DEUTSCHE
BUNDESBANK
EUROSYSTEM

Discussion Paper

Deutsche Bundesbank
No 34/2019

Expectations formation, sticky prices, and the ZLB

Betsy Bersson
(Duke University)
Patrick Hürtgen
(Deutsche Bundesbank)
Matthias Paustian
(Federal Reserve Board of Governors)

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet http://www.bundesbank.de

ISBN 978–3–95729–625–2 (Printversion)
ISBN 978–3–95729–626–9 (Internetversion)

Non-technical summary

Research Question

In the aftermath of the financial crisis, central banks have widely used communication about the future path of monetary policy — forward guidance (FG) — as an instrument to compensate for the inability to move the policy rate in the near term. Against this background, this paper addresses three key policy questions. First, how should monetary policy optimally be conducted under commitment when an adverse demand shock drives the economy to the zero lower bound (ZLB) on nominal interest rates for an extended period? Second, what are the macro effects of delaying the liftoff from the ZLB as in a time-dependent FG policy? Third, how large are fiscal multipliers when monetary policy is at the lower bound?

Contribution

We introduce level-k thinking that is consistent with the micro evidence on expectations formation into a sticky-price model. In that framework, individuals iteratively update their beliefs about the future. At level-1, agents calculate the direct effect of a particular policy intervention holding fixed beliefs about future variables. At higher levels, those beliefs are updated from the sequence of economic outcomes that would be obtained period by period given the direct effect of the intervention and the previous level of beliefs. In contrast to the empirical evidence that has found low levels of k, most model-based analysis typically considers the case that this iterative process goes to infinity ("rational expectations").

Results

The key qualitative prescription from optimal monetary policy under commitment at the ZLB derived in rational expectations models continues to hold with expectations formation according to level-k thinking albeit with a reduced macroeconomic benefit. Under rational expectations two puzzling features arise: (i) a time-dependent FG policy can have implausibly large effects on inflation and output and (ii) if such a delay lasts for too long, the policy can contract rather than stimulate the economy. We resolve both of these puzzles when using level-k thinking. Similarly, we also show that the macroeconomic stimulus from increases in government expenditures under constant interest rates is smaller under level-k thinking than with fully rational expectations. Thus, our main finding of more attenuated macroeconomic effects of policy stimulus with bounded rationality carries over to fiscal policy.

Nichttechnische Zusammenfassung

Fragestellung

Nach der Finanzkrise sind viele Zentralbanken dazu übergegangen Kommunikation über die zukünftige Ausrichtung der Geldpolitik – forward guidance (FG) – als Instrument zu verwenden, insbesondere wenn der Politikzins kurzfristig nicht verwendet werden kann. Vor diesem Hintergrund adressiert das vorliegende Papier drei zentrale Politikfragen. Erstens, wie soll optimale Geldpolitik unter commitment gestaltet werden, wenn ein adverser Nachfrageschock die Ökonomie für längere Zeit an die Nullzinsgrenze bindet? Zweitens, was sind die makroökonomischen Effekte einer Verzögerung der Zinsanhebung von der Nullzinsgrenze, wie bei einer zeitpfadabhängigen FG-Politik? Drittens, wie hoch sind die Fiskalmultiplikatoren, wenn die Geldpolitik an der Nullzinsgrenze ist?

Beitrag

In einem neukeynesianischen Modell führen wir Level-k Denken ein. Dies ist konsistent mit der Mikrodatenevidenz zur Erwartungsbildung. Level-k Denken bedeutet, dass Individuen ihre Erwartungen über die Zukunft iterativ bilden. Bei Level-1 berechnen Individuen den direkten Effekt einer Politikintervention und halten ihre Erwartungen über die zukünftigen Auswirkungen konstant. Bei weiteren Iterationen werden die Erwartungen basierend auf dem direkten Effekt der Intervention und der vorangegangenen Runde der Erwartungsbildung formiert. Entgegen der empirischen Evidenz, die geringe Werte für k findet, wird in der modellbasierten Analyse in der Regel der Fall unterstellt, dass der iterative Prozess gegen unendlich strebt ("Rationale Erwartungen").

Ergebnisse

Die Politikempfehlung für optimale verpflichtende Geldpolitik an der Nullzinsgrenze in rationalen Erwartungsmodellen hält auch unter Level-k Denken, allerdings mit einem abgeschwächten makroökonomischen Nutzen. Unter rationalen Erwartungen treten zwei unplausible Effekte auf: (i) eine zeitpfadabhängige FG Politik hat außergewöhnlich starke Effekte auf Inflation und das BIP und (ii) wenn eine Nullzinspolitik lange verzögert wird, kann die Politik eher kontraktiv als expansiv auf die Ökonomie wirken. Mit Level-k Denken treten beide Effekte nicht auf. Zudem ist eine expansive Fiskalpolitik mit konstanten Zinsen geringer unter Level-k Denken als mit rationalen Erwartungen. Daher überträgt sich unser Ergebnis von abgeschwächten makroökonomischen Effekten einer expansiven Politik auch auf die Fiskalpolitik.

Bundesbank Discussion Paper No 34/2019

Expectations formation, sticky prices, and the ZLB*

Betsy Bersson[†], Patrick Hürtgen[‡], and Matthias Paustian[§]

Abstract

At the zero lower bound (ZLB), expectations about the future path of monetary or fiscal policy are crucial. We model expectations formation under level-k thinking, a form of bounded rationality introduced by García-Schmidt and Woodford (2019) and Farhi and Werning (2017), consistent with experimental evidence. This process does not lead to a number of puzzling features from rational expectations models, such as the forward guidance and the reversal puzzle, or implausible large fiscal multipliers. Optimal monetary policy at the ZLB under level-k thinking prescribes keeping the nominal rate lower for longer, but short-run macroeconomic stabilization is less powerful compared to rational expectations.

Keywords: expectations formation, optimal monetary policy, New Keynesian model, zero lower bound, forward guidance puzzle, reversal puzzle, fiscal multiplier.

JEL classification: E32.

*The views expressed in this paper are those of the authors, and not necessarily those of the Board of Governors of the Federal Reserve System, its staff or the Deutsche Bundesbank. We would like to thank for comments Klaus Adam, Jeff Campbell, Tom Holden, David Lebow, Taisuke Nakata, Alexander Scheer, Sebastian Schmidt, Todd Walker, as well as seminar and conference participants at the Bank of Canada, the European Central Bank, the Federal Reserve Bank of Cleveland 2019 Inflation Conference, the Deutsche Bundesbank, the Federal Reserve Board, the 2019 Conference on Macroeconomic Modelling and Model Comparison in Frankfurt, and the EEA 2019 Annual Congress. All errors are our own.

[†]Duke University, betsybersson@gmail.com.

[‡]Deutsche Bundesbank patrick.huertgen@bundesbank.de.

[§]Federal Reserve Board of Governors, matthias.o.paustian@frb.gov.

"With nominal short-term interest rates at or close to their effective lower bound in many countries, the broader question of how expectations are formed has taken on heightened importance." — Janet Yellen (2016)

1 Introduction

Most model-based analysis of monetary policy assumes that agents have rational expectations. However, the empirical literature documents that the expectations formation of firms and households is hard to reconcile with rational expectations; see, for instance, Coibion and Gorodnichenko (2012, 2015); Coibion, Gorodnichenko, and Kamdar (2018). Motivated by this discrepancy, this paper studies macroeconomic dynamics in a model of bounded rationality. Our focus is on the zero lower bound (ZLB) in sticky-price models. It is well known that under rational expectations in a ZLB environment, expectations about outcomes in the distant future can be crucial for economic dynamics along the entire path. Furthermore, in the aftermath of the financial crisis, central banks have widely used communication about the future path of monetary policy as an instrument to compensate for the inability to move the policy rate in the near term. Any assessment of the effects of such communication should take into account how expectations are actually formed as stressed in the quote above by former Federal Reserve Chair Yellen (2016).

Our analysis focuses on ZLB scenarios and policy interventions with which households and firms have little experience. In these scenarios, forming model-consistent beliefs about the macroeconomic outcomes is naturally much more challenging than in normal times. We assume that expectations are formed according to level-k thinking as in García-Schmidt and Woodford (2015, 2019) and Farhi and Werning (2017). In that framework, individuals iteratively update their beliefs about the future starting out from a baseline belief (referred to as level-0). At level-1, agents then calculate the direct effect of a particular policy intervention holding fixed beliefs about future variables at the baseline. At higher levels, those beliefs are then updated from the sequence of economic outcomes that would obtain period by period given the direct effect of the intervention and the previous level of beliefs about the future.[1] When this updating process is stable, the equilibrium converges to rational expectations.

For policy advice in a temporary ZLB environment, this type of bounded rationality analysis is likely to be more appropriate than rational expectations, because one cannot merely assume that agents' expectations will approximate rational expectations. The typical justification that rational expectations equilibria are learnable need not apply when agents have spent little time in a liquidity trap and are facing new policy interventions

[1] For any given set of beliefs, outcomes are a temporary equilibrium in the sense of Grandmont (1977).

such as forward guidance. In this paper, we mainly focus on low levels of belief revisions, in line with the experimental evidence (see, for example, Nagel, 1995; Camerer, Ho, and Chong, 2004; Mauersberger and Nagel, 2018; Coibion, Gorodnichenko, Kumar, and Ryngaert, 2018). We are primarily interested in the short-run effects of policy interventions. Consequently, we abstract from the question of how agents adjust expectations in response to forecast errors, which is clearly a key question for the longer run effects of permanent changes in policy.

We address three key policy questions for which communication about the future path of monetary or fiscal policy, and hence expectations formation, is crucial. These questions have received a large amount of attention in the context of rational expectations models.[2] First, how should monetary policy optimally be conducted under commitment when an adverse demand shock drives the economy to the ZLB for an extended period? Second, what are the macroeconomic effects of delaying the liftoff from the ZLB as in a time-dependent forward-guidance policy? Third, how large are fiscal multipliers when monetary policy is at the lower bound? Our main findings are summarized below.

1. Optimal monetary policy under commitment at the ZLB:

How should monetary policy optimally be conducted when the private sector forms beliefs according to level-k thinking? To answer this question, we study the optimal commitment policy in the context of a demand shock that brings the economy to the ZLB. Our core finding is that the central bank optimally chooses to lift off from the ZLB later under level-k thinking than with rational expectations and hence provides more forward guidance. Our analysis therefore shows that the pure fact that agents are not fully rational does not imply that forward-guidance policies are inappropriate. On the contrary, in the context of our model, the central bank optimally chooses to deliver more stimulus in the future with level-k thinking than with rational expectations and hence lifts off from the ZLB later, because that future stimulus has a smaller impact on near-term economic outcomes than under rational expectations. Quantitatively, we find that the central bank is able to deliver less macroeconomic stabilization at the ZLB under optimal commitment policy with level-k thinking than with rational expectations.

We also examine the case where a central bank minimizes a loss function with equal weights on inflation and the output gap as opposed to the welfare-based loss function in

[2]or related analysis regarding optimal monetary policy at the ZLB under commitment, see Eggertsson and Woodford (2003). For related analysis regarding the effects of a delayed liftoff (including the so-called reversal puzzle), see Carlstrom, Fuerst, and Paustian (2015) or Del Negro, Giannoni, and Patterson (2012). For related analysis regarding the fiscal multiplier at the ZLB, see Christiano, Eichenbaum, and Rebelo (2011).

our baseline case. This choice is motivated by a dual mandate. Under rational expectations, we find that the central bank optimally raises the nominal interest rate off the ZLB earlier than under a simple Taylor rule and prescribes a higher path for nominal interest rates thereafter. Despite higher nominal interest rates, this policy implies lower real interest rates than under a Taylor rule and hence improves macro outcomes. Clearly, this policy prescription crucially relies on rational expectations, and one would expect it to be fragile once bounded rationality is introduced. We show that this is indeed the case. In particular, the path for the nominal rate that is optimal if the private sector forms expectations according to level-k in fact delays the liftoff relative to the Taylor rule and prescribes a lower path for nominal interest rates thereafter. This arises naturally, because inflation expectations adjust sluggishly with level-k thinking in response to changes in the nominal rates. Hence, the central bank cannot rely on the immediate adjustment of inflation expectations when announcing its path for the nominal rate.

2. The effects of time-dependent forward guidance at the ZLB:

We study the New Keynesian model with endogenous inflation persistence and consider an adverse demand shock that brings the economy to the ZLB for an extended period. In that context, we examine the effects of time-dependent forward guidance, which we model as a policy of delaying the liftoff from the ZLB by an exogenous number of quarters before returning to the Taylor rule. As shown in Carlstrom et al. (2015), such a policy can be extremely powerful in rational expectations models. That is, delaying the liftoff by only a few quarters can more than offset the effect of the adverse demand shock and the ZLB. This is one manifestation of the so-called forward-guidance puzzle. Under rational expectations, the delayed liftoff creates a powerful feedback loop between expectations of high inflation in the future, the resulting improvement in inflation during the spell at the ZLB, and even higher inflation in the future due to inflation persistence. Under level-k thinking, this feedback loop is much more muted because agents' limited degree of reflection dampens the process. Consequently, actual inflation, ex-ante real interest rates, and real activity are improved to a much smaller extent by forward guidance.

As pointed out by Carlstrom et al. (2015), under some conditions, a policy of delaying the liftoff can result in the so-called reversal puzzle under rational expectations. The reversal puzzle is the implication that extending the length of the interest rate peg beyond some critical duration results in a fall in inflation and activity instead of an increase. Our key finding is that the reversal puzzle does not arise with level-k thinking. That is, while a delay in the liftoff by too many quarters results in a counterintuitive deterioration of the macroeconomy under rational expectations, the same policy results in an improvement under level-k thinking. Furthermore, in this setting, the solution under level-k thinking does not converge to the counterintuitive rational expectations solutions as the level of thinking k grows. Convergence is achieved for small enough delays for which forward guidance has the conventional effects. Hence, our analysis provides a basis for discounting the reversal puzzle based on implausible expectations formation.

3. The fiscal multiplier under constant interest rates:

The literature has pointed out that the fiscal multiplier under constant interest rates can be large.[3] As in Christiano et al. (2011), we consider a coordinated policy experiment: a fiscal stimulus and an interest rate peg, both of which extend into the future jointly with a certain probability p. As shown in Carlstrom, Fuerst, and Paustian (2014), rational expectations about events in the far future are a key factor contributing to large fiscal multipliers in this context. Hence, it is natural to examine how sensitive the fiscal multiplier is to deviations from the rational expectations assumption.

We consider a baseline calibration where the rational expectations fiscal multiplier is very large. We find that the level-k multiplier for low levels of k is only modestly larger than unity. Convergence to the rational expectations multipliers is extremely slow.[4] However, for an alternative parameterization where the fiscal multiplier under rational expectations is more modest, convergence of level-k thinking to rational expectations is fairly fast.[5]

Recent work by Mertens and Ravn (2014) has shown that the effect of government spending is different in a non-fundamental liquidity trap than in a fundamentals-driven one. In the non-fundamentals-driven ZLB episode, government spending multipliers are always smaller than unity, and inflation falls rather than rises with additional spending. That prediction from the rational expectations model is certainly puzzling, as it implies that agents sharply change even the sign of their inflation expectations in response to small parameter changes that put the model into the indeterminacy region. We compute fiscal multipliers under level-k thinking in the full parameter space covering the determinacy and indeterminacy region from the rational expectations model and show, remarkably, that there is no discontinuity in the behavior of inflation expectations and hence actual inflation as under rational expectations.

[3] See Christiano et al. (2011), and Woodford (2011).

[4] For instance, in our baseline calibration where the rational expectations multiplier is 4.9, even with level-100 thinking, the fiscal multiplier is only about 85 percent of its rational expectations counterpart and full convergence requires level-500 thinking.

[5] For instance, when the expected duration of the fiscal expansion is only a little less than five quarters, the fiscal multiplier under rational expectations drops to 1.3 and the level-5 thinking multiplier is already 1.13, with full convergence achieved at level-20 thinking.

Related literature

This paper contributes to the literature that relaxes the arguably strong assumptions of full information and rational expectations. Learning models as surveyed in Evans and Honkapohja (2001) assume that agents do not know the full economic structure and run regression on past data when forming expectations. For a recent contribution in the context of the ZLB in a New Keynesian model, see Evans and McGough (2018). In contrast, the expectations formation process applied in this paper follows earlier work by García-Schmidt and Woodford (2015, 2019) and Farhi and Werning (2017), which assumes that agents know the structure of the economy, but cannot compute the fully model-consistent forecasts.[6]

Compared to the present paper, García-Schmidt and Woodford (2019) employ their framework to a different question — namely, whether low nominal interest rates result in deflation in the context of a linear sticky-price model with an exogenous interest rate peg. They impose a continuous model of belief revisions. For simplicity, we assume a discrete belief revision model, and we show in the Appendix that qualitatively our results hold when applying continuous belief revisions. Farhi and Werning (2017) study the combination of level-k thinking, incomplete markets, and occasionally binding borrowing constraints and find that all features together provide a solution to the forward-guidance puzzle, but they abstract their analysis from the ZLB.[7] Recent work by Woodford (2018) and Woodford and Xie (2019) studies a different aspect of limitations to fully rational expectations formation — namely, finite planning horizons. A related, but conceptually different, literature addresses higher-order beliefs. The work by Wiederholt (2015) and Angeletos and Lian (2018) also show the fragility of predictions that rest on long series of forward-looking, general equilibrium feedback loops, as in parts of this paper. However, agents in those models do form rational expectations, albeit under imperfect common knowledge.

The rest of the paper is organized as follows. Section 2 presents the core model setup and introduces bounded rationality. Section 3 examines optimal monetary policy at the ZLB rational expectations and level-k thinking. Section 4 discusses the implications of bounded rationality on the forward guidance and the reversal puzzle. Section 5 revisits fiscal multipliers at the ZLB. Section 6 concludes.

[6] See also the related earlier work by Guesnerie (1992, 2008).

[7] Sergeyev and Iovino (2019) apply level-k thinking to analyze the effects of quantitative easing and foreign exchange rate interventions.

2 Level-k reasoning and a simple example

Section 2.1 formally defines a level-k equilibrium. In section 2.2 we then provide a simple example in the context of the purely forward-looking baseline New Keynesian model.

2.1 Equilibrium definitions under level-k thinking

We define an equilibrium with level-k thinking, which builds on the concept of temporary equilibrium defined below.

Definition 1. *(Temporary Equilibrium)*
A temporary equilibrium in period t is an outcome for endogenous variables X_t such that given a sequence of beliefs $\{B_{t+j}\}_{j=1}^{\infty}$ about relevant future variables $\{X_{t+j}\}_{j=1}^{\infty}$, all period t equilibrium conditions (e.g., market clearing, feasibility, and optimal choice given beliefs) are satisfied.

Formally, a temporary equilibrium in period t is an outcome for prices and quantities generated by a mapping from beliefs $\{B_{t+j}\}_{j=1}^{\infty}$ about future relevant variables $\{X_{t+j}\}_{j=1}^{\infty}$ to current period equilibrium values X_t that satisfies the assumptions above. We denote this mapping by Φ, which may depend on predetermined variables X_{t-1}:

$$X_t = \Phi\left(\{B_{t+j}\}_{j=1}^{\infty}, X_{t-1}\right). \tag{1}$$

Level-k thinking, following Farhi and Werning (2017), specifies how agents form beliefs.[8]

Definition 2. *(Level-k equilibrium)*
A level-k equilibrium is a temporary equilibrium where beliefs $\{B_{t+j}\}_{j=1}^{\infty}$ are given by the level-$(k-1)$ equilibrium sequences. These are generated recursively given an initial belief $\{X_{t+j}^{0}\}_{j=1}^{\infty}$:

$$X_t = \Phi\left(\{X_{t+j}^{k-1}\}_{j=1}^{\infty}, X_{t-1}^{k}\right). \tag{2}$$

A key parameter for quantitative results is the level of belief revision, k. We find it reasonable that boundedly-rational agents will perform at most a small number of iterations when forming their beliefs. That assumption seems to be confirmed in experimental studies, see Mauersberger and Nagel (2018) and Coibion et al. (2018), who find a typical level of reasoning no higher than three or four, respectively. One of the core questions in

[8] In infinite-horizon models, updating beliefs about an infinite sequence is infeasible in practice, and we need to truncate the problem at a large integer T. We choose T large enough such that the economy has converged back to steady-state well before T.

this paper, then, is whether the rational expectations outcome is quantitatively or even qualitatively similar to level-k thinking for reasonably small levels of k.

Our analysis abstracts from the important question of how agents learn over longer time periods from their forecast errors. As such, the analysis is clearly not well suited — nor intended — to study permanent changes in policy. Instead, the iterative process of level-k thinking is one that takes place in meta time and is therefore useful for studying the impact response of the economy to policy changes and the resulting dynamics over the near to medium term. Finally, we also explore the possibility that beliefs are updated continuously as proposed in García-Schmidt and Woodford (2019) instead of in a discrete manner (see Appendix D).

2.2 A simple example in the purely forward-looking model

Our baseline model is the three-equation New Keynesian model. Firms set prices in a staggered fashion as in the model of Calvo (1983) and wages are flexible. A representative household supplies labor, consumes and demands bonds in zero net supply. For arbitrary expectations, one can derive the following equilibrium conditions from the log-linear optimality conditions of individual firms and households (see, for example, Preston, 2005):

$$y_t = E_t \sum_{s=0}^{\infty} \beta^s [(1-\beta) y_{t+1+s} - \frac{1}{\sigma}(i_{t+s} - nr_{t+s} - \pi_{t+1+s})], \tag{3}$$

$$\pi_t = E_t \sum_{s=0}^{\infty} (\beta\varphi)^s \left[\beta(1-\varphi)\pi_{t+1+s} + \kappa y_{t+s}\right], \tag{4}$$

Here, π_t, y_t, and i_t, denote inflation, the output gap, and the nominal rate, respectively, all measured as log-deviations from the steady state. The exogenous variable nr_t represents the natural rate of interest that acts as a demand shock in this framework. The first equation describes aggregate demand, derived by integrating over all individual consumption functions, imposing market clearing $C_t = Y_t$, and noting that bonds are in zero net supply. Here, β is the discount factor, and σ is the inverse of the intertemporal elasticity of substitution. The second equation describes aggregate supply and is derived from the optimality condition for the firm's price-setting problem, where φ is the Calvo probability that the firm will not have a chance to reset its price. The slope of the Phillips curve, κ, is defined by

$$\kappa \equiv \frac{(1-\beta\varphi)(1-\varphi)}{\varphi}(\sigma + \omega^{-1}), \tag{5}$$

where ω is the Frisch elasticity of labor supply.

In the context of our purely forward-looking model, the level-k equilibrium is described by the following recursions given a baseline belief about the sequences for future endogenous variables at level-$k = 0$:

$$y_t^k = \sum_{s=0}^{\infty} \beta^s[(1-\beta)y_{t+1+s}^{k-1} - \frac{1}{\sigma}(i_{t+s} - nr_{t+s} - \pi_{t+1+s}^{k-1})] \tag{6}$$

$$\pi_t^k = \kappa y_t^k + \sum_{s=0}^{\infty}(\beta\varphi)^s \left[\beta(1-\varphi)\pi_{t+1+s}^{k-1} + \kappa\beta\varphi y_{t+1+s}^{k-1}\right] \ . \tag{7}$$

It is useful to begin with a simple example that illustrates how the belief formation under level-k thinking works. Suppose in period 1 the central bank announces it will raise the nominal interest rate by Δ in period 2. We assume it is understood that the nominal interest rate is fixed in period 1 and that the central bank operates under a Taylor rule from period 3 onwards. The economy is in steady state prior to the announcement. For simplicity we focus on $\sigma = 1$.

Under level-1 thinking, expectations about future variables are at the baseline (corresponding to the naive belief that all other agents in the future will not react to the policy change) and only the direct effect of the interest rate change affects household choices. Output falls by Δ in period 2 and due to discounting in the Euler equation by $\beta\Delta$ in period 1. Since firms hold expectations about future variables at the baseline, inflation falls by κ times the respective output movements in each period. Note that the level-1 equilibrium in period 2 coincides with the rational expectations solution, because the model is purely forward looking and there is no policy change from period 3 onwards.

Level-2 thinking now updates expectations about future aggregate variables from the sequence of temporary equilibria under level-1 thinking. Households' choices in period 1 now reflect the direct effect of the nominal interest rate change, $-\beta\Delta$, expectations about their lower future income,$-(1-\beta)\Delta$, and the effect of lower expected inflation, $-\kappa\Delta$. These effects equal the sum of the ex-ante real interest rates, $-(1+\kappa)\Delta$, as under rational expectations. Higher levels of belief formation would not yield any change in outcomes as the beliefs about future variables have already converged at level-2. It is easy to see from Table 1 that the level-2 equilibrium is the rational expectations equilibrium in this example. Inflation equals the discounted sum of current and future output gaps and aggregate output is given by the sum of real interest rates as under rational expectations.

More generally, in this purely forward-looking model where the baseline belief corresponds to the rational expectations equilibrium prior to some policy change at most T periods in the future, the level-$k = T$ equilibrium is the rational expectations equilibrium. Clearly, this will no longer be true in a model with endogenous state variables, which we also consider in this paper.

		$t=1$	$t=2$	$t=3,\ldots$
level-1	y_t	$-\beta\Delta$	$-\Delta$	0
	π_t	$-\kappa\beta\Delta$	$-\kappa\Delta$	0
level-2	y_t	$-(1+\kappa)\Delta$	$-\Delta$	0
	π_t	$-\kappa(1+\kappa+\beta)\Delta$	$-\kappa\Delta$	0

Table 1: Level-1 and level-2 equilibrium in simple example

3 Optimal monetary policy under commitment

This section asks how optimal monetary policy at the ZLB should be conducted in the context of the simple forward-looking model outlined above. In Section 3.1 we outline the policy experiment and the calibration. In Section 3.2 we provide results when the policymaker minimizes a welfare-based loss function. In Section 3.3 we consider a loss function with equal weights on inflation and output as in a dual mandate.

3.1 Policy experiment and calibration

We consider the purely-forward looking model outlined in section 2.2. Under rational expectations, the equilibrium conditions of the model are reduced to the familiar IS and Phillips curves.

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma}\left(i_t - nr_t - \mathbb{E}_t \pi_{t+1}\right) \tag{8}$$

$$\pi_t = \kappa y_t + \beta \mathbb{E}_t \pi_{t+1} \tag{9}$$

Under level-k thinking, the equations outlined earlier in (6) and (7) apply. We consider a recession scenario where a large adverse natural rate shock drives the economy to the ZLB for an extended period when monetary policy operates under the following rule

$$i_t = \max(nr_t + 1.5\pi_t + 0.5y_t, ZLB). \tag{10}$$

We refer to outcomes under this simple rule as the "baseline". The rule prescribes that the nominal interest rate follows the natural rate one-for-one (subject to the ZLB constraint and a response to endogenous variables to ensure determinacy), which is exactly the discretionary policy in rational expectations models as in Levin, Lopez-Salido, Nelson, and Yun (2010). Hence, it is a good point of comparison for the optimal commitment policy.

The policy problem under commitment is set up as follows. The private sector forms expectations in line with level-k, and the central bank knows the level-k of belief formation and is choosing its instrument optimally to minimize the discounted sum of variances of realized inflation and realized output. In other words, only the private sector is subject to bounded rationality. We assume the central bank can credibly announce any interest rate path that respects the ZLB, and the private-sector expectations about monetary policy align with it. To provide a fair comparison, the initial beliefs in this exercise are set to the rational expectations solution under the Taylor rule subject to the ZLB, in line with the assumption in Farhi and Werning (2017). In other words, under the Taylor rule, the economies with rational expectations and with level-k thinking result in the same economic outcomes by construction.[9] We believe this assumption is the most useful one when advising a policymaker who treats the underlying shocks as unobservable, but who is presented with a baseline outlook for macro variables under a simple rule. The policymaker would like to know how alternative assumptions about expectations formation of the private sector affect her optimal policy design. Our assumption of holding a baseline outlook constant across model specifications is familiar in the literature (see, for example, Boneva, Braun, and Waki, 2016).

The optimal commitment policy is computed numerically by solving a quadratic programming problem subject to an inequality constraint on the equilibrium interest rate, which is the policymaker's choice variable.[10] The algorithm is sketched in Appendix B; it requires that the infinite horizon problem be approximated by an arbitrarily long but finite horizon problem. To create the baseline, we assume that the natural rate in the initial period falls to -0.033, after which it follows an autoregressive process with persistence of $\rho = 0.9$. Under the Taylor rule, the ZLB is binding for 12 quarters.

3.2 A welfare-based loss function

We examine the results of optimal policy with a welfare-based loss function. That is, the relative weight on output gap stabilization is $\frac{\kappa}{\epsilon} = \frac{.0255}{6} = 0.00425$, where ϵ is the elasticity

[9]This is isomorphic to assuming that agents' initial beliefs under bounded rationality come from the level-k solution, but with a different set of underlying anticipated shocks (in each of the three structural equations), such that the macro observables follow the rational expectations path.

[10]Formally, we chose the path for the nominal interest rate via adding anticipated monetary policy shocks to an interest rate rule as in Laseen and Svensson (2011). Under rational expectations, the particular rule chosen to calculate the impulse response to anticipated policy shocks is irrelevant as long as it generates a determinate equilibrium. Under level-k thinking, it is important that the rule does not include an endogenous feedback to equilibrium variables, but instead be purely exogenous. That is required so that both the actual and the expected path of the nominal interest rate obey the ZLB when those expectations are non-rational. An interest rate rule that is purely exogenous does not generate an issue of determinacy under level-k thinking as outlined in Farhi and Werning (2017). Appendix A shows the impulse response to anticipated policy shocks under level-k thinking that are at the core of our algorithm.

of substitution between goods varieties.

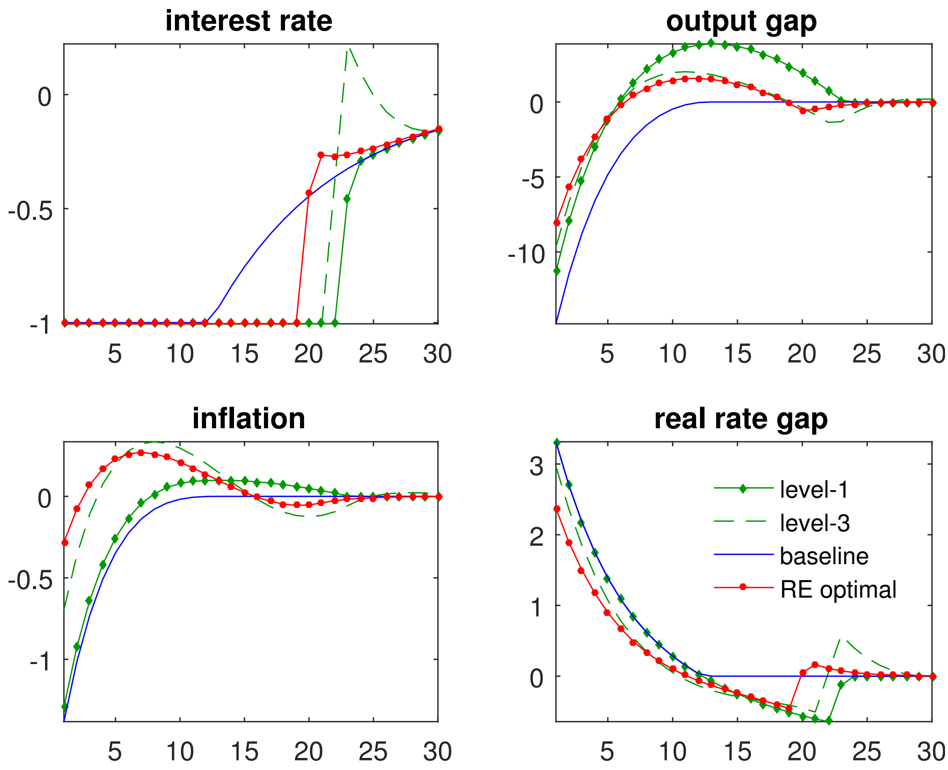

Figure 1: Optimal policy with level-1, and level-3 thinking and a welfare-based loss function.

Figure 1 shows the outcomes under the optimal commitment policies with rational expectations as well as with level-k thinking for levels 1 and 3 (level-2 is not shown to keep the figure readable). In all of these cases, the optimal policy delays liftoff relative to the Taylor rule both under rational expectations and under level-k thinking. That is, the pure fact that agents are boundedly rational about future inflation and future output does not imply that forward-guidance policies are inappropriate. Note that monetary policy optimally stays at the ZLB for longer under level-k thinking than under rational expectations. Focusing on level-1 thinking, the policymaker optimally chooses to stay at the ZLB for almost a year longer than under rational expectations and thus creates stronger overshooting of output and inflation in those future periods. Despite the stronger stimulus in the future, the near-term effects of monetary policy are less stimulative than under rational expectations, because agents to do not adjust their expectations about future income and inflation at all under level-1 thinking. With higher levels of thinking, here shown for level-3 thinking, expectations about the future do adjust, but qualitatively the findings are similar to level-1. The date of liftoff is delayed relative to rational expectations and yet monetary policy achieves smaller improvements in inflation and output in the initial periods than under rational expectations.[11]

3.3 Equal weights in the loss function

We now assume that the central bank has a loss function with equal weights on inflation and the output gap, which we motivate with the dual mandate of many central banks in practice. Figure 2 shows the resulting outcomes under the optimal policies for rational expectations and level-k thinking.

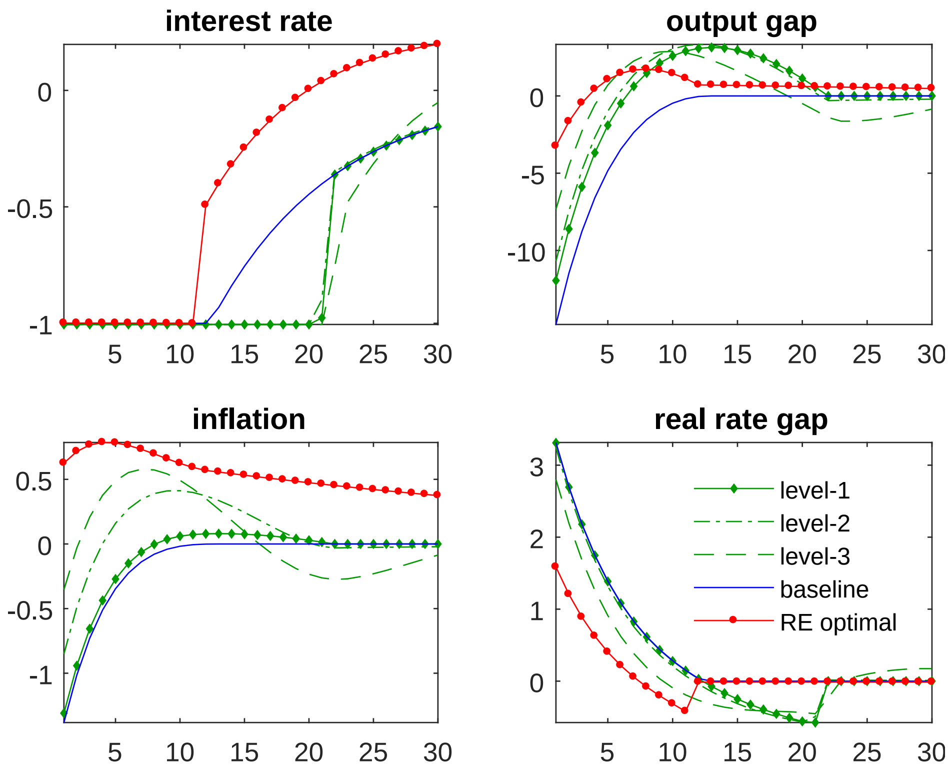

Figure 2: Optimal policy with level-1, level-2, and level-3 thinking and rational expectations under an equal-weights loss function.

Remarkably, under rational expectations, the optimal commitment policy lifts off from the ZLB earlier than under the Taylor rule and has a higher path for the nominal interest rate after the liftoff. Despite this higher path for the policy instrument, macroeconomic outcomes are much better. This arises because the real interest rate is lower under optimal policy due to the higher expected inflation. This can be seen in the lower-right panel, which plots the real rate gap — defined as the gap between the ex-ante real interest rate

[11]These results are consistent with those in Nakata, Ogaki, Schmidt, and Yoo (2018), who study optimal commitment policy based on ad-hoc models with discounting in the Euler equation and the Phillips curve. They also find that optimal policy is less powerful due to discounting, and that the central bank optimally extends the ZLB spell relative to the liftoff date in the absence of discounting in the standard NK model.

and the natural rate shock. Clearly, this prescription for optimal monetary relies on the ability to influence inflation expectations in line with the rational expectations path. One would therefore suspect that the optimal policy prescription is fragile to the introduction of bounded rationality, which turns out to be the case.

Figure 2 shows that with the level-k process for expectations formation, rather than raising the nominal interest rate from the lower bound earlier than under the Taylor rule prescribes as under rational expectations, the nominal rate lifts off much later. In fact, liftoff is delayed to period 20, a full two years later than under the Taylor rule. Macroeconomic outcomes do improve, but only mildly so compared to the remarkable improvement under optimal commitment with rational expectations.

We now discuss expectations formation and the associated macroeconomic outcomes for the different levels of belief revision k in more detail. Under level-1 thinking, inflation and output gap expectations are at the baseline, and only the direct effect of the delayed liftoff from the ZLB affects macroeconomic outcomes. Since the path for the interest rate under optimal policy is assumed to be known, and it enters the Euler equation directly (albeit with small discounting), the output gap is improved noticeably. Inflation barely moves with level-1 thinking, because expectations about future inflation and the future output gap are at the baseline and the impact of the current period output gap on inflation is small. Under level-2 thinking, the path for inflation is substantially improved as price setters update their expectations of the future output gap (and, much less importantly also update their expectations of future inflation). The path for the output gap under level-2 thinking is little changed from that of level-1, because the level-1 path for inflation that is used in expectations formation for the real interest rate is little changed from the baseline, and the path for the output gap that *is* changed does not carry much weight in the Euler equation.[12] Under level-3 thinking, the main change relative to level-2 thinking is in inflation expectations that are now revised substantially and are in line with the actual paths for inflation under level-2 thinking. Hence, the ex-ante real interest rate gap (the gap between the ex-ante real interest rate computed using households' subjective inflation expectations and the natural rate shock) is smaller than in the baseline and hence the output gap improves further compared to level-2 thinking.

Nevertheless, even under level-3 thinking, output falls about 7 percent below the steady state on impact – almost twice as much as under rational expectations – and inflation

[12]It is important to recognize that Figure 2 recomputes optimal monetary policy for any given level of thinking of the private sector. Hence, the path for inflation expectations under level-2 thinking that is used in the computations of optimal policy under level-2 thinking is not the one plotted in this figure for level-1 thinking, because the latter is based on a different optimal interest rate path. Nevertheless, for the purpose of discussion and since the difference in optimal interest rate paths as k varies is small, one can gain intuition by reading beliefs under level-2 thinking for inflation (or output) off of the path under level-1 thinking shown in this figure for those variables.

remains below the baseline on impact against a rise above the baseline under rational expectations. Hence, we conclude that macroeconomic stabilization is more difficult for the central bank in the context of level-k thinking than with rational expectations.

It is natural to ask whether the policy that a central bank would optimally choose knowing that the private sector forms expectations according to level-k would converge to the optimal policy under rational expectations for high level of k. We find that, numerically, it does converge. At first glance, this is surprising. After all, the optimal commitment policy under rational expectations shown in Figure 2 is higher than that of the Taylor rule over the 30 quarters that are plotted in the figure. It is easy to see from the discussion in section 2.2 that a uniformly higher path for interest rates leads to uniformly worse outcomes for inflation and output for any level of k. However, under rational expectations and for a high level of k, the nominal interest rate in the last period of the finite horizon policy problem is slightly below that of the Taylor rule. It turns out that this final period stimulus is crucial for the dynamics along the entire path. This is merely another manifestation of the forward-guidance puzzle under rational expectations and under the near-rational expectations solution for high levels of k.[13]

The analysis makes it clear that the optimal path for the nominal interest rate under rational expectations relies crucially on model-consistent expectations about future inflation and output. In practice, one may hold the view that these expectations will not instantaneously adjust perfectly in line with rational expectations and that small changes to interest rates at far distant horizons have little to no effect on current period outcomes. But if that is so, the prescriptions from the bounded rationality models for low levels of k examined here will likely give more suitable policy advice. These prescriptions turned out to be in line with actual forward guidance policies chosen in the aftermath of the financial crisis.

We close this section with a discussion of a possible caveat — namely, the fact that agents are assumed not to learn from their forecast errors (as is standard in the level-k literature). Incorporating learning into the level-k framework is conceptually difficult for a number of reasons. First, much of the learning literature operates in a linear economy with constant coefficients. However, in the context of a quasi-linear model with the ZLB, the law of motion has time-varying coefficients, as illustrated, for example, in Guerrieri and Iacoviello (2015). As such, it is less clear how one should adjust beliefs in response to forecast errors compared with a time-invariant law of motion. Second, like García-Schmidt and Woodford (2019) and Farhi and Werning (2017), we do not model the ultimate reasons why agents do not update beliefs arbitrarily often, — in other words, k is not endogenous.

[13] In a rational expectations model and without imposing the ZLB, the importance of small changes in the last period in a finite horizon model of optimal monetary policy has been pointed out by Campbell and Weber (2018).

But one may suspect that if agents were to make costly and persistent forecast errors, they would eventually adjust their level of belief formation k. Therefore, we find it more natural to provide some sensitivity analysis with respect to k. Finally, we stress that the framework is not intended to give advice about the long-run consequences of permanent policy changes for which the treatment of forecast errors is clearly more important.

4 Delayed liftoff and the reversal puzzle

In this section, we consider the effect of a liftoff delay from the ZLB by a fixed number of quarters relative to the liftoff date under a benchmark Taylor rule. Such "lower for longer policies" have been widely used in the aftermath of the financial crisis by central banks. A number of authors have shown that such a time-dependent forward-guidance policy can have implausibly large effects on initial inflation and output; see, for instance, Carlstrom et al. (2015) or Del Negro et al. (2012).

We add endogenous inflation persistence to the basic New Keynesian model. This will allow us to contrast results under rational expectations and level-k thinking for both the forward-guidance puzzle (which does not require endogenous states) and the less well known "reversal puzzle" outlined in the introduction (which does require endogenous states).

4.1 Model setup and policy experiment

We assume that those firms that do not receive a signal to update their price fully index to last period's inflation rate. For arbitrary expectations, the equilibrium conditions can now be expressed as:

$$y_t = E_t \sum_{s=0}^{\infty} \beta^s [(1-\beta) y_{t+1+s} - \frac{1}{\sigma}(i_{t+s} - nr_{t+s} - \pi_{t+1+s})] \tag{11}$$

$$p_t^* + \pi_t = (1-\beta\varphi) E_t \sum_{s=0}^{\infty} (\beta\varphi)^s [\pi_{t+s} + (\omega + \sigma^{-1}) y_{t+s}] \tag{12}$$

$$p_t^* = \frac{\varphi}{1-\varphi}(\pi_t - \pi_{t-1}) \ . \tag{13}$$

The second equation is the first-order condition for firms that optimally adjust their price, where p_t^* denotes the price of adjusting firms relative to the aggregate price index. The third equation follows from the recursion for the aggregate price index. These equations can be reduced under rational expectations to the familiar system:

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma}\left(i_t - nr_t - \mathbb{E}_t \pi_{t+1}\right) \tag{14}$$

$$\pi_t = \frac{1}{1+\beta}\pi_{t-1} + \frac{\beta}{1+\beta}\mathbb{E}_t \hat{\pi}_{t+1} + \frac{1}{1+\beta}\kappa y_t \ . \tag{15}$$

Under level-k thinking, the equilibrium conditions are given by the following recursions:

$$y_t^k = \sum_{s=0}^{\infty} \beta^s [(1-\beta) y_{t+1+s}^{k-1} - \frac{1}{\sigma}(i_{t+s}^{k-1} - nr_{t+s} - \pi_{t+1+s}^{k-1})] \tag{16}$$

$$(p_t^*)^k + \pi_t^k = (1-\beta\varphi)\left[\pi_t^k + (\omega+\sigma^{-1}) y_t^k + \sum_{s=1}^{\infty} (\beta\varphi)^s \left[\pi_{t+s}^{k-1} + (\omega+\sigma^{-1}) y_{t+s}^{k-1}\right]\right] \tag{17}$$

$$(p_t^*)^k = \frac{\varphi}{1-\varphi}(\pi_t^k - \pi_{t-1}^k) \ . \tag{18}$$

We consider an environment where the central bank communicates that the nominal interest rate will stay at the ZLB for an extended period before returning to an interest rate rule. This occurs against the background of an adverse demand shock that drives the economy endogenously to the ZLB. In particular, we assume the central bank announces an interest rate rule given by:

$$i_t = \begin{cases} ZLB & t = 1, 2, ..., t^*, t^{*+1}, ..., t^{*+j} \\ \max(ZLB, r_t + \phi_\pi \pi_t + \phi_y y_t) & t \geq t^{*+j+1} \end{cases}$$

Here, t^* is the period prior to liftoff under the baseline Taylor rule. This policy delays liftoff by j periods relative to the Taylor rule. Post-liftoff, the baseline rule applies again, subject to the ZLB. This policy experiment models in a simple way key elements of forward guidance implemented by several central banks in the aftermath of the financial crisis.

We assume a standard calibration: $\sigma = 1$, $\omega = 0$, $\beta = 0.99$. The probability of keeping the price fixed is $\varphi = 0.8564$ such that the slope of the Phillips curve is $\kappa = 0.0255$. The interest rate reaction coefficients are given by $\phi_\pi = 1.5$ and $\phi_y = 0.5$. The initial innovation into the natural rate is set to -0.015, and we assume the natural rate follows an AR(1) process with persistence of 0.9.

4.2 Comparing rational expectations and level-k thinking

Figure 3 shows the solution under rational expectations under the standard Taylor rule in the blue solid line and with a delayed liftoff by one, two, and three quarters. Note that a delay of one quarter improves outcomes very little. However, a delay of two quarters results in a dramatic improvement. Output rises above the steady state instead of contracting by nearly 2 percent on impact under the baseline and inflation is persistently above the steady state rather than below under the baseline rule. Under rational expectations, macro economic outcomes are extremely sensitive to very small variations in the liftoff date in this economy. This arises because there is a strong feedback between expectations of future inflation, improvements in inflation during the spell at the ZLB and the resulting rise in future inflation due to indexation.

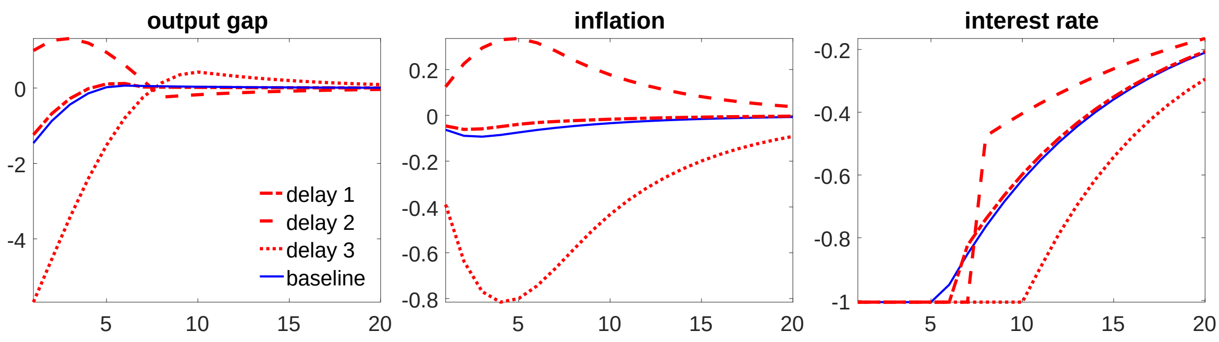

Figure 3: Delayed liftoff under rational expectations

One would suspect that a delay of three quarters would result in further increases in output and inflation. However, in the context of this calibrated model, a delayed liftoff policy under rational expectations results in a so-called *reversal puzzle* as pointed out in Carlstrom et al. (2015). Rather than further stimulating the economy, a delay in the liftoff date of three quarters results in a contraction in output and inflation. The weak economic outcomes then cause the ZLB to bind further and the ZLB binds endogenously for a total of 10 quarters. It is difficult to provide an intuition for why the reversal occurs under rational expectations. Note, however, that the feedback loop described above naturally becomes stronger as the duration of the peg is extended. The reversal then occurs when the feedback process becomes so strong that inflation would rise without bounds unless the initial response were to switch the sign. These perverse movements of inflation and output in response to a time-dependent delay in the liftoff have been discussed in detail in Carlstrom et al. (2015), and they cast a doubt on the rational expectations assumption. One may question why households and firms should be expected to revise their expectations in response to a delayed liftoff of three periods in a qualitatively completely different way than for a delay of two periods. In particular, Figure 3 makes it clear that under rational expectations, the general equilibrium effects of monetary policy must be working in the opposite direction to the expansionary partial equilibrium effect from the interest rate change alone whenever the so-called reversal puzzle arises.

We next examine the same delayed liftoff policy under level-k thinking. We again assume that the initial expectations are given by the baseline outlook under rational expectations and the Taylor rule. Figure 4 shows the solution to the model with a delayed liftoff by the same one, two, and three quarters under level-2 thinking. (Level-3 or level-4 thinking give rise to qualitatively similar results.)

Note first that a delay by one and two quarters results in substantially smaller macroeconomic improvement than under rational expectations. In other words, the forward guidance policy is substantially less powerful with boundedly rational agents. This arises, in part, because the strong endogenous feedback loop described earlier is muted with level-2 thinking. Furthermore, there is no reversal puzzle under level-k thinking. Keeping interest rates lower for longer improves outcomes by more, the longer the interest rate stays at the ZLB. Numerical results (not shown here) confirm that this is the case for further delays in the liftoff date beyond the 3three quarters that are plotted in the figure. It is also worth mentioning that under continuous belief revision the results are qualitatively similar (see Figure 8 in the Appendix).

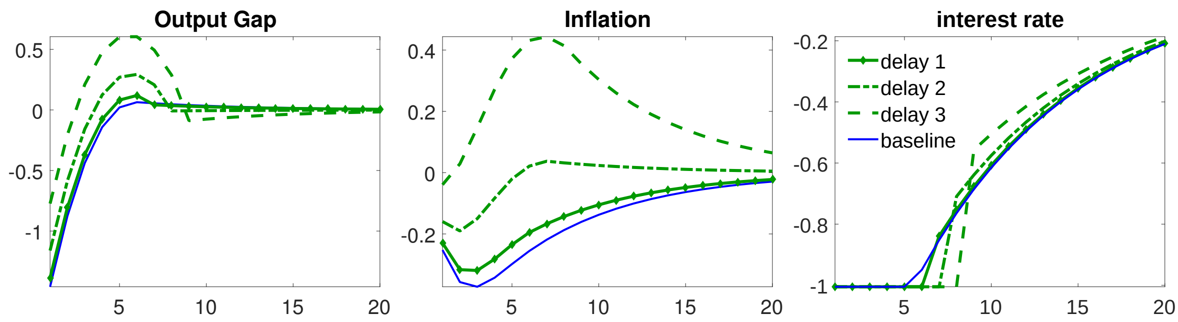

Figure 4: Level-2 solution with delayed liftoff by one, two, and three quarters.

We close this section by discussing whether the iterative process of belief revision for large k converges to the rational expectations solution. We find numerically that it does so when the delay in the liftoff is small (a delay of one or two periods) such that the policy is expansionary for output and inflation under both rational expectations and level-k thinking. However, it does not converge to the rational expectations solution for a delay in the liftoff that results in a (counterintuitive) contraction in real activity and inflation under rational expectations, for instance, for a delay of three periods, as the level-k increases the effects on initial inflation and output diverge to arbitrarily large values.[14]

[14] When we study convergence, we assume that updates are only partial in the direction of the sequence of new temporary equilibria as in the work of García-Schmidt and Woodford (2019). It is well known from the literature on the convergence properties of the Fair-Taylor algorithm (which is closely related to belief revision as in level-k thinking), that the algorithm often does not converge when other algorithms such as the stacked time algorithm does. Partial updates or dampening often helps to achieve convergence

In our judgment, non-convergence provides a formal basis for discounting this particular prescription from the rational expectations framework as not relevant in practice.

5 Fiscal multipliers at the ZLB

This section considers fiscal multipliers at the ZLB under level-k thinking and compares the results with those under rational expectations. We begin with the case of an equilibrium at the ZLB driven by fundamentals and then proceed to a so-called expectations-driven liquidity trap similar to Mertens and Ravn (2014).

5.1 A fundamentals-driven equilibrium

Much of the literature has pointed out that the fiscal multiplier under constant interest rates in a ZLB episode can be large. The mechanism is well understood: higher government spending raises inflation expectations, which reduces ex-ante real interest rates and hence crowds in private expenditures, thus raising the multiplier above unity. What is surprising is that the effect can be quantitatively very large. For instance, in the well-known paper by Christiano et al. (2011), the fiscal multiplier is 3.7 on impact for an expansion with an expected duration of five quarters in the context of a simple, purely forward-looking model. Carlstrom et al. (2014) point out that the mechanism in this model relies on rational expectations about small-probability events in which the fiscal expansion lasts for a long time and the expected macro outcomes are huge. Similarly, Dupor and Li (2015) present empirical evidence that the large inflation expectations channel associated with government spending highlighted in the theory did not hold in the 2009 Recovery Act period in the U.S. data. It is thus natural to examine the sensitivity of the fiscal multiplier under constant interest rates to an explicit model of belief formation such as level-k thinking.

For simplicity, we consider an environment similar to that of section 3 of Christiano et al. (2011). We consider a coordinated policy experiment in which (i) government spending is set above steady state $g_t = \overline{g} > 0$, and simultaneously (ii) the central bank announces an interest rate peg $i_t = 0$. Each period there is probability p that this policy will continue so that the expected duration of the expansion is $T = \frac{1}{1-p}$. With probability $(1 - p)$ the fiscal expansion ends, at which point fiscal policy returns to steady state, $g_t = 0$, and monetary policy reverts to a typical Taylor rule:

$$i_t = \phi_\pi \pi_t + \phi_y y_t. \tag{19}$$

in this context.

Under standard assumptions on ϕ_π and ϕ_y, there is a unique stable equilibrium after the period of the peg. Since there are no state variables or exogenous shocks during these subsequent periods, the unique equilibrium after the policy experiment is given by $\pi_t = y_t = 0$.

Under arbitrary expectations about the outcomes during a fiscal expansion, the equilibrium conditions are given by:

$$c_t = E_t \sum_{s=0}^{\infty} \beta^s p^{s+1} [(1-\beta) c_{t+1+s} - \frac{1}{\sigma}(p^{-1}\bar{i} - \pi_{t+1+s})] \tag{20}$$

$$\pi_t = E_t \sum_{s=0}^{\infty} (\beta\varphi)^s \left[p^{s+1}\beta(1-\varphi)\pi_{t+1+s} + p^s \kappa(\sigma c_{t+s} + \omega^{-1} y_{t+s})\right] \tag{21}$$

$$y_t = (1-s)\, c_t + s g_t \; . \tag{22}$$

Here, $\bar{i}$ is the constant value of the natural rate of interest, which we set to zero without loss of generality. The constant $s = \frac{G}{Y_{ss}}$ is the share of government spending in the steady state and (with a slight change of notation) κ is now defined as $\kappa \equiv (1-\beta\varphi)(1-\varphi)\varphi^{-1}$.

Under rational expectations the model is given by:

$$c_t = \mathbb{E}_t c_{t+1} - \frac{1}{\sigma}(i_t - \mathbb{E}_t \pi_{t+1}) \tag{23}$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa m c_t \tag{24}$$

where marginal cost is given by $mc_t = \sigma c_t + \omega^{-1} y_t$.

Under rational expectations, the fiscal multiplier during the interest rate peg is given by:

$$\frac{dY}{dG} \equiv \left(\frac{1}{s}\right)\frac{dy_t}{dg_t} = \left[\frac{\sigma\left[(1-p)(1-\beta p) - \kappa p\right]}{\Delta}\right] , \tag{25}$$

where

$$\Delta \equiv \sigma(1-p)(1-\beta p) - \kappa\left[\sigma + \omega^{-1}(1-s)\right] p. \tag{26}$$

The model has a stable, unique equilibrium whenever $\Delta > 0$. We restrict attention to that case for now. We use the same baseline parameter values as in Carlstrom et al. (2014): $\beta = 0.99$, $\kappa = 0.028$, $\omega^{-1} = 0.5$, $\sigma = 2$, $s = 0.2$ and $p = 5/6$. Under this calibration, the rational expectations fiscal multiplier is large. In fact, it is 4.9.

Our system of equations under level-k thinking is,

$$c_t^k = \sum_{s=0}^{\infty} p^{s+1}\beta^s[(1-\beta)c_{t+1+s}^{k-1} - \frac{1}{\sigma}(p^{-1}\bar{i} - \pi_{t+1+s}^{k-1})] \quad (27)$$

$$\pi_t^k = \kappa(\sigma c_t^k + \omega^{-1}y_t^k) + \sum_{s=0}^{\infty}(\beta\varphi)^s p^{s+1}\left[\beta(1-\varphi)\pi_{t+1+s}^{k-1} + \kappa\beta\varphi(\sigma c_{t+1+s}^{k-1} + \omega^{-1}y_{t+1+s}^{k-1})\right] \quad (28)$$

$$y_t^k = (1-s)\, c_t^k + s g_t^k\ . \quad (29)$$

We assume that agents have boundedly rational beliefs about an equilibrium where all future variables during the fiscal expansion take on constant values; that is $\{\pi_{t+1+s}^k\}_{s=0}^{\infty} = \bar{\pi}^k$, $\{c_{t+1+s}^k\}_{s=0}^{\infty} = \bar{c}^k$. Equations (27) and (28) provide updating formulas to revise these beliefs iteratively. The assumption of constant values during the expansion is consistent with rational expectations. For this exercise as well as the one in the next subsection, we initialize beliefs at the rational expectations equilibrium prior to the change in policy. This turns out to be the steady state. Formally, the level-k equilibrium iterates on the following recursions for the vector $X = [y, \pi]'$:

$$X_{(k)} = CX_{(k-1)} + G. \quad (30)$$

Here, G is a vector of constants related to the exogenous level of government spending and the matrix C is given by

$$C = \begin{bmatrix} \tilde{p} & \frac{\tilde{p}}{\sigma(1-\beta)} \\ \tilde{k}(\tilde{p} + \beta\varphi\hat{p}) & \frac{\tilde{k}\tilde{p}}{\sigma(1-\beta)} + \hat{p}\beta(1-\varphi) \end{bmatrix},$$

where we define:

$$\tilde{k} \equiv \frac{(1-\beta\varphi)(1-\varphi)}{\varphi}(\sigma + \omega^{-1}(1-s)), \quad \tilde{p} \equiv \frac{p(1-\beta)}{1-p\beta}, \quad \hat{p} \equiv \frac{p}{1-p\beta\varphi}.$$

We begin discussing convergence of the fiscal multiplier under level-k to the rational expectations multiplier, which is summarized in Proposition 1.

Proposition 1. *(Convergence to rational expectations equilibrium)*
The iterations in (30) converge if $\Delta \equiv \sigma(1-p)(1-\beta p) - \kappa[\sigma + \omega^{-1}(1-s)]\, p > 0$.

Note that this parameter restriction is the same as the determinacy condition for the rational expectations version of the model. Hence, if the model has a unique stable equilibrium under rational expectations, the level-k equilibrium converges to that rational expectations equilibrium as k increases.

The proof is outlined in Appendix C. We now turn to the quantitative implications of level-k thinking for the size of the fiscal multiplier under our baseline calibration.

level-k	1	2	5	10	20	30	40	50	100	200	500
multiplier	1	1.03	1.23	1.53	2.05	2.5	2.87	3.19	4.16	4.76	4.9
(% of RE)	(20)	(21)	(25)	(31)	(42)	(51)	(59)	(65)	(85)	(97)	(100)

Table 2: Fiscal multiplier under level-k thinking

Table 2 shows that to achieve beliefs consistent with the rational expectations fiscal multiplier of 4.9, agents are required to iteratively update their beliefs to an extraordinarily high level. For example, level-5 beliefs result in a fiscal multiplier that is only 1.23. Even level-100 beliefs only generate a multiplier that is 85% of its rational expectations counterpart, while full convergence is only achieved at level-500 thinking.

Is this slow convergence to rational expectations under level-k thinking a generic feature of fiscal multipliers in this framework, or is convergence only slow when the multiplier is unusually large under rational expectations? To answer this, we reduce the probability of staying in the expansionary fiscal policy regime from $p = 5/6 \approx 0.83$ to $p = 0.8$. Under rational expectations, this lowers the fiscal multiplier from 4.9 to 1.3. Under level-k thinking, convergence to this smaller fiscal multiplier is relatively rapid. In particular, level-5 thinking implies a multiplier of 1.13, level-10 a multiplier of 1.2 and at level-20 the model has almost converged to rational expectations producing a multiplier of 1.29.

We conclude that the huge fiscal multipliers which can occur in this very simple model under rational expectations require beliefs that can be approximated in this framework only by an unusually high level of k. When fiscal multipliers are smaller and arguably more reasonable, the approach of bounded rationality taken here approximates rational expectations closely for relatively low levels of k.

5.2 An expectations-driven liquidity trap

Recent work by Mertens and Ravn (2014) has shown that the effect of government spending is different in an expectations-driven liquidity trap than in a fundamentals driven one. In particular, their analysis restricts attention to a minimum state variable solution (MSV) augmented with a sunspot shock that follows a Markov process. Whenever the sunspot shock is persistent enough, the augmented MSV solution leads to self-fulfilling spells at the ZLB in which government spending multipliers are always smaller than unity rather than bigger than unity and inflation falls rather than rises with additional spending.

In line with their approach, we can calculate fiscal multipliers under the MSV solution even for $p > p^*$. Here, p^* is the critical value for which $\Delta = 0$ in equation (26), which is the boundary of the determinacy region. Another interpretation of this multiplier is that it is the fiscal multiplier in an expectations-driven liquidity trap driven by a sunspot shock that persists with probability p. As discussed in Farhi and Werning (2017), with level-k thinking, indeterminacy of equilibrium in a linear model never arises.

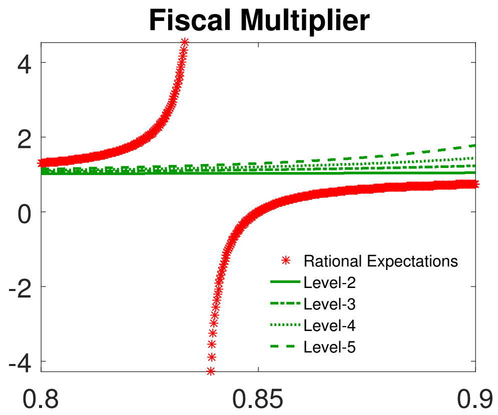

Figure 5: Fiscal multiplier as a function of the probability p

Figure 5 plots the fiscal multiplier across both regions of p under rational expectations and for level-k thinking. The multipliers are broadly similar for $p \sim 0.8$. Under rational expectations, the fiscal multiplier then grows rapidly for small increases of p. It asymptotes at the determinacy region to unboundedly large positive values, before collapsing to unboundedly negative values and then rises again towards unity from below. In contrast, the fiscal multiplier under level-k thinking is a smooth function of p and only mildly increasing in p. Furthermore, Proposition 1 implies that the fiscal multiplier in the indeterminacy region under rational expectations is not the limit of the fiscal multiplier under bounded rationality as the level of belief revision k increases. This casts some doubt on the practical relevance of the rational-expectations dynamics under indeterminacy. The sign flip in the fiscal multiplier under rational expectations reflects the fact that inflation expectations rise with additional government spending for $p < p^*$, but fall for $p > p^*$. Under level-k thinking, none of these unusual switches in inflation expectations occur locally around p^*. These results are reminiscent of the findings in Wieland (2018), who points out that the typical assumption of the MSV solution in the indeterminacy region implies that an equilibrium (out of the many possible equilibria) is picked whose dynamics are qualitatively very different than in the determinacy region. Our analysis makes clear that an explicit model of belief formation results in dynamics that do not resemble those from the MSV solution under indeterminacy at all. Furthermore, the fact that the MSV solution is not the limit of higher levels of k in the indeterminacy region gives reason to doubt the plausibility of the resulting dynamics as implausible.

6 Conclusion

Many puzzling features of rational expectations model at the ZLB disappear once an empirically relevant process for expectations formation is embedded in an otherwise standard macroeconomic model. Overall, level-k thinking mitigates the large general equilibrium effects of announced future paths of monetary and fiscal policy. Level-k expectation formation appears particularly reasonable in a policy environment where households and firms have little experience with regard to how these policy interventions affect macroeconomic outcomes.

More research is needed to understand how households and firms form their expectations in practice. Recent work by Coibion, Gorodnichenko, and Weber (2019) is a welcome contribution in that regard. Another fruitful avenue for future research is an evaluation of the empirical fit of large-scale DSGE models in a level-k thinking environment.

References

Angeletos, G.-M. and C. Lian (2018). Forward guidance without common knowledge. *American Economic Review 108*(9), 2477–2512.

Boneva, L. M., R. A. Braun, and Y. Waki (2016). Some unpleasant properties of loglinearized solutions when the nominal rate is zero. *Journal of Monetary Economics 84*, 216 – 232.

Calvo, G. A. (1983). Staggered price setting in a utility-maximizing framework. *Journal of Monetary Economics 12*, 383–398.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics 119*(3), 861–898.

Campbell, J. R. and J. P. Weber (2018). Open mouth operations. Federal Reserve Bank of Chicago WP No. 2018-03.

Carlstrom, C. T., T. S. Fuerst, and M. Paustian (2014). Fiscal multipliers under an interest rate peg of deterministic versus stochastic duration. *Journal of Money, Credit and Banking 46*(6), 1293–1312.

Carlstrom, C. T., T. S. Fuerst, and M. Paustian (2015). Inflation and output in new keynesian models with a transient interest rate peg. *Journal of Monetary Economics 76*, 230 – 243.

Christiano, L. J., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy 119*, 78–121.

Coibion, O. and Y. Gorodnichenko (2012). What Can Survey Forecasts Tell Us about Information Rigidities? *Journal of Political Economy 120*(1), 116 – 159.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105*(8), 2644–78.

Coibion, O., Y. Gorodnichenko, and R. Kamdar (2018). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature 56*(4), 1447–91.

Coibion, O., Y. Gorodnichenko, S. Kumar, and J. Ryngaert (2018). Do you know that i know that you know...? higher-order beliefs in survey data. NBER Working Papers 24987, National Bureau of Economic Research, Inc.

Coibion, O., Y. Gorodnichenko, and M. Weber (2019). Monetary policy communications and their effects on household inflation expectations. NBER Working Papers 25482, National Bureau of Economic Research, Inc.

Del Negro, M., M. Giannoni, and C. Patterson (2012). The forward guidance puzzle. Staff Reports 574, Federal Reserve Bank of New York.

Dupor, B. and R. Li (2015). The expected inflation channel of government spending in the postwar u.s. *European Economic Review 74*, 36 – 56.

Eggertsson, G. B. and M. Woodford (2003). The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity 34*(1), 139–235.

Evans, G. W. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics.* Princeton University Press.

Evans, G. W. and B. McGough (2018). Interest-rate pegs in new keynesian models. *Journal of Money, Credit and Banking 50*, 939–65.

Farhi, E. and I. Werning (2017). Monetary policy, bounded rationality, and incomplete markets. Revise and resubmit, American Economic Review.

García-Schmidt, M. and M. Woodford (2015). Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis. NBER Working Papers 21614, National Bureau of Economic Research, Inc.

García-Schmidt, M. and M. Woodford (2019). Are low interest rates deflationary? a paradox of perfect-foresight analysis. *American Economic Review 109*, 86–120.

Grandmont, J. M. (1977). Temporary general equilibrium theory. *Econometrica 45*(3), 535–572.

Guerrieri, L. and M. Iacoviello (2015). Occbin: A toolkit to solve models with occasionally binding constraints easily. *Journal of Monetary Economics 70*, 22–38.

Guesnerie, R. (1992). An exploration of the eductive justifications of the rational-expectations hypothesis. *American Economic Review 82*, 1254–78.

Guesnerie, R. (2008). Macroeconomic and monetary policies from the eductive viewpoint. In K. Schmidt-Hebbel and C. Walsh (Eds.), *Monetary Policy under Uncertainty and Learning*, pp. 171–202. Central Bank of Chile.

Laseen, S. and L. Svensson (2011). Anticipated alternative policy rate paths in policy simulations. *International Journal of Central Banking 7*(3), 1–35.

Levin, A., D. Lopez-Salido, E. Nelson, and T. Yun (2010). Limitations on the effectiveness of forward guidance at the zero lower bound. *International Journal of Central Banking 6*, 143–185.

Mauersberger, F. and R. Nagel (2018). Chapter 10 - levels of reasoning in keynesian beauty contests: A generative framework. In C. Hommes and B. LeBaron (Eds.), *Handbook of Computational Economics*, Volume 4 of *Handbook of Computational Economics*, pp. 541 – 634. Elsevier.

Mertens, K. and M. Ravn (2014). Fiscal policy in an expectations-driven liquidity trap. *Review of Economic Studies 81*(4), 1637–1667.

Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *American Economic Review 85*(5), 1313–1326.

Nakata, T., R. Ogaki, S. Schmidt, and P. Yoo (2018). Attenuating the Forward Guidance Puzzle : Implications for Optimal Monetary Policy. Finance and Economics Discussion Series 2018-049, Board of Governors of the Federal Reserve System (US).

Preston, B. (2005). Learning about monetary policy rules when long-horizon expectations matter. *International Journal of Central Banking 1*(2).

Sergeyev, D. and L. Iovino (2019). Quantitative Easing without Rational Expectations. Technical report, Working Paper.

Wiederholt, M. (2015). Empirical properties of inflation expectations and the zero lower bound. Technical report, R&R JPE.

Wieland, J. (2018). State-dependence of the zero lower bound government spending multiplier. Working paper.

Woodford, M. (2011). Simple analytics of the government expenditure multiplier. *American Economic Journal: Macroeconomics 3*, 1–35.

Woodford, M. (2018). Monetary Policy Analysis When Planning Horizons Are Finite. In *NBER Macroeconomics Annual 2018, volume 33*, NBER Chapters. National Bureau of Economic Research, Inc.

Woodford, M. and Y. Xie (2019). Policy Options at the Zero Lower Bound When Foresight is Limited. *AEA Papers and Proceedings 109*, 433–437.

Yellen, J. (2016). Macroeconomic research after the crisis. Speech at "The Elusive 'Great' Recovery: Causes and Implications for Future Business Cycle Dynamics" 60th annual economic conference sponsored by the Federal Reserve Bank of Boston.

Online Appendix A: IRFs to anticipated monetary policy shocks

A useful starting point for understanding the transmission of monetary policy is the impulse response to an anticipated loosening of monetary policy in the near and distant future. We assume further that the nominal interest rate is unchanged in all periods leading up to the future monetary loosening. This experiment is directly relevant for understanding optimal policy, because these impulse responses to anticipated monetary policy shocks are the key input into the computational algorithm we are using when solving for optimal monetary policy.

Figure 10 shows the responses to an anticipated shock one quarter and 100 years in the future. Note that the experiment in the top panel corresponds to the simple example presented in Section 2.2. Because the model is purely forward looking one can also read the bottom panel to show the response to any anticipated shock at any horizon earlier than 100 years in the future. For instance, the effect of a policy loosening 25 years in the future on current inflation and output can be read off the chart by considering quarter 300 (25 years prior to the loosening in quarter 400 assumed in the experiment plotted in the chart).

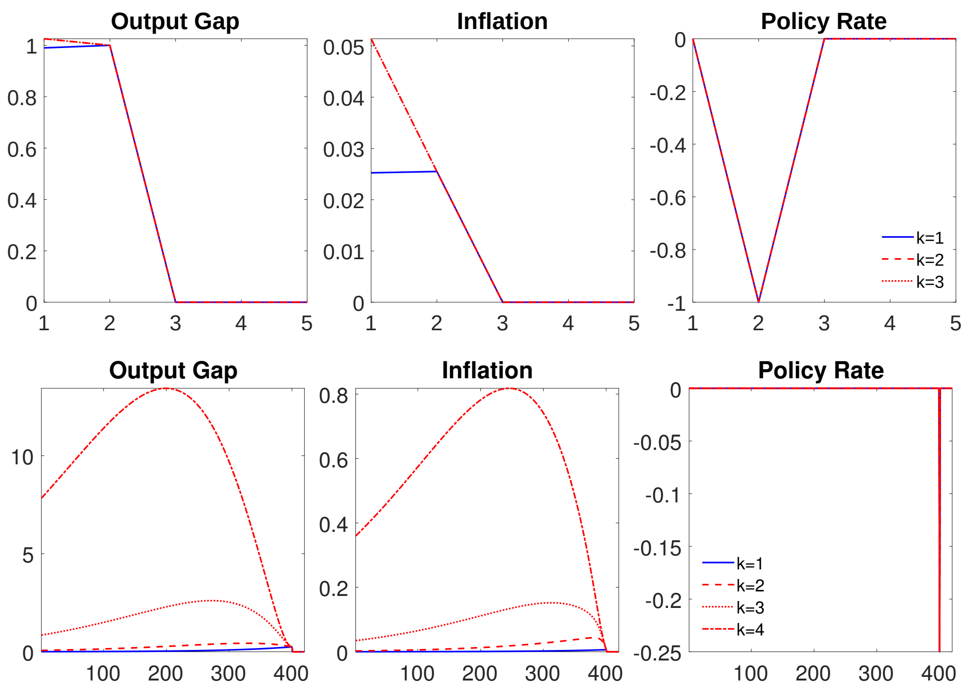

Figure 6: Anticipated monetary policy shock 1 quarter and 400 quarters ahead (variables are in percentage deviations from steady state)

At level-1 thinking, all future endogenous variables are held at a baseline path equal to steady state. Since the interest rate is discounted in the Euler equation, the effect on inflation and output is smaller the further in the future the loosening occurs. Farhi and Werning (2017) call this the horizon effect. At higher levels, the response is hump shaped for the following reasons. Beliefs are now formed from the paths for output and inflation from previous levels which are above the baseline. The higher income and inflation expectations accumulate for longer the further in the future the movement in the interest rate is and this creates an anti-horizon effect. But beliefs about future inflation and income are also discounted. With β close to unity, the discounting matters little for relatively modest horizons and hence monetary policy loosening is *more* powerful the further in future it occurs. For longer horizons, the effect of discounting becomes large, and the anti horizon effect dominates; monetary policy loosening is *less* powerful the further in future it occurs. For comparison, under rational expectations, the impact of the monetary policy loosening 100 years in the future is incredibly large in this linear model: The output response under rational expectations is on the order of 10^{25} percent above steady state for reasons that have been pointed out in Carlstrom et al. (2015), whereas it is only a couple of percentage points above steady state with low levels of k.

Online Appendix B: Computing Optimal Policy

We compute optimal monetary policy via the choice of anticipated monetary policy shocks. The key idea behind the approach is that in a linear model, impulse responses are additive in the exogenous disturbances. So we can always write that a vector of endogenous variables is equal to the baseline projection plus the response to anticipated monetary policy deviations from the rule. Because these are anticipated, the approach assumes commitment. The deviations will be chosen optimally given a loss function. Let us stack the impulse responses of our n endogenous variables over the T^{**} periods in an nT^{**} vector Y:[15]

$$Y = B + DX \,. \tag{31}$$

Here, B is the nT^{**} vector containing the baseline projection given the rule and some (non-policy) shocks. The matrix D is nT^{**} by T^{**} containing the response of observables to the anticipated deviations, which are denoted by X. Specifically, in the first row, column j collects the response of the first endogenous variable in period 1 of the planning

[15] For the purpose of constructing optimal policy, we do not need to include in Y the responses of all model variables, just those that feature the loss function. Once the optimal deviations X have been constructed, the implied response of those variables not in the loss function can be constructed ex-post.

horizon to an anticipated deviation j periods out. The second row contains the response of the first variable in period two of the planning problem to those deviations, etc. The row $n+1$ contains the response of the second variable in the first period of the planning problem to the anticipated deviations, etc.

The welfare weights are contained in a nT^{**} by nT^{**} matrix W. Note that we must include discounting here. Rows $1, ..., T^{**}$ are multiplied by β^{t-1} for $t = 1, 2, ..., T^{**}$ and the same for each block of rows $(j-1)T^{**} + 1 : jT^{**}$ for $k = 1, 2, ..., n$.

The optimal policy problem is then simply to minimize the quadratic form:

$$\min_X \frac{1}{2} Y'WY = \min_X \frac{1}{2} X'D'WDX + B'WDX + t.i.p. \tag{32}$$

The optimal weights X^* are given by:

$$(X^*)' = -B'WD(D'WD)^{+}?;. \tag{33}$$

Here, $^+$ denotes the pseudo inverse. Once the weights have been determined, the path of all model variables under optimal policy is just the sum of the baseline responses plus the responses to anticipated deviations of size X^*, aka $Y^* = B + DX^*$

Incorporating the ZLB in this algorithm amounts to adding a nonlinear constraint in the optimization problem:

$$\min_X \frac{1}{2} X'D'WDX + B'WDX \tag{34}$$

subject to:

$$-D_r X \leq B_r - c \tag{35}$$

Here D_r is the matrix containing the response of the nominal interest rates to the anticipated deviations. B_r is the vector containing the baseline projection for the nominal interest rate, and c is a vector of constants containing the maximum loglinear deviation of the nominal rate from steady state admissible so that the ZLB is just reached. One can use the MATLAB function `quadprog` to solve this problem.

Online Appendix C: Proof of conditions of convergence

The iterations converge if all roots of C are within the unit circle. It can be shown that

$$\det(C) = \frac{p^2\beta\left[(1-\beta)(1-\varphi) - \frac{\tilde{\kappa}}{\sigma}\varphi\right]}{(1-p\beta)(1-p\beta\varphi)} < 0, \tag{36}$$

$$\operatorname{tr}(C) = \frac{p(1-p\beta\varphi)\left[\frac{\tilde{\kappa}}{\sigma} + (1-\beta)\right] + p\beta(1-\varphi)(1-p\beta)}{(1-p\beta)(1-p\beta\varphi)} > 0. \tag{37}$$

Given the signs of the determinant and the trace above, all roots are within the unit circle if $1 - \operatorname{tr}(C) + \det(C) > 0$. Simple, but tedious, algebra shows that this is satisfied if

$$\sigma(1-p)(1-\beta p) - p\tilde{\kappa} > 0 \tag{38}$$

This is the same condition as the parameter restriction for which the rational expectations model has a unique equilibrium, aka (26) in the main text. Hence, if the model has unique stable rational expectations equilibrium then the level-k iterations converge to that equilibrium. Otherwise, they diverge.

Online Appendix D: Comparison with a continuous model of belief revision

So far we have applied a discrete updating of beliefs. In this section we explore how our results are affected under continuous belief revisions. Overall, while there are some quantitative differences, the qualitative results are robust to this alternative model of belief revisions.

Following the algorithm outlined in García-Schmidt and Woodford (2019), we assume that the continuous updating procedure is approximated by a discrete procedure. Specifically, the belief update is given a (small) weight of γ, whereas the fraction $1-\gamma$ is the weight on the belief from the previous round of belief revision. In most of their applications, García-Schmidt and Woodford (2019) set $\gamma = 0.02$. Using their notation it holds that $n \approx k\gamma$. Hence, for $n = 2$, the average individual in the economy is expected to iterate two times, i.e., the average level-2 thinking equilibrium. Since continuous updating naturally involves more rounds of (partial) belief revisions compared to discrete belief revisions, the aggregate macroeconomic outcomes under the former procedure converge somewhat faster to the rational expectations equilibrium.

Figure 7 presents results for $n = 0.05, 0.075, 0.1$ using $\gamma = 0.02$ for the optimal monetary policy experiment. As the model is purely forward-looking we choose lower values for the belief revision n, as under this type of belief revision the model converges faster to the rational expectations solution. Qualitatively, the macroeconomic dynamics are the same as under discrete belief formation (see Figure 2).

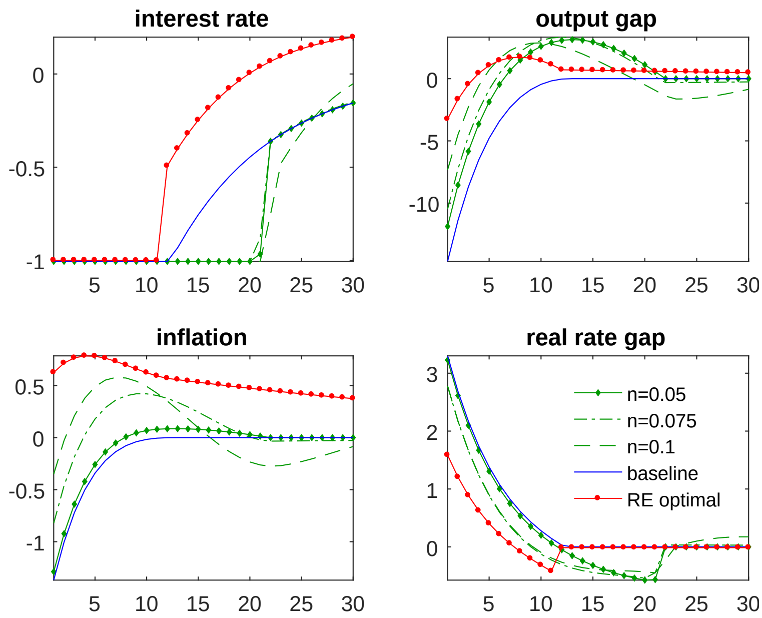

Figure 7: Optimal policy with continuous belief revisions and rational expectations (variables are in percentage-deviations from steady state). The welfare weights are normalized to one.

Figure 8 presents results for $n = 2.0$ using $\gamma = 0.02$ for the delayed lift-off experiment. With continuous belief revision, agents have revised their beliefs 100 times and this is commonly labeled as the average level-2 equilibrium outcome. As expected, compared to discrete updating (see Figure 2), the quantitative results are slightly larger under continuous updating of beliefs.

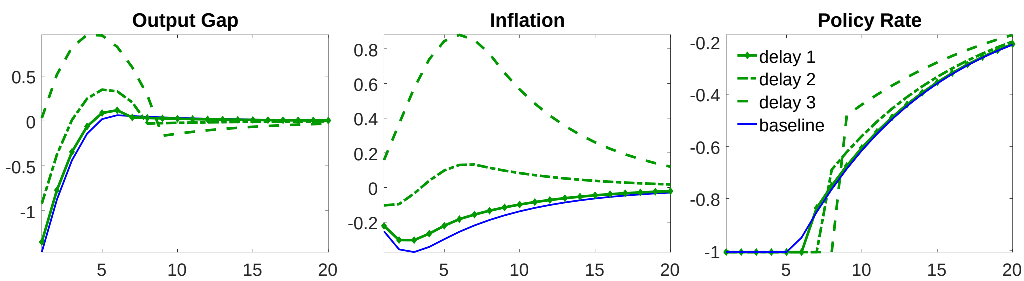

Figure 8: Continuous belief revision for $n = 2.0$ with delayed liftoff by one, two, and three quarters (variables are in percentage-deviations from steady state).

Finally, in Figure 9, we provide results for the fiscal multiplier experiment (see Figure 5 for results with discrete belief revisions). Reassuringly, the results are robust to using continuous belief revision. Interestingly, even for $n = 4$, the multipliers are still far away from the rational expectations solution indicating that in this application very high levels of cognitive belief revisions are required to match the rational expectations solution. We confirm that both with discrete and continuous belief revisions, equilibrium indeterminacy does not arise.

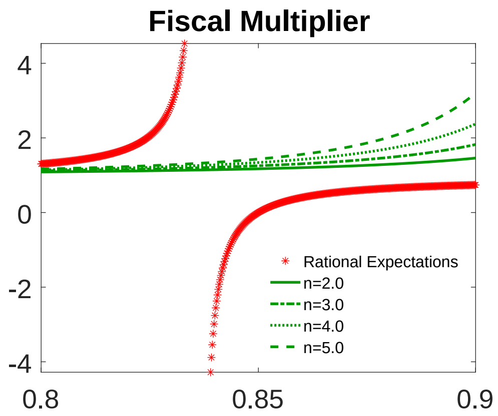

Figure 9: Fiscal multiplier as a function of the probability p