**Finance and Economics Discussion Series**
**Divisions of Research & Statistics and Monetary Affairs**
**Federal Reserve Board, Washington, D.C.**

# Learning and Misperception: Implications for Price-Level Targeting

**Martin Bodenstein, James Hebden, and Fabian Winkler**

**2019-078**

Please cite this paper as:
Bodenstein, Martin, James Hebden, and Fabian Winkler (2019). "Learning and Misperception: Implications for Price-Level Targeting," Finance and Economics Discussion Series 2019-078. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2019.078.

# Learning and Misperception: Implications for Price-Level Targeting

Martin Bodenstein James Hebden Fabian Winkler*

October 25, 2019

**Abstract**

Monetary policy strategies that target the price level have been advocated as a more effective way to provide economic stimulus in a deep recession when conventional monetary policy is limited by the zero lower bound on nominal interest rates. Yet, the effectiveness of these strategies depends on a central bank's ability to steer agents' expectations about the future path of the policy rate. We develop a flexible method of learning about the central bank's policy rule from observed interest rates that takes into account the limited informational content at the zero lower bound. When agents learn, switching from an inflation targeting to a price-level targeting strategy at the onset of a recession does not yield the desired stabilization benefits. These benefits only materialize after the policy rule has been in place for a sufficiently long time. Temporary price-level targeting strategies are likely to be much less effective than their permanent counterparts.

*Federal Reserve Board, 20th St and Constitution Ave NW, Washington DC 20551. Our email addresses are martin.r.bodenstein@frb.gov, james.s.hebden@frb.gov and fabian.winkler@frb.gov. We thank Etienne Gagnon, Christopher Gust, David Lopez-Salido, Dirk Niepelt, Bob Tetlow, and seminar participants at the European Central Bank for helpful comments. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System.

# 1 Introduction

A number of ongoing structural developments have tested the ability of central banks around the globe to achieve their goals. The neutral real interest rate has likely fallen, implying less leeway to lower interest rates in the event of a recession because of the zero lower bound (ZLB). This development poses a challenge for the predominant monetary policy approach of flexible inflation targeting to manage future recessions and has led some central banks, including the Federal Reserve, to review their existing monetary policy frameworks.[1] Academics and policymakers have explored "makeup strategies" that aim to offset, at least in part, past misses of inflation from its long-run target, in contrast to flexible inflation targeting where the history of past deviations of inflation from this target is irrelevant. Makeup strategies are thought to have large stabilization benefits that stem from their effect on expectations: A commitment to make up for inflation shortfalls through lower interest rates in the future raises near-term inflation expectations and lowers real interest rates which then stimulate the economy even when the short-term nominal interest rate is constrained by the ZLB. A particular type of makeup strategies that has received most of the attention is price-level targeting.[2]

When assessing the effects of price-level targeting, it is commonly assumed that private sector agents know the strategy pursued by the central bank and fully believe policymakers' commitment to this strategy. While these assumptions may be reasonable when the monetary policy strategy has been stable over time, they are questionable when the central bank changes to a strategy without historical precedent and the public has no experience with the new strategy. A switch to price-level targeting or a similar makeup strategy would constitute a radical departure from current practice and the public would require time to learn and to trust this new approach to monetary policy.

[1] Compare the Minutes of the Federal Open Market Committee, September 17-18, 2019 available at `www.federalreserve.gov/monetarypolicy/files/fomcminutes20190918.pdf`.

[2] Eggertsson and Woodford (2004) show that price-level targeting is the optimal commitment policy in the textbook New Keynesian model subject to the ZLB. Reifschneider and Williams (2000), Hebden and López-Salido (2018), Bernanke et al. (2019), and Mertens and Williams (2019) discuss the stabilization properties of various strategies that seek to stabilize the price level in quantitative models. See also Svensson (2019) for a recent review. For a discussion of other benefits of price-level targeting, see Svensson (1999), Vestin (2006); Ambler (2009) and Hatcher and Minford (2016) offer extensive surveys.

After all, even in the comparatively benign case of the shift in monetary policy strategy by Federal Reserve Chairmen Volcker and Greenspan towards a low inflation target, it took considerable time for the public to understand the new strategy, as argued in Erceg and Levin (2003). Since the effectiveness of price-level targeting and other makeup strategies hinges crucially on expectations, we think that a complete evaluation of the merits of these strategies must consider deviations from the full information, rational expectations paradigm, taking into account the public's uncertainty about the central bank's strategy.

There exists an active emerging literature on the effect of agents' cognitive limitations on the effectiveness of monetary policy, including Gabaix (2016), Farhi and Werning (2017), Woodford (2018), and Angeletos and Lian (2018). These studies have focused on mitigating the so-called "forward guidance puzzle" by which the effect of announcements of future monetary policy on inflation and output today increase with the horizon at which the changes are expected to occur.[3] However, these cognitive frictions impair the effectiveness of price-level targeting strategies as well, since these frictions generally limit the effect of future monetary policy actions on expectations. Eusepi and Preston (2018) study the effectiveness of price-level targeting strategies under adaptive learning and find that it retains its stabilization properties.[4] All of these studies, however, consider cases in which a policy strategy is in place indefinitely, and misperceptions caused by agents' cognitive limitations also persist indefinitely. The literature therefore falls short on the question of how expectations can affect the *transition* between policy strategies.

In this paper, we develop a method of learning about the central bank's policy strategy that can be applied to models with a ZLB constraint. Agents hold subjective beliefs about the parameters in the central bank's rule for the short-term nominal interest rate—the policy rate—and they update these beliefs solely from observations of the policy rate and the inputs to the policy rule such as inflation and the output gap. Contrary

[3] Del Negro et al. (2012) coined the term "forward guidance puzzle" in work showing that standard medium-scale DSGE models grossly overestimate the impact of forward guidance on the macroeconomy.

[4] Mele et al. (2018) argue that price-level targeting need not be optimal when a rational central bank interacts strategically with a learning private sector.

to the learning environments featured in Evans and Honkapohja (2001) and Eusepi and Preston (2018), the entire structure of the economy is common knowledge and agents make rational forecasts conditional on their perceived policy rule. If the true policy rule stays in place over a sufficiently long time horizon, then the learning equilibrium converges to the full information, rational expectations equilibrium. However, when the central bank switches to a new strategy, it takes time for agents to learn this new strategy. During this transition, misperception of the policy strategy by the private sector can have unintended consequences for economic outcomes. The convergence of beliefs can be hampered in particular when the economy is at the ZLB, because the observed policy rate carries little information about the true rule parameters in this case.

To assess the implications of our learning framework for monetary policy, we consider the textbook New Keynesian model. The central bank sets the policy rate $i_t$ according to a simple rule that can react to an inflation gap and/or a price-level gap. We focus on a switch from a (flexible) inflation targeting to a (flexible) price-level targeting strategy. Our analysis suggests that under learning the switch to price-level targeting falls short of delivering the stabilization benefits that are found under full information in a demand-driven recession with a binding ZLB.

When the central bank switches to price-level targeting at the onset of a demand-driven recession, the switch mitigates the loss in output and the shortfall in inflation under rational expectations and full information. By contrast, when agents are learning, output and inflation initially fall just as much as under the inflation targeting strategy despite the switch. Because agents do not immediately understand the switch in the policy rule, they initially attribute the differences in the policy rate resulting under the price-level targeting strategy to a series of discretionary policy shocks rather than the switch in strategy. Agents therefore fail to anticipate the more accommodative policy associated with the price-level targeting strategy. The learning problem is further complicated by the fact that the policy rate quickly reaches the ZLB. At the ZLB the interest rate prescribed by the policy rule is censored as the actual policy rate cannot fall below zero and the private sector agents receive little information about the true rule parameters. As a result, under learning, the central bank is left with a much larger negative price-level gap than under full information, and thus has to allow for substantial overshooting of inflation after the recession to deliver on its promise of price level stabilization. The costs of this overshooting are incurred without having accrued any stabilization benefits in the midst of the recession.

In order for the stabilization benefits of price-level targeting to materialize, price-level targeting should be introduced in relatively calm times—that is, when inflation is not persistently undershooting its long-run target and the federal funds rate is not constrained by the ZLB—to give agents the opportunity to learn the new policy strategy. When put in place for a sufficiently long time, systematic price-level targeting then becomes superior to inflation targeting, just as under full information.

We also show that a price-level targeting strategy that is permanently in operation is preferable to a temporary price-level targeting strategy of the type suggested by Evans (2012) and Bernanke et al. (2019). Under temporary price-level targeting, monetary policy falls into two regimes: The central bank targets the price level when the ZLB binds but switches back to targeting inflation once the price-level gap accumulated during the ZLB episode has been closed. Consequently, the price-level targeting regime is active precisely when it is difficult for agents to infer changes in monetary policy. Agents are unlikely to anticipate that monetary policy will be more accommodative in this regime, a failure that renders the strategy ineffective. This result echoes the concern voiced by Svensson (2019) that the effectiveness of makeup strategies "probably requires that economic agents need to see the policy practiced and its principles obeyed for some time, in order to believe that it will be maintained and be successful in the future."

Our behavioral learning framework builds on Tetlow and von zur Muehlen (2001) and Cogley et al. (2015). We extend their approach to take into account the limited informational content of interest rate observations at the ZLB by explicitly modeling the associated non-linearity in the observation equation of agents' Bayesian state-space system. In related work, Gust et al. (2018) assume that agents are only uncertain about the value of the intercept term in the policy rule where the intercept term can take on values from a small, finite, and publicly known set. All other rule parameters are known and fixed at all times. Price-level targeting strategies are not considered in their work.

Limited information and learning also play a role in the literature on imperfect credibility, which is often modeled as agents not observing parts of the central bank's policy strategy. For example, Erceg and Levin (2003) and Schorfheide (2005) interpret the shifts in monetary policy in the 1980s and 1990s through the lens of learning and limited credibility. Bodenstein et al. (2012) show that, under imperfect credibility, private sector agents may doubt that the central bank will honor its announcement to keep interest rates low for longer. As a result, the interest rate path expected by the private sector lies above the path announced by the central bank. In our paper, the private sector projects future monetary policy to be tighter than intended by the central bank under price-level targeting, leading to outcomes that resemble those under a lack of credibility. Finally, Kryvtsov et al. (2008) model the switch from an inflation targeting strategy to a price-level targeting strategy under imperfect credibility. They model credibility to be independent of the observed interest rate path and they do not impose a ZLB constraint.

The remainder of this paper is structured as follows. Section 2 describes our generalized model of learning, while Section 3 contains our application to the introduction of price-level targeting. In Section 4 we discuss the best timing for a central bank to switch from inflation targeting to price-level targeting. Section 5 analyzes a temporary price-level targeting strategy. Section 6 concludes.

## 2 A Model of Learning the Monetary Policy Strategy

For the purpose of our analysis, a monetary policy strategy is fully described by a simple policy rule that describes how the central bank maps inflation, the output gap, and possibly other variables into a value of the short-term nominal interest rate, the policy rate. We assume that the private sector agents do not know with certainty the policy rule of the central bank, and we therefore distinguish between the *actual* policy rule followed by the central bank and the *perceived* policy rule that agents use to form expectations. Agents infer the parameters of the perceived policy rule solely from observations of economic data as described in the following.

## 2.1 The Economy

We conduct the analysis in a discrete-time model that is linear except for the ZLB. Excluding the description of monetary policy, the equilibrium conditions are summarized by a linear forward-looking system of equations of the form:

$$0 = F_2 \mathbb{E}_t [x_{t+1}] + F_1 x_t + F_0 x_{t-1} + F_i i_t + F_u u_t \tag{1}$$

All variables are expressed in deviations from the deterministic steady state of the model. The endogenous variables $x_t$ enter with their current, past, and expected future values into the model. Exogenous disturbances enter through the random vector $u_t$. The shocks are iid and have mean zero. Finally, the policy instrument $i_t$ is determined by the central bank as described below.

## 2.2 The Central Bank

The actual policy rule for the policy rate (in deviations from its deterministic steady state) $i_t$ is of the form:

$$i_t = \max \{\underline{i}, i_t^*\}, \; i_t^* = \Psi(\beta_t) x_t + e_t \tag{2}$$

where $i_t^*$ is the notional interest rate and $\underline{i}$ is the lower bound on the policy rate. The notional rate is set according to a linear rule with parameters $\Psi(\beta_t)$ that are a function of a small set of parameters $\beta_t$. These parameters can vary over time to accommodate changes in the central bank's systematic response to economic outcomes. The policy rate is also affected by a white noise process $e_t$. This shock represents one-time discretionary adjustments to the policy rate.[5]

[5] It is straightforward to extend our framework to include persistent monetary policy shocks to model time-varying changes in the intercept term of the rule as in Gust et al. (2018)

## 2.3 Full information rational expectations equilibrium

In the full information rational expectations equilibrium, private sector agents observe the sequences of past and current realizations of the endogenous variables $x_t$, the exogenous disturbances $u_t$ and the policy rate $i_t$. Agents also know the linear economic model given in equation (1). We assume that the central bank commits to the values of the parameters in the policy rule $\beta_t$ in equation (2). Private sector agents know these parameters and the form of the policy rule and hence have a complete understanding of the monetary policy strategy. Agents also observe the current value of the policy shock $e_t$. At every point in time, the private sector agents know the correct policy rate path that the central bank intends to implement contingent on the state of the economy. We solve for the full information rational expectations equilibrium with an occasionally binding ZLB constraint using the algorithm of Holden (2016).

## 2.4 Learning Equilibrium

Under learning, private sector agents also observe the sequences of past and current realizations of the endogenous variables $x_t$, the exogenous disturbances $u_t$ and the policy rate $i_t$. They also know the linear economic model given in equation (1). While agents know the general form of the actual policy rule in equation (2), they do not observe the values of the parameters $\beta_t$. Neither do they observe the realizations of the monetary policy shock $e_t$. Instead, agents believe that the transitory shock and changes to the parameters at time $t$ are normally distributed white noise with:

$$\begin{pmatrix} e_t \\ \beta_t - \beta_{t-1} \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} \sigma_{et}^2 & 0 \\ 0 & \Sigma_{\beta t} \end{pmatrix}\right). \tag{3}$$

The variances $\sigma_{et}^2 > 0$ and $\Sigma_{\beta t} > 0$ are subjective and are an exogenous input to the learning process. The assumption that policy parameters change over time is common in empirical work. Notably, Boivin (2006) assumes that policy rule parameters follow random walks to assess how the conduct of U.S. monetary policy has changed over time. His estimates suggest that the rule parameters evolve gradually and feature wide error bands. Strictly speaking, these beliefs render agents boundedly rational in our model since the true rule parameters are constant except for a one-time discrete jump. We view this setup as a simplified representation of an environment with fundamental uncertainty about the actual policy rule.

Our formulation of the learning equilibrium follows Cogley et al. (2015) with the important difference that we include the informational constraints arising from the ZLB. Private sector agents enter period $t$ with beliefs about the policy rule parameters inherited from $t-1$. In formulating decisions plans, agents treat the mean parameter estimates as if known with certainty. Then period $t$ shocks are realized. Agents observe the realizations of the private-sector shocks and the central bank's policy action and infer a perceived policy shock $\tilde{e}_t$. Outcomes are determined in accordance with the beginning-of-period plans. Finally, agents update their estimates of the rule parameters.

In providing a detailed description of the learning framework, we start by defining the mean of the posterior distribution of $\beta_{t-1}$ at the end of period $t-1$ with $\hat{\beta}_{t-1}$. Following Kreps (1998), agents plan under anticipated utility and view the rule parameters as fixed at $\hat{\beta}_{t-1}$.[6] These plans also depend on the agents' perceived value $\tilde{e}_t$ of the actual policy shock. Agents solve for state-contingent paths starting in period $t$ denoted by $\left\{x_s^{(t)}, s \geq t\right\}$ that satisfy the system of equations (1) and the policy rule (2) with $\beta_s = \hat{\beta}_{t-1}$ for all $s \geq t$. Merging the two conditions, the solution needs to satisfy

$$0 = F_2 \mathbb{E}_s \left[x_{s+1}^{(t)}\right] + F_1 x_s^{(t)} + F_0 x_{s-1}^{(t)} + F_i \max\left(\underline{i}, \Psi\left(\hat{\beta}_{t-1}\right) x_s^{(t)} + e_s^{(t)}\right) + F_u u_s \tag{4}$$

for all $s \geq t$ with the initial condition $x_{t-1}^{(t)} = x_{t-1}$. In solving this problem, agents take the perceived policy shock sequence $e_s^{(t)}$ as distributed iid $\mathcal{N}\left(0, \sigma_e^2\right)$ with $e_t^{(t)} = \tilde{e}_t$. Thus, the solution $x_s^{(t)}$ represents the agents' expectations about the future evolution of the economy at time $t$. Again, we rely on the computationally efficient algorithm of Holden

[6] Anticipated utility refers to the widely used assumption in the learning literature that agents derive their decisions and expectations about future developments under the assumption that their current perception of the economic environment, in our case the policy rule parameters, persists indefinitely. This simplifying assumption ignores that, at the same time, the public treats the parameters in the policy rule as random variables in the learning problem. See (Cogley and Sargent, 2008) on interpreting anticipated utility as an approximation to Bayesian optimal learning.

(2016) in this step.

The perceived value of the policy shock for the current period $\tilde{e}_t$ reflects the observed value of the policy rate:

$$\tilde{e}_t = \mathbb{E}\left[e_t \mid i_t = \max\left(\underline{i}, \Psi\left(\hat{\beta}_{t-1}\right)x_t + e_t\right), x_t\right] = \begin{cases} i_t - \Psi(\hat{\beta}_{t-1})x_t & \text{if } i_t > \underline{i} \\ -\sigma_{et}\frac{\phi\left(\frac{\underline{i}-\Psi(\hat{\beta}_{t-1})x_t}{\sigma_{et}}\right)}{\Phi\left(\frac{\underline{i}-\Psi(\hat{\beta}_{t-1})x_t}{\sigma_{et}}\right)} & \text{if } i_t = \underline{i} \end{cases} \tag{5}$$

where $\phi$ and $\Phi$ are the standard normal density and cumulative distribution functions, respectively. If the observed policy rate is above the ZLB, $\tilde{e}_t$ equals the difference between the observed policy rate $i_t$ (derived from the actual policy rule with parameters $\beta_t$) and the policy rate projected by the private sector (derived from the perceived policy rule with parameters $\hat{\beta}_{t-1}$). If the observed policy rate is at the ZLB, $\tilde{e}_t$ equals the conditional expectation of the policy shock $e_t$ when the notional rate $i_t^*$ is below the lower bound $\bar{i}$, i.e., the mean of a truncated normal distribution.

To obtain for the equilibrium in period $t$, we solve simultaneously for $x_t = x_t^{(t)}$, $i_t$, and $\tilde{e}_t$ using equations (4), (5), and the actual policy rule in (2). The appendix provides details on the solution algorithm.

Having observed the equilibrium outcomes $x_t$, agents update their beliefs about the rule parameters $\beta_t$ by solving the following Bayesian filtering problem:

$$i_t = \max\left\{\underline{i}, i_t^*\right\} \tag{6}$$

$$i_t^* = \Psi\left(\beta_t\right)x_t + e_t,\ e_t \sim \mathcal{N}\left(0, \sigma_{et}^2\right) \tag{7}$$

$$\beta_t = \beta_{t-1} + \epsilon_{\beta t},\ \sim \mathcal{N}\left(0, \Sigma_{\beta t}\right). \tag{8}$$

Using the posterior distribution of beliefs about $\beta_{t-1}$ as the prior distribution in this filtering problem, agents derive a new posterior distribution of beliefs about $\beta_t$ given the observations of $i_t$ and $x_t$. As in Cogley et al. (2015), agents treat $x_t$ as exogenous and thus as independent of the shocks $e_t$ and $\epsilon_{\beta t}$. By ignoring the correlation between the policy shocks and the economic outcomes, agents do not make use of all the available information. However, form an analytical perspective, this exogeneity assumption greatly simplifies the agents' filtering problem because the problem reduces to a Bayesian regression with truncation in this case. In addition, we make two computational approximations to the filtering problem: First, we replace the posterior distribution by its Laplace approximation, i.e. we approximate the posterior distribution of $\beta_t$ with a normal distribution. Second, we approximate the potentially non-linear mapping $\Phi(\beta_t)$ with a first-order Taylor expansion as in extended Kalman filtering, but keep the non-linearity arising from the ZLB. Again, we refer to the appendix for details on the solution algorithm.

Several of the stated behavioral assumptions imply that the private sector agents are boundedly rational in our model. First, agents behave as anticipated utility modelers and treat the current estimates of the rule parameters $\hat{\beta}_{t-1}$ as if known with certainty when deriving their economic decisions and when computing the perceived policy shock $\tilde{e}_t$. Second, agents update the estimates of the rule parameters through a filtering problem that, although Bayesian, ignores the endogeneity of the policy shocks $e_t$ to the model outcomes $x_t$. Third, agents take a perceived law of motion for policy innovations in equation (3) as given (in particular the variances $\sigma_{et}$ and $\Sigma_{\beta t}$), even though this law of motion may not coincide with the central bank's actual formulation of monetary policy.

# 3 Learning a Price-Level Targeting Strategy

We apply this learning framework to a situation in which the central bank switches from an inflation targeting strategy to a price-level targeting strategy. A price-level targeting strategy actively seeks to offset passed misses of inflation from the central bank's inflation goal. As laid out in the introduction, price-level targeting is considered to be better suited than inflation targeting to stabilize the macroeconomy, particularly when the policy rate is at the ZLB. Most of the subsequent analysis also applies for other makeup strategies such as average inflation targeting or shadow rate rules.[7]

While our learning framework can be embedded into any linear model of the form

[7] For a summary of makeup strategies see, Bernanke et al. (2019) and Hebden and López-Salido (2018).

in equation (1), we adopt the textbook New Keynesian model to characterize the underlying economic environment in our illustration of the impact of learning on the effectiveness of price-level targeting. This choice of model seems well suited given that most theoretical arguments about the benefits of price-level targeting are formulated within this model.

### 3.1 Economic Model

The textbook New Keynesian model features two equations:

$$\pi_t = \kappa y_t^{gap} + \beta \mathbb{E}_t \left[\pi_{t+1}\right] + \upsilon_t \tag{9}$$

$$y_t^{gap} = \mathbb{E}_t \left[y_{t+1}^{gap}|\right] - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t \pi_{t+1} - g_t\right). \tag{10}$$

All variables are expressed in deviation from their non-stochastic steady state values. The New Keynesian Phillips Curve in (9) links inflation (measured relative to its long-run target $\pi^*$) $\pi_t$ to its expected value and the output gap $y_t^{gap}$. The output gap is defined as the difference between actual output and the natural level of output, $y_t^{gap} = y_t - y_t^*$. $\upsilon_t$ is an inefficient cost-push shock. The Aggregate Demand Curve in (10) provides the connection between the output gap, inflation, the policy rate $i_t$ and the natural rate of interest $g_t$. $\mathbb{E}_t$ denotes the subjective expectations of the private sector conditional on its information set $\mathbb{I}_t$.

We set the discount factor $\beta$ equal to $0.9956$ to imply a steady state real interest rate of 1.75 percent, and set the intertemporal elasticity of substitution $\sigma$ equal to $1$. The slope of the Phillips Curve $\kappa$ is fixed at $0.1$. The demand and supply shocks $g_t$ and $\upsilon_t$ follow first-order autoregressive processes with autocorrelations $\rho_g = \rho_u = 0.9$. The standard deviations of the innovations are $\sigma_g = 0.3$ and $\sigma_u = 0.03$, respectively, in order to match the volatility of inflation and the output gap in quarterly U.S. data from 1984Q1 to 2007Q4.

To highlight the importance of the private sector's expectations about monetary policy for the evolution of the economy—the expectations channel of monetary policy—we iterate equations (9) and (10) forward. The Aggregate Demand Curve can thus be written as

$$y_t^{gap} = -\frac{1}{\sigma}\mathbb{E}_t \sum_{j=s}^{\infty} [i_{t+s} - \pi_{t+s+1}] - \frac{1}{\sigma}\mathbb{E}_t \sum_{j=s}^{\infty} [g_{t+j}] \tag{11}$$

which reveals the dependence of the output gap on the expected path for the real interest rate. Similarly, the Phillips Curve implies that inflation equals the discounted sum of current and future expected output gaps

$$\pi_t = \kappa \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t [y_{t+s}^{gap}] + \mathbb{E}_t \sum_{s=0}^{\infty} \beta^s [v_{t+s}] . \tag{12}$$

As a result, inflation also depends on the expected path of the real interest rate. Thus, the ability of the central bank to steer inflation and the output gap into any desired direction depends importantly on the ability of the central bank to steer the private sector's expectations about current and future interest rates.

If the monetary policy strategy—the state-contingent path of the policy rate $i_t$—is known to the private sector, then the central bank can successfully steer the economy in the desired direction. However, if the monetary policy strategy is not known with certainty or, alternatively, not perceived as credible by the private sector, perceived and actual monetary policies may differ importantly from each other. As a result, the realized economic outcomes may end up differing substantially from those intended by the central bank.

## 3.2 Monetary Policy Strategies

We assume that, at every point in time, the strategy of the central bank is fully described by the interest rate rule

$$i_t^* = \rho_i i_{t-1} + (1-\rho_i)\left( (1+\phi_{\pi t})\,\pi_t + \phi_{pt}\frac{p_t^{gap}}{4} + \phi_y \frac{y_t^{gap}}{4} \right) + e_t \tag{13}$$

$$i_t = \max\{i_t^*, \underline{i}\} \tag{14}$$

which is consistent with the steady state form given in equation (2). The lower bound $\underline{i}$ on the policy rate is expressed relative to the steady state and equals $\underline{i} = -\frac{\pi^*}{\beta}$ where the steady state inflation rate $\pi^*$ equals the central bank's inflation target.

The central bank arrives at a value for the notional interest rate $i_t^*$ from the current values of inflation, the output gap, and, possibly, the lagged realized value of the policy rate. Under price-level targeting, the central bank also responds to a price-level gap. The price-level gap $p_t^{gap}$ records the cumulative departure of inflation from its target value from a fixed date in the past to the present period $t$ and evolves according to $p_t^{gap} = p_{t-1}^{gap} + \pi_t$. Given the ZLB constraint, the actual policy rate $i_t$ equals the notional rate $i_t^*$ if the latter is above $\underline{i}$, and equal to $\underline{i}$ otherwise.

The subsequent analysis assumes that before some period $\tau_0$, the central bank follows an inflation targeting strategy. More precisely, we characterize inflation targeting by the inertial version of the Taylor (1999) rule which has been featured in public documents of the Federal Reserve such as the Monetary Policy Report to Congress and Brayton et al. (2014) and accounts for the empirical observation that central banks adjust rates sluggishly in response to economic conditions, see English et al. (2003).[8]

The parameters in the inertial Taylor rule assume the values $\rho_i = 0.85$, $\phi_y = 1$, $\phi_\pi = 0.5$ and $\phi_p = 0$. The inflation target $\pi^*$ is set to two percent. The central bank always adheres to the prescriptions from the rule, and in particular the true monetary policy shock is always zero. However, as stated above, the private sector does not necessarily understand this feature of policymaking and, at times, will perceive the monetary policy shock to differ from zero.

Given these assumptions, we consider a switch in monetary policy in period $\tau_0$ to a price-level targeting rule with $\phi_p = 1$ and $\phi_\pi = 0$. We set the reference date for the price-level gap to be period $\tau_0 - 1$ so that upon switching the policy strategy the price-level

[8]Because of interest rate inertia, the inertial Taylor rule allows for some history dependence when the economy is not at the ZLB. Repeated substitutions of the lagged policy rate term in the rule imply that the current value of the notional rate responds to averages of current and past inflation and current and past output gaps with lower weights on observations further in the past. However, while the economy is at the ZLB, this history dependence comes to a partial halt as the actual policy rate no longer records fully the deviations of inflation and the output gap from their long-run target values. Nevertheless, for the same values of inflation and the output gap, the inertial Taylor rule prescribes a shallower rate path than its non-inertial counterpart with $\rho_i = 0$.

gap is closed, i.e., $p^{gap}_{\tau_0-1} = 0$. Without loss in generality, we assume the parameters on the lagged policy rate $\rho_i$ and the output gap $\phi_y$ to remain unchanged and that the private sector understands this to be the case.

### 3.3 Beliefs

While private sector agents do not understand how the parameters $\phi_\pi$ and $\phi_p$ have changed in period $\tau_0$, agents can infer the new parameter values from the data by solving the Bayesian regression problem described above. This problem depends on the subjective beliefs about the rule parameters. Initially, agents know the parameters of the inertial Taylor rule with certainty, as would be the case if that rule had been in place for a sufficiently long period for the beliefs to have converged to the true values. To update their estimates of the rule parameters $\phi_\pi$ and $\phi_p$ after the switch to price-level targeting, agents apply the filtering problem in equations (6)–(8) with

$$\beta_t = \begin{pmatrix} \phi_{\pi t} \\ \phi_{pt} \end{pmatrix}, \; \Phi\left(\beta_t\right) x_t = \rho_i i_{t-1} + \left(1 - \rho_i\right) \left( \left(1 + \phi_{\pi t}\right) \pi_t + \frac{\phi_{pt}}{4} p^{gap}_t + \frac{\phi_y}{4} y^{gap}_t \right)$$

given the rule described in equation (13). Since the parameters on the lagged value of the policy rate $\rho_i$ and the output gap $\phi_y$ remain unchanged throughout the analysis, we assume for simplicity that private sector agents do not consider the possibility of changes in these parameters either. In addition, agents set the initial price-level gap to zero.

Over time, the parameter beliefs converge to the true parameter values. The speed of convergence depends on the subjective prior uncertainty about the rule parameters $\Sigma_{\beta t}$ and the subjective variance of policy shocks $\sigma^2_{et}$. The ratio $\Sigma_{\beta t}/\sigma^2_{et}$ can be understood as a signal-to-noise ratio in the Bayesian regression problem. If the entry in $\Sigma_{\beta t}$ associated with a specific parameter is larger, then a given-size forecast error in the policy rate carries more information about this specific parameter and the resulting update in this parameter will be larger.

We think it is unclear how to judge empirically how agents would adjust their expectations in the wake of the adoption of price-level targeting that is without precedent in recent history.[9] In our benchmark specification, we choose $\sigma_{et}^2$ and $\Sigma_{\beta t}$ such that beliefs converge about half-way to the truth in 20 quarters after the strategy switch. In detail, we set

$$\sigma_{et}^2 = 0.01, \; \frac{\Sigma_{\beta t}}{\sigma_{et}^2} = \begin{pmatrix} 0.4 & -0.11 \\ -0.11 & 0.05 \end{pmatrix}$$

implying to a perceived standard deviation of the policy shocks of 10 basis points and perceived standard deviations of the innovations in the parameter $\phi_\pi$ and $\phi_p$ of 0.06 and 0.02, respectively, and a perceived correlation between innovations to $\phi_\pi$ and $\phi_p$ of -0.25. In other words, agents think that the central bank will tend to pursue either inflation targeting or price-level targeting, but they do not think that they are mutually exclusive.

We also consider a "slower learning" case implying less willingness of the agents to adjust their views about monetary policy. In this case, we reduce the parameter innovation matrix $\Sigma_{\beta t}$ by a factor of ten which yields much slower convergence of the parameters to the truth.

### 3.4 Learning in Normal Times

We illustrate the learning mechanism using stochastic simulations. Each simulation is initialized by a draw from the ergodic distribution of the variables generated under the inflation targeting rule (with parameters $\phi_\pi = 0.5$ and $\phi_p = 0$. Given these initial conditions, policymakers switch to the new rule in period $\tau_0 = 0$ with parameters $\phi_\pi = 0$ and $\phi_p = 1$.

Figure 1 plots how the private sector's parameter beliefs evolve under full information and under learning, respectively. Under full information (red lines), private sector beliefs immediately adjust to their new true values.

By contrast, in our benchmark learning case (yellow lines), the parameter beliefs ad-

[9] One could follow the approach in Boivin (2006) and estimate a policy rule with time-varying parameters from historical data and use the resulting estimates to discipline the belief process in our learning model. However, such an analysis is unlikely to yield much information on the past evolution of the parameter on the price-level gap $\phi_{pt}$ since no central bank has pursued a price-level target, and even less on its future evolution.

Figure 1: Beliefs of rule parameters in normal times.

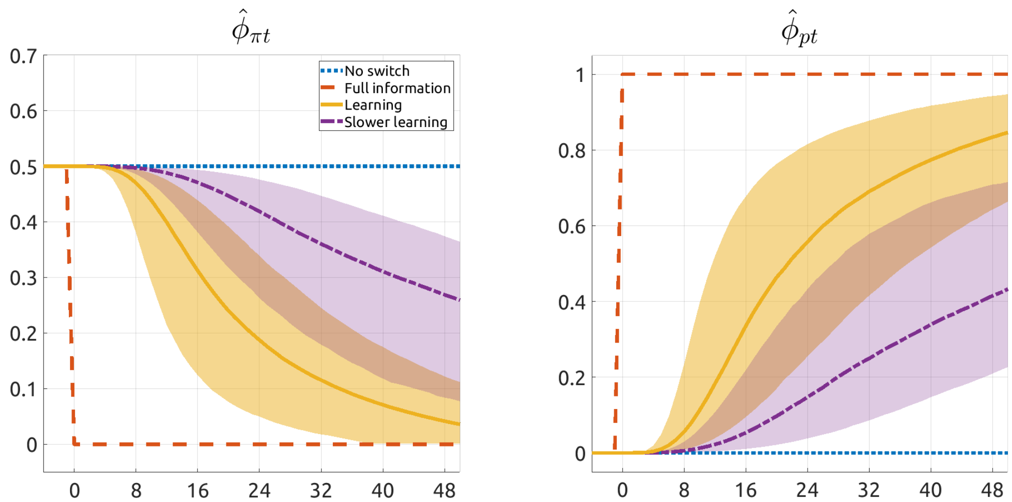


Note: Solid lines show median beliefs $\hat{\phi}_{\pi t}$ and $\hat{\phi}_{pt}$ across 1,000 simulations. Shaded areas show 10th and 90th percentiles.

just slowly, but converge over time. The paths of the parameter beliefs along a particular simulation depend on the realizations of the underlying economic shocks. The solid lines in Figure 1 show the median beliefs over the rule parameters. Across simulations, these beliefs can vary considerably as evidenced by the shaded area representing the 10th and 90th percentiles of the distribution of beliefs. Initially, the beliefs move strongly into the right direction because, as we discuss further below, there is more variation in observed outcomes to learn from just after the rule has switched. The speed of learning depends importantly on the subjective uncertainty that agents place on changes in the rule parameters. In our benchmark case, the median belief has almost converged to the true parameters after 15 years. Beliefs about the parameter on inflation converge somewhat faster than those on the price-level gap parameter. Under the "slower learning" specification (purple lines), agents apportion the differences between the observed policy rate and the rate prescribed under the perceived rule more to the discretionary policy shock $e_t$ than to the changes in the rule parameters $\beta_t$. Hence, the updating steps in the parameters are smaller.

Figure 2 plots the joint distribution of beliefs about the parameters $\phi_\pi$ and $\phi_p$ in our benchmark learning case at three points in time. As agents update their beliefs, a larger estimate of the parameter on the price-level gap term $\hat{\phi}_{pt}$ is associated with a lower esti-

Figure 2: Joint distribution of rule parameter beliefs in normal times.

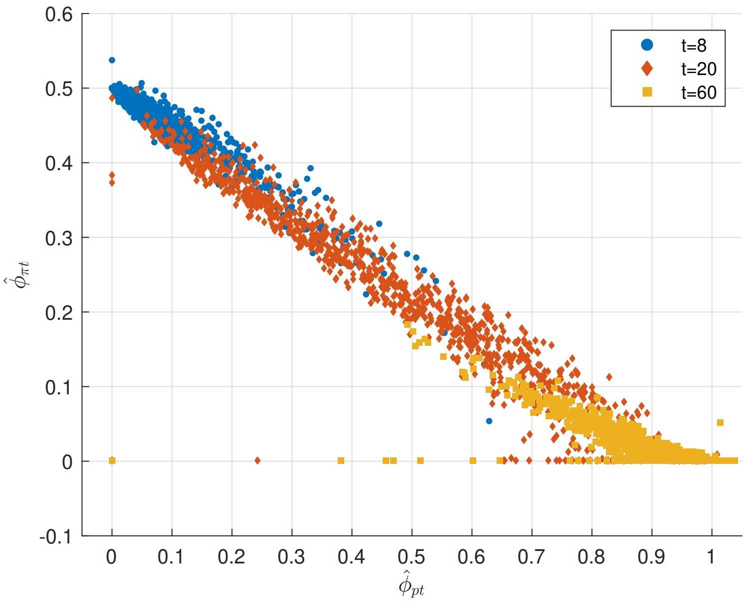


Note: For each time period $t$ shown, a dot corresponds to a belief $\left(\hat{\phi}_{\pi t}, \hat{\phi}_{pt}\right)$ in one of 1,000 simulations under the benchmark learning parameterization.

mate of the parameter on the inflation term $\hat{\phi}_{\pi t}$.

The stabilization benefits of switching to price-level targeting are illustrated in Figure 3. We show the volatilities of inflation and the output gap as measured by the interquartile range of realizations across simulations. By excluding the tails of the distribution the figure captures volatilities in "normal times."

Under full information, price-level targeting reduces inflation volatility immediately by half relative to inflation targeting, while the interquartile range for the output gap drops by a small amount. This feature highlights the stabilizing properties of price-level targeting in forward-looking models: When inflation is low, price-level targeting calls for bringing about higher future inflation to stabilize the price level. Provided that the policy is understood and credible, inflation expectations rise, through the forward-looking Phillips curve (9), the initial shortfall in inflation is mitigated.

By contrast, under learning the volatility of inflation falls slowly, while output gap volatility even rises initially after the switch. The central bank's ability to stabilize the economy through the expectations channel of monetary policy deteriorates initially for two reasons. First, because agents perceive the policy rule to be different from the ac-

Figure 3: Inflation and output gap volatility in normal times.

(a) Annualized inflation $4\pi_t$.

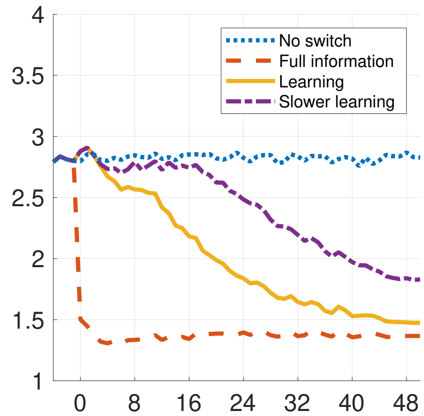


(b) Output gap $y_t^{gap}$.

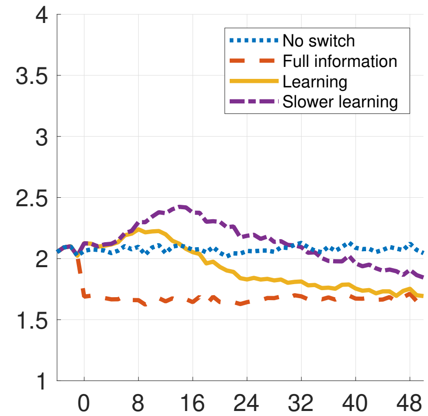


Note: Lines show the interquartile range, i.e. the 75th minus the 25th percentile, of outcomes across 1,000 simulations at each time period $t$.

tual one, agents expect a lower or higher interest rate path in a given simulation than policymakers intend to pursue under the actual policy rule. Mechanically, such misperception has similar effects on the economy as a monetary policy shock. Second, belief updating is itself a source of volatility. Because agents' beliefs about the policy rule parameters change so do their implied expectations about the future interest rate path. Over time, however, as beliefs have moved close enough to the true parameter values, the volatilities fall below their pre-switch levels and approach the levels under full information.

The better agents understand the switch of the central bank to price-level targeting, the smaller will be the variation in inflation and the output gap across simulations. Consequently, once agents have made sufficient learning progress towards the true parameter values, the data provides less identifying variation, and parameter learning slows down. This nexus between the agents' beliefs and the variability of economic variables explains the slowdown in the pace of learning in the later part of the simulations shown in Figure 1.

### 3.5 Learning During a Deep Recession

Proponents of price-level targeting (or other makeup) strategies have emphasized the stabilizing features of this strategy in particular when the policy rate is at the ZLB, although, as illustrated above, the benefits may apply more broadly.[10] If, in a deep demand-driven recession, the policy rate reaches the ZLB, a price-level targeting central bank provides automatically additional monetary accommodation. In order to make up for the shortfall of inflation from its long-run target, the central bank will need to keep the interest rate path sufficiently low to induce future catch-up inflation. In many models, the anticipation of higher future inflation and low future nominal interest rates lowers the expected path for real interest rates which in turn can stimulate the economy up front without any contemporaneous interest rate adjustment (because of the ZLB).[11]

For price-level targeting strategies to stabilize the economy through this expectations channel, private sector agents must understand the policy strategy and consider it credible. Under learning, the central bank may not achieve the desired outcomes from adopting the price-level targeting strategy because the central bank cannot reveal its commitment to the new strategy through the observed policy rate as long as the policy rate is at the ZLB. Agents receive little information about the new policy rule while at the ZLB and they fail to anticipate that the central bank will keep the policy rate path low in the future. As a result, the switch in policy strategy does not provide further monetary stimulus. If, nevertheless, the central bank remains committed to its new strategy, it will subsequently have to allow for higher inflation to undo the accumulated price-level gap.

We illustrate these challenges to price-level targeting by considering a switch from inflation targeting to price-level targeting in a demand-driven recession during which the ZLB binds. Specifically, we choose a combination of demand shocks $g_t$ and supply shocks $u_t$ that, under inflation targeting, induce inflation to fall by roughly 1.5 percent-

[10] Compare also to Svensson (1999) and Vestin (2006).

[11] Under full information, pairing inflation targeting with forward guidance on the path of the policy rate could yield similar effects as permanent price-level targeting. However, under credible permanent price-level targeting future policy accommodation is automatic at the ZLB; absent an explicit rule for forward guidance the adhoc nature of this approach may delay the communication of future policy accommodation. See also the discussion in Section 5 on temporary price-level targeting.

age points and the output gap to fall by roughly 5 percentage points, magnitudes of declines that are comparable to those during the Great Recession. The innovations to the shocks start in period $t = 0$ and end in period $t = 11$; afterwards $g_t$ and $u_t$ converge back to zero at speeds dictated by the auto-regressive parameters $\rho_g$ and $\rho_u$.

Figure 4 shows the median outcomes across simulations conditional on this sequence of shocks (i.e., shocks are sampled randomly except for periods $t = 0, \ldots, 11$). The blue lines (labeled "no switch") depict the outcomes in the recession scenario when the central bank follows the inflation-targeting rule throughout the entire simulation. The presence of the effective lower bound exacerbates the effects of the recessionary shocks and leads to a large drop in inflation and the output gap. Notably, under the inflation targeting strategy the central bank does not makeup for any deviations of inflation from its long-run target during the ZLB episode; there is no inflation overshooting.

When the central bank adopts a price-level targeting rule at the onset of the recession ($\tau_0 = 0$), this new strategy is very effective in mitigating the adverse effects of the recession under full information (red lines): Inflation only drops half as much as under inflation targeting and the fall in the output gap is reduced. Later in the simulation, inflation overshoots its target, as the price-level targeting rule keeps policy rates low for longer to close the price-level gap. In fact, it is precisely the expectation of this more accommodative policy stance and the accompanying inflation overshoot that, through the expectational channels of the New Keynesian Phillips Curve (9) and the Aggregate Demand Curve (10), prevents inflation from falling during the recession.

By contrast, under learning (yellow lines in Figure 4), agents fail to anticipate the full extent of future policy accommodation under price-level targeting. As a result, inflation expectations are lower than under full information, the drop in the real interest rate is restrained, and little buffer is provided against the declines in inflation and the output gap. The large and persistent drop in inflation accumulates to a sizable price-level gap over time and, consequently, the central bank must keep the policy rate lower for longer than under full information to close this gap. A sizable overshoot of inflation results from the attempts to close the gap later in the simulation.

Figure 5 shows the evolution of the private sector beliefs that correspond to these

Figure 4: Outcomes during a deep recession.

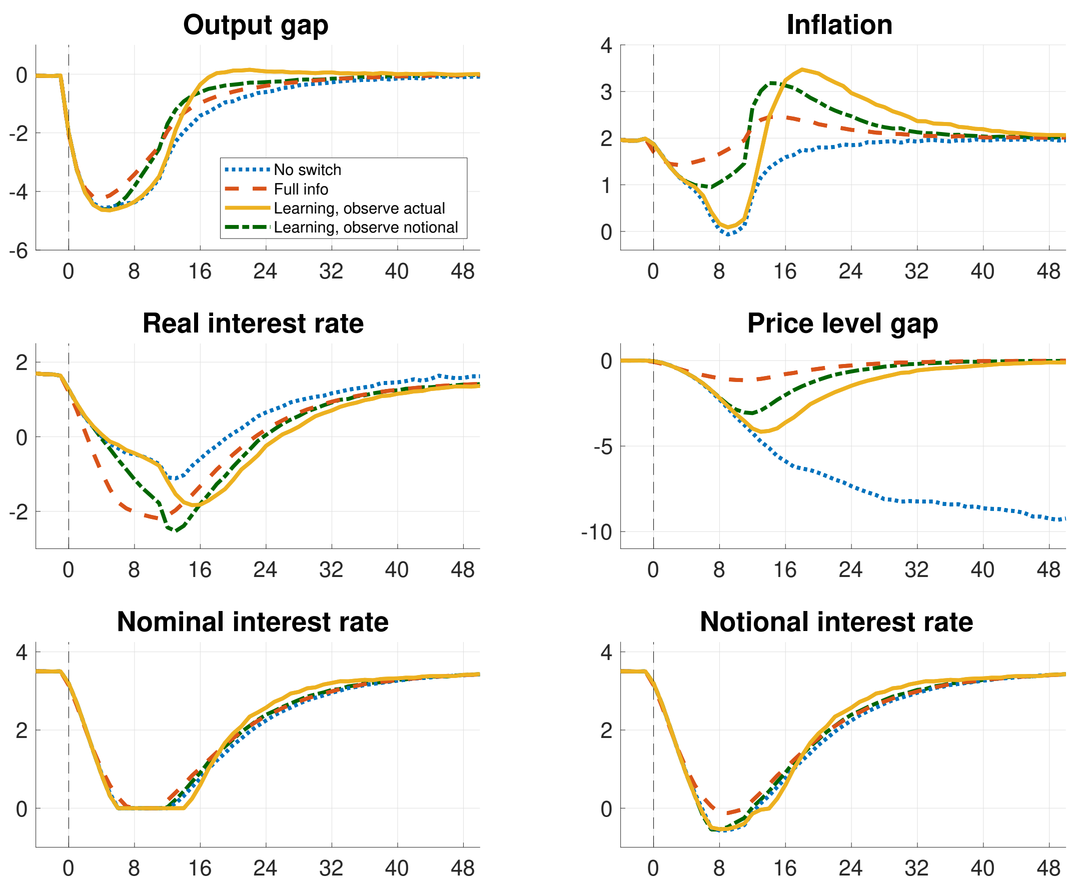


Note: Solid lines show median outcomes and beliefs across 1,000 simulations. Shocks are sampled randomly except for periods $t = 0, \ldots, 11$ during which they are fixed at the sequence described in the text. The rule followed by the central bank switches from inflation to price-level targeting in period $\tau_0 = 0$. Variable definitions are as follows: "Output gap" is $y_t^{gap}$, "Inflation" is $2 + 4\pi_t$, "Real Interest Rate" is $4\left(1/\beta + i_t - \mathbb{E}_t\pi_{t+1}^{(t)}\right)$, "Nominal Interest Rate" is $2{+}4\left(1/\beta + i_t\right)$, "Notional Interest Rate" is $2{+}4\left(1/\beta + i_t^*\right)$, "Price Level Gap" is $p_t^{gap}$.

Figure 5: Beliefs during a deep recession.

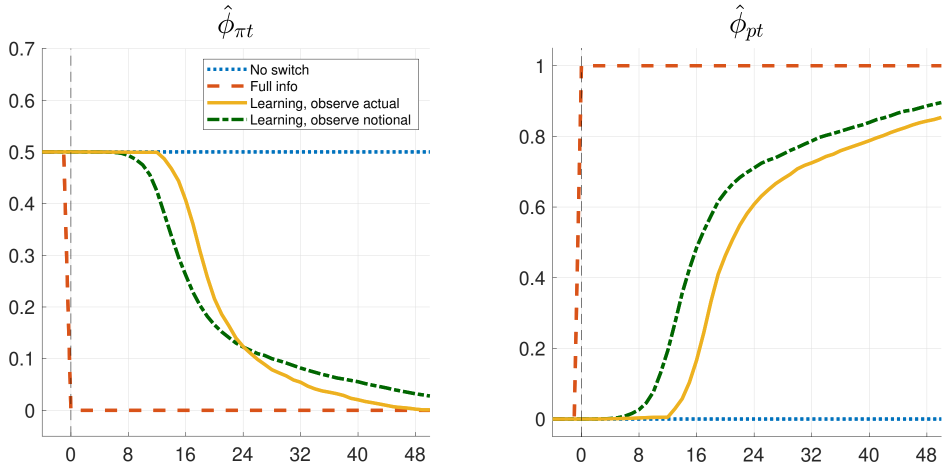


Note: Solid lines show median outcomes and beliefs across 1,000 simulations. Shocks are sampled randomly except for periods $t = 0, \ldots, 11$ during which they are fixed at the sequence described in the text. The rule followed by the central bank switches from inflation to price-level targeting in period $\tau_0 = 0$.

simulations. Compared to the evolution of beliefs in normal times (shown in Figure 1), the median parameter beliefs hardly move towards the true rule parameters in the first 14 quarters after the shock as learning is particularly hampered by the presence of the ZLB. To elaborate on this finding, Figures 4 and 5 include the outcomes and beliefs when agents observe the notional rate $i_t^*$ instead of the policy rate $i_t$ (green lines). As the notional interest rate is not censored and can assume negative values, agents receive more information about the true rule parameters while the actual policy rate is at the ZLB. As a result, beliefs adjust earlier than in the case that agents observe the actual policy rate only. This earlier adjustment of beliefs is sufficient to noticeably stabilize inflation. In other words, the loss of stabilization benefits under learning is greatly amplified by the limited informational content of the actual policy rates when the ZLB is binding.

Overall, these simulations highlight that the effectiveness of flexible price-level targeting depends importantly on the formation of expectations. When agents do not understand the future effects of current policy changes, the commitment to stabilize the price level requires a prolonged period of policy accommodation and high inflation later on without the benefits of closer-to-target inflation and output during the recession. The simulations provide an example of a commitment-based policy that is designed to achieve sizable stabilization benefits by steering expectations, yet may turn out to be undesirable if expectations fail to respond as intended.

## 4 When to Adopt New Policy Strategies?

If a central bank intends to switch from an inflation to a price-level targeting strategy, we advise to do so as early as possible, but not around an episode in which the ZLB becomes or is already binding. To substantiate this recommendation, we vary the timing of the adoption of price-level targeting—before, during, or after the onset of the recession—while keeping fixed the recession scenario introduced above and rank the resulting economic outcomes according to the central bank's loss function. Neither the central bank nor private sector agents are assumed to have any advance information about the recession prior to its realization.

We specify the loss function of the central bank to be:

$$\mathcal{L}_{T_0}^{T_1} = \sum_{s=T_0}^{T_1} \beta^{s-T_0} \left(\pi_s^2 + (y_s^{gap})^2\right). \tag{15}$$

This loss function places equal weights on squared deviations of the inflation rate from its long-run target and of output from its natural level. Period losses are discounted by the factor $\beta$. We measure the discounted loss that occurs between periods $T_0$ and $T_1$, where we set $T_0 = -20$. The sequence of recession shocks starts in period $0$.

Figure 6 plots the value of the loss function under full information and under learning as a function of the timing of adopting price-level targeting $t_0$. The loss under inflation targeting is normalized to $1$. The left panel reports losses that accrue into the infinite future ($T_1 = \infty$) and the right panel considers the losses that accumulate just around the recession scenario ($T_1 = 20$).

Given the parameterization of the policy rules and the loss function, price-level targeting is always preferred to inflation targeting under full information. The benefits of price-level targeting are larger the further in advance of the recession the central bank adopts its new strategy. These benefits are diminished when price-level targeting is adopted after the economy has already fallen into recession.

Figure 6: Expected losses and the timing of introducing price-level targeting.

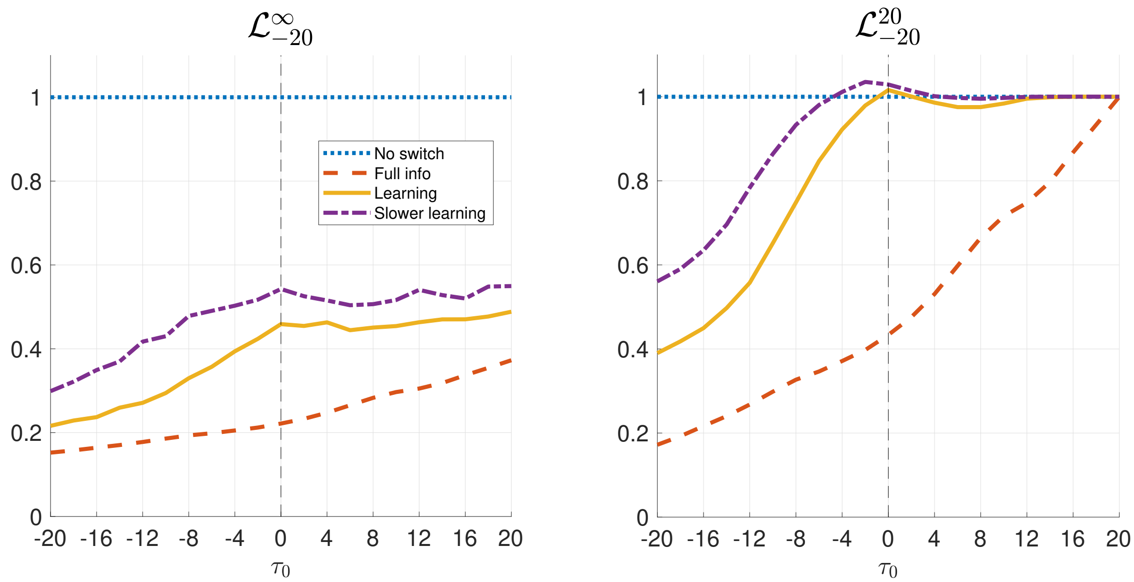


Note: Lines show simulated values of $\mathcal{L}_{-20}^{\infty}$ and $\mathcal{L}_{-20}^{20}$, conditional on the recession scenario starting in $t = 0$ and the rule switching at $t = \tau_0$. 1,000 simulations for each rule switch period $\tau_0$. Losses are normalized to one for the "no switch" case.

Under learning, the same considerations apply, but with the additional challenge that private sector agents do not lower their expectations about future nominal and real interest rates as quickly as under full information. The switch to price-level targeting still reduces the expected loss relative to inflation targeting over the long term regardless of the time of adoption. However, in the near term—spanning periods $T_0 = -20$ to $T_1 = 20$ shown in the right panel of Figure 6—we find that late adoption of the price-level targeting strategy has no advantage over inflation targeting. Hence, in the case of late adoption the advantage of price-level targeting over the long term simply reflects the optimality of price-level targeting at and away from the ZLB given the central bank's loss function. In line with our earlier discussion, during the recession, policy is perceived to be less aggressive in stabilizing prices than it actually is under the new rule and inflation is more volatile. The slower beliefs adjust, the more do the potential benefits of price-level targeting evaporate.

Figure 7: Beliefs and outcomes with earlier adoption of price-level targeting.

(a) Outcomes.

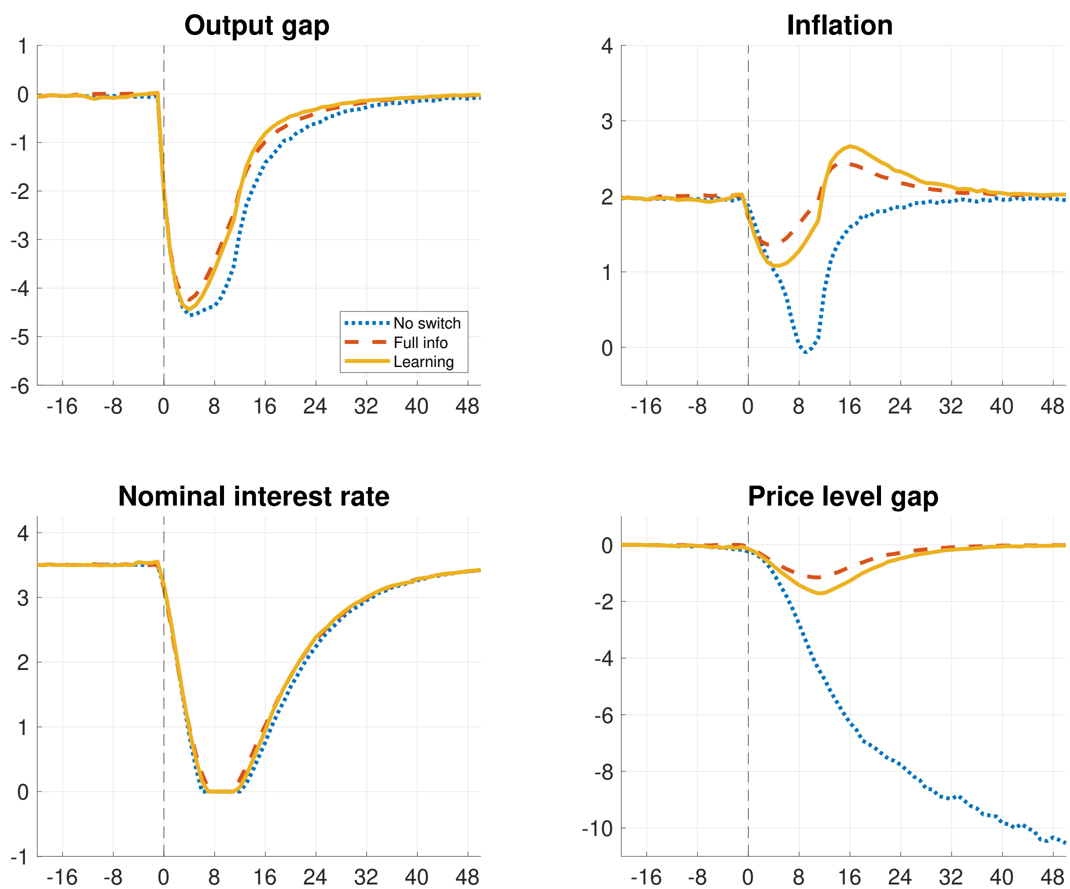


(b) Beliefs.

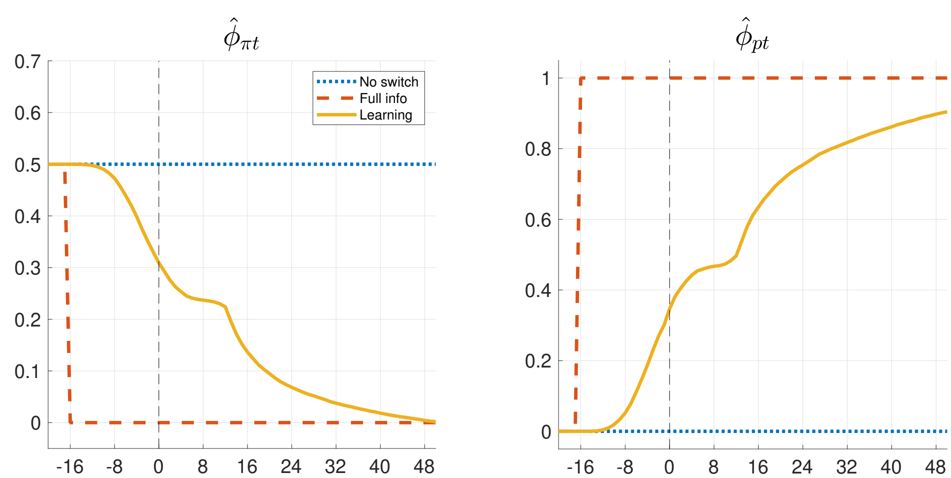


Note: Lines show median outcomes and beliefs across 1,000 simulations. Shocks are sampled randomly except for periods $t = 0, \ldots, 11$ during which they are fixed at the sequence described in the text. The rule followed by the central bank switches from inflation to price-level targeting in period $\tau_0 = -16$. Variable definitions as in Figure 4.

Table 1: Stabilization gains.

| | | (1) ZLB binding | (2) Mean $y^{gap}$ | (3) Mean $\pi$ | (4) s.d. $y^{gap}$ | (5) s.d. $\pi$ | (6) Loss |
|---|---|---|---|---|---|---|---|
| $\mathcal{L}_{-20}^{\infty}$ | $\tau_0 = -\infty$ | 3.19% | -0.23 | 2.00 | 1.23 | 0.99 | 0.151 |
| | $\tau_0 = -16$ | 4.49% | -0.24 | 1.99 | 1.41 | 1.37 | 0.236 |
| | $\tau_0 = 0$ | 5.76% | -0.27 | 1.96 | 1.72 | 1.96 | 0.449 |
| | $\tau_0 = 8$ | 5.86% | -0.30 | 1.90 | 1.73 | 1.95 | 0.452 |
| | $\tau_0 = \infty$ | 10.05% | -0.61 | 1.50 | 2.59 | 3.48 | 1.000 |
| $\mathcal{L}_{-20}^{20}$ | $\tau_0 = -\infty$ | 13.10% | -1.23 | 1.94 | 1.04 | 0.83 | 0.169 |
| | $\tau_0 = -16$ | 19.66% | -1.35 | 1.83 | 1.91 | 2.53 | 0.440 |
| | $\tau_0 = 0$ | 25.29% | -1.75 | 1.27 | 2.91 | 4.03 | 1.013 |
| | $\tau_0 = 8$ | 25.10% | -1.91 | 0.99 | 2.90 | 3.91 | 0.970 |
| | $\tau_0 = \infty$ | 23.16% | -2.05 | 0.79 | 2.90 | 3.93 | 1.000 |

Note: Results based on 1,000 simulations for each rule switch period $\tau_0$ shown. Recession periods are $t = 0, \ldots, 8$. ZLB periods are the fraction of periods across simulations and across time during which the ZLB is binding. $\tau_0 = -\infty$ refers to the full information case in which price-level targeting is in place from the start of each simulation, while $\tau_0 = \infty$ refers to the case in which inflation targeting is in place indefinitely. Simulations are based on the benchmark parameterization of subjective belief uncertainty. Loss function values are normalized to one for the case $\tau_0 = \infty$.

Price-level targeting is more beneficial under learning if adopted well in advance of the recession. When price-level targeting has been in place sufficiently long, private sector agents have had the opportunity to learn the new policy strategy before the recession begins, so that the stabilizing benefits of this strategy come to fruition. Figure 7 shows the evolution of beliefs and outcomes with learning for the case in which the central bank switches in $\tau_0 = -16$. In this case, beliefs have partially adjusted towards the new rule parameters by the onset of the recession. Even with this partial understanding of the policy rule, policymakers already achieve similar outcomes of inflation and output as under full information.

Table 1 reinforces our message about the timing of adoption by showing additional statistics for the benchmark learning case. Early adoption of price-level targeting ($\tau_0 = -16$) yields similar outcomes (means and standard deviations) of inflation and the output gap and losses as the full information case ($\tau_0 = -\infty$), in particular when we consider the long horizon. By contrast, adopting price-level targeting at the onset of a recession ($\tau_0 = 0$) results in greater volatility of inflation and the output gap.

Overall, our results suggest that a central bank planning to switch to a price-level targeting strategy should do so as early as possible, unless it attaches a high probability to a deep recession in the near future. In that case, it can be beneficial to postpone announcing price-level targeting until after the recession is over to avoid being stuck with a commitment to make up a large price-level gap but little additional stabilization of the economy during the recession.

# 5 Temporary Price-Level Targeting

So far, we have focused on the adoption of a permanent price-level targeting strategy, under which the central bank seeks to close the price-level gap regardless of the gap sign and the economic conditions. We now turn to the more state-contingent variant of temporary price-level targeting (TPLT). Under TPLT the central bank only seeks to close the negative price-level gap that stems from a ZLB episode; once this negative gap has been eliminated, the strategy switches back to inflation targeting. Evans (2012) and Bernanke (2017) argue that TPLT can provide the full stabilization benefits of permanent price-level targeting during steep declines of aggregate demand while, at the same time, can help avoiding the potential difficulties associated with communicating to the public that tighter monetary policy is needed to reduce a positive price-level gap. Both these studies assume that the private sector has full information and that the strategy is perfectly credible.

However, the validity of these assumptions seems to be even more questionable in the case of a TPLT than a permanent price-level targeting strategy. Svensson (2019) articulates these concerns by stating that, if price-level targeting strategies "are only applied occasionally and temporarily, economic agents will not be very used to them, and considerable explanation and communication may be necessary. But this may still not be sufficient for the temporary price-level target to be credible, in which case the favorable effect of raised inflation expectations may not occur. Credibility normally needs to be earned, meaning that economic agents need to see the policy put into practice and its principles obeyed, in order to believe that it will be maintained and be successful in the future." Our learning framework directly speaks to Svensson's concern, as agents understand and believe a price-level targeting strategy only once the strategy can be inferred from the observations of the policy rate.

A TPLT strategy differs from a permanent price-level targeting strategy along two dimensions: the definition of the makeup measure and the state-contingent rule parameters. In our formulation of TPLT, the makeup measure accumulates past deviations of inflation from its target since a state-contingent reference period $\tau_0(t)$:

$$z_t = \sum_{s=\tau_0(t)}^{t} \pi_s. \tag{16}$$

The reference period evolves according to:

$$\tau_0(t) = \begin{cases} t & \text{if } i_{t-1} = \underline{i} \text{ and } \max_{\tau_0(t-1)\leq s\leq t-1} z_s \geq 0 \\ \tau_0(t-1) & \text{if } i_{t-1} > \underline{i} \text{ or } \max_{\tau_0(t-1)\leq s\leq t-1} z_s < 0 \end{cases}. \tag{17}$$

Intuitively, the reference period is the last time that the policy rate reached the ZLB. The reference period, and therefore the makeup measure $z_t$, are reset when the policy rate is at the ZLB and the makeup measure has ever turned positive since the previous reference period. By contrast, under permanent price-level targeting the makeup measure is given by the accumulated (positive or negative) price-level gap since a fixed reference period.

We now turn to the weight that the central bank assigns to the makeup measure in its interest rate rule. The policy rule continues to be of the form in equation (13):

$$i_t^* = \rho_i i_{t-1} + (1-\rho_i)\left((1+\phi_{\pi t})\,\pi_t + \phi_{zt}\frac{z_t}{4} + \phi_y \frac{y_t^{gap}}{4}\right) + e_t. \tag{18}$$

The parameters $\phi_{\pi t}$ on inflation and $\phi_{zt}$ on the makeup measure are state-contingent to split the TPLT strategy de facto into an inflation targeting regime and a price-level targeting regime:

$$(\phi_{\pi t}, \phi_{zt}) = \begin{cases} (1,0) & \text{if } i_{t-1} = \underline{i} \text{ or } \max_{\tau_0(t) \leq s \leq t-1} z_s < 0 \\ (0, 0.5) & \text{otherwise} \end{cases}. \quad (19)$$

The central bank assigns positive weight to the makeup measure only in the price-level targeting regime which gets triggered when the policy rate first reaches the ZLB. The regime stays in place until the makeup measure has been made up for. After that, the central bank switches back to the inflation targeting regime.

In line with our previous formulation of beliefs under learning, we assume that agents perfectly observe the makeup measure $z_t$, but do not observe the parameters $\phi_{zt}$ and $\phi_{\pi t}$. The beliefs about the evolution of these parameters are parameterized in the same way as in Section 3.3. In particular, agents think of $\phi_{zt}$ and $\phi_{\pi t}$ as time-varying, but they do not have the knowledge that the parameters follow two discrete regimes.

As for the permanent price-level targeting strategy, we consider the severe demand-driven recession scenario for a TPLT strategy to contrast the performance of the economy under learning with its performance under full information (agents observe the true rule parameters and understand their dependence on the economic conditions). The outcomes and beliefs under learning and full information are shown in Figure 8.

Over the course of the recession, the economic outcomes under TPLT are virtually the same as under the permanent price-level targeting strategy. Initially, the anticipated stabilization benefits of the strategy do not materialize because agents require time to learn the new strategy, in particular while the policy rate is constrained by the ZLB. Consequently agents fail to anticipate the more accommodative path of monetary policy in the future. The cumulative shortfall in inflation is larger under learning and, as a result thereof, the central bank stays in the price-level targeting regime much longer than under full information.

Once the economy has recovered sufficiently, the central bank returns to its inflation targeting regime and the rule parameters switch yet again as shown in the bottom panels of Figure 8. Under learning this change in parameters initiates a new adjust-

Figure 8: Beliefs and outcomes with temporary price-level targeting.

(a) Outcomes.

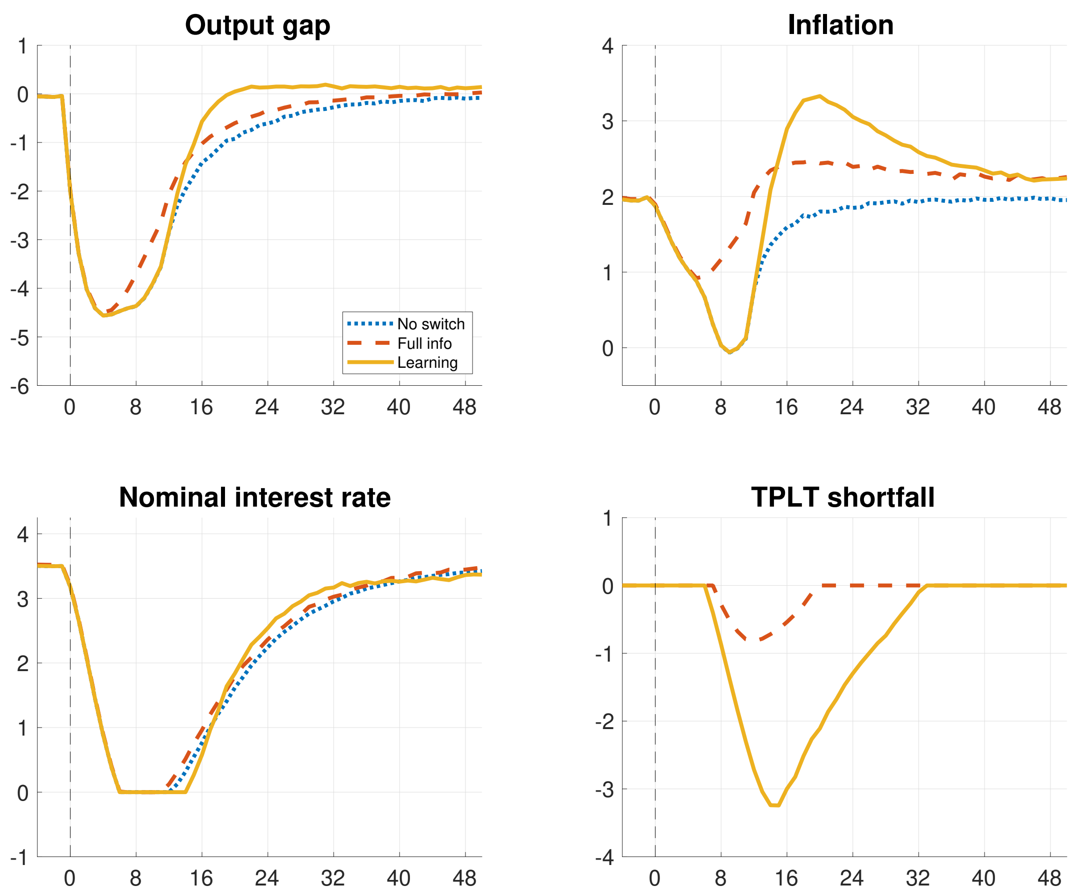


(b) Beliefs.

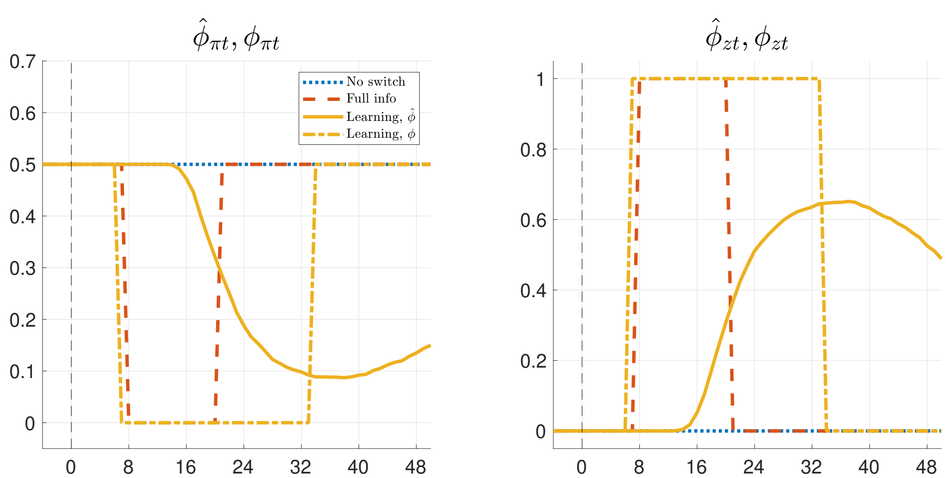


Note: Solid lines show median outcomes and beliefs across 1,000 simulations. The dash-dotted line in the lower panel additionally plots the median actual rule parameters $\phi_{\pi t}$ and $\phi_{zt}$ under learning. The simulations are initialized at the ergodic distribution of outcomes obtained under the inertial Taylor rule, and in period $t = 0$ the central bank starts following the TPLT strategy (see text). The starting values for the rule inputs are $z_t = 0$ and $\tau_0(0) = -1$. Variable definitions as in Figure 4, except for "TPLT shortfall" defined as $\min\{z_t, 0\}$.

ment process for the agents. Over the course of the price-level targeting regime, agents' beliefs partially adjusted to the true parameters under the price-level targeting regime. But upon the central bank's return to the inflation targeting regime of the TPLT strategy, agents will reverse their beliefs and move their parameter estimates to be yet again closer to the parameters of the inflation targeting regime. The beliefs and the direction in which the parameter estimates move over time will again be reversed on the next occasion the central bank is in the price-level targeting regime during a ZLB episode. In our learning formulation, agents will never understand the state-contingent nature of the TPLT strategy. As a result, and in contrast to permanent price-level targeting, agents will never be in the position to correctly anticipate the central bank's policy actions. In particular at the ZLB, the TPLT strategy will never be as effective in stabilizing the economy as the permanent price-level targeting strategy (which agents will come to fully understand over time).

The ineffectiveness of TPLT is in part the result of our assumptions about the beliefs that agents can entertain. In particular, agents cannot entertain the idea of regime switches in the policy rules embedded in the TPLT strategy. Yet, even if we allowed agents to consider the possibility of switches between two regimes, it would still be difficult for agents to learn the TPLT strategy. The simple reason is that agents cannot infer anything about the regime the economy is currently not in. Before a ZLB episode occurs, agents have no opportunity to learn about the central bank's likely behavior during that episode; at the ZLB, there is virtually no information that allows to discriminate between different rules; and the period of the recovery during which the central bank still follows price-level targeting is short. Depending on the speed of learning, it would presumably take several zero-lower bound episodes before agents would fully understand the contingent behavior of the central bank. Moreover, agents would also have to learn the conditions that trigger the switch from one regime to the other, further complicating the inference problem relative to a permanent price-level targeting strategy.

# 6 Conclusion

We have developed a method of learning about the central bank's policy strategy from observed policy rates that explicitly takes into account the limited informational content of observed policy rates at the ZLB. We have applied this method to a simple New-Keynesian model in which the central bank can pursue either an inflation targeting or price-level targeting strategy.

When the central bank switches to price-level targeting at the onset of a deep recession, this switch mitigates the loss in output and the shortfall in inflation under rational expectations and full information, as is well known. But when agents are learning, the benefits of price-level targeting do not materialize because agents do not understand the new policy regime immediately. The learning problem is further complicated by the fact that the policy rate quickly hits the ZLB, at which point agents receive little information about the true parameters of the policy rule. As a result, under learning, the central bank is left with a much larger negative price-level gap than under full information, and thus has to allow for substantial overshooting of inflation after the recession without having accrued any stabilization benefits in the midst of the recession. In order for these benefits to materialize, price-level targeting should be introduced in calm times to give agents the opportunity to learn this new policy strategy rather than being deployed as a policy tool in a deep recession. Temporary price-level targeting strategies are likely to be much less effective than their permanent counterparts.

# References

**Ambler, Steve**, "Price-Level Targeting And Stabilisation Policy: A Survey," *Journal of Economic Surveys*, December 2009, *23* (5), 974–997.

**Angeletos, George-Marios and Chen Lian**, "Forward guidance without common knowledge," *American Economic Review*, 2018, *108* (9), 2477–2512.

**Bernanke, Ben S.**, "Monetary Policy in a New Era," Technical Report, prepared for the conference on Rethinking Macroeconomic Policy, Peterson Institute, Washington, DC, October 12-13, 2017, Brookings Institution 2017.

_ **, Michael T. Kiley, and John M. Roberts**, "Monetary Policy Strategies for a Low-Rate Environment," Finance and Economics Discussion Series 2019-009, Board of Governors of the Federal Reserve System (US) February 2019.

**Bodenstein, Martin, James Hebden, and Ricardo Nunes**, "Imperfect credibility and the zero lower bound," *Journal of Monetary Economics*, 2012, *59* (2), 135–149.

**Boivin, Jean**, "Has U.S. Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data," *Journal of Money, Credit and Banking*, August 2006, *38* (5), 1149–1173.

**Brayton, Flint, Thomas Laubach, and David L. Reifschneider**, "The FRB/US Model : A Tool for Macroeconomic Policy Analysis," FEDS Notes 2014-04-03, Board of Governors of the Federal Reserve System (U.S.) April 2014.

**Cogley, Timothy and Thomas J. Sargent**, "Anticipated Utility And Rational Expectations As Approximations Of Bayesian Decision Making," *International Economic Review*, February 2008, *49* (1), 185–221.

_ **, Christian Matthes, and Argia M. Sbordone**, "Optimized Taylor rules for disinflation when agents are learning," *Journal of Monetary Economics*, 2015, *72* (C), 131–147.

**Eggertsson, Gauti B and Michael Woodford**, "Policy options in a liquidity trap," *American Economic Review*, 2004, *94* (2), 76–79.

**English, William B., William R. Nelson, and Brian P. Sack**, "Interpreting the Significance of the Lagged Interest Rate in Estimated Monetary Policy Rules," *The B.E. Journal of Macroeconomics*, April 2003, *3* (1), 1–18.

**Erceg, Christopher J. and Andrew T. Levin**, "Imperfect credibility and inflation persistence," *Journal of Monetary Economics*, May 2003, *50* (4), 915–944.

**Eusepi, Stefano and Bruce Preston**, "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, March 2018, *56* (1), 3–59.

**Evans, Charles L.**, "Monetary Policy in a Low-Inflation Environment: Developing a State-Contingent Price-Level Target," *Journal of Money, Credit and Banking*, 2012, *44* (s1), 147–155.

**Evans, George W. and Seppo Honkapohja**, *Learning and Expectations in Macroeconomics* Princeton University Press, Princeton: Princeton University Press, 2001.

**Farhi, Emmanuel and Iván Werning**, "Monetary Policy, Bounded Rationality, and Incomplete Markets," NBER Working Papers 23281, National Bureau of Economic Research, Inc 2017.

**Gabaix, Xavier**, "A Behavioral New Keynesian Model," CEPR Discussion Papers 11729, C.E.P.R. Discussion Papers December 2016.

**Gust, Christopher J., Edward Herbst, and David López-Salido**, "Forward Guidance with Bayesian Learning and Estimation," Finance and Economics Discussion Series 2018-072, Board of Governors of the Federal Reserve System (US) October 2018.

**Hatcher, Michael and Patrick Minford**, "Stabilisation Policy, Rational Expectations And Price-Level Versus Inflation Targeting: A Survey," *Journal of Economic Surveys*, April 2016, *30* (2), 327–355.

**Hebden, James and David López-Salido**, "From Taylor's Rule to Bernanke's Temporary Price Level Targeting," Finance and Economics Discussion Series 2018-051, Board of Governors of the Federal Reserve System (US) July 2018.

**Holden, Tom**, "Computation of solutions to dynamic models with occasionally binding constraints," Working paper 2016.

**Kreps, David**, "Anticipated Utility and Dynamic Choice, 1997 Schwartz Lecture," in Donald P. Jacobs, Ehud Kalai, and Morton Kamien, eds., *Frontiers of Research in Economic Theory*, Cambridge, England: Cambridge University Press, 1998.

**Kryvtsov, Oleksiy, Malik Shukayev, and Alexander Ueberfeldt**, "Adopting Price-Level Targeting under Imperfect Credibility: An Update," Staff Working Papers 08-37, Bank of Canada 2008.

**Mele, Antonio, Krisztina Molnar, and Sergio Santoro**, "On the perils of stabilizing prices when agents are learning," Discussion Paper Series in Economics 22/2018, Norwegian School of Economics, Department of Economics October 2018.

**Mertens, Thomas M. and John C. Williams**, "Monetary Policy Frameworks and the Effective Lower Bound on Interest Rates," *AEA Papers and Proceedings*, May 2019, *109*, 427–432.

**Negro, Marco Del, Marc Giannoni, and Christina Patterson**, "The forward guidance puzzle," Staff Reports 574, Federal Reserve Bank of New York 2012.

**Reifschneider, David and John C. Williams**, "Three Lessons for Monetary Policy in a Low-Inflation Era," *Journal of Money, Credit and Banking*, 2000, *32* (4), 936–966.

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

**Svensson, Lars E O**, "Price-Level Targeting versus Inflation Targeting: A Free Lunch?," *Journal of Money, Credit and Banking*, August 1999, *31* (3), 277–295.

**Svensson, Lars E.0.**, "Monetary Policy Strategies for the Federal Reserve," Technical Report, Federal Reserve Bank of Chicago June 4-5 2019. prepared for the conference Monetary Policy Strategy, Tools, and Communication Practices-A Fed Listens Event.

**Tetlow, Robert and Peter von zur Muehlen**, "Simplicity versus optimality: The choice of monetary policy rules when agents must learn," *Journal of Economic Dynamics and Control*, 2001, *25* (1-2), 245–279.

**Vestin, David**, "Price-level versus inflation targeting," *Journal of Monetary Economics*, October 2006, *53* (7), 1361–1376.

**Woodford, Michael**, “Monetary Policy Analysis When Planning Horizons Are Finite,” in “NBER Macroeconomics Annual 2018, volume 33” NBER Chapters, National Bureau of Economic Research, Inc, June 2018.

# A Details on the solution algorithm

This appendix describes the numerical algorithm used to compute the learning equilibrium in Section (2).

1. Start with a prior for $\beta_{t-1}$ that is normally distributed as $\mathcal{N}\left(\hat{\beta}_{t-1}, P_{t-1}\right)$.

2. We compute $x_t$ as a function of $\tilde{e}_t$, $u_t$, $x_{t-1}$ and $\hat{\beta}_{t-1}$: $x_t = f\left(x_{t-1}, u_t, \tilde{e}_t, \hat{\beta}_{t-1}\right)$. In particular, we augment equation (1) with anticipated shocks to the policy rule equation. Following Holden (2016), we use a mixed-integer linear programming solver to determine the sequence of anticipated shocks such that the max operator in equation (2)will hold period-by-period as projected under perfect foresight.

3. Find $(x_t, i_t, \tilde{e}_t)$ as the solution to the system of equations:

$$\tilde{e}_t = \begin{cases} i_t - \Psi(\hat{\beta}_{t-1})x_t & \text{if } i_t > \underline{i} \\ -\sigma_{et}\dfrac{\phi\left(\frac{\underline{i}-\Psi(\hat{\beta}_{t-1})x_t}{\sigma_{et}}\right)}{\Phi\left(\frac{\underline{i}-\Psi(\hat{\beta}_{t-1})x_t}{\sigma_{et}}\right)} & \text{if } i_t = \underline{i} \end{cases} \tag{A.1}$$

$$x_t = f\left(x_{t-1}, u_t, \tilde{e}_t, \hat{\beta}_{t-1}\right) \tag{A.2}$$

$$i_t = \max\left\{\underline{i}, \Psi(\hat{\beta}_{t-1})x_t + e_t\right\}. \tag{A.3}$$

4. Obtain a posterior for $\beta_t$ through the filtering problem (6)–(8). Even though we take $x_t$ as exogenous, the non-linearity stemming from the ZLB and from the potential non-linearity of $\Psi(\cdot)$ make this a non-linear filtering problem. To avoid having to use a numerically expensive particle filter, we make some numerical approximations. First, we approximate the non-linearity from $\Psi(\cdot)$ by taking a first-order Taylor expansion of the notional rate around $\hat{\beta}_{t-1}$:

$$i_t = \max\{\underline{i}, i_t^*\} \tag{A.4}$$

$$i_t^* \approx \Psi(\hat{\beta}_{t-1})x_t + \beta_t'\frac{\partial\Psi}{\partial\beta}(\hat{\beta}_{t-1})x_t + e_t,\ e_t \sim \mathcal{N}\left(0, \sigma_{et}^2\right) \tag{A.5}$$

$$\beta_t = \beta_{t-1} + \epsilon_{\beta t},\ \beta_{t-1} \sim \mathcal{N}\left(\hat{\beta}_{t-1}, P_{t-1}\right),\ \epsilon_{\beta t} \sim \mathcal{N}\left(0, \Sigma_{\beta t}\right) \tag{A.6}$$

Note that, in our application to price-level targeting in the paper, $\Psi(\cdot)$ is already linear, so the above is an equality rather than an approximation.

We will work with the systematic part of the notional rate, which we denote by $s_t = i_t^* - e_t$. The prior of $s_t$ given $x_t$ is normally distributed:

$$\mathbb{E}[s_t \mid x_t] = m_t = \Psi(\hat{\beta}_{t-1})x_t \tag{A.7}$$

$$\mathbb{V}[s_t \mid x_t] = S_t = H_t' (P_{t-1} + \Sigma_{\beta t}) H_t \tag{A.8}$$

$$\text{where } H_t = \frac{\partial \Psi}{\partial \beta}(\hat{\beta}_{t-1})x_t. \tag{A.9}$$

To get to the posterior of $\beta_t$ after observing $i_t$, we have to distinguish whether the ZLB is binding or not.

(a) If $i_t > \underline{i}$, the filtering problem (A.4)–(A.6) reduces to the extended Kalman filter (EKF) and the posterior is normally distributed as $\mathcal{N}\left(\hat{\beta}_t, P_t\right)$. The filtering equations are standard:

$$K_t = \frac{(P_{t-1} + \Sigma_{\beta t}) H_t}{S_t + \sigma_{et}^2} \tag{A.10}$$

$$\hat{\beta}_t = \hat{\beta}_{t-1} + K_t (i_t - m_t) \tag{A.11}$$

$$P_t = P_{t-1} + \Sigma_{\beta t} - K_t \left(S_t + \sigma_{et}^2\right) K_t'. \tag{A.12}$$

If $\Psi$ is linear, then the EKF is just the standard Kalman filter and we have found an exact solution to the posterior.

(b) If $i_t = \underline{i}$, we compute the mean and the variance of the posterior of $s_t$ given $x_t$ and the observation that $s_t + e_t \leq \underline{i}$. For an arbitrary integrable function $g$, we have that the posterior mean of $g(s_t)$ is given by:

$$
\begin{aligned}
\mathbb{E}\left[g(s_t) \mid i_t = \underline{i}, x_t\right] &= \mathbb{E}\left[g(s_t) \mid x_t, s_t \leq \underline{i} - e_t\right] \\
&= \frac{\mathbb{E}\left[g(s_t)\,\mathbb{1}\left\{s_t \leq \underline{i} - e_t\right\} \mid x_t\right]}{\mathbb{P}\left(s_t \leq \underline{i} - e_t \mid x_t\right)} \\
&= \frac{\mathbb{E}\left[g(s_t)\,\mathbb{E}\left[\mathbb{1}\left\{s_t \leq \underline{i} - e_t\right\} \mid s_t, x_t\right] \mid x_t\right]}{\mathbb{P}\left(s_t \leq \underline{i} - e_t \mid x_t\right)} \\
&= \mathbb{E}\left[g(s_t)\,\frac{\mathbb{P}\left(e_t \leq \underline{i} - s_t \mid s_t, x_t\right)}{\mathbb{P}\left(s_t + e_t \leq \underline{i} \mid x_t\right)} \mid x_t\right] \\
&= \int_{-\infty}^{\infty} g(s)\,\frac{\Phi\left(\frac{\underline{i}-s}{\sigma_{et}^2}\right)}{\Phi\left(\frac{\underline{i}-m_t}{S_t+\sigma_{et}^2}\right)}\,\frac{1}{\sqrt{2\pi}S_t}e^{\frac{(s-m_t)^2}{2S_t^2}}\,ds. \qquad \text{(A.13)}
\end{aligned}
$$

We compute these expressions using Gaussian quadrature for $g(s) = s$ and $g(s) = s^2$ to obtain the posterior mean and variance of $s_t$, which we denote by $\tilde{m}_t$ and $\tilde{S}_t$. We now approximate the posterior distribution of $s_t$ as $\mathcal{N}\left(\tilde{m}_t, \tilde{S}_t\right)$. With this approximation, the posterior for $\beta_t$ given $x_t$ and $i_t = \underline{i}$ is normally distributed as $\mathcal{N}\left(\hat{\beta}_t, P_t\right)$, with the updating formula:

$$K_t = \frac{\left(P_{t-1} + \Sigma_{\beta t}\right) H_t}{S_t} \qquad \text{(A.14)}$$

$$\hat{\beta}_t = \hat{\beta}_{t-1} + K_t\left(\tilde{m}_t - m_t\right) \qquad \text{(A.15)}$$

$$P_t = P_{t-1} + \Sigma_{\beta t} - K_t\left(S_t - \tilde{S}_t\right) K_t'. \qquad \text{(A.16)}$$