omega} \frac{r_t}{1+\pi_{t+1}} \right], \tag{3}$$

$$q_t = 1 + \gamma \left(\frac{i_t}{k_{t-1}} - \delta \right), \tag{4}$$

$$q_t = \beta \mathrm{E}_t \left\{ \left(\frac{c_t}{c_{t+1}} \right)^{\omega} \left[r_{t+1}^k - \frac{\gamma}{2} \left(\frac{i_{t+1}}{k_t} - \delta \right)^2 + \gamma \left(\frac{i_{t+1}}{k_t} - \delta \right) \frac{i_{t+1}}{k_t} + q_{t+1}(1-\delta) \right] \right\}, \tag{5}$$

Eq. (2) pins down the supply of labor by equating the marginal cost of work, in consumption units, with the marginal benefit, the real wage. Eq. (3), the standard Euler equation for bonds, balances the marginal cost of purchasing a bond with the corresponding expected benefit. Eq. (4) defines an investment demand function relating net investment (i.e. $i_t - \delta k_{t-1}$) to Tobin's q. More precisely, net investment will be positive if and only if Tobin's q exceeds unity. Lastly, eq. (5) defines Tobin's q as the present discounted value of the marginal profits from an extra unit of capital, measured in terms of future marginal products and reductions in installation costs.

2.1.2. *Firms.* The production sector consists of intermediate goods firms producing a continuum of differentiated inputs under monopolistic competition, and a representative final goods firm.

The final goods firm produces gross output, Y_t, using a Dixit-Stiglitz aggregator:

$$Y_t = \left[\int_0^1 y_{i,t}^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}},$$

where $y_{i,t}$ is the quantity of differentiated good sold by intermediate firm $i \in [0,1]$, and $\epsilon > 1$ is the elasticity of substitution between the different inputs. Since the final goods producer maximises profits in a perfectly competitive environment, the demand function facing each intermediate firm i is:

$$y_{i,t} = \left[\frac{p_{i,t}}{P_t}\right]^{-\epsilon} Y_t,$$

where $p_{i,t}$ is firm i's sale price and P_t is the price of the composite good, which is defined by:

$$P_t = \left[\int_0^1 p_{i,t}^{1-\epsilon} di\right]^{\frac{1}{1-\epsilon}}.$$

Each intermediate firm i operates an identical technology function given by $y_t^i = h_{i,t}^{1-\alpha} k_{i,t-1}^{\alpha}$, where $\alpha \in [0,1]$. Variables $h_{i,t}$ and $k_{i,t-1}$ are the levels of employment and capital used by firm i. In this context, firm i's real marginal cost, $mc_{i,t}$, evolves according to:

$$mc_{i,t} = \left[\frac{w_t}{1-\alpha}\right]^{1-\alpha} \left[\frac{r_t^k}{\alpha}\right]^{\alpha}. \tag{6}$$

In equilibrium, firm i balances the relative price of its inputs with the corresponding ratio of marginal productivities:

$$\frac{w_t}{r_t^k} = \frac{1-\alpha}{\alpha} \frac{k_{i,t-1}}{h_{i,t}}. \tag{7}$$

Every firm relies on the same combination of inputs.

The model features price-setting frictions along the lines proposed by Rotemberg (1982). Each firm faces a real cost to adjust its price, which captures the negative effects that price changes have on customer-firm relationships. This menu cost is given by:

$$\frac{\psi}{2}\left[\frac{p_{i,t}}{p_{i,t-1}} - 1\right]^2 y_t,$$

where $\psi > 0$ determines the size of the adjustment cost. Each firm i chooses its price to maximise expected real profits, given by:

$$\left[\frac{p_{i,t}}{P_t}\right] y_{i,t} - \left[\frac{mc_{i,t}}{P_t}\right] y_{i,t} - \frac{\psi}{2}\left[\frac{p_{i,t}}{p_{i,t-1}} - 1\right]^2 y_t$$
$$+ \beta\Lambda_{t,t+1}\left\{\left[\frac{p_{i,t+1}}{P_{t+1}}\right] y_{i,t+1} - \left[\frac{mc_{i,t+1}}{P_{t+1}}\right] y_{i,t+1} - \frac{\psi}{2}\left[\frac{p_{i,t+1}}{p_{i,t}} - 1\right]^2 y_{t+1}\right\},$$

where $\Lambda_{t,t+1}$ is the household intertemporal marginal rate of substitution. Differentiating the objective function with respect to $p_{i,t}$ and imposing symmetry among all retail goods yields a standard New Keynesian Phillips Curve:

$$\psi\pi_t(1+\pi_t) = 1 + \epsilon\left[mc_t - 1\right] + \beta \mathrm{E}_t\left[\left(\frac{c_t}{c_{t+1}}\right)^{\omega}\psi\pi_{t+1}(1+\pi_{t+1})\frac{y_{t+1}}{y_t}\right]. \tag{8}$$

As usual, eq. (8) links current inflation to next period's expected inflation and output gap. If there were no price-setting frictions, the real marginal cost of producing a unit of output would be constant and equal to $\frac{\epsilon-1}{\epsilon}$.

2.1.3. *Aggregate resources and private sector equilibrium conditions.* Rotemberg pricing ensures that labor and capital inputs are equally distributed among the various intermediate goods firms. Aggregate output is thus given by:

$$y_t = h_t^{1-\alpha}k_{t-1}^{\alpha}, \tag{9}$$

where $h_t = \sum_i h_{i,t}$ and $k_{t-1} = \sum_i k_{i,t-1}$. The economy-wide resource constraint is:

$$\left[1 - \frac{\psi}{2}\pi_t^2\right] y_t = c_t + i_t + \Gamma(i_t/k_{t-1})k_{t-1}, \tag{10}$$

where the left hand side represents aggregate production net of price adjustment costs, and the right hand side corresponds to aggregate use of resources.

The private sector equilibrium conditions are equations (1) through (10). These are 10 equations in the following 11 endogenous unknowns:

$$c_t, h_t, i_t, k_t, w_t, r_t^k, q_t, mc_t, y_t, r_t, \pi_t.$$

As it stands, the system is undetermined, because monetary policy is not yet defined. I turn to this in the following subsection.

2.2. **Monetary policy and expectations formation.**

2.2.1. *The Taylor rule.* The monetary authority sets the gross nominal interest rate, r_t, according to the truncated Taylor rule:

$$r_t = \max\left[1, \rho r_{t-1} + (1-\rho)(\bar{r} + \phi_t\pi_{t-1}) + u_t\right], \tag{11}$$

where $\bar{r}$ is the steady-state gross nominal rate, $\rho \in [0,1[$ is a smoothing parameter, the shock, $u_t \sim \mathcal{N}(0,\sigma_u^2)$, is a proxy for discretionary monetary policy, and ϕ_t is the central bank's response to inflation. The latter evolves according to a 2-state Markov chain with transition matrix, P. For row i and column j of P, element $p_{i,j} = P(\phi_t = \phi_j | \phi_{t-1} = \phi_i)$ for $i, j \in \{L, H\}$, where $0 \leq p_{i,j} < 1$ and $\sum_{j=1}^{2} p_{i,j} = 1$ for all i.

The Markov process models changes in the central bank's reaction function. Under Regime 1, henceforth the *passive regime*, the central bank de-emphasises inflation stabilisation by violating the Taylor principle, i.e., $\phi_L \leq 1$. In contrast, under Regime 2, henceforth the *active regime*, the central bank emphasises inflation stabilisation by respecting the Taylor principle, i.e., $\phi_H > 1$.[5] This regime changing device thus captures the systematic deviations from inflation stabilisation observed in the US in the post-war period.[6]

Just to be clear: movements in ϕ_t represent persistent changes in the central bank's inflationary stance, whereas u_t represents transitory deviations from the systematic component of the policy rule.[7] My regime switching framework bears a close resemblance to the ones used by Leeper and Zha (2003), Schorfheide (2005), and Bianchi and Melosi (2018). Hence, it does not explain why monetary policy shifts over time. Instead, it simply assumes that there are two regimes subject to constant transition probabilities.

2.2.2. *Expectations formation.* Bayesian learning is introduced as in Richter and Throckmorton (2015), who builds on Hamilton (1989). As mentioned previously, households and firms do not observe the central bank's inflation parameter, ϕ_t, for it is obscured by the discretionary shock, u_t. Instead, they make an inference about its value based on observed realisations of $\{r_t, \pi_{t-1}\}$.

The expectations formation process summarizes the model as:

$$E_t \left[F(\mathbf{x}_{t+1}, \mathbf{x}_t) \right] = 0,$$

[5]The choice of words *passive regime* and *active regime* is taken from Bianchi and Melosi (2018). The passive regime temporarily weakens the central bank's commitment to keeping inflation -and hence, inflation expectations- under control. Thus, this regime changing device resembles the constrained discretion strategy [Bernanke and Mishkin (1997)], which allows the central bank some flexibility to temporarily de-emphasise inflation stabilisation. To be sure, this regime changing device is not the only way of modelling temporary deviations from inflation targeting. For example, modelling the shock u_t as an autoregressive process would also give rise to persistent deviations from inflation stabilisation. The reason why I favour the Markov chain approach is twofold. First, it is closer to previous studies [e.g. Leeper and Zha (2003) and Bianchi and Melosi (2018)]. Second, it allows me to easily consider limited information via the Hamilton filter.

[6]Please refer to Taylor (2012), Bianchi (2013) and Mishkin (2018) for empirical and anecdotal evidence on how the Federal Reserve has regularly de-emphasised inflation stabilisation for both short and long periods of time.

[7]In reality, these intra-regime shifts in policy are either deliberate or due to implementation error [Schorfheide (2005)].

where F is a vector value function representing the economic environment and the first order conditions of the model, and **x** is a vector containing the observed variables in the model:

$$\mathbf{x} = \begin{cases} c_t, h_t, i_t, k_t, w_t, r_t^k, q_t, mc_t, y_t, r_t, \pi_t, \phi_t, & \text{under Rational Expectations,} \\ c_t, h_t, i_t, k_t, w_t, r_t^k, q_t, mc_t, y_t, r_t, \pi_t, \lambda_t, & \text{under Learning.} \end{cases}$$

I introduce the Rational Expectations equilibrium of the model, in which ϕ_t is perfectly observed, to have a benchmark. In the Learning model, λ_t is a vector of conditional probabilities that $\phi_t = \phi_i$ for $i \in \{L, H\}$, which agents update to form expectations. Specifically, agents use the nonlinear filter in Hamilton (1989) to compute the conditional probabilities:

$$\lambda_t^i = \mathrm{P}(\phi_t = \phi_i | \Omega_t, \Theta).$$

Here $\Omega_t = \{r_t, \cdots, r_0, \pi_{t-1}, \cdots, \pi_0\}$ is the history of the nominal interest rate and inflation, and $\Theta = \{P, \phi_1, \phi_2, \sigma^2\}$ is the vector of relevant parameters. With each additional observation of $\{r_t, \pi_{t-1}\}$, agents update their inference about ϕ_t.[8]

Under Rational Expectations, ϕ_t is known. Given $\phi_t = \phi_i$, forming expectations is simple:

$$\mathrm{E}_t\left[\mathrm{F}(\mathbf{x}_{t+1}, \mathbf{x}_t)\right] = \sum_{j=L}^{H} p_{i,j} \int_{-\infty}^{\infty} \mathrm{F}(\mathbf{x}_{t+1}, \mathbf{x}_t)\vartheta(u_{t+1})du_{t+1} = 0, \tag{12}$$

where $\vartheta(.)$ is the normal probability density function. In words, the sum operates across the two monetary policy regimes, while the integral operates across realisations of the discretionary shock.

Under learning ϕ_t is unknown, creating a signal extraction problem. Given λ_t, the expectation formation process is:

$$\mathrm{E}_t\left[\mathrm{F}(\mathbf{x}_{t+1}, \mathbf{x}_t)\right] = \sum_{i=L}^{H} \lambda_t^i \sum_{j=L}^{H} p_{i,j} \int_{-\infty}^{\infty} \mathrm{F}(\mathbf{x}_{t+1}, \mathbf{x}_t)\vartheta(u_{t+1})du_{t+1} = 0. \tag{13}$$

According to the first sum operator, agents account for monetary policy uncertainty by weighting each realisation of ϕ by its conditional likelihood, λ_t^i.

Comparing eq. (12) and eq. (13) provides three useful insights:

[8] Subsection 3.2 explains how to compute λ_t.

(i) The unknown policy regime generates forecast errors with respect to the Rational Expectations benchmark. These errors distort current decision-making via the three intertemporal optimality conditions.

(ii) As agents' beliefs become increasingly accurate, the learning economy converges to the Rational Expectations benchmark. In other words, deviations from Rational Expectations are a function of how far agents' beliefs are from the truth.

(iii) Learning smooths expectations, because it is unlikely that $\lambda_t^i = 1$ for any i. That is, agents are aware that they do not known ϕ_t, so they put some weight on choices they would make in either monetary regime. As will become clear, this smoothing of expectations makes the learning economy less sensitive to monetary policy.

2.2.3. *Measuring MPU.* This section concludes by presenting the model-based MPU index I use throughout the paper. Each period, the Hamilton filter provides two conditional probabilities: $\lambda_t^i = \mathrm{P}(\phi_t = \phi_i | \Omega_t, \Theta)$ for $i \in \{L, H\}$. Periods with $\lambda_t^i = 0.5$ feature absolute uncertainty: agents consider the two possible policy regimes equally likely. In contrast, periods with $\lambda_t^i = 1$ for either i feature absolute certainty: agents are convinced they know ϕ_t. Following Richter and Throckmorton (2015), I measure uncertainty as:

$$\zeta_t = \frac{\sqrt{0.5} - \sqrt{\sum_{i=L}^{H} (\lambda_t^i - 0.5)^2}}{\sqrt{0.5}}, \tag{14}$$

which ranges from 0 (absolute certainty) to 1 (absolute uncertainty). Importantly, ζ_t reflects how confident agents are about their inferences, not the accuracy of these inferences.

One final remark. Eq. (11) mean that both the nominal interest rate, r_t, and the vector of conditional probabilities, λ_t, are part of the state space. Consequently, the index ζ_t does not respond to current realisations of control variables such as output and inflation, because it is predetermined.

3. Calibration, Solution Technique and Beliefs

3.1. **Calibration.** Table 1 lists the values assigned to the model parameters. I choose conventional values for the taste and technology parameters. The intertemporal elasticity of substitution, $1/\omega$, is set to 1. The quarterly discount factor, β, is set to 0.99, corresponding to a 4% annual real interest rate. The Frisch elasticity of labor supply, $1/\eta$, is set to 1, which is in line with Christiano et al.

(2010). The leisure-preference parameter, χ, is set to 6.88, implying that one third of time is spent working in the deterministic steady state.

Also following convention, the elasticity of substitution between intermediate goods, ϵ, is set to 10, while the price adjustment cost parameter, ψ, is set to 105. This value mimics a Calvo price setting specification where the average duration between price changes is four quarters. Capital's share of output, α, is set to 0.33, and the depreciation rate, δ, is set to 0.025. Finally, the capital adjustment cost, γ, is set to 6, which follows Erceg and Levin (2003) and Gavin et al. (2015).

Concerning monetary policy, the smoothing parameter, ρ, is set to 0.6, which is consistent with Clarida et al. (2000). The four parameters governing the 2-state Markov chain for ϕ_t, $\{\phi_L, \phi_H, p_{L,L}, p_{H,H}\}$, are chosen to be consistent with the following three outcomes. First, deviations from the Taylor principle are short, yet pronounced.[9] Second, the ergodic mean of ϕ is in the neighbourhood of 1.5, which is the standard value found in the literature. Third, the calibration is in line with Bianchi and Melosi (2018). The resulting values are: $\phi_L = 0.2$, $\phi_H = 2$, $p_{L,L} = 0.7$, and $p_{H,H} = 0.85$.[10]

I finally parameterise the standard deviation of the discretionary shock, σ_u. This is an important parameter. When σ_u is very small, u_t has no significant effects on the policy rate; any additional observation of $\{r_t, \pi_{t-1}\}$ fully reveals ϕ_t. When σ_u is very large however, u_t has huge effects on the policy rate; any additional observation of $\{r_t, \pi_{t-1}\}$ adds little information about ϕ_t. I am interested in neither of these cases. Instead, I want limited information to play a meaningful role, so I move away from these extremes and set $\sigma_u = 0.001$. Though this choice is somewhat arbitrary, it does ensure that my numerical simulations satisfy two requirements: (i) beliefs are often close to reality; and (ii) movements in the MPU index, ζ, roughly resemble those observed in US data. Subsection 3.3 offers more details on these two conditions, while Subsection 4.1.3 tests alternative calibrations for σ_u.

3.2. **Model solution.** I solve the model using a textbook projection method. I want this paper to be self-contained, so let me only sketch the basic features of this solution method.[11]

[9] This is in line with Bernanke and Mishkin (1997).

[10] As long as $\phi_L < 1$, $\phi_H > 1$, and $p_{L,L} < p_{H,H}$ (all of them reasonable assumptions), the main qualitative conclusions are unaffected.

[11] For a detailed review on this technique, see Fernandez-Villaverde et al. (2016).

TABLE 1. Parameter values.

	Symbol	Value		Symbol	Value
Taste and technology			**Monetary policy**		
Discount factor	β	0.99	Smoothness parameter	ρ	0.60
Inv. Inter. elasticity	ω	1.00	Passive regime	ϕ_L	0.20
Inverse Frisch elasticity	η	1.00	Active regime	ϕ_H	2.00
Leisure parameter	χ	6.88	Trans. probability matrix	$p_{L,L}$	0.70
Elasticity of substitution	ϵ	10.0	Trans. probability matrix	$p_{H,H}$	0.85
Price adjustment	ψ	105	STD discretionary shock (%)	σ_u	0.10
Capital's share of output	α	0.33			
Capital depreciation (%)	δ	2.50			
Capital Adjustment cost	γ	6.00			

Any general equilibrium model can be expressed as a nonlinear first-order system of expectational difference equations:

$$\Gamma(E_t\mathbf{z}_{t+1}, \mathbf{z}_t, \mathbf{v}_{t+1}) = \mathbf{0}.$$

Here $\mathbf{z}$ is a $n\ x\ 1$ vector of stationary variables, $\mathbf{v}$ is a $m\ x\ 1$ vector of structural shocks, and $\mathbf{0}$ is a $n\ x\ 1$ vector of 0.

The solution I seek is of the form:

$$\mathbf{c}_t = d(\mathbf{s}_t),$$

$$\mathbf{s}_t = g(\mathbf{s}_{t-1}, \mathbf{v}_t).$$

The first set of equations, $d(\cdot)$, represents the optimal mapping from the state variables, $\mathbf{s}_t$, to the control variables, $\mathbf{c}_t$. This set of equations is known as policy functions. In turn, the second set of equations, $g(\cdot)$, describe the evolution of the state variables.

The set of policy functions is the solution of:

$$G(d(\mathbf{s}_t)) = 0, \tag{15}$$

where $G(\cdot)$ is an operator defined over function spaces.

Projection methods solve (15) in three steps. First, they specify a linear combination:

$$d^j(\mathbf{s}|\Xi) = \sum_{i=0}^{j} \xi_i \Psi_i(\mathbf{s})$$

of basis functions $\Psi_i(\mathbf{s})$ given coefficients $\Xi = \{\xi_0, ..., \xi_j\}$. Second, they define a residual function:

$$R(\mathbf{s}|\Xi) = G(d^j(\mathbf{s}|\Xi)).$$

Third, they select the values of Ξ that minimize this residual given some metric.

Hence, I first have to select a basis $\sum_{i=0}^{j} \xi_i \Psi_i(\mathbf{s})$. I then need to "project" $G(\cdot)$ against that basis to find Ξ. Different choices for these two inputs lead to slightly different projection methods. In this work, I opt for Chebyshev polynomials as basis, and I choose the vector of parameters Ξ to make the residual function $R(\mathbf{s}|\Xi)$ zero at a set of Chebyshev nodes. This approach is known as orthogonal collocation.[12]

I now turn to the computational aspect of expectations formation. Recall from Subsection 2.2.2 that under limited information agents observe r_t directly but can only make an inference about the value of ϕ_t based on what they see happening with $\{r_t, \pi_{t-1}\}$. This inference takes the form of the vector of conditional probabilities λ_t.

The Hamilton filter [Hamilton (1989)] performs such an inference iteratively for $t = 1, 2, 3, ...$ with step t taking as input λ_{t-1}. The crucial magnitudes agents use to perform iteration t are the probability densities under the two regimes:

$$\eta_{j,t} = f(r_t|\phi_t = \phi_j, \Omega_t, \Theta),$$

for $j = H, L$. Here Ω_t is the history of the nominal interest rate and inflation, and Θ is the vector of relevant parameters (see Subsection 2.2.2). The density functions $f(.)$ are specified below.

Given $\{\lambda_{t-1}, \eta_{j,t}\}$, the conditional density of the tth observation is given by:

$$h(r_t|\Omega_t, \Theta) = \sum_{i=1}^{2} \sum_{j=1}^{2} p_{i,j} \lambda_{t-1}^{i} \eta_{j,t}.$$

The desired output is:

$$\lambda_t^j = \frac{\sum_{i=1}^{2} p_{i,j} \lambda_{t-1}^{i} \eta_{j,t}}{h(r_t|\Omega_t, \Theta)}.$$

Crucially, the elements of η_t need not be Gaussian densities or belong to the same family of distributions. This property is essential, because the max operator in eq. (11) defines two completely different scenarios.

When $r_t > 0$ the elements of η_t are normal densities:

$$\eta_{j,t} = f(r_t|\phi_t = \phi_j, \Omega_t, \Theta) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(r_t - \rho r_{t-1} - (1-\rho)(\bar{r} + \phi_j \pi_{t-1}))^2}{2\sigma^2}},$$

[12] Appendix A details how I apply this general procedure to the model presented in Section 2.

for $j = H, L$. In these instances, the latest observation of the nominal interest rate and inflation contain relevant information about ϕ_t, which agents insert into the Hamilton filter.

However, when the economy is against the Zero Lower Bound ($r_t = 0$), agents face a degenerate distribution, which entails a probability mass function $\eta_{j,t} = 1$ for $j = 1, 2$. It is trivial to verify that under $r_t = 0$ agents update their beliefs according to:

$$\lambda_t^i = p_{i,i}\lambda_{t-1}^i + p_{i,j}\lambda_{t-1}^j.$$

Iterating the above expression forward shows that λ_t^i converges to the unconditional probability $P(\phi = \phi_i)$. This result is intuitive. At the Zero Lower Bound (ZLB), neither the nominal interest rate nor the inflation rate conveys useful information about the monetary policy regime. Agents' best inference thus dictates that ϕ_t converges to its ergodic distribution.

A point here deserves further comment. For simplicity, the model has a single source of exogenous fluctuations: monetary policy. Consequently, in the simulations shown in Section 4, movements in the inflation rate are not large enough to bring the economy to the ZLB. Yes, hitting the ZLB would require additional shocks that increase macroeconomic volatility. However, more shocks would also increase the number of state variables, and thus complicate the numerical solution of the model. I therefore leave the quantitative assessment of how MPU interacts with the ZLB for future research.

3.3. **Beliefs and uncertainty.** Agents' beliefs should not be arbitrary; people should not make systematic errors. Instead, theoretical models should feature a close match between beliefs and reality [Woodford (2013)].

The calibration discussed above ensures my model is consistent with this argument. To make my point clear, I assess the density function of the absolute-value norm between the true policy regime and private sector beliefs: $\theta_t = |\phi_t - (\lambda_t^L\phi_L + (1 - \lambda_t^L)\phi_H)|$, with $\theta_t \in [0, |\phi_H - \phi_L|]$, on a 100,000-period simulation of the model. As seen in Figure 1, most of the time, beliefs are close to the truth. In other words, on average, the signal extraction problem does not imply a substantial deviation from the Rational Expectation equilibrium of the model.

My model has monetary policy as the only source of exogenous fluctuations. As a result, it cannot possibly fit the data with a high degree of accuracy. Yet, movements in my MPU index, ζ_t, do resemble those observed in the US (see Table 2). As with the MPU indexes for the US developed by Baker et al. (2016) and Husted et al. (2019), in my model MPU is positively skewed

FIGURE 1. Accuracy of agents' beliefs: density function.

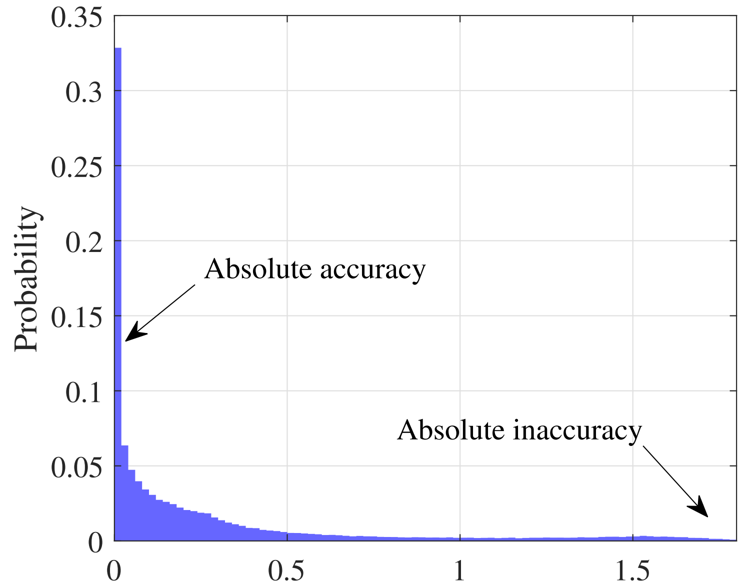

Notes. Density function of the absolute-value norm between the true policy regime and agents beliefs: $\theta_t = |\phi_t - (\lambda_t^L \phi_L + (1-\lambda_t^L)\phi_H)|$, with $\theta_t \in [0, |\phi_H - \phi_L|]$. Based on a 100,000-period stochastic simulation of the model.

TABLE 2. MPU index: model versus data.

	Model	US Data	
	ζ	Baker-Bloom-Davis	Husted-Rogers-Sun
Standard deviation to mean ratio	1.23	0.52	0.42
Skewness	1.41	1.25	1.92
Kurtosis	4.17	4.72	8.92

Notes. Results in the first column are based on a 100,000-period stochastic simulation of the model. The second and third columns refer to US quarterly data from 1985Q1 to 2021Q1. Baker-Bloom-Davis stands from the MPU index developed by Baker et al. (2016). Husted-Rogers-Sun stands for the MPU index developed by Husted et al. (2019).

and leptokurtic. That is, both in my model and in the data, MPU has a long right tail and features more outliers than the normal distribution. The model, however, generates too much volatility: the relative standard deviation is above 1 in the model, but well below 1 in the US data.

On the whole, the baseline calibration allows the learning model to be close to the Rational Expectation equilibrium, and to generate movements in MPU not dissimilar from those observed in the US data.

4. CENTRAL RESULTS

4.1. **The real effects of monetary policy uncertainty.** This section reproduces the contractionary effects of increases in MPU documented by Creal and Wu (2017) and Husted et al. (2019).

FIGURE 2. Impulse responses to a one-standard deviation surprise increase in MPU.

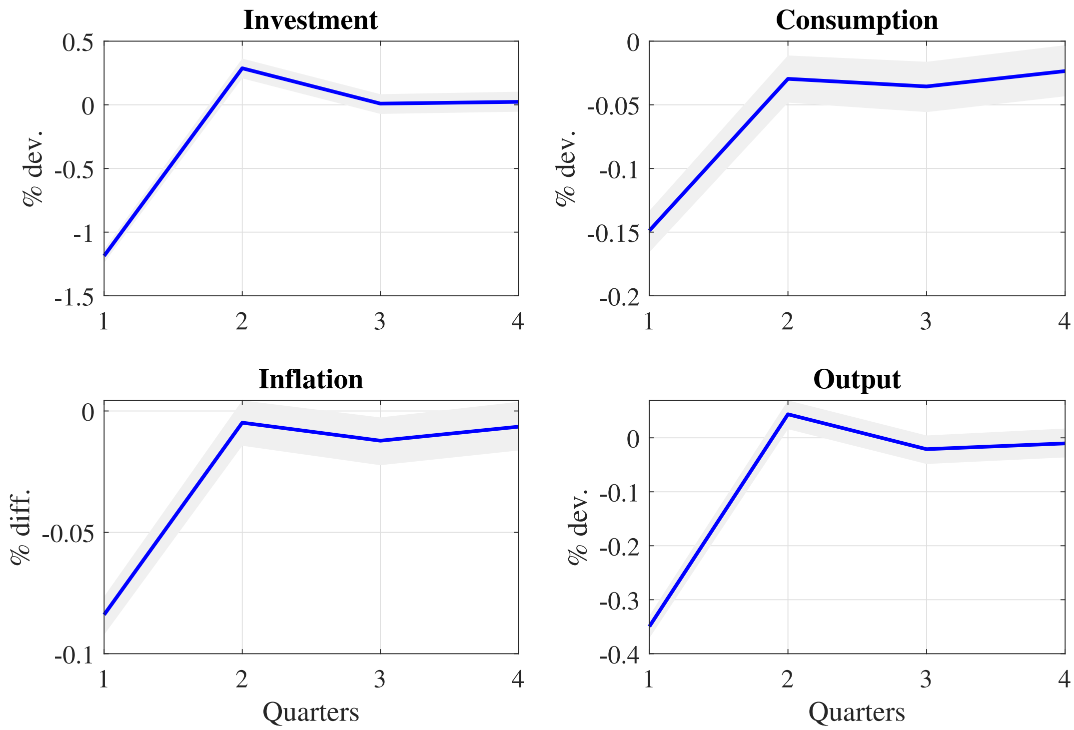

Notes. Impulse response computed via Jorda (2005) Local Projection methods. Shaded areas are the 95% confidence intervals computed using Monte Carlo methods.

4.1.1. *Understanding the mechanism.* I estimate the predictive relation between the MPU index and investment, consumption, inflation and output using Local Projection methods [Jorda (2005)]. The empirical specification is:

$$\hat{Z}_{t+h} = \mu_h + \beta_h \zeta_t + \gamma_h X_t + \epsilon_{t+h}. \tag{16}$$

Here $\hat{Z}$ is either investment, consumption or output (all in percentage deviations from their steady state values), or inflation (in percentage differences from its steady state value). Also, ζ_t is the MPU index, X_t is a set of control variables[13], ϵ_{t+h} is the projection residual, and μ_h, β_h, γ_h are the projection coefficients. The Local Projection impulse response function of $\hat{Z}$ with respect to ζ is given by $\{\beta_h\}_{h\geq 0}$.[14]

Figure 2 reports the Local Projection impulse responses to a one-standard-deviation increase in ζ_t. The message is clear: MPU lowers aggregate demand, thereby hindering both output and

[13] I include the current stock of capital, k_{t-1}, the current nominal interest rate, r_t, and agents' beliefs, λ_t. These three variables constitute the state space of the model. The results are robust to alternative controls.

[14] Since the dependent variable is in % deviations from steady state (% differences for inflation), the coefficient of MPU implies that variable Z deviates from its long run equilibrium by β_h% for every extra unit of MPU.

inflation.[15] These effects are short-lived: one quarter after the shock, all variables are already close to their steady state values.

The rationale behind the negative effects of higher MPU is as follows. Beliefs about the monetary regime have an impact on aggregate demand via their effects on the next period's expected nominal interest rate. Specifically, whenever current inflation is below (above) the central bank target, aggregate demand benefits from agents believing that monetary policy is active (passive). This occurs because with inflation below (above) target, an active (passive) Taylor principle signals a protracted period of low nominal interest rates.

Uncertainty kills this signalling mechanism. That is, uncertainty blurs the likely future path of interest rates, thereby smoothing changes in agent behaviour.

Let me stress what the previous paragraph says and does not say. It says that spikes in MPU make consumption and investment a weighted average of their hypothetical values if agents had known the monetary regime with certainty. It does not say, however, that spikes in MPU *always* depress the economy. For example, if inflation is below (above) target, then agents will consume and invest more when MPU is high than when they believe that monetary policy is passive (active).

But, why does Figure 2 then report clear contractionary effects from increases in MPU? Because the model is non-linear: the positive and negative effects of MPU do not cancel out in the long run. The net effect is not zero, but negative. Below, I shall perform a set of comparative statics to assess the importance of convex adjustment costs and concave preferences in obtaining these contractionary effects.

Still in Figure 2, note that spikes in MPU trigger, first and foremost, a strong decline in investment (roughly 1.2%). Indeed, by obscuring the likely future path of interest rates, uncertainty blurs the expected marginal product of capital[16], making it rational, on average, to defer investment. This outcome recalls the seminal work by Bernanke (1983) and Abel and Eberly (1994), emphasising the importance of waiting until economic uncertainty is resolved. Husted et al. (2019) also stress that uncertainty about monetary policy hinders investment.

[15]This simultaneous fall in investment, consumption, and output following a spike in uncertainty is in agreement with Basu and Bundick (2017), who stress the importance of sticky prices to produce this co-movement.

[16]The nominal interest rate and the rental rate of capital are tightly linked: their correlation is roughly -0.95 in a long run simulation. By construction, the partial derivative of the marginal product of capital with respect to hours worked is positive. Therefore, a decline in the nominal interest rate raises the rental rate of capital by boosting aggregate demand and labor input.

Consumption falls too, for MPU distorts intertemporal saving decisions via the Euler equation for bonds. However, the decline is relatively mild, as households attempt to smooth consumption. Lastly, to accommodate the ensuing fall in aggregate demand, firms cut labor input, thus lowering output and inflation.[17]

As mentioned earlier, the negative effects of increases in MPU are short-lived. The reason is twofold. First, spikes in MPU are themselves short-lived, for agents quickly learn about the monetary policy reaction function.[18] Indeed, the auto-correlation coefficient of ζ is close to -0.1. Second, the basic New Keynesian model has weak endogenous propagation mechanisms. Nonetheless, investment does rebound slightly, as the resolution of uncertainty leads households to bring the capital stock back to its steady state level. As for output, the lack of a strong rebound does not contradict Husted et al. (2019), who find that industrial production has a hump-shaped response to MPU shocks.

4.1.2. *Sensitivity analysis.* The previous subsection conveyed one key message: on average, spikes in MPU dampen aggregate demand, and thus lower output. I now analyse how changes in crucial parameters affect this key finding. My goal is twofold: to show its robustness, and to shed light on the forces driving it.

To be more precise, I use eq. (16) to compute the impact response to a one-standard deviation increase in MPU[19] under different calibrations of the capital adjustment cost parameter, the price adjustment cost parameter, and the elasticity of intertemporal substitution parameter. Table 3 reports the results.

Panel A shows that for larger values of the capital adjustment cost parameter, γ, investment falls less after a spike in MPU. The logic is straightforward. Larger γ means a higher cost of deviating from steady state investment. Said differently, larger values of γ reduce the sensitivity of investment to changes in both current and expected prices.[20] It follows that the effects of spikes in MPU on output decrease with the size of γ.

The main point from Panel B is clear too: a low price adjustment cost ($\psi = 70$) makes the economy less responsive to changes in MPU. This is intuitive, for in the absence of price rigidities

[17]The employment response to a one-standard deviation surprise increase in MPU (not shown) is similar to that of output. This occurs because the stock of capital is pre-determined.

[18]This logic is consistent with the small deviations from the Rational Expectations equilibrium discussed in Subsection 3.3.

[19]Recall that β_0 in eq. (16) measures the impact response of variable $\hat{Z}$ to a unit increase in MPU.

[20]The right hand side of eq. (5) shows that capital adjustment costs lessen the link between interest rates and the marginal profit of accumulating an extra unit of capital.

TABLE 3. Impact response to a one-standard deviation surprise increase in MPU.

	Capital adjustment cost		
	$\gamma = 2$	$\gamma = 6$ (baseline)	$\gamma = 10$
Panel A			
Investment	-1.70	-1.19	-0.14
Consumption	-0.12	-0.15	-0.04
Inflation	-0.10	-0.08	-0.02
Output	-0.43	-0.34	-0.06
	Price adjustment cost		
	$\psi = 70$	$\psi = 105$ (baseline)	$\psi = 140$
Panel B			
Investment	-0.49	-1.19	-1.19
Consumption	-0.07	-0.15	-0.15
Inflation	-0.02	-0.08	-0.07
Output	-0.16	-0.34	-0.35
	Inverse elasticity of intertemporal substitution		
	$\omega = 1$ (baseline)	$\omega = 3$	$\omega = 8$
Panel C			
Investment	-1.19	-0.92	-0.48
Consumption	-0.15	-0.05	0.00
Inflation	-0.08	-0.08	0.02
Output	-0.34	-0.23	-0.10

Notes. Impact response to a one-standard deviation surprise increase in MPU, computed using eq. (16). Investment, consumption and output are in percentage deviations from their steady state values. Inflation is in percentage differences from its steady state value. The impact responses under the baseline calibration are the same than in Figure 2.

the model would feature money neutrality. That is, if ψ were 0, real variables would not respond to monetary shocks, and the classical dichotomy would hold.

Lastly, in Panel C I vary the elasticity of intertemporal substitution, which is also the risk aversion parameter. Quite naturally, as ω goes up, consumption responds less to changes in the expected interest rate. In the extreme when $\omega \to \infty$, households would be infinitely unwilling to substitute consumption over time, making consumption independent of the monetary policy regime. It then comes as no surprise that as the utility function becomes more concave, increases in MPU have smaller overall effects.

In sum, this subsection has established that the negative effects of MPU do not hinge on a narrow set of parameter values. Below I explore one more dimension before turning to the effect of MPU on the transmission of monetary policy.

4.1.3. *An additional remark.* As stated in Subsection 3.1, the standard deviation of the transitory shock, σ_u, is a crucial parameter. Let σ_u approach 0, and the evolution of $\{r_t, \pi_{t-1}\}$ will convey perfect information about the monetary policy regime. Agents' inferences will be utterly precise,

TABLE 4. Dynamic implications of σ_u.

	Mean	Correlation with ζ_t			
σ_u (%)	θ_t	Y_t	i_t	c_t	π_t
0.005	0.00	0.00	0.00	0.00	0.00
0.070	0.19	-0.07	-0.12	0.05	0.00
0.100	0.27	-0.54	-0.77	-0.12	-0.20
0.130	0.36	-0.17	-0.24	0.01	-0.05
0.500	0.72	-0.01	-0.01	-0.01	-0.02

Notes. Based on a 100,000-period stochastic simulation of the model. The second column reports the mean of the inference error, θ. Columns 3 to 6 report the correlation between the MPU index, ζ_t, and output, investment, consumption, and inflation, respectively.

thus making the signal extraction problem trivial. The link between MPU and economic activity will therefore disappear.

In contrast, let σ_u go to ∞, and the monetary regime will become irrelevant, for r_t will only be driven by the discretionary shock, u_t. Furthermore, $\{r_t, \pi_{t-1}\}$ will no longer convey useful information about the monetary regime. Agents' beliefs will thus converge to their ergodic distribution: $\lambda_t^i \to \mathrm{P}(\phi_t = \phi_i)$ for $i \in \{L, H\}$. In this case, the relevance of MPU will also vanish.

The link between MPU and economic activity therefore occurs for values of σ_u for which the signal extraction problem is neither trivial nor impossible. Table 4 provides a numerical illustration of this argument. The statistics are based on 100,000-period simulations of the model.

Not surprisingly, the mean of the inference error, $\bar{\theta}$, increases monotonically with σ_u, as the signal extraction problem becomes increasingly challenging. More importantly, low (high) values of σ_u make learning too easy (difficult), thus reducing the relevance of MPU. That the correlation between macroeconomic variables and σ_u is hump-shaped follows.

4.2. **Learning and the transmission mechanism.** This second section evaluates the impact of learning on the monetary policy transmission mechanism. The key insight is clear: learning makes the economy less sensitive to the central bank reaction function. Spikes in MPU intensify this link, for agents are forced to base their actions on the choices they would make conditional on *both* monetary regimes.

4.2.1. *Correlations.* I first compare the links between the exogenous policy processes (ϕ_t and u_t) and output and inflation. Table 5 reports the contemporaneous correlations in a 100,000-period simulation of the model under learning and under Rational Expectations.

TABLE 5. Correlation coefficients.

	Correlation with ϕ_t		Correlation with u_t	
	Learning	Rational Expectations	Learning	Rational Expectations
Y_t	0.59	0.81	-0.68	-0.45
π_t	0.77	0.87	-0.54	-0.38

Notes. Based on a 100,000-period stochastic simulation of the model. Learning stands for the model with limited information. Rational Expectations stands for the model with perfect information.

The table conveys two messages. First, learning renders prices and quantities *less* sensitive to the central bank inflation response. The correlations with ϕ_t are significantly closer to 0 in the learning economy. The intuition is straightforward: since agents know that their beliefs can be inaccurate, they put equal weight in expectation on the choices they would make in either regime (see Subsection 2.2.2). Their behaviour is thus tempered by uncertainty, making them less responsive to the monetary regime.

Second, learning makes the economy *more* sensitive to discretionary shocks. In absolute value, correlations with u_t are larger in the learning economy. The logic is as follows. Under learning, transitory shocks affect agents' beliefs, λ_t. As a result, they have stronger effects on the expected path of future interest rates than in the Rational Expectations model (where agents perfectly observe the transitory disturbances). These stronger effects, combined with the three inter-temporal optimality conditions, account for the amplified responses to transitory shocks.

4.2.2. *Growth rates and regime changes.* This final subsection digs deeper into the effects of uncertainty on the responses to regime changes. I first identify all quarters featuring a regime change in the simulation. Then I compute the absolute value of the growth rate of output and inflation in those quarters, and compare them with their unconditional means. I focus on the absolute value of the growth rates, because I am interested in the magnitude of the changes, not their sign. Formally, I first identify all periods t^* in which $\phi_{t^*} = \phi_i$ and $\phi_{t^*-1} = \phi_j$ for $i \neq j$, where $i, j \in \{L, H\}$. Next, I compute:

$$\eta_x = \frac{\text{mean}\left[|\log(x_{t^*}) - \log(x_{t^*-1})|\right]}{\text{mean}\left[|\log(x_t) - \log(x_{t-1})|\right]},$$

where $x = \{Y, \pi\}$. In words, the numerator is the mean of the absolute value of the growth rate of x conditional on a regime change. The denominator is the mean of absolute value of the growth rate of x in the entire 100,000-period simulation.

Columns 2 and 3 in Table 6 support previous findings. The Rational Expectations economy reacts much more strongly to changes in ϕ_t. For example, conditional on a regime change, the

TABLE 6. Growth rates and regime changes.

	Rational Expectations	Learning		
		Unconditional	Conditional on	
			High ζ_t	Low ζ_t
Output, η_y	2.45	1.49	1.31	1.69
Inflation, η_π	2.74	1.84	1.35	2.40

Notes. Based on a 100,000-period stochastic simulation of the model. Learning stands for the model with limited information. Rational Expectations stands for the model with perfect information. In column 4 (5), the numerator of η_x is conditioned on ζ_t being higher (lower) than its 0.8 (0.2) percentile.

absolute growth rate of output is 2.45 times larger than its mean. In the learning model, that number is 1.49. Inflation features similar dynamics.[21]

Because uncertainty tempers agents' behaviour, regime changes conditional on high MPU must feature lower η_x statistics. Column 4 confirms this intuition. When the numerator of η_x is conditioned on ζ_t being higher than its 0.8 percentile, the resulting statistics are far below those in column 3. In contrast, when MPU is low (ζ_t lower than its 0.2 percentile), responses in the learning economy are closer to those in the Rational Expectations economy (see columns 2 and 5).

5. Conclusion

I present a New Keynesian model in which the central bank response to inflation varies over time. Agents do not observe the current monetary policy regime, but learn about it using Bayes theorem. I solve the model using a textbook projection method to account for the endogenous non-linearity introduced by imperfect information.

The model reproduces the contractionary effects of monetary policy uncertainty (MPU) shocks documented by Husted et al. (2019). The mechanism is straightforward: spikes in MPU obscure the likely future path of interest rates and, hence, of capital returns. This unpredictability makes it rational to defer investment and consumption until uncertainty is resolved.

I also explore how MPU affects the monetary policy transmission mechanism. The key insight is intuitive: by softening the link between fundamentals and equilibrium prices and allocations, learning renders the economy less responsive to the monetary policy inflation response.

Further research is needed in the following areas: (i) the implications of uncertainty for the effectiveness of unconventional monetary measures, and (ii) the implications of uncertainty for

[21] Bloom et al. (2018) also found that the presence of uncertainty significantly lowers the effects of policy changes.

the rules-versus-discretion debate. The study of effectiveness could benefit from a positive analysis using a framework similar to the one presented here. On the other hand, the rules-versus-discretion debate may require a normative analysis using a framework with endogenous regime changes.

References

Abel, A. B. and Eberly, J. C. (1994). A Unified Model of Investment under Uncertainty. *American Economic Review*, 84(5):1369–1384.

Andreasen, M. M., Fernández-Villaverde, J., and Rubio-Ramírez, J. F. (2018). The Pruned State-Space System for Non-Linear DSGE Models: Theory and Empirical Applications. *Review of Economic Studies*, 85(1):1–49.

Azqueta-Gavaldon, A., Hirschuhl, D., Onorante, L., and Saiz, L. (2020). Economic policy uncertainty in the euro area: an unsupervised machine learning approach. Working Paper Series 2359, European Central Bank.

Bachmann, R. and Bayer, C. (2013). Wait-and-See business cycles? *Journal of Monetary Economics*, 60(6):704–719.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.

Basu, S. and Bundick, B. (2017). Uncertainty Shocks in a Model of Effective Demand. *Econometrica*, 85:937–958.

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

Bernanke, B. S. and Mishkin, F. S. (1997). Inflation targeting: A new framework for monetary policy? *Journal of Economic Perspectives*, 11(2):97–116.

Bi, H., Leeper, E. M., and Leith, C. (2013). Uncertain Fiscal Consolidations. *Economic Journal*, 0:31–63.

Bianchi, F. (2013). Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics. *Review of Economic Studies*, 80(2):463–490.

Bianchi, F. and Melosi, L. (2018). Constrained Discretion and Central Bank Transparency. *The Review of Economics and Statistics*, 100(1):187–202.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.

Bloom, N., Floetotto, M., Jaimovich, N., Eksten, I. S., and Terry, S. J. (2018). Really Uncertain Business Cycles. *Econometrica*, 86(3):1031–1065.

Candia, B., Coibion, O., and Gorodnichenko, Y. (2021). The Inflation Expectations of U.S. Firms: Evidence from a new survey. NBER Working Papers 28836, National Bureau of Economic Research, Inc.

Carvalho, C. and Nechio, F. (2014). Do people understand monetary policy? *Journal of Monetary Economics*, 66(C):108–123.

Christiano, L. J., Trabandt, M., and Walentin, K. (2010). DSGE Models for Monetary Policy Analysis. In Friedman, B. M. and Woodford, M., editors, *Handbook of Monetary Economics*, volume 3 of *Handbook of Monetary Economics*, chapter 7, pages 285–367. Elsevier.

Clarida, R., Galí, J., and Gertler, M. (2000). Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *The Quarterly Journal of Economics*, 115(1):147–180.

Cogley, T., Matthes, C., and Sbordone, A. M. (2015). Optimized Taylor rules for disinflation when agents are learning. *Journal of Monetary Economics*, 72(C):131–147.

Creal, D. D. and Wu, J. C. (2017). Monetary Policy Uncertainty And Economic Fluctuations. *International Economic Review*, 58(4):1317–1354.

Davig, T. (2004). Regime-switching debt and taxation. *Journal of Monetary Economics*, 51(4):837–859.

Davig, T. and Foerster, A. (2019). Uncertainty and Fiscal Cliffs. *Journal of Money, Credit and Banking*, 51(7):1857–1887.

Erceg, C. J. and Levin, A. T. (2003). Imperfect credibility and inflation persistence. *Journal of Monetary Economics*, 50(4):915–944.

Eusepi, S. and Preston, B. (2010). Central Bank Communication and Expectations Stabilization. *American Economic Journal: Macroeconomics*, 2(3):235–271.

Eusepi, S. and Preston, B. (2011). Expectations, Learning, and Business Cycle Fluctuations. *American Economic Review*, 101(6):2844–2872.

Evans, G. W. and Honkapohja, S. (2009). Learning and Macroeconomics. *Annual Review of Economics*, 1(1):421–451.

Fernandez-Villaverde, J. and Levintal, O. (2018). Solution methods for models with rare disasters. *Quantitative Economics*, 9(2):903–944.

Fernandez-Villaverde, J., Rubio-Ramirez, J., and Schorfheide, F. (2016). Solution and Estimation Methods for DSGE Models. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2 of *Handbook of Macroeconomics*, chapter 0, pages 527–724. Elsevier.

Gavin, W. T., Keen, B. D., Richter, A. W., and Throckmorton, N. A. (2015). The zero lower bound, the dual mandate, and unconventional dynamics. *Journal of Economic Dynamics and Control*, 55(C):14–38.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–384.

Husted, L., Rogers, J., and Sun, B. (2019). Monetary policy uncertainty. *Journal of Monetary Economics*.

Jorda, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Leeper, E. M. and Zha, T. (2003). Modest policy interventions. *Journal of Monetary Economics*, 50(8):1673–1700.

Mishkin, F. S. (2018). Improving the use of discretion in monetary policy. *International Finance*, 21(3):224–238.

Richter, A. W. and Throckmorton, N. A. (2015). The consequences of an unknown debt target. *European Economic Review*, 78(C):76–96.

Rotemberg, J. (1982). Monopolistic price adjustment and aggregate output. *Review of Economic Studies*, 49(4):517–531.

Schorfheide, F. (2005). Learning and Monetary Policy Shifts. *Review of Economic Dynamics*, 8(2):392–419.

Taylor, J. B. (2012). Monetary policy rules work and discretion doesn't a tale of two eras. *Journal of Money, Credit and Banking*, 44(6):1017–1032.

Terasvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, 89(425):208–218.

van der Cruijsen, C., Jansen, D.-J., and de Haan, J. (2015). How Much Does the Public Know about the ECB's Monetary Policy? Evidence from a Survey of Dutch Households. *International Journal of Central Banking*, 11(4):169–218.

Woodford, M. (2013). Macroeconomic Analysis Without the Rational Expectations Hypothesis. *Annual Review of Economics*, 5(1):303–346.

APPENDIX A. MODEL SOLUTION AND STABILITY

I solve the model using a textbook projection method with Chebyshev polynomials and orthogonal collocation. For a detailed review on this technique, see Fernandez-Villaverde et al. (2016). I first describe the solution under Rational Expectations, then move to the introduction of limited information. I then discuss accuracy. Lastly, I show the stability of the model.

A.1. **The Full Information economy.** The state space of the model consists of two continuous variables, $\{k_{t-1}, r_t\}$, and one discrete variable, ϕ_t. I approximate the decision rules for labor, inflation, and the price of capital for all possible values of ϕ_t using 3 Chebyshev polynomials for both k_{t-1} and r_t. I need to estimate 54 parameters ($3x3x3x2$).

Evaluating the residuals of eq. (3), eq. (5) and eq. (8) at each of the 3 zeros of the Chebyshev of order 3 for k and r (i.e., collocation points), and the 2 levels of ϕ gives 54 equations to solve for those 54 coefficients. Given a good initial guess (for which I use a linear solution of a simpler model) a Newton solver easily deals with this system.

A.2. **The Limited Information economy.** In Case 2, variable $\lambda_t^L = \mathrm{P}(\phi_t = \phi_L | \Omega_t, \Theta)$ substitutes ϕ_t in the state space. I still approximate the decision rules for labor, inflation, and the price of capital, but I now do it using 3 Chebyshev polynomials for $\{k_{t-1}, r_t, \lambda_t^L\}$. As before, I choose the resulting 81 parameters to minimise both residual functions at the 81 collocation points. To update λ_t^L, agents use the filter in Hamilton (1989).

A.3. **Accuracy.** Following Fernandez-Villaverde et al. (2016) and Fernandez-Villaverde and Levintal (2018), I assess accuracy by computing the mean and maximum unit-free Euler error across the ergodic set of the model. Overall, the solutions for both models are accurate. The mean Euler error in log10 units is -3.2 under Rational Expectations, and -2.8 under limited information, while the maximum Euler error are -2.9 and -2.1, respectively. To put these numbers into perspective, a value of -2 means $1 mistake for each $100, a value of -3 means $1 mistake for each $1000, and so on.

A.4. **Stability.** I now show that the model generates stable time series with unique statistical properties regardless initial conditions and histories of shocks. That is, even though the Taylor principle is only satisfied in the active monetary policy regime, the model converges to a unique ergodic distribution.

FIGURE 3. Stochastic paths of log output.

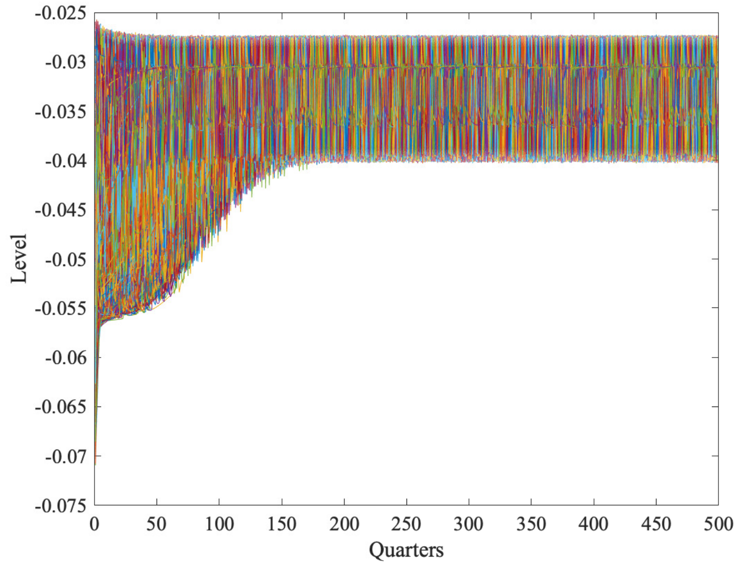

Notes. Each line represents the path of log output in stochastic simulation n for $n = 1, 2, ..., 2000$. I draw the initial conditions from the following distributions: $q_0 \sim U(0, 1)$, $k_0 \sim U(0.9, 1.1)k_{ss}$, $r_0 \sim U(0, 0.02)$. Here k_{ss} is the deterministic steady state value of capital, and $U(\cdot)$ stands for the uniform distribution.

Dynamic systems of non-linear equations cannot, in general, be shown to have finite unconditional moments analytically [see e.g. Terasvirta (1994) and Andreasen et al. (2018)]. Therefore, I show the stability of the model numerically.

To this end, I perform a simple exercise. I will simulate the model N times for T periods conditional on N different sequences of shocks, $\Xi = \{\phi_t, u_t\}_{t=1}^{t=T}$, and N different initial conditions, $\Omega = \{q_0, r_0, k_0\}$. I will then check that the resulting ergodic distributions in all N simulations coincide.

Figure 3, which sets $N = 2,000$ and $T = 500$, reports the paths of log output[22] in each simulation. Importantly, there is not a single path that explodes off to either plus or minus infinity. Instead, all paths converge to the same distribution.

In light of Figure 3, I argue that the model does have a unique ergodic distribution.

This result was not entirely unexpected. In the model, the ergodic mean of ϕ_t is given by:

$$\bar{\phi} = \frac{1 - p_{L,L}}{2 - p_{L,L} - p_{H,H}}\phi_H + \frac{1 - p_{H,H}}{2 - p_{L,L} - p_{H,H}}\phi_L.$$

[22] I could have reported the paths of any endogenous variable; the message would be the same.

As seen in subsection 3.1, I calibrate the model so that $\bar{\phi} = 1.4$. That is, on average the Taylor principle is satisfied, for, on average, a 1% rise in inflation is actually met with a greater than 1% rise in the nominal interest rate.