The Central Bank's Dilemma: Look Through Supply Shocks or Control Inflation Expectations?*

Paul Beaudry
University of British Columbia
and NBER

Thomas J. Carter
IMF

Amartya Lahiri
University of British Columbia

July 8, 2025

Abstract

When countries are hit by supply shocks, central banks often face the dilemma of either looking through such shocks or reacting to them to ensure that inflation expectations remain anchored. In this paper, we propose a tractable framework to capture this dilemma and explore optimal policy under a range of assumptions on how expectations are formed, including a form of bounded rationality involving level-k thinking (LKT). Despite modelling LKT in a way that nests both adaptive and rational expectations as special cases, we show that the optimal policy under LKT is qualitatively different and involves abrupt pivots in the policy stance. In particular, it is optimal for the central bank to initially look through supply shocks until a threshold is reached, then pivot discontinuously to a more hawkish anti-inflationary stance. We find that such pivots can, if optimally executed, be compatible with soft landings in the sense that most (or even all) of the reduction in inflation occurs through re-anchoring of expectations rather than economic slack. We also discuss risks and why policy errors in terms of tightening too late or too slowly can be especially costly in such an environment.

Keywords: Inflation expectations, supply shocks, monetary policy

JEL Classification: E3, E5

*We would like to thank Roc Armenter, Jenny Chan, José Dorich, Frank Smets, Ludwig Straub and George-Marios Angeletos for comments. Paul Beaudry and Amartya Lahiri would like to acknowledge financial support from SSHRC Canada. The views expressed are those of the authors and do not necessarily reflect the views of the IMF, its Executive Board, or IMF management.

1 Introduction

How best to respond to supply shocks is a classic topic in monetary economics, and one that advanced economies' recent experiences with inflation have brought back to the forefront. Starting in mid-2021, a series of supply shocks led central banks in many countries to ask similar questions:[1] should they look through supply shocks – even temporarily – at the risk of de-anchoring inflation expectations; or should they immediately tighten monetary policy to keep expectations anchored, albeit at the risk of hindering the recovery from the COVID pandemic? More recently, the possibility of large shifts in the global trade landscape has rekindled discussion on the best way to manage inflation expectations in the presence of large supply shocks, and how policymakers should weigh this priority against other objectives.[2]

To adequately address the question of the appropriate policy response to supply shocks, it is important to tackle them in a framework rich enough to capture both the costs and benefits of looking through these shocks. In our view, this requires two key modelling ingredients: (i) an explicit mechanism for why looking through supply shocks can be attractive to policy-makers under certain circumstances; and (ii) a flexible approach to modelling inflation expectations that allows for the possibility that looking through supply shocks could lead to de-anchoring. While each of these elements is present in the literature, the aim of this paper is to bring them together in a setup that is sufficiently tractable to allow for an explicit characterization of the relevant policy trade-offs and their optimal resolution under a range of assumptions on the expectation formation process.

To preview our model's main features, the two key ingredients that we just described will respectively take the form of wage rigidities and a model of bounded rationality that involves level-k thinking but nests both adaptive and rational expectations as special cases. With regard to the first of these key ingredients, one might be tempted to instead turn to the canonical three-equation New Keynesian (NK) model, which assumes a combination of sticky prices and flexible wages and plays a central role in many monetary policy discussions. However, as we are reminded in Reis (2022), an incentive to look through supply shocks only arises in this context when these shocks are conflated with mark-up shocks, which we view as inappropriate given that many of the shocks contributing to the recent inflation surge were "true" supply shocks with direct bearing on the productive capacities of many economies (e.g., supply chain disruptions, commodity price

[1] Recent work in di Giovanni et al. (2022) provides empirical evidence for the significant role of supply shocks in accounting for the recent inflation surges in the euro area and the United States.

[2] See, e.g., Macklem (2025).

shocks).[3] When facing such shocks, the prescription of the canonical NK model is instead that policy-makers should *never* look through them, even if they are temporary, since a policy that adjusts output one-for-one with its natural level will achieve a "divine coincidence" under which the inflation and output gaps are simultaneously closed. As a result, we view the canonical model as an inadequate laboratory for our analysis but recognize that wage rigidities are well known to overturn its stark prescriptions.[4] This makes a model with sticky wages (or more generally with stickier wages than prices) a more promising laboratory – in addition to offering more realism to the extent that empirical estimates generally favour higher degrees of wage rigidity than price rigidity.

To get more specific, we model wage rigidity by considering an economy subject to productivity shocks and assuming that wages are set before observing the current shock, while prices are set after. This creates potential for wage-price spirals in the sense that wage-setters set wages based on the prices they expect firms to set, and these wages then become part of the costs that firms factor into their pricing decisions. This configuration also gives rise to a Phillips curve in which productivity shocks enter as a residual, and policy-makers will generally have an incentive to look through these shocks to avoid variation in employment and output.

However, the cost of looking through supply shocks will be that it could lead to a de-anchoring in inflation expectations, and this brings us to the second of the model's key ingredients. Here our goal is to understand how outcomes and optimal policy are shaped by the degree of sophistication with which agents form expectations – especially with regard to the way that they account for the central bank's announced policies in their inflation expectations. To do so, we rely on level-k thinking (LKT), which is a form of bounded rationality that, as we detail below, enjoys significant support in the experimental literature and has recently been applied to a range of topics in macroeconomics. We specifically model LKT in such a way that expectations are adaptive at level $k = 0$, while the "maximum sophistication" limit $k \to \infty$ corresponds to the case of rational expectations, where agents can fully account for the effects of monetary policy. In contrast, intermediate values of k lead to a situation where agents' expectations are based on a combination of backward- and forward-looking elements, with the latter reflecting an imperfect understanding of the way that monetary policy impacts the economy.

With these ingredients in place, we examine how monetary policy should respond to supply

[3] In the model, supply shocks are specifically assumed to take the form of productivity shocks, but these can be interpreted as subsuming a range of potential supply shocks. For example, in the Online Appendix, we establish a formal equivalence between productivity shocks and oil price shocks in the context of a small, oil-importing economy.

[4] This is an example of how models with sticky prices and sticky wages often deliver similar prescriptions with regard to the management of demand shocks, but can differ considerably in the case of supply shocks.

shocks when policy-makers aim to minimize a loss function that penalizes deviations of inflation from target and employment from its natural level. To set the stage, we first show that under rational expectations (RE), it is optimal for central banks to fully look through supply shocks. In contrast, the optimal policy under adaptive expectations (AE) is generally tighter but involves a constant degree of "look-through" in the precise sense that the rate at which policy-makers make trade-offs between the inflation and employment gaps never changes over time.

Our main analysis focuses on the intermediate case where agents form expectations under LKT. In this case, the optimal policy that emerges is qualitatively very different from those emerging under RE and AE. In particular, one of our main results is to identify conditions under which the optimal policy under LKT involves a type of "pivoting" behaviour in the sense that policy-makers initially opt to mostly look through supply shocks, but then switch suddenly to a very aggressive policy response if shocks cumulate beyond some threshold level.

These abrupt pivots occur because LKT gives rise to a non-convexity in the central bank's loss function. The context for this non-convexity is that the employment impacts of policy tightening work through two off-setting channels in the model economy: (i) a "direct" channel whereby tighter policy reduces economic activity at a given level of private-sector expectations; and (ii) an "indirect" channel whereby tighter policy helps to re-anchor inflation expectations, thus reducing the amount of *actual* economic slack that policy-makers have to engineer to stabilize realized inflation. All else equal, the indirect channel will tend to be stronger during periods when the economy has recently experienced large inflation overshoots, since in this case expectations will already be elevated due to the backward-looking elements on which they are partially based under LKT. However, it is also possible for the indirect channel to get stronger *as policy-makers tighten*, since we find that a key property of private-sector expectations under LKT is that they are more sensitive to changes in monetary policy if policy is already tight to begin with.[5] This is the underlying source of non-convexity, and it contributes to a situation where there are benefits to tightening by a sufficiently wide margin, but doing so only after the economy has experienced sufficiently large overshoots.

Another key result is that with the direct and indirect channels above both in play, pivots need not be associated with hard landings in the sense of triggering large contractions. In particular, if the indirect channel is strong enough, then pivots could help to re-anchor expectations without requiring a significant amount of slack. In this respect, we highlight a special case where the two

[5] In the Online Appendix, we show that this key property also holds under models of bounded rationality involving reflective expectations (García-Schmidt and Woodford, 2019) and cognitive hierarchies (Camerer et al., 2004).

channels offset one another exactly in expectation, implying that pivots have no impact on the expected level of employment. Outside this special case, we also find that expected impacts on employment are consistently small in numerical exercises. In this sense, the model predicts that optimally-executed pivots should be compatible with soft landings, at least in expectation.

However, this is not to say that pivots are costless, since they also increase the variance of employment by making monetary policy more responsive to supply shocks. As a result, pivots are risky and involve accepting some probability that poor draws of the supply shock could lead to hard landings. In simulations, we further show that the costs of mis-timing pivots can be substantial, as can the costs of attempting to achieve similar degrees of tightening in more gradual ways.

Three points regarding our approach and results are in order at this stage. First, to make maximal use of the LKT framework as a way of exploring the risk of de-anchoring, our approach downplays the intertemporal coordination frictions that typically arise in NK models due to staggered price- and wage-setting, instead putting the focus on *intratemporal* coordination frictions. In particular, we assume that prices and wages reset each period, though in the case of wages this occurs before agents have an opportunity to observe the current productivity shock. As a result, the key coordination issue in our model is that wage-setters have to form expectations on prices, which will then be set by firms following a mark-up rule that partly depends on the average wage. Wage-setters are thus effectively forced to forecast each other's expectations. This gives rise to a static coordination game to which LKT can readily be applied and within which monetary policy serves as a coordination mechanism for nominal choices.

Second, in order to focus on the policy dilemmas presented by supply shocks, we mostly disregard the role of demand shocks in our analysis. We adopt this focus not to be dismissive of the likelihood that demand forces played an important role in the recent inflation surges experienced in many countries. Rather, the choice reflects the fact that our approach has nothing novel to contribute regarding the demand side of the economy. Accordingly, one should view our results as describing the optimal adjustments that supply shocks dictate, over and above the central bank's underlying response to demand shocks. We discuss and formalize this interpretation in Section 6.

Third, it is important to bear in mind that the issue of how to respond to supply-driven inflation shocks is not unique to recent times, nor to industrial countries. Arguably, the question is of even greater importance in emerging economies, where food and fuel expenditures comprise a much larger share of consumption expenditure. Given that food and energy prices are volatile and often driven by local or global supply shocks, inflation management becomes a much trickier exercise for

central banks in these countries. Questions related to whether a central bank should look through or react to inflation movements that are driven by supply-side developments are thus recurrent and highly germane in emerging economies.[6]

Our work connects with four distinct but related strands of the literature. The first two of these strands are the voluminous literature on NK macroeconomic models and the nascent literature on the recent global inflation episode. Comprehensive summaries of the former can be found in Galí (2015) and Woodford (2004), amongst many others, while Reis (2022) provides an excellent overview of the various hypotheses regarding the recent episode, the role of the inflation anchor and the associated challenges for central banks.

Another related literature focuses on bounded rationality, LKT, and their implications for macroeconomics. In experimental game theory, an extensive body of laboratory experiments offer evidence for LKT, with most estimates pointing to low k values that place agents far from the RE benchmark (among many others, see Arad and Rubinstein, 2012; Costa-Gomes et al., 2013; Kneeland, 2015). Moreover, LKT and the closely related concept of "reflective expectations" have now been applied to a wide range of macroeconomic questions in both monetary contexts (e.g., Farhi and Werning, 2019; García-Schmidt and Woodford, 2019; Iovino and Sergeyev, 2023) and fiscal ones (Bianchi Vimercati, Eichenbaum, and Guerriro, 2021).[7]

Finally, our findings regarding soft landings connect with the literature on factors that contribute to successful disinflations (e.g., Hazell et al., 2022; Sargent, 1982). Within this literature, a common theme is that effective expectation management can make it possible for disinflations to occur faster than slack measures would suggest,[8] with the expectational mechanisms commonly emphasized having to do with perceptions around the long-term credibility of policy-makers. Though such mechanisms clearly contributed to past disinflations, their applicability to the recent episode is less clear given that long-run inflation expectations were relatively stable over this episode. In this sense, one of our contributions is to show how LKT gives rise to a complementary mechanism for how disinflations can be achieved without large shifts in *both* slack *and* long-term expectations.[9]

[6] As noted in a previous footnote, the productivity shocks on which the model focuses can be interpreted as the mirror image of oil price shocks, which are one of the more volatile and frequent shocks hitting both emerging and advanced economies. See the Online Appendix for details.

[7] In the Online Appendix, we confirm that our main results hold under reflective expectations and also do so under the "cognitive hierarchy" framework proposed by Camerer et al. (2004).

[8] See also Reis (2021) on this issue.

[9] Our paper is also related to some recent work by Lorenzoni and Werning (2023), who study the phenomenon of wage-price spirals in a general setting and explore implications for policy. One way in which our work is distinct from theirs is that we focus on an environment with boundedly rational agents.

The rest of the paper is organized as follows. Section 2 presents the model, while Section 3 formalizes the optimal policy problem and presents special cases of the problem that can be solved analytically. Section 4 presents the full dynamic policy problem and numerical simulations of the optimal policy. Section 5 explores the consequences of policy errors, while Section 6 discusses demand shocks. Section 7 then elaborates on the model's interpretation of the recent inflation surge, while Section 8 concludes. Most of our proofs are collected in an Appendix, and additional material is provided in an Online Appendix.

2 Model

We consider a simple economy in which the key players are: (i) a set of households that supply labour services on a monopolistically competitive basis; (ii) a set of firms that supply goods on a monopolistically competitive basis; and (iii) a central bank that sets monetary policy with an eye toward minimizing an ad-hoc loss function.

As we stressed in the introduction, an important element will be that wages are set one period in advance, while prices are flexible, and this will be key to delivering a Phillips curve that differs from canonical NK models and helps to rationalize why looking through supply shocks can sometimes be warranted. The absence of multi-period nominal rigidities will also allow us to focus on the cross-sectional coordination problems that arise among price- and wage-setters at each point in time. In this sense, the model can be seen as being closer in spirit to the original environment of Blanchard and Kiyotaki (1987).

Before we turn to the model's components, it is important to note that expectation operators should not be assumed to be rational unless this is specifically indicated, since our ultimate goal is to understand the policy implications of different expectation formation processes. To be as clear as possible in this area, we will attach superscripts to all expectations to convey relevant information about *whose* expectations they represent, or *how* they are being formed.

2.1 Firms

We assume a two-stage production process, with distinct final and intermediate goods. The details on each stage are as follows.

2.1.1 Final good producers

Final goods are produced from a unit measure of intermediate goods using a technology of the form

$$Y_t = \left(\int_0^1 Y_{ft}^{\frac{\gamma-1}{\gamma}} df \right)^{\frac{\gamma}{\gamma-1}},$$

where Y_t denotes output of final goods; Y_{ft} denotes purchases of intermediate good $f \in [0,1]$; and $\gamma > 1$ denotes the elasticity of substitution among intermediate goods. Assuming that final goods producers are perfectly competitive and make zero profits, their demand for intermediate good f is then given by

$$Y_{ft} = Y_t \left(\frac{P_{ft}}{P_t} \right)^{-\gamma}, \tag{2.1}$$

where

$$P_t \equiv \left(\int_0^1 P_{ft}^{1-\gamma} df \right)^{\frac{1}{1-\gamma}}$$

measures the aggregate price level.

2.1.2 Intermediate good producers

Each intermediate good is generated from a unit measure of labour services using a technology of the form

$$Y_{ft} = \theta_t \left(\int_0^1 N_{sft}^{\frac{\rho-1}{\rho}} ds \right)^{\frac{\rho}{\rho-1}},$$

where N_{sft} denotes hiring of service $s \in [0,1]$; $\rho > 1$ denotes the elasticity of substitution among services; and θ_t denotes an aggregate productivity level, which we assume to be stochastic and the only source of randomness in the economy. While the particular form of the process governing θ_t is not critical to our results, it will sometimes be convenient to assume that θ_t follows a random walk in logarithms – specifically,

$$\ln \theta_t = \ln \theta_{t-1} + \epsilon_t,$$

where ϵ_t is an i.i.d. normal shock with mean zero and variance σ_θ^2.

We assume that intermediate good producers take wages as given but set their own prices on a monopolistically competitive basis. As a result of the wage-taking assumption, a given producer's

demand for labour service s is given by

$$N_{sft} = \left(\frac{Y_{ft}}{\theta_t}\right)\left(\frac{W_{st}}{W_t}\right)^{-\rho},$$

where

$$W_t \equiv \left(\int_0^1 W_{st}^{1-\rho} ds\right)^{\frac{1}{1-\rho}} \tag{2.2}$$

measures the aggregate wage level. The producer's implied profits are in turn given by

$$P_{ft}Y_{ft} - \int_0^1 W_{st}N_{sft}ds = \left(P_{ft} - \frac{W_t}{\theta_t}\right)Y_{ft}. \tag{2.3}$$

Turning to the pricing behaviour of intermediate good producers, our key assumption here is that prices are flexible in the sense that firms can adjust them every period and do so after observing the current level of productivity. As a result, the optimization problem facing a given firm involves choosing (Y_{ft}, P_{ft}) to maximize the profit function on line 2.3, subject to the demand curve on line 2.1. This leads to a standard mark-up rule of the form

$$P_{ft} = \frac{\gamma}{\gamma - 1}\frac{W_t}{\theta_t}, \tag{2.4}$$

implying an inverse relationship between prices and productivity that will carry over to an inverse relationship between productivity and inflation at the aggregate level. Another corollary of this mark-up rule is that all intermediate good producers behave identically in the sense that $(P_{ft}, Y_{ft}) = (P_t, Y_t)\ \forall f \in [0,1]$, so overall demand for labour service s is given by

$$\int_0^1 \left(\frac{Y_{ft}}{\theta_t}\right)\left(\frac{W_{st}}{W_t}\right)^{-\rho} df = \left(\frac{Y_t}{\theta_t}\right)\left(\frac{W_{st}}{W_t}\right)^{-\rho}. \tag{2.5}$$

2.2 Households

We assume a unit measure of households, which are indexed by $h \in [0,1]$ and each supply a range of labour services. A given household's utility is given by

$$\mathbb{E}_t^{HH(h)}\left[\sum_{\tau=0}^{\infty}\beta^\tau\left(\ln C_{h,t+\tau} - \eta\int_0^1 \tilde{N}_{s,h,t+\tau}ds\right)\right], \tag{2.6}$$

where $\mathbb{E}_t^{HH(h)}(\cdot)$ denotes the expectations of the particular household in question; $\beta \in (0,1)$ is a discount factor; $\eta > 0$ is a parameter scaling the disutility of labour; C_{ht} denotes the household's consumption; and $\tilde{N}_{sht}$ denotes the total hours of work that the household supplies under labour service s. The tilde above this last variable is thus meant to distinguish a "supply-side" labour quantity from the "demand-side" quantities considered in our characterization of firms' behaviour, and this notational convention will be maintained throughout the remainder of this section.

When optimizing, households face a sequence of budget constraints of the form

$$\int_0^1 W_{st}\tilde{N}_{sht}ds + D_t + T_t + B_{h,t-1} = P_t C_{ht} + \frac{B_{ht}}{1+i_t}, \tag{2.7}$$

where D_t denotes dividends received from firms (which are assumed to be distributed equally across households); T_t denotes transfers received from the government (which are distributed on a similar basis); B_{ht} denotes the face value of the household's purchases of risk-free, one-period bonds; and i_t denotes the nominal interest rate on those bonds.

We assume that supply decisions for each labour service are made by service-specific unions, while the household's overall consumption-saving decision is delegated to a household head. The details on these decisions are as follows.

2.2.1 Consumption-saving decision

The consumption-saving decision is assumed to take place after observing the current level of productivity in the economy. This leads to an Euler equation of the form

$$\beta \mathbb{E}_t^{HH(h)} \left(\frac{C_{ht}}{C_{h,t+1}} \frac{1+i_t}{P_{t+1}/P_t} \right) = 1, \tag{2.8}$$

where the subscript in the expectation $\mathbb{E}_t^{HH(h)}(\cdot)$ indicates that the current productivity level θ_t is part of the information set to which the household head had access when forming this expectation.

2.2.2 Labour-supply decision

Turning to labour unions, we assume that they set the wages for their respective services on a monopolistically competitive basis, taking as given the behaviour of all other unions, along with that of household heads. Unlike our treatment of prices, we also assume a modest degree of wage stickiness in the sense that unions can adjust their wages every period but must do so before

observing the current level of productivity.

For the union representing a given service $s \in [0,1]$, the relevant program involves choosing a nominal wage rate W_{st}, which the labour demand curve on line 2.5 will then associate with a certain number of hours $\tilde{N}_{st}$ to be worked by each of the union's members – i.e., $\tilde{N}_{hst} = \tilde{N}_{st} = (Y_t/\theta_t)(W_{st}/W_t)^{-\rho}\ \forall h \in [0,1]$. When making this decision, the union is assumed to value wages and the disutility of labour based on the marginal preferences of its members. This leads to an objective of the form

$$\mathbb{E}_{t-1}^{WS(s)}\left(W_{st}\tilde{N}_{st}\int_0^1 \frac{1}{P_t C_{ht}}dh - \eta\tilde{N}_{st}\right), \tag{2.9}$$

where $\mathbb{E}_{t-1}^{WS(s)}(\cdot)$ denotes the expectation of the union setting wages for service s, conditional on an information set that excludes θ_t; and the integral $\int_0^1[1/(P_tC_{ht})]dh$ measures the average household's marginal utility of nominal income. The implied wage-setting rule can then be written as

$$W_{st} = \frac{\rho\eta}{\rho-1}\frac{\mathbb{E}_{t-1}^{WS(s)}\left(\frac{Y_t}{\theta_t}\right)}{\mathbb{E}_{t-1}^{WS(s)}\left(\frac{Y_t}{\theta_t}\int_0^1 \frac{1}{P_tC_{ht}}dh\right)}. \tag{2.10}$$

The expectations on the right-hand side of this last expression will be key objects in our analysis, and we will entertain a range of assumptions on how they are formed. However, a key feature common to all of the models of expectation formation that will be considered in the main text is that they each predict that all wage-setters should share common, *though potentially incorrect*, expectations on all aggregate outcomes in the economy[10] – i.e., $\mathbb{E}_{t-1}^{WS(s)}(X_t) = \mathbb{E}_{t-1}^{WS(\tilde{s})}(X_t)$ for an arbitrary aggregate variable X_t and any pair of wage-setters $(s,\tilde{s}) \in [0,1]^2$. Given the form of the expectations on the right-hand side of line 2.10, this implies that all wage-setters should behave identically in the sense of selecting a common wage rate – i.e., $W_{st} = W_t\ \forall s \in [0,1]$. This allows us to re-express the wage-setting condition on line 2.10 in terms of the aggregate wage rate and a set of common expectations $\mathbb{E}_{t-1}^{WS}(\cdot)$ that describe the subjective beliefs of all wage-setters in the economy:

$$W_t = \frac{\rho\eta}{\rho-1}\frac{\mathbb{E}_{t-1}^{WS}\left(\frac{Y_t}{\theta_t}\right)}{\mathbb{E}_{t-1}^{WS}\left(\frac{Y_t}{\theta_t}\int_0^1 \frac{1}{P_tC_{ht}}dh\right)}. \tag{2.11}$$

[10] At the same time, we stress that common expectations are not critical to our results. For example, we show in the Online Appendix that our main results carry over to a version of the model in which we allow for heterogeneity in wage-setters' expectations using the "cognitive hierarchy" framework of Camerer et al. (2004).

2.3 Phillips and IS curves

Before turning to the model's policy side, it is now convenient to derive the Phillips and IS curves associated with the private-sector outcomes above. As a first step, we also solve for the "natural" levels of output and employment that would prevail in a counterfactual economy with flexible wages.

2.3.1 A symmetric, flexible-wage benchmark

Consider a counterfactual economy in which wages are flexible in the sense of being set after observing the current level of productivity, and in which consumption is equally distributed across households. This is equivalent to replacing expectation operators with identity maps on the right-hand side of the wage-setting condition on line 2.11, while also setting $C_{ht} = Y_t \ \forall h \in [0,1]$. This leads to a new, flexible wage-setting condition of the form

$$W_t = \frac{\rho\eta}{\rho - 1} P_t Y_t.$$

At the same time, the mark-up rule on line 2.4 can be written as follows after multiplying both sides by total output and defining $N_t \equiv Y_t/\theta_t$:

$$P_t Y_t = \frac{\gamma}{\gamma - 1} W_t N_t.$$

Combining these last two expressions then allows us to pin down the natural level of employment:

$$N_t = \frac{\rho - 1}{\rho\eta} \frac{\gamma - 1}{\gamma} \equiv \bar{N}.$$

The natural level of employment is thus constant and intuitively decreasing in the disutility of labour, as well as the degrees of monopoly distortion assumed in the goods and labour markets. The implied level of potential output is $\bar{Y}_t \equiv \theta_t \bar{N}$.

2.3.2 Phillips curve

We now turn our attention back to the "true" wage-setting condition on line 2.11. By combining this expression with the mark-up rule on line 2.4 and then linearizing around the symmetric, flexible-wage benchmark derived above, we reach a Phillips curve of the form

$$\hat{\pi}_t = \mathbb{E}_{t-1}^{WS}(\hat{N}_t) + \mathbb{E}_{t-1}^{WS}(\hat{\pi}_t) - \hat{\theta}_t, \tag{2.12}$$

where $\hat{\pi}_t \equiv \ln P_t - \ln P_{t-1} - \pi^*$ denotes the deviation of inflation from some target value π^*; $\hat{N}_t \equiv \ln N_t - \ln \bar{N}$ denotes the employment gap; and $\hat{\theta}_t \equiv \ln \theta_t - \mathbb{E}_{t-1}^{WS}(\ln \theta_t)$ can be interpreted as a supply shock.

Some important features of this Phillips curve should be noted. The first is that the residual in equation 2.12 is a "true" supply shock rather than a mark-up shock. This is a key dimension in which this Phillips curve differs from the canonical NK Phillips curve and is a consequence of our focus on wage rigidities; had we instead assumed sticky prices but flexible wages, then it can easily be confirmed that this term would not appear in the Phillips curve. At the same time, the fact that the expression on line 2.12 can be derived without taking a specific stand on the way that wage-setters form their expectations is a consequence of our focusing on a setting where nominal frictions arise from a Lucas (1972)-style assumption that wages are set one period in advance. As explained in Preston (2005), this is a convenient property of Lucas-style nominal rigidities, whereas the Calvo (1983)-style rigidities typically assumed in the NK literature make the precise form of the implied supply curve more sensitive to the details of how private-sector expectations are formed.

2.3.3 IS curve

Turning to the model's demand side, the same arguments presented in Preston (2005) make it difficult to derive a single IS curve that holds across a range of potential choices on the way that one could specify the household expectations entering into the Euler equation on line 2.8. As a result, we pursue a two-pronged strategy under which we: (i) first focus on the standard case where household expectations are fully rational for the purposes of the main text; then (ii) use Section A of the Online Appendix to show how our results generalize to a broader range of expectational assumptions. For the purposes of the main text, this leads to a standard IS curve of the following form after linearizing equation (2.8) and accounting for market clearing:

$$\hat{N}_t = \mathbb{E}_t^{RE}(\hat{N}_{t+1}) - [\hat{\imath}_t - \mathbb{E}_t^{RE}(\hat{\pi}_{t+1})], \tag{2.13}$$

where $\mathbb{E}_t^{RE}(\cdot)$ denotes a rational expectation; and $\hat{\imath}_t \equiv i_t - (\pi^* - \ln \beta)$ denotes the deviation of the nominal rate from its steady-state value. However, we stress from the outset that the particular assumptions made with regard to household expectations have no bearing on our main results (and will shortly expand on this point in Subsection 2.5).

2.4 Monetary and fiscal policy

We now close the model by specifying the behaviour of the central bank and fiscal authority. The details on each are as follows.

2.4.1 Central bank

On the monetary policy side, we could immediately specify an interest rate rule and later optimize on the parameters in that rule. However, it will be more convenient to think of the central bank as directly controlling the level of employment and setting it as a (potentially time-varying) function of the prevailing rate of inflation.

This will involve a feedback rule of the form

$$\hat{N}_t = -\phi_t \hat{\pi}_t, \tag{2.14}$$

where ϕ_t is a policy stance reflecting the rate at which policy-makers are prepared to tighten in response to off-target inflation. As we detail in our next section, policy-makers are assumed to update this policy stance each period, announcing its new value before observing the current level of productivity, and doing so with an eye toward minimizing an ad-hoc loss function. By allowing for time variation in the policy stance and the implied level of feedback from $\hat{\pi}_t$ to $\hat{N}_t$, an approach along these lines will ultimately enable us to identify the conditions under which policy-makers find it optimal to "pivot" in the sense of shifting from low values of ϕ_t at which they do not respond much to inflation to high values associated with much stronger reactions.

One way to interpret the feedback rule above is by forward-solving the IS curve – i.e.,

$$\hat{N}_t = -\mathbb{E}_t^{RE}\left[\sum_{\tau=0}^{\infty}(\hat{\imath}_{t+\tau} - \hat{\pi}_{t+\tau+1})\right].$$

This equation allows us to express the current employment gap as a negative function of the cumulative sum of all expected future deviations of the real interest rate from its steady-state value. As a result, this cumulative sum can be viewed as a measure of the total stimulus being provided by monetary policy, and equation (2.14) can then be interpreted as the product of a rule according to which policy-makers adjust this measure of stimulus as a time-varying function of the

prevailing inflation rate – that is,

$$\phi_t \hat{\pi}_t = \mathbb{E}_t^{RE} \left[\sum_{\tau=0}^{\infty} (\hat{\imath}_{t+\tau} - \hat{\pi}_{t+\tau+1}) \right]. \tag{2.15}$$

Given the importance that real-world policy-makers place on communicating a policy stance that goes beyond the current value of the policy rate, an interpretation along these lines can be viewed as more realistic and encompassing relative to more standard policy rules.

2.4.2 Fiscal authority

Turning finally to fiscal policy, we assume for concreteness that the government operates a balanced budget each period and keeps bonds in zero net supply. However, none of these assumptions are essential to our results, which go through for a wide range of Ricardian fiscal policies.

2.5 Summary of equilibrium conditions

To summarize, the model's key variables are the inflation gap $\hat{\pi}_t$, employment gap $\hat{N}_t$, nominal rate gap $\hat{\imath}_t$, and policy stance ϕ_t. In equilibrium, these four unknowns are jointly pinned down by the Phillips and IS curves on lines 2.12 and 2.13, the policy feedback rule on line 2.14, and the solution to a policy problem on which we elaborate in the next section.

A key property of this system is that the model can be solved recursively in the sense that one can first solve for $(\hat{\pi}_t, \hat{N}_t, \phi_t)$ independent of the IS curve, then use the IS curve to back out a value for $\hat{\imath}_t$. This is the approach that we take below, and it importantly makes all of our main results independent of the details of the IS curve, including the way that households are assumed to form expectations at the time they make their consumption-saving decisions. Instead, the key private-sector expectations in our analysis will be those of wage-setters, and much of our attention will be on how these are influenced by monetary policy under different models of expectation formation.

3 Expectations and optimal policy

We now consider the design of optimal monetary policy in the model economy, and how it is shaped by different assumptions on the expectation formation process in the private sector. However, before tackling these issues, it will be useful to lay out the theories of expectation formation that we will consider and the outcomes to which they would lead for arbitrary values of the policy stance

ϕ_t. Several additional models of expectations formation are considered in the Online Appendix, including models involving hybrid Phillips curve, econometric learning, and various forms of bounded rationality beyond LKT.

3.1 Expectations

3.1.1 Rational expectations (RE)

If we assume that wage-setters' expectations are rational, then taking $(t-1)$-dated expectations on both sides of the Phillips curve on line 2.12 gives $\mathbb{E}_{t-1}^{WS}(\hat{N}_t) = 0$. Combining this result with the feedback rule on line 2.14 then further yields $\mathbb{E}_{t-1}^{WS}(\hat{\pi}_t) = 0$ so long that $\phi_t > 0$. Substituting these results back into equations (2.12) and (2.14), we see that

$$(\hat{\pi}_t, \hat{N}_t) = (-\hat{\theta}_t, \phi_t \hat{\theta}_t),$$

so inflation collapses to an i.i.d. process and lies above (below) its target value whenever the supply shock is negative (positive). Any $\phi_t > 0$ thus suffices to fully anchor inflation expectations at target, and employment is also fully stabilized in the limit $\phi_t \to 0$.

3.1.2 Adaptive expectations (AE)

For this case, we consider the simplest possible version of adaptive expectations, with

$$[\mathbb{E}_{t-1}^{WS}(\hat{\pi}_t), \mathbb{E}_{t-1}^{WS}(\hat{N}_t)] = (\hat{\pi}_{t-1}, \hat{N}_{t-1}). \tag{3.16}$$

Under this assumption, the inflation and employment gaps are respectively given by the following:

$$\begin{aligned} \hat{\pi}_t &= \hat{N}_{t-1} + \hat{\pi}_{t-1} - \hat{\theta}_t \\ \hat{N}_t &= -\phi_t(\hat{N}_{t-1} + \hat{\pi}_{t-1} - \hat{\theta}_t). \end{aligned}$$

As a result, inflation collapses to a random walk in the special case where policy-makers opt to fully look through inflation pressures in the sense of setting $\phi_t = 0$. Contrasting this result with the rational case above illustrates how the mapping between policy and inflation outcomes can vary drastically depending on the nature of private-sector expectations.[11]

[11] Given such a stark contrast with the fully rational case, it is natural to explore several potential extensions of this simple adaptive benchmark. In this respect, Section D of the Online Appendix focuses on an extension where a

3.1.3 Level-k thinking (LKT)

Turning finally to LKT, our goal here is to consider a more general framework for expectation formation that can accommodate RE and AE within a broader range of possibilities. LKT provides one such framework.

To motivate LKT, it is useful to note from the wage-setting condition on line 2.10 that wage-setters in this economy face an important coordination problem. This is because they each need to forecast aggregate outcomes that depend on the choices of all other wage-setters via the aggregate wage W_t entering into firms' mark-up rule. As a result, each wage-setter must "forecast the forecasts" of all other wage-setters.

Under LKT, "mutual forecasting" problems of this sort are formalized as iterative processes. At the first stage $\tilde{k} = 1$, agents posit some initial guess on the expectations of all other agents and compute the aggregate outcomes that would emerge in this case. At the next stage $\tilde{k} = 2$, the guess on others' expectations is updated to reflect the aggregate outcomes implied by the previous guess and used as the initial guess of a new iteration. This process is then repeated up to some finite $\tilde{k} = k$, with each iteration thus accounting for a new layer of higher-order expectations. The restriction to a finite number of iterations reflects some bounded computing power on the part of individual agents on account of limited resources or capacity for forecasting others' behaviour.

In the case of our model economy, the first stage of this process involves positing some initial (level-zero) guesses $[\mathbb{E}^0_{t-1}(\hat{\pi}_t), \mathbb{E}^0_{t-1}(\hat{N}_t)]$, which are then updated via the Phillips curve and the central bank's feedback rule – i.e.,

$$\begin{aligned}\mathbb{E}^1_{t-1}(\hat{\pi}_t) &= \mathbb{E}^0_{t-1}(\hat{N}_t) + \mathbb{E}^0_{t-1}(\hat{\pi}_t)\\ \mathbb{E}^1_{t-1}(\hat{N}_t) &= -\phi_t[\mathbb{E}^0_{t-1}(\hat{N}_t) + \mathbb{E}^0_{t-1}(\hat{\pi}_t)],\end{aligned}$$

where we have assumed that LKT has no impact on agents' expectations on purely exogenous variables, so $\mathbb{E}^0_{t-1}(\hat{\theta}_t) = \mathbb{E}^1_{t-1}(\hat{\theta}_t) = 0$. Iterating on this recursion up to the final level k then yields expectations of the following form:

$$\mathbb{E}^k_{t-1}(\hat{\pi}_t) = (1-\phi_t)^{k-1}[\mathbb{E}^0_{t-1}(\hat{N}_t) + \mathbb{E}^0_{t-1}(\hat{\pi}_t)] \tag{3.17}$$

$$\mathbb{E}^k_{t-1}(\hat{N}_t) = -\phi_t(1-\phi_t)^{k-1}[\mathbb{E}^0_{t-1}(\hat{N}_t) + \mathbb{E}^0_{t-1}(\hat{\pi}_t)]. \tag{3.18}$$

mix of adaptive and rational wage-setters give rise to a hybrid Phillips curve. On the other hand, Section E shows how our results generalize to cases where we replace the simple backward-looking expectations on line 3.16 with more sophisticated backward-looking rules that embody various forms of econometric learning.

Throughout the main text, we will assume that the initial guesses seeding this process take the form $\mathbb{E}^0_{t-1}(\hat{\pi}_t) = \hat{\pi}_{t-1}$ and $\mathbb{E}^0_{t-1}(\hat{N}_t) = \hat{N}_{t-1}$ – i.e., agents use a simple adaptive benchmark as the starting point for their reasoning.[12] Under this assumption, feeding the final, level-k expectations above into the model's "true" Phillips curve and feedback rule leaves us with the following equilibrium system:

$$\hat{\pi}_t = (1-\phi_t)^k(\hat{N}_{t-1} + \hat{\pi}_{t-1}) - \hat{\theta}_t$$
$$\hat{N}_t = -\phi_t[(1-\phi_t)^k(\hat{N}_{t-1} + \hat{\pi}_{t-1}) - \hat{\theta}_t].$$

Comparing this system with the equilibrium system that was shown to emerge under AE, we see that the two coincide when $k = 0$. On the other hand, when k goes to infinity, the system collapses to $(\hat{\pi}_t, \hat{N}_t) = (-\hat{\theta}_t, \phi_t\hat{\theta}_t)$ so long that $\phi_t \in (0, 1]$, and we thus recover the RE solution that we described earlier. In these senses, LKT can be viewed as nesting AE and RE as opposite ends of a common iterative spectrum.[13]

3.2 Optimal policy

We now turn our attention to the policy problem that the central bank faces when selecting the policy stance ϕ_t. Here we assume that policy-makers announce their stance before observing the current productivity level θ_t and do so with eye toward minimizing an ad-hoc, quadratic loss function of the form[14]

$$\mathbb{E}^{RE}_{t-1}\left[\sum_{\tau=0}^{\infty}\beta^{\tau}_{CB}\left(\hat{\pi}^2_{t+\tau} + \mu\hat{N}^2_{t+\tau}\right)\right], \tag{3.19}$$

where $\beta_{CB} \in [0, 1)$ is a discount factor that may differ from that of the private sector; $\mu > 0$ is a parameter reflecting policy-makers' relative weight on employment stabilization; and the rational

[12] See Sections E and F in the Online Appendix for robustness checks exploring a range of potential alternative choices on agents' initial guesses.

[13] We note here that a key property of LKT is that each wage-setter's inflation expectation will generally fail to coincide with their expectation of the inflation expectations of all other wage-setters – i.e., $\mathbb{E}^{WS(s)}_{t-1}(\hat{\pi}_t) \neq \mathbb{E}^{WS(s)}_{t-1}[\overline{\mathbb{E}}^{WS}_{t-1}(\hat{\pi}_t)]$, where $\overline{\mathbb{E}}^{WS}_{t-1}(\hat{\pi}_t) \equiv \int_0^1 \mathbb{E}^{WS(\tilde{s})}_{t-1}(\hat{\pi}_t)d\tilde{s}$ denotes the average expectation in the economy. This discrepancy arises because each wage-setter believes that they are more sophisticated than all other wage-setters under LKT. A key advantage of LKT is that it offers a simple, tractable way to capture such tensions in agents' expectations relative to many alternative models of higher-order expectations.

[14] By "ad-hoc", we mean that the loss function being optimized has not been derived from the model's underlying welfare function. Our interest in a loss function that penalizes the squared inflation and employment gaps instead stems from the fact that such functions provide a natural (and commonly used) way to capture the fact that real-world central banks place weight on both price stability and stabilization of the real economy, whether the latter is formalized as part of an explicit dual mandate or simply treated as a secondary objective in the context of a flexible inflation targeting framework.

expectation $\mathbb{E}_{t-1}^{RE}(\cdot)$ reflects our assumption that policy-makers are not subject to the expectational frictions that may arise in the private sector. The optimization is subject to the Phillips curve on line 2.12 and feedback rule on line 2.14 and thus implicitly assumes that the monetary authority can credibly commit to carrying out the policy actions prescribed by ϕ_t even if it may not be optimal to do so once $\hat{\theta}_t$ has been revealed.[15]

This problem can be simplified by defining $x_t \equiv \hat{\pi}_t + \hat{N}_t$ as a measure of the total overheating that the economy experiences in a given period. In this case, the policy problem collapses to choosing ϕ_t to minimize

$$\mathbb{E}_{t-1}^{RE}\left[\sum_{\tau=0}^{\infty}\beta_{CB}^{\tau}(1+\mu\phi_{t+\tau}^2)\left[(1-\phi_{t+\tau})^{2k}x_{t+\tau-1}^2+\sigma_\theta^2\right]\right], \tag{3.20}$$

subject to the following law of motion for the overheating measure:

$$x_t = (1-\phi_t)[(1-\phi_t)^k x_{t-1} - \hat{\theta}_t]. \tag{3.21}$$

With this simplified policy problem in hand, we will now characterize policy-makers' optimal choice on ϕ_t under RE ($k \to \infty$), AE ($k = 0$), and LKT ($k \in \{1, 2, \ldots\}$).

3.2.1 Optimal policy under RE ($k \to \infty$)

For any sequence of positive policy stances, the loss function under RE reads as

$$\mathbb{E}_{t-1}^{RE}\left[\sum_{\tau=0}^{\infty}\beta_{CB}^{\tau}(1+\mu\phi_{t+\tau}^2)\sigma_\theta^2\right],$$

while the law of motion for the overheating measure reads as $x_t = -(1-\phi_t)\hat{\theta}_t$.

With losses thus independent of x_t and strictly increasing in ϕ_t^2, it is clearly optimal for the central bank to make the latter as small as possible. However, since uniqueness issues arise when $\phi_t = 0$ exactly,[16] we interpret the RE economy as one in which it is optimal for policy-makers to set ϕ_t at a positive but "vanishingly low" level, and stress that this leads to full anchoring of inflation

[15] Note that the realized values of inflation, employment, and the supply shock are all assumed to be revealed to policymakers and the private sector at the end of a given period, so policymakers' commitment to ϕ_t can immediately be verified ex-post.

[16] Another way to see the uniqueness issues that arise under RE when $\phi_t = 0$ is that in this case the nominal rate setting consistent with the central bank's preferred "zero gaps" equilibrium would be constant at $\hat{\imath}_t = 0$, but merely announcing this would lead to indeterminacy. When $\phi_t = 0$, these issues could be overcome by supplementing the policy feedback rule with additional information regarding the central bank's intended off-equilibrium behaviour.

expectations while keeping the variance of employment also vanishingly low.

In this sense, RE effectively allows the central bank to fully look through any deviations of inflation from target when setting monetary policy. Intuitively, this reflects the fact that rational agents fully understand that all deviations from the inflation target are transient for even vanishingly positive values of ϕ_t, so their inflation expectations are always anchored at target – even if the economy has recently experienced a long sequence of negative supply shocks. Put differently, even a very weak commitment to keeping inflation at target is enough to keep *expected* inflation at target when the private sector is fully rational.

3.2.2 Optimal policy under AE ($k = 0$)

Under AE, the policy problem can be represented as a Bellman equation of the form

$$V(x_{t-1}) = \min_{\phi_t} \left\{ (1 + \mu\phi_t^2) \left(x_{t-1}^2 + \sigma_\theta^2 \right) + \beta_{CB} E_{t-1}^{RE} \left[V(x_t) \right] \right\},$$

where the minimization is subject to the following law of motion for the overheating measure: $x_t = (1 - \phi_t)(x_{t-1} - \hat{\theta}_t)$.

We solve this problem using a guess-and-verify approach. Specifically, we conjecture a value function of the form $V(x) = a_1 x^2 + a_2 \sigma_\theta^2$. This leads to a solution of the following form:

$$\begin{aligned} \phi_t &= \frac{\beta_{CB} a_1}{\mu + \beta_{CB} a_1} \\ a_2 &= \frac{1 + \mu\phi_t^2 + \beta_{CB} a_1 (1 - \phi_t)^2}{1 - \beta_{CB}}, \end{aligned}$$

where a_1 corresponds to the unique positive solution of the following equation:

$$a_1 = 1 + \frac{\mu \beta_{CB} a_1}{\mu + \beta_{CB} a_1}.$$

Several features of this solution are noteworthy. First, as under RE, the optimal policy involves a constant stance and thus precludes any pivoting. Moreover, in the special case of a myopic central bank with $\beta_{CB} = 0$, it is optimal for policy-makers to fully look through supply shocks in the sense of setting $\phi_t = 0$, much as under the RE case. In contrast, a forward-looking central bank will opt for a more hawkish policy under which $\phi_t > 0$. This last result reflects the fact that policy-makers cannot influence the private sector's *contemporaneous* expectations under AE, but they can

influence the expectations that agents will carry into *future* periods, and a forward-looking central bank will find it optimal to exploit this channel to at least some extent.

3.2.3 Optimal policy under LKT ($k \in \{1, 2, \ldots\}$)

Under LKT, the policy problem does not admit a closed-form solution, nor easy qualitative characterizations. As a result, we will proceed in two steps. First, we will use the remainder of this section to examine the policy problem when the central bank is assumed to be myopic in the sense that $\beta_{CB} = 0$. Although this is a special case, it will allow us to derive several important features of policy-makers' optimal choice on ϕ_t when viewed as a function of key parameters and the state variable x_{t-1}. In the following section, we then turn our attention to the case of a forward-looking central bank with $\beta_{CB} > 0$. This involves solving the full dynamic version of the policy problem using numerical methods and will allow us to establish the robustness of the qualitative results derived under the myopic special case.

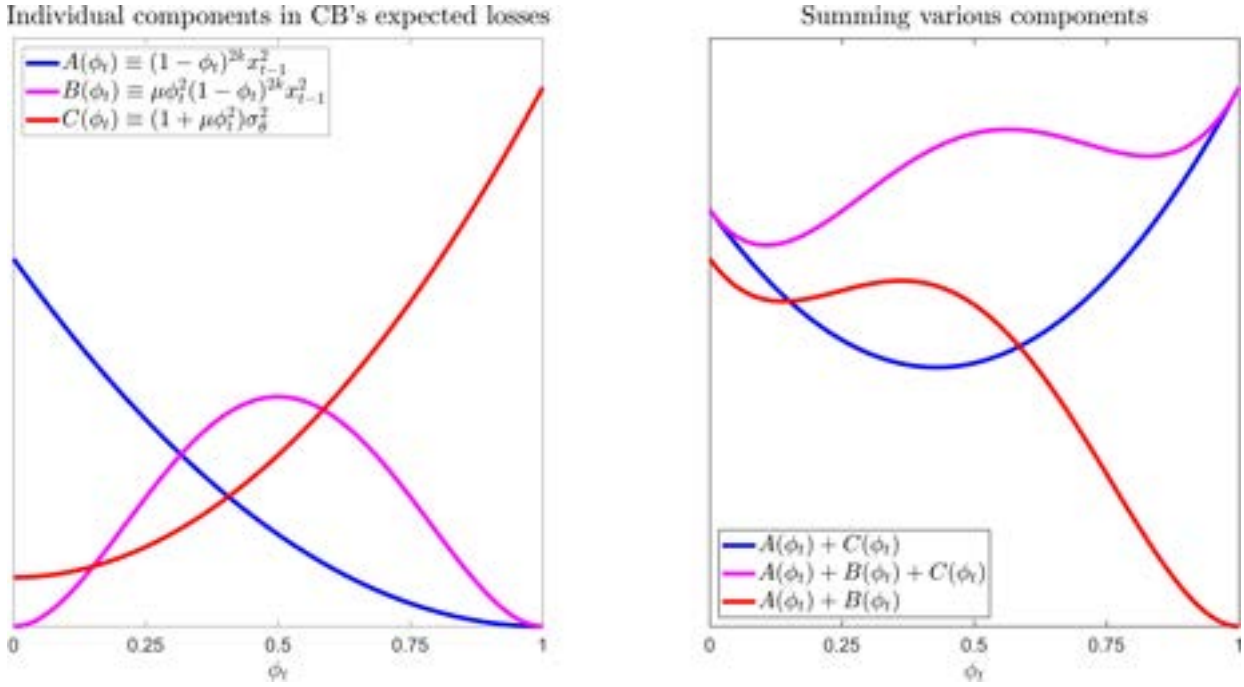
To understand why it is difficult to solve the policy problem even in the myopic case, it is useful to decompose the central bank's expected losses in a given period as follows:

$$(1+\mu\phi_t^2)\left[(1-\phi_t)^{2k}x_{t-1}^2+\sigma_\theta^2\right] = \overbrace{(1-\phi_t)^{2k}x_{t-1}^2+\sigma_\theta^2}^{=\mathbb{E}_{t-1}^{RE}(\hat{\pi}_t^2)} + \underbrace{\mu\phi_t^2(1-\phi_t)^{2k}x_{t-1}^2+\mu\phi_t^2\sigma_\theta^2}_{=\mu\mathbb{E}_{t-1}^{RE}(\hat{N}_t^2)}. \tag{3.22}$$

Since the solution for private-sector expectations on lines 3.17 and 3.18 implies that $\mathbb{E}_{t-1}^{WS}(x_t) = (1-\phi_t)^k x_{t-1}$, the term $(1-\phi_t)^{2k}x_{t-1}^2$ in the decomposition above can be interpreted as the cost of inflation gaps driven by poor anchoring of private-sector expectations, while the term σ_θ^2 can be interpreted as the cost of inflation gaps driven by supply shocks. By extension, the term $\mu\phi_t^2(1-\phi_t)^{2k}x_{t-1}^2$ can be interpreted as the employment cost of central bank tightening *in response to inflation gaps driven by poor anchoring*, while the final term $\mu\phi_t^2\sigma_\theta^2$ can be interpreted as the employment cost of tightening *in response to inflation gaps driven by supply shocks*.

As we illustrate in Figure 1, this decomposition is useful because it identifies the "cross-term" $\mu\phi_t^2(1-\phi_t)^{2k}x_{t-1}^2$ as the only obstacle to a simple solution, since suppressing this term would lead to a situation where total expected losses are a single-troughed function of the policy stance. In contrast, the cross-term itself is hump-shaped as a function of the policy stance, since the costs that it captures are best mitigated either by setting ϕ_t very low (and thus committing not to tighten much in the first place), or by setting ϕ_t very high (and thus ensuring that expectations are very

Figure 1: Decomposition of the central bank's expected losses



Notes. The left-hand panel breaks the decomposition on line 3.22 into three components, while the right-hand panel sums these components in various combinations. The specific values assumed in the panels are $(\mu, k, x_{t-1}^2/\sigma_\theta^2) = (10, 1, 7.5)$.

well-anchored). In cases where the economy has recently experienced significant overheating (i.e., $x_{t-1}^2 >> 0$) and the policy preference parameter μ is sufficiently large, this cross-term can get large enough to leave a significant signature on the shape of the central bank's overall losses, leading to a W-shaped profile along the lines illustrated by the magenta line in the right-hand panel of Figure 1. At the same time, the red line in the right-hand panel provides some insight into the role that LKT plays in shaping the central bank's losses, namely by collecting the terms associated with poor anchoring (both of which would fall to zero in the rational limit $k \to \infty$).

With the central bank's expected losses potentially W-shaped as a function of the policy stance, a key challenge when solving for optimal policy under LKT will be to account for the possibility of multiple local extrema. With this challenge in mind, we proceed as follows. Letting $\mathcal{L}(\phi_t, x_{t-1}^2)$ denote the expected losses on line 3.22, the first-order condition for a myopic central bank can be written as

$$0.5\mathcal{L}_1(\phi_t, x_{t-1}^2) = \mu\phi_t \left[(1-\phi_t)^{2k} x_{t-1}^2 + \sigma_\theta^2\right] - k(1-\phi_t)^{2k-1} x_{t-1}^2 (1+\mu\phi_t^2) = 0. \tag{3.23}$$

The richness of the policy problem is then immediately evident: even in the simplest case where $k = 1$, the first-order condition will be cubic in the policy stance and may thus admit multiple solutions. In other words, this equation will generally take the form of a correspondence, rather than a function, when viewed as a mapping from the (squared) overheating measure x_{t-1}^2 to the policy stance ϕ_t. For these reasons, it will be more convenient to focus on the inverse mapping that the first-order condition implies from ϕ_t back to x_{t-1}^2, since this mapping will take the form of a function:

$$\frac{\mu\phi_t}{(1-\phi_t)^{2k-1}\left[k(1+\mu\phi_t^2)-\mu\phi_t(1-\phi_t)\right]} = \frac{x_{t-1}^2}{\sigma_\theta^2}. \tag{3.24}$$

Letting $f(\phi_t)$ denote the left-hand side of this expression, most of our analysis will therefore focus on exploring the shape of the function $f(\cdot)$.

However, before doing so, it is useful to establish some additional restrictions that the optimal choice on ϕ_t should respect, apart from satisfying equation (3.24):

Lemma 3.1. *The optimal policy stance should satisfy both $\phi_t \in [0,1]$ and $f'(\phi_t) > 0$.*

Proof. To place the optimal stance in the unit interval, we note that all $\phi_t \in (1,\infty)$ satisfy $\mathcal{L}(\phi_t, x_{t-1}^2) > (1+\mu)\sigma_\theta^2 = \mathcal{L}(1, x_{t-1}^2)$, while all $\phi_t \in (-\infty, 0)$ satisfy $\mathcal{L}(\phi_t, x_{t-1}^2) > x_{t-1}^2 + \sigma_\theta^2 = \mathcal{L}(0, x_{t-1}^2)$.

To further see that the optimal stance should satisfy $f'(\phi_t) > 0$, note that equation (3.24) can only hold at points where the denominator on its left-hand side is positive, which is equivalent to $\mathcal{L}_{12}(\phi_t, x_{t-1}^2) < 0$. Totally differentiating equations (3.23) and (3.24) around any such point then yields $d\phi_t/dx_{t-1}^2 = 1/\left[\sigma_\theta^2 f'(\phi_t)\right] = (-1)\mathcal{L}_{12}(\phi_t, x_{t-1}^2)/\mathcal{L}_{11}(\phi_t, x_{t-1}^2)$, so $f'(\phi_t)$ and $\mathcal{L}_{11}(\phi_t, x_{t-1}^2)$ must share a common sign. As a result, the central bank's second-order condition holds only if $f'(\phi_t) > 0$. ∎

With these restrictions in mind, we will now use a pair of lemmata to build up a characterization of the function $f(\cdot)$, relegating all remaining proofs to the Online Appendix. As a first step, we define $d(\phi_t)$ as the term in square brackets in the denominator on the left-hand side of equation (3.24) – i.e., $f(\phi_t) = \mu\phi_t/[(1-\phi_t)^{2k-1}d(\phi_t)]$. $d(\cdot)$ is a quadratic function whose behaviour over the unit interval can be summarized as follows:

Lemma 3.2. *If policy-makers' relative weight on employment is sufficiently low in the sense that*

$\mu < 4k(1+k)$, *then* $d(\phi_t) > 0\ \forall \phi_t \in [0,1]$. *Otherwise, the function* $d(\cdot)$ *has two real roots,*

$$\phi_1^d = \frac{1}{2(1+k)} - \sqrt{\frac{1}{4(1+k)^2} - \frac{k}{\mu(1+k)}} > 0, \text{ and}$$
$$\phi_2^d = \frac{1}{2(1+k)} + \sqrt{\frac{1}{4(1+k)^2} - \frac{k}{\mu(1+k)}} \in [\phi_1^d, 1),$$

with $d(\phi_t) < 0$ *i.f.f.* $\phi_t \in (\phi_1^d, \phi_2^d)$.

Our next step is to turn our attention to the derivative

$$f'(\phi_t) = \frac{\mu k \left[2\mu(k+1)\phi_t^3 - 3\mu\phi_t^2 + 2(k-1)\phi_t + 1\right]}{(1-\phi_t)^{2k}\left[d(\phi_t)\right]^2}, \tag{3.25}$$

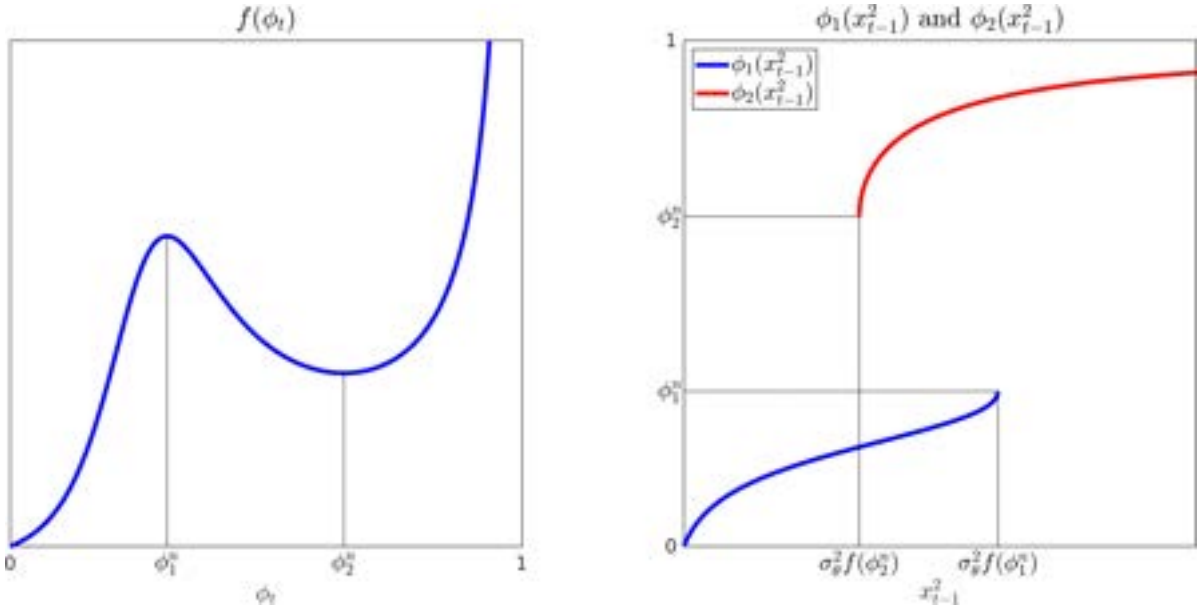
which shares its sign with the cubic function in its numerator, $n(\phi_t) \equiv 2\mu(k+1)\phi_t^3 - 3\mu\phi_t^2 + 2(k-1)\phi_t + 1$. The key features of this cubic can be summarized as follows:

Lemma 3.3. *There exists some* $\overline{\mu} \in (0, 4k(1+k))$ *with the following properties:*

- *If* $\mu < \overline{\mu}$, *then* $n(\phi_t) > 0\ \forall \phi_t \in [0,1]$;
- *If instead* $\mu \geq \overline{\mu}$, *then the function* $n(\cdot)$ *has two roots in the unit interval,* $\phi_1^n > 0$ *and* $\phi_2^n \in [\phi_1^n, 1)$. *These roots have the property that* $n(\phi_t) < 0\ \forall \phi_t \in (\phi_1^n, \phi_2^n)$, *with* $n(\phi_t) > 0$ $\forall \phi_t \in [0, \phi_1^n) \cup (\phi_2^n, 1]$, *and are distinct when* $\mu > \overline{\mu}$;
- *Finally, if* $\mu \geq 4k(1+k)$ *as well, then the roots described above can be ordered as follows relative to those described in Lemma 3.2:* $\phi_1^d \leq \phi_1^n \leq \phi_2^d < \phi_2^n$.

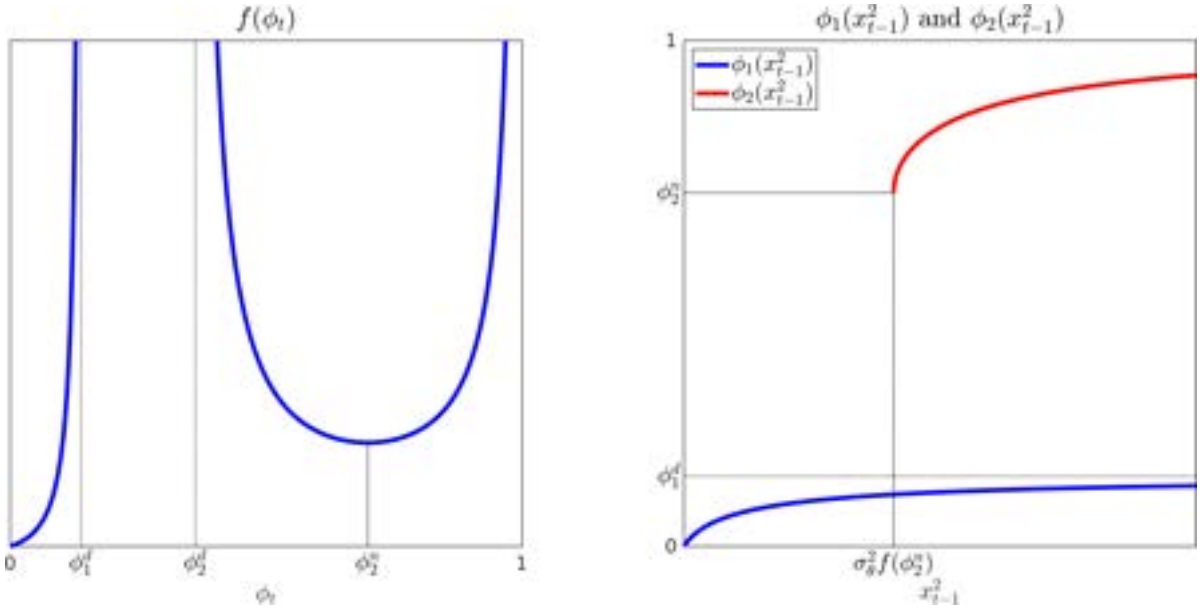
Based on these lemmata and the limiting values $f(0) = 0$ and $\lim_{\phi_t \nearrow 1}\{f(\phi_t)\} = \infty$, we can conclude that one of three scenarios must obtain with regard to the shape of the function $f(\cdot)$, whose domain we will now restrict to the unit interval. If $\mu \leq \overline{\mu}$, then $f(\cdot)$ is continuous and strictly increasing, and the uniquely optimal policy stance at any level of past overheating x_{t-1} is given by the inverse $f^{-1}(x_{t-1}^2/\sigma_\theta^2)$. If instead $\mu \in (\overline{\mu}, 4k(1+k))$, then $f(\cdot)$ is continuous but non-monotonic, and the optimal policy stance will lie on one of the two branches along which this function is strictly increasing: either a lower branch mapping from $[0, \phi_1^n]$ to $[0, f(\phi_1^n)]$, or an upper branch that maps from $[\phi_2^n, 1)$ to $[f(\phi_2^n), \infty)$ and thus involves strictly higher values for the policy stance. A similar situation obtains if $\mu \geq 4k(1+k)$, except in this case $f(\cdot)$ is discontinuous, and the lower branch now maps from $[0, \phi_1^d)$ to $[0, \infty)$. The latter two scenarios are illustrated in

Figure 2: Case $\mu \in (\overline{\mu}, 4k(1+k))$



Notes. The left-hand panel depicts the function on the left-hand side of equation (3.24). The right-hand panel depicts the functions $\phi_1(\cdot)$ and $\phi_2(\cdot)$ described in the main text, each of which corresponds to a local optimum. The parameter values assumed in the panels are $(\mu, k) = (6, 1)$.

Figure 3: Case $\mu \geq 4k(1+k)$



Notes. The left-hand panel depicts the function on the left-hand side of equation (3.24). The right-hand panel depicts the functions $\phi_1(\cdot)$ and $\phi_2(\cdot)$ described in the main text, each of which corresponds to a local optimum. The parameter values assumed in the panels are $(\mu, k) = (10, 1)$.

Figures 2 and 3, with $\phi_1(\cdot)$ denoting the implied one-to-one mapping from x_{t-1}^2 to ϕ_t along the lower branch, while $\phi_2(\cdot)$ denotes the analogous mapping along the upper branch.

When $\mu > \overline{\mu}$, this characterization opens up the key possibility that the globally optimal policy stance could respond to changes in x_{t-1}^2 by jumping discretely between the local minima represented by the two branches that we just described. In this respect, it is straightforward to show which branches will be selected when x_{t-1}^2 is either very small or very large. For example, policy-makers will strictly prefer the upper branch when x_{t-1}^2 is very large, since in this case the terms $(1-\phi_t)^{2k}x_{t-1}^2$ and $\mu\phi_t^2(1-\phi_t)^{2k}x_{t-1}^2$ will dominate all others in the loss function on line 3.22. If $\phi_*(x_{t-1}^2)$ denotes the globally optimal policy stance at an arbitrary value of x_{t-1}^2, we can thus conclude that $\phi_*(x_{t-1}^2) = \phi_2(x_{t-1}^2)$ when x_{t-1}^2 is sufficiently large. At the same time, we must have $\phi_*(x_{t-1}^2) = \phi_1(x_{t-1}^2)$ when $x_{t-1}^2 \in [0, \sigma_\theta^2 f(\phi_2^n))$, since the upper branch is inaccessible in this range.

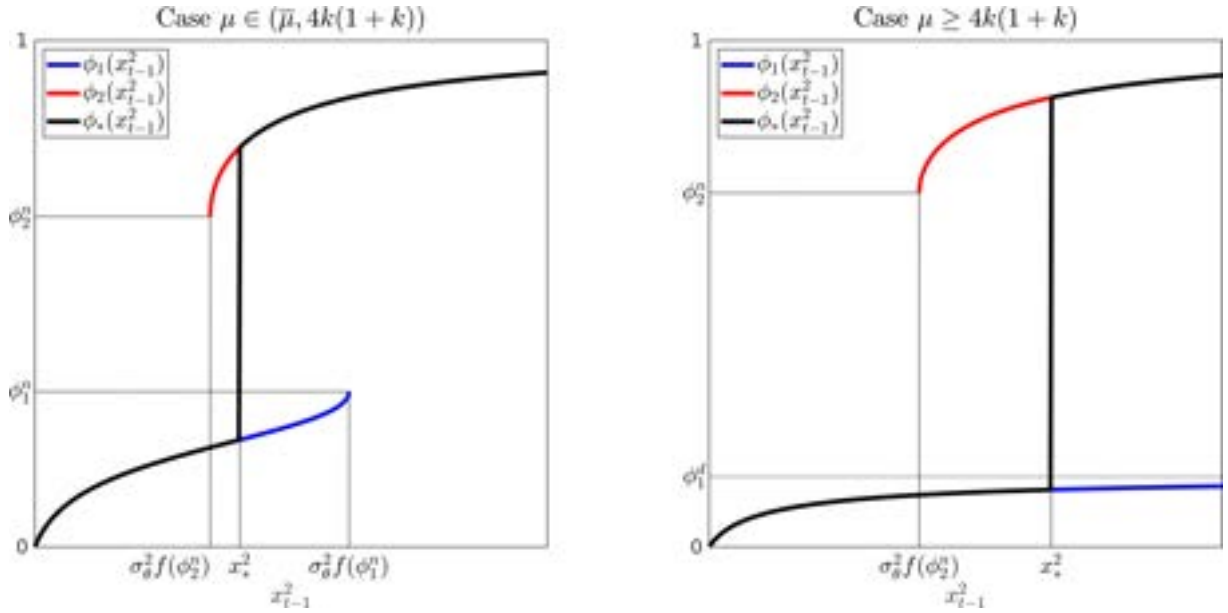
Based on these limiting outcomes, we can conclude that when $\mu > \overline{\mu}$ there must exist at least one value for x_{t-1}^2 around which it is optimal for policy-makers to pivot in the sense of jumping from the lower branch to the upper branch. In the Appendix, we further establish a single-crossing property that makes this pivot point unique, leading to our main result:

Proposition 3.1. *The globally optimal policy stance $\phi_*(x_{t-1}^2)$ is strictly increasing in x_{t-1}^2. If $\mu \leq \overline{\mu}$, then it is also continuous. Otherwise, there exists some cutoff $x_*^2 \in [\sigma_\theta^2 f(\phi_2^n), \infty)$ around which the optimal policy stance jumps discontinuously in the sense that $\phi_*(x_{t-1}^2) = \phi_1(x_{t-1}^2)$ $\forall x_{t-1}^2 \in [0, x_*^2)$, but $\phi_*(x_{t-1}^2) = \phi_2(x_{t-1}^2)$ $\forall x_{t-1}^2 \in (x_*^2, \infty)$.*

Figure 4 illustrates the implied policy profile when $\mu > \overline{\mu}$, with policy-makers initially opting for low values of ϕ_t at which they are mostly looking through supply shocks, then suddenly jumping to a much more aggressive stance if shocks cumulate up to a point that drives the overheating measure beyond some threshold level.

The intuition for Proposition 3.1 is that as x_{t-1}^2 rises and puts more pressure on private-sector expectations, it is optimal for policy-makers to compensate by increasing ϕ_t to at least some extent. When μ is relatively low, this can be done in a smooth way. However, when μ is high, a key concern for policy-makers is that smooth approaches will eventually run afoul of the "cross-term" that we emphasized in our discussion of Figure 1 – that is, the term $\mu\phi_t^2(1-\phi_t)^{2k}x_{t-1}^2$ capturing the employment costs of tightening in response to poorly anchored expectations. This is because smoothly ramping up the policy stance will eventually lead to situations where ϕ_t takes intermediate values that are still too low to effectively re-anchor private-sector expectations, but

Figure 4: Optimal pivots around $x_{t-1}^2 = x_*^2$



Notes. The panels depict the globally optimal policy stance, $\phi_*(x_{t-1}^2)$, for a range of values for the state variable x_{t-1}^2 and two cases involving different values for the policy preference parameter μ. The locally optimal stances $\phi_1(x_{t-1}^2)$ and $\phi_2(x_{t-1}^2)$ are also depicted. The parameter values assumed in the left-hand panel are the same as in Figure 2, while those assumed in the right-hand panel are the same as in Figure 3.

high enough that policy-makers have to engineer significant amounts of slack to offset the impact of those expectations on realized inflation outcomes. Pivots give central banks with high μ values a way to avoid such costly scenarios, and one of the key benefits of the more decisive re-anchoring with which they are associated is precisely that this allows policy-makers to re-stabilize realized inflation without necessitating large employment gaps.

With these points in mind, we use the remainder of this section to elaborate on some additional key features of our solution. In particular, we start with our attention on the sensitivity of private-sector expectations to the policy stance, as measured by the (unsigned) elasticity

$$\left| \frac{d\mathbb{E}_{t-1}^{WS}(x_t)}{d\phi_t} \frac{\phi_t}{\mathbb{E}_{t-1}^{WS}(x_t)} \right| = \frac{k\phi_t}{1-\phi_t} \equiv \zeta(\phi_t), \tag{3.26}$$

and elaborate on the important role that this object plays in driving our results, along with the related role being played by the parameter k controlling the sophistication of private-sector expectations. We then turn our attention to the way that the employment gap behaves around the pivot point and whether pivots can be compatible with soft landings.

Role of the elasticity $\zeta(\phi_t)$ and parameter k. To better understand the role being played by the elasticity $\zeta(\phi_t)$, it is useful to re-express the central bank's first-order condition in a way that allows us to distinguish between the "direct" channels via which the policy stance impacts the economy at a given level of private-sector expectations, versus the "indirect" channels that are intermediated by changes in these expectations. Dividing line 3.23 by $\phi_t^{-1}\mathcal{L}(\phi_t, x_{t-1}^2)$ and re-arranging terms, we get

$$\frac{\mu\phi_t^2}{1+\mu\phi_t^2} = \zeta(\phi_t)\frac{(1-\phi_t)^{2k}x_{t-1}^2}{(1-\phi_t)^{2k}x_{t-1}^2+\sigma_\theta^2}, \tag{3.27}$$

where the left-hand ratio $\mu\phi_t^2/(1+\mu\phi_t^2)$ can be interpreted as the elasticity of the central bank's expected losses with respect to the policy stance, holding private-sector expectations constant, while the right-hand ratio $(1-\phi_t)^{2k}x_{t-1}^2/[(1-\phi_t)^{2k}x_{t-1}^2+\sigma_\theta^2] \equiv r(\phi_t, x_{t-1}^2)$ can be interpreted as the elasticity of expected losses with respect to expectations, now holding direct effects constant. The right-hand side of line 3.27 thus captures the marginal benefits of tightening in terms of better anchoring expectations, while the left-hand side captures the marginal costs in terms of making employment more volatile at a given level of expectations.

What is key here is that while marginal costs and the ratio $r(\phi_t, x_{t-1}^2)$ are respectively increasing and decreasing in ϕ_t, marginal benefits are generally non-monotonic due to the fact the elasticity $\zeta(\phi_t)$ is increasing in ϕ_t – i.e., expectations tend to be more responsive to changes in the policy stance when ϕ_t is high. This critical property of private-sector expectations under LKT is the underlying source of non-convexity in the economy and leads to a situation where there are benefits to tightening by a sufficiently wide margin. Moreover, since the ratio $r(\phi_t, x_{t-1}^2)$ is increasing in x_{t-1}^2, the marginal benefits of tightening are amplified when the economy has recently experienced significant overheating, and this gives policy-makers an incentive to delay large increases in ϕ_t until the overheating measure has reached sufficiently high levels.

At the same time, the fact that the elasticity on line 3.26 is increasing in the "levels of thinking" parameter k helps to explain why monetary policy is so powerful in the model's RE limit. It also suggests that policy-makers facing very large but finite values for k may be able to make do without pivots, while others facing lower values for k may find pivots highly valuable as a way of engineering large changes in private-sector expectations. In the Appendix, we are able to verify and formalize this intuition as follows:

Proposition 3.2. *The threshold $\overline{\mu}$ described in Lemma 3.3 is strictly increasing in the parameter*

k and satisfies $\lim_{k\to\infty}\{\overline{\mu}\} = \infty$ *in the RE limit. As a result, increasing k narrows the range of values for* μ *within which pivots emerge as a feature of optimal policy.*

Pivots and soft landings. We finally turn to the way that employment behaves around the pivot point. While pivots are associated with large increases in the policy stance, their actual impact on employment is unclear and depends on the balance of the direct and indirect channels that we emphasized when discussing equation (3.27): if the latter channel were very strong, then the adjustment in inflation expectations that pivoting induces could reduce inflation pressures enough that the amount of *actual* slack that policy-makers have to engineer would be relatively low. This tension is reflected in the fact that the (rationally) expected employment gap is given by

$$\mathbb{E}_{t-1}^{RE}(\hat{N}_t) = -\phi_t(1-\phi_t)^k x_{t-1}, \tag{3.28}$$

so its behaviour around the pivot point is driven by how the jump to a higher policy stance impacts the product $\phi_t(1-\phi_t)^k$.

While this impact is difficult to characterize as a general matter, we show in the Appendix that the special case with $k = 1$ level of thinking has the property that the direct and indirect channels offset one another exactly in expectation, leading to no change in $\phi_t(1-\phi_t)^k$ or the expected level of employment around the pivot point. In other words, pivots are compatible with soft landings when $k = 1$, *at least in expectation*, and essentially amount to a "tough talk" strategy that re-anchors expectations in a way that eliminates the need for actual slack.

However, if pivots are expected to lead to soft landings, then why wait to implement them? The answer has to do with the impact of pivots on the variance of the employment gap,

$$\text{var}_{t-1}^{RE}(\hat{N}_t) = \phi_t^2\sigma_\theta^2, \tag{3.29}$$

which increases discontinuously around the pivot point. This reflects the fact that an aggressive policy stance makes employment more sensitive to supply shocks by committing the central bank to respond strongly to off-target inflation, regardless of whether that inflation is driven by expectations or shocks. Optimally-timed pivots under $k = 1$ are thus associated with "risky" or "narrow" soft landings in the sense that they re-stabilize inflation without triggering changes in the *expected* level of employment, but do so at the cost of increasing the uncertainty around employment outcomes.

To summarize:

Proposition 3.3. *If $\mu > \overline{\mu}$, then the pivots described in Proposition 3.1 are associated with discontinuous increases in the variance of the employment gap, as given by equation* (3.29) *above. However, if $k = 1$ as well, then they have no effect on the expected level of employment, as given by equation* (3.28) *above.*

4 Dynamic policy problem: numerical solution

Having shown in our previous section how pivots can emerge as a feature of optimal policy when policy-makers are myopic, we now ask if pivots are robust to extensions in which policy-makers are assumed to be forward-looking. This will involve using numerical methods to solve the full dynamic problem facing a central bank with a non-zero discount factor $\beta_{CB} > 0$.

However, before doing this, it will be useful to generalize the Phillips curve derived in our previous section to allow for an arbitrary slope on the expected employment gap, rather than the slope of one implicit in equation (2.12). This will be helpful in calibrating the numerical model and interpreting its results, in addition to allowing for greater generality.

4.1 Generalizing the Phillips curve

To allow for a non-unit slope in the Phillips curve, we replace the household preferences in the previous section with ones in line with Greenwood, Hercowitz, and Huffman (1988). Specifically, we now assume that household h has a utility function of the form

$$\mathbb{E}_t^{HH(h)}\left[\sum_{\tau=0}^{\infty}\beta^{\tau}\ln\left(C_{h,t+\tau}-\eta\theta_{t+\tau}\int_0^1\tilde{N}_{h,s,t+\tau}^{1+\chi}ds\right)\right],$$

where $\chi > 0$ is the inverse Frisch elasticity of labour supply. In the Online Appendix, we show that this leads to a modified Phillips curve of the following form, where $\lambda \equiv \chi/(1+\rho\chi)$:

$$\hat{\pi}_t = \lambda\mathbb{E}_{t-1}^{WS}(\hat{N}_t) + \mathbb{E}_{t-1}^{WS}(\hat{\pi}_t) - \hat{\theta}_t,$$

In the same appendix, we further show how non-unit values for the slope parameter λ impact the problem facing policy-makers. This involves re-defining our overheating measure as $x_t \equiv \hat{\pi}_t + \lambda\hat{N}_t$, in addition to defining $\tilde{\mu} = \mu/\lambda^2$ as policy-makers' "effective weight" on stabilizing employment, and $\tilde{\phi}_t \equiv \lambda\phi_t$ as their "effective policy stance" in a given period. Under these definitions, the

central bank's problem amounts to choosing $\tilde{\phi}_t$ to minimize

$$\mathbb{E}_{t-1}^{RE}\left[\sum_{\tau=0}^{\infty}\beta_{CB}^{\tau}(1+\tilde{\mu}\tilde{\phi}_{t+\tau}^{2})\left[(1-\tilde{\phi}_{t+\tau})^{2k}x_{t+\tau-1}^{2}+\sigma_{\theta}^{2}\right]\right], \tag{4.30}$$

subject to the following law of motion:

$$x_t = (1-\tilde{\phi}_t)[(1-\tilde{\phi}_t)^k x_{t-1} - \hat{\theta}_t]. \tag{4.31}$$

Comparing this program with the original program on lines 3.20 and 3.21 of the previous section, we see that the two coincide when we replace (μ, ϕ_t) in the latter program with $(\tilde{\mu}, \tilde{\phi}_t)$. In this sense, solving the policy problem in the general case of a non-unit slope is equivalent to solving our original program with a modified weight on the central bank's employment objective, though one must be mindful to infer the true policy stance via $\tilde{\phi}_t = \lambda\phi_t$.

This equivalence has several key implications. Most importantly, it implies that all of the qualitative results established in the previous section readily carry over to the case of an arbitrarily sloped Phillips curve. Moreover, since the literature favours values for λ much less than one, it implies a significant widening in the range of values for the "true weight" μ within which the pivoting behaviour described in Proposition 3.1 emerges. To the extent that λ takes values less than one, pivots should also be associated with wider swings in the true policy stance $\phi_t = \tilde{\phi}_t/\lambda$.

4.2 Baseline parameterization and solution method

We are now ready to solve the generalized policy problem on lines F.63 and B.45 using numerical methods. Given the model's emphasis on wage-setting, coupled with its property that wages are sticky within periods but fully flexible across periods, it is natural to focus on parameterizations that identify model periods with years. We therefore set policy-makers' discount factor to $\beta_{CB} = 0.995$, aligning with the midpoint of Bank of Canada staff's current assessed range of 0 to 1% for the real neutral rate in Canada (Champagne et al., 2023).

As for the model's other parameters, our baseline parameterization will set the relative weight on employment in the central bank's loss function to $\mu = 1$, a natural benchmark consistent with policy-makers placing equal weight on deviations of inflation from target and employment from its natural level. At the same time, we set the Phillips curve slope to $\lambda = 0.092$, a value broadly in line with Canadian estimates in Djeutem et al. (2022) and Wagner et al. (2022). We also assume $k = 2$

levels of thinking in the private sector, placing us in the middle of the range of $k = 1$ to 4 considered in Farhi and Werning (2019) and in line with the experimental literature (which generally estimates k values in the low single digits). The only remaining parameter is then the supply-shock variance, σ_θ^2, which we calibrate to match the variance of CPI inflation in Canada, computed using a data sample that begins when the country's current 2% inflation target first formally started applying in the mid-1990s.

For a given set of parameters, we solve the model by value function iteration. Details on our solution algorithm are given in the Online Appendix.

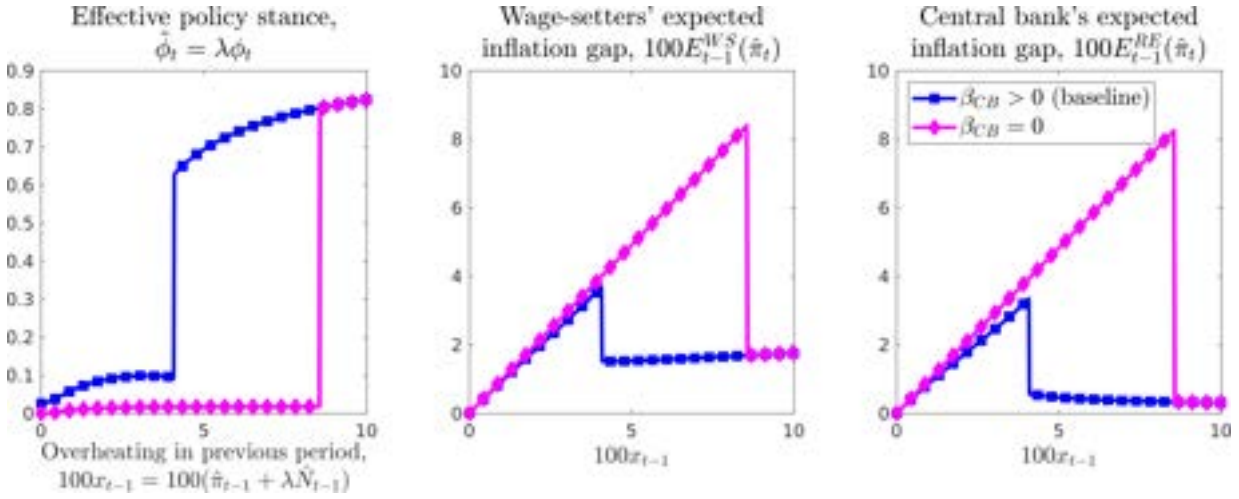
4.3 Results

Results for the baseline parameterization are given by the blue lines in Figure 5, with the left-hand panel depicting the optimal policy stance as a function of the past overheating measure x_{t-1}. Pivots clearly remain a key feature of optimal policy, with a threshold value for x_{t-1} around 4% triggering a sudden increase in the policy stance. As shown in the middle and right-hand panels, which respectively plot the inflation expectations of wage-setters and policy-makers as a function of past overheating, pivots are also associated with sharp re-normalizations in both rational and LKT-based inflation expectations.

Figure 5 also reports results for a myopic parameterization under which we set $\beta_{CB} = 0$ while keeping all other parameters at their baseline values. Those results are given by the magenta lines in the figure. Comparing those results to the non-myopic base case in blue, we see that allowing for a realistic degree of forward-looking behaviour on the part of the central bank leads to a significant "pull forward" of the threshold around which policy-makers are prepared to pivot (left-hand panel), setting the scene for earlier (and thus smaller) re-normalizations in both rational and LKT-based inflation expectations (right-hand and middle panels, respectively). This willingness to pivot earlier is a natural consequence of the fact that a forward-looking central bank internalizes the benefits that stabilizing inflation in the current period will generate in future periods through better-anchored expectations. In other words, our results suggest that pivoting behaviour is not merely *robust* to allowing for a forward-looking central bank: there is a precise sense in which this behaviour is *actually strengthened* by non-zero values for the central bank's discount factor.

With these points in mind, we use the remainder of this section to explore the roles that two additional parameters play in driving our numerical results, namely: (i) the "levels of thinking" parameter k; and (ii) the ratio $\tilde{\mu} = \mu/\lambda^2$ representing the effective weight that policy-makers place

Figure 5: Optimal pivots under the baseline parameterization and $\beta_{CB} = 0$



on their employment objective.

Role of the parameter k. Figure 6 reports results from varying k between one and five while keeping all other parameters at their baseline values. The dashed red lines in the figure also report illustrative results for an extreme parameterization under which we set $k = 40$, well outside the range supported by experimental studies – and high enough to deliver results relatively close to those associated with the rational expectations limit $k \to \infty$ (for which results are reported in dashed blue). A key pattern to note is that while the cases involving $k \in \{1, 2, 3\}$ all lead to pivots around thresholds in the 3-5% range, higher-k cases have no pivots occurring even when x_{t-1} is as high as 10%. This is consistent with the message from Proposition 3.2 that policy-makers can make do without pivots when facing a sufficiently sophisticated private sector.

We also note that within the low single-digit range of k values typically supported by the experimental literature, Figure 6 tends to associate higher values for k with smaller jumps in the policy stance around the pivot point (left-hand panel), along with less sharp adjustments in the expectations of wage-setters (middle panel) and policy-makers (right-hand panel). These patterns reflect the fact that policy-makers facing higher values for k have greater leverage over private-sector expectations and thus do not need to tighten as aggressively to re-stabilize inflation.

These points are important partly because a smaller jump in the policy stance around a given pivot point should lead to a situation where employment outcomes are less sensitive to supply shocks following pivots, all else being equal. This suggests that the "risky soft landing" problem that we discussed earlier should be less of an issue when k is high. To confirm this intuition, we use

Figure 6: Impact of changing the levels of private-sector thinking, k

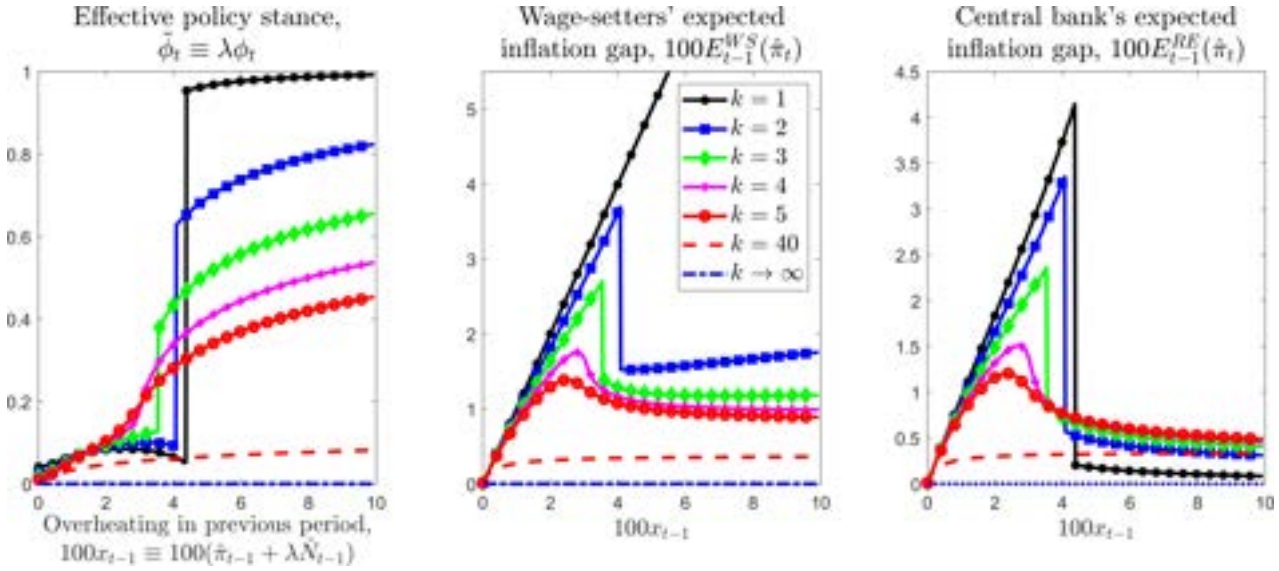


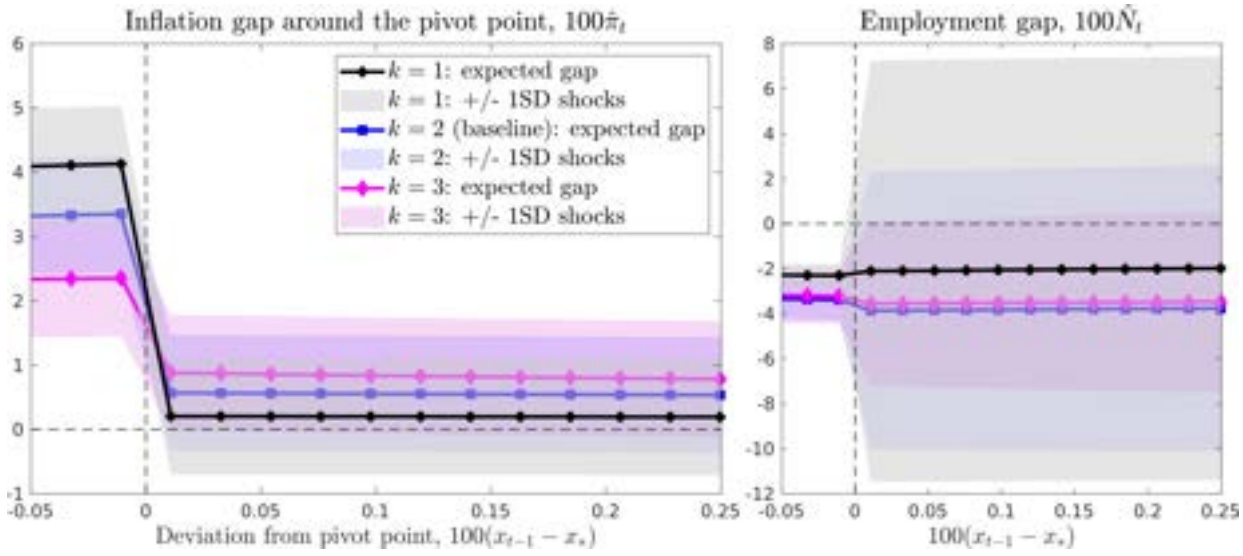
Figure 7 to zoom in on the model's behaviour around the pivot points associated with the $k = 1$ through 3 parameterizations described earlier. In this figure, solid lines depict (rationally) expected inflation and employment outcomes before observing the supply shock $\hat{\theta}_t$, while the surrounding bands depict realized outcomes for a range of potential shock values. Since the horizontal axes in both panels have been normalized so that pivots occur at a value of zero for all parameterizations, this makes it easy to compare parameterizations in terms of how the first and second moments of the employment gap respond to pivots.

In particular, we note that pivots have negligible impacts on expected employment for all parameterizations considered, and that the main employment impact of pivots is instead to increase the uncertainty surrounding employment outcomes – especially when k is small. This indicates that risky soft landings remain a key feature of the model economy, and that higher values for k open up a better policy trade-off between stabilizing expectations and mitigating the variance of employment and associated risk of hard landings.

Role of the parameter $\tilde{\mu}$. We now turn our attention to the effective weight that policy-makers place on minimizing deviations of employment from its natural level, as measured by the ratio $\tilde{\mu} = \mu/\lambda^2$. In Figure 8, we report results for a range of values for this ratio, holding all other parameters at their baseline values.

Consistent with our analytical results, the optimal policy stance is discontinuous when $\tilde{\mu}$ is large

Figure 7: Risky soft landing under different levels of private-sector thinking (k)



but takes on a smooth profile for small values of $\tilde{\mu}$, suggesting that $\tilde{\mu}$ must exceed some threshold in order for pivoting to occur. Moreover, conditional on being above this threshold, higher values for $\tilde{\mu}$ are associated with higher pivot points, reflecting a greater willingness to look through periods of high inflation when policy-makers place greater weight on their employment objectives.

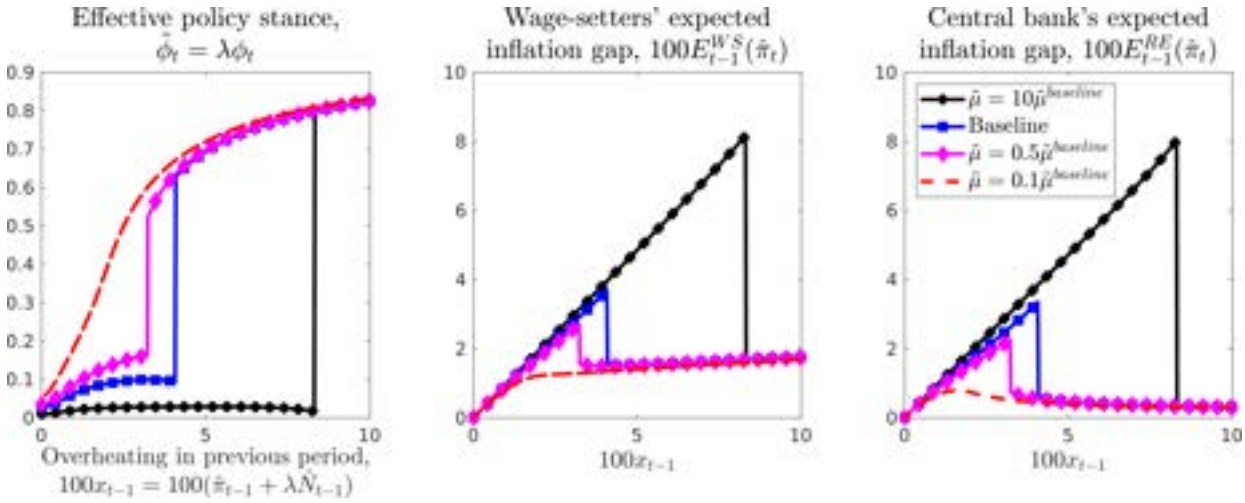
While policy-makers thus tend to delay pivots when $\tilde{\mu}$ is high, it is important to note from Figure 8 that in this case pivots will be associated with larger swings in the policy stance should x_{t-1} eventually reach levels high enough to trigger them. In this sense, pivots "come later", but they are also more pronounced.[17]

5 Costs of deviating from the optimal policy

While the analysis in the last two sections identified conditions under which discontinuous pivots represent an optimal response to sufficient levels of overheating, such pivots are clearly difficult to time as a practical matter, and concerns that an abrupt tightening could trigger economic disruptions often lead policy-makers to take more gradual approaches. While gradual approaches are generally sub-optimal when viewed through the lens of the model, we use this section to get

[17] In this respect, we note that recent pivots in advanced economies occurred at a time when (i) some central banks had recently announced framework changes placing greater emphasis on employment outcomes; and (ii) a range of studies suggested that Phillips curves in advanced economies had flattened significantly. To the extent that these factors respectively contributed to a higher value for μ and lower value for λ, they would have increased the ratio $\tilde{\mu} = \mu/\lambda^2$, setting the scene for a later but more pronounced pivot when viewed through the lens of this framework.

Figure 8: Impact of changing policy-makers' effective weight on employment, $\tilde{\mu} = \mu/\lambda^2$



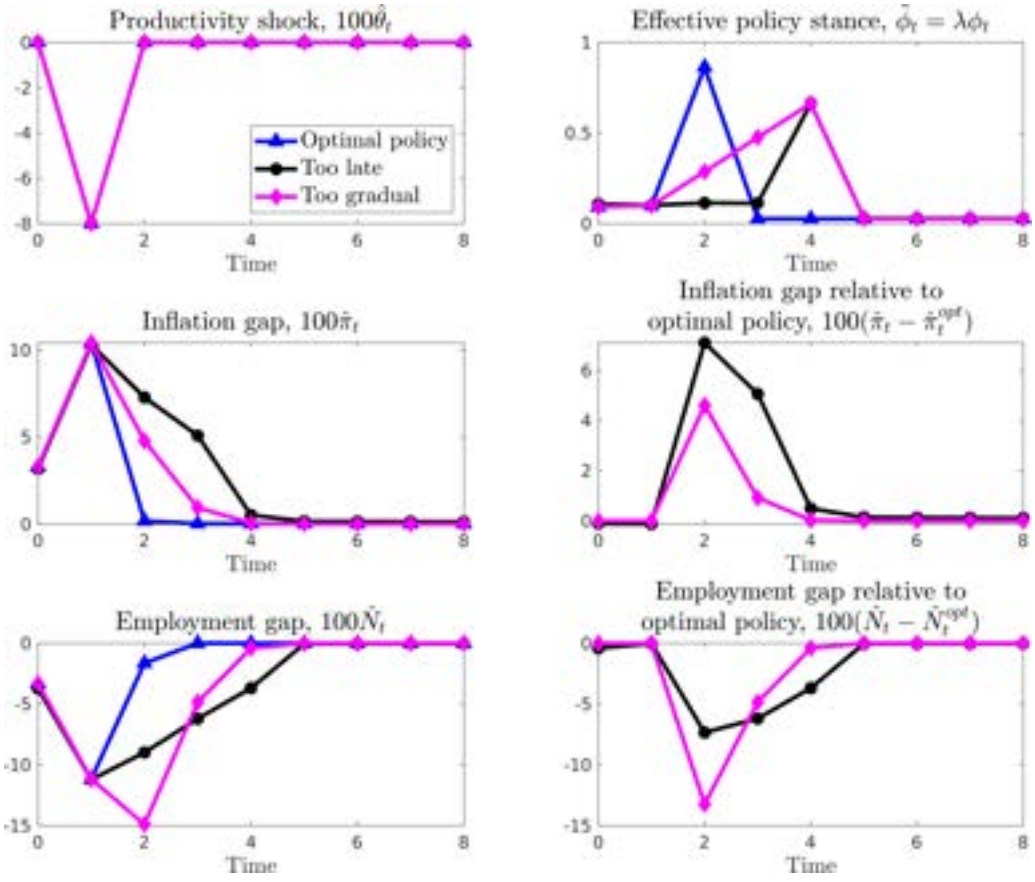
more specific about their costs, as well of those of poorly timed pivots.

Of course, there are many ways in which the monetary authority could deviate from the optimal policy in our model. We will focus on two possible deviations in simulations where we initialize the baseline parameterization of the model at a modest level of recent overheating and then assume a supply shock large enough to trigger pivoting under the optimal policy. Figure 9 depicts the implied time paths for key variables under the optimal policy, as well as the two deviant policies.

The first of these deviant policies aims to capture a scenario where the policy pivot is mis-timed and occurs too late. To construct this "too late" scenario, we make use of the fact that the central bank's loss function generally admits multiple, locally optimal values for the policy stance, with pivots involving jumps from the lowest of these optima to a higher one. In particular, we assume that policy-makers (i) initially stay on their lowest local optimum after the supply shocks arrives, even though it is no longer globally optimal to do so; and (ii) only recognize and correct this mistake after three periods. As shown in the middle and bottom panels of Figure 9, this leads to significantly worse inflation and employment outcomes relative to the globally optimal policy.

The second deviation that we consider aims to capture a scenario where policy-makers no longer pivot and instead tighten gradually in response to the supply shock. For this scenario, we assume that the central bank smoothly increases its stance in the three periods following the supply shock, eventually reaching the same peak value as in the "too late" scenario; once they reach this peak, policy-makers are assumed to revert back to the globally optimal policy. As shown in Figure 9, a key feature of this "too gradual" scenario is that it leads to worse employment outcomes relative

Figure 9: Simulating an optimal pivot and two possible deviations



to *both* the optimal policy and the "too late" scenario. This illustrates the risk of hard landings associated with intermediate values for the policy stance and more generally highlights the rich, non-monotonic mapping that the model admits between the policy stance and aggregate outcomes.

6 What about demand shocks?

So far, we have ignored demand shocks, despite the key role that they played in the recent inflation surge and other real-world inflation episodes. In this section, we briefly discuss the consequences of introducing demand shocks and how our results can be interpreted in this broader context. We will focus on the standard case where demand shocks arise from shocks to the household discount factor, leading the IS curve on line 2.13 to include a stochastic intercept d_t. Forward-solving then yields

$$\hat{N}_t = -\mathbb{E}_t^{RE}\left[\sum_{\tau=0}^{\infty}(\hat{\imath}_{t+\tau} - \hat{\pi}_{t+\tau+1})\right] + \mathbb{E}_t^{RE}\left(\sum_{\tau=0}^{\infty} d_{t+\tau}\right), \tag{6.32}$$

so the appropriate measure of the stimulus being provided by policymakers is now the cumulative sum of expected future differences between the real rate and a time-varying natural rate linked to the demand shock. In special cases where the central bank is able to ensure an average difference of zero, this would imply full stabilization of the employment gap.

With these points in mind, we can still interpret the feedback rule on line 2.14 as the product of a situation where policy-makers set the level of stimulus as a function of the prevailing rate of inflation, though the appropriate measure of stimulus should now be based on equation (6.32):

$$\phi_t \hat{\pi}_t = \mathbb{E}_t^{RE}\left[\sum_{\tau=0}^{\infty}\left(\hat{\imath}_{t+\tau} - \hat{\pi}_{t+\tau+1}\right)\right] + \mathbb{E}_t^{RE}\left(\sum_{\tau=0}^{\infty} d_{t+\tau}\right).$$

Under this interpretation, the core policy problem laid out in Section 3 is unaltered relative to its form in an economy without demand shocks, and the only impact of demand shocks is thus to shift the nominal rate settings that the solution to this problem would imply.

Put differently, when the economy is subject to both demand and supply shocks, the policy-setting process can be viewed as a two-stage process where policy-makers first gauge the rate path needed to fully insulate employment from demand shocks, and then judge if adjustments are needed to address supply shocks or poor anchoring of expectations. All of our results thus carry over to this context but should be interpreted as informing the second of these two stages.

Of course, this assumes that policy-makers are able to distinguish between demand and supply shocks, and this is a non-trivial assumption. If, for example, a central bank mistakenly interpreted some demand shocks as supply shocks, then this could lead to an overly stimulative stance to the extent that the two-stage process above will generally involve fully offsetting demand shocks but looking through supply shocks to at least some extent. Similar issues would arise if policy-makers simply mis-judged the strength of demand. "Policy errors" of this sort could be incorporated into our analysis but would involve a signal-extraction problem that we leave for future work.

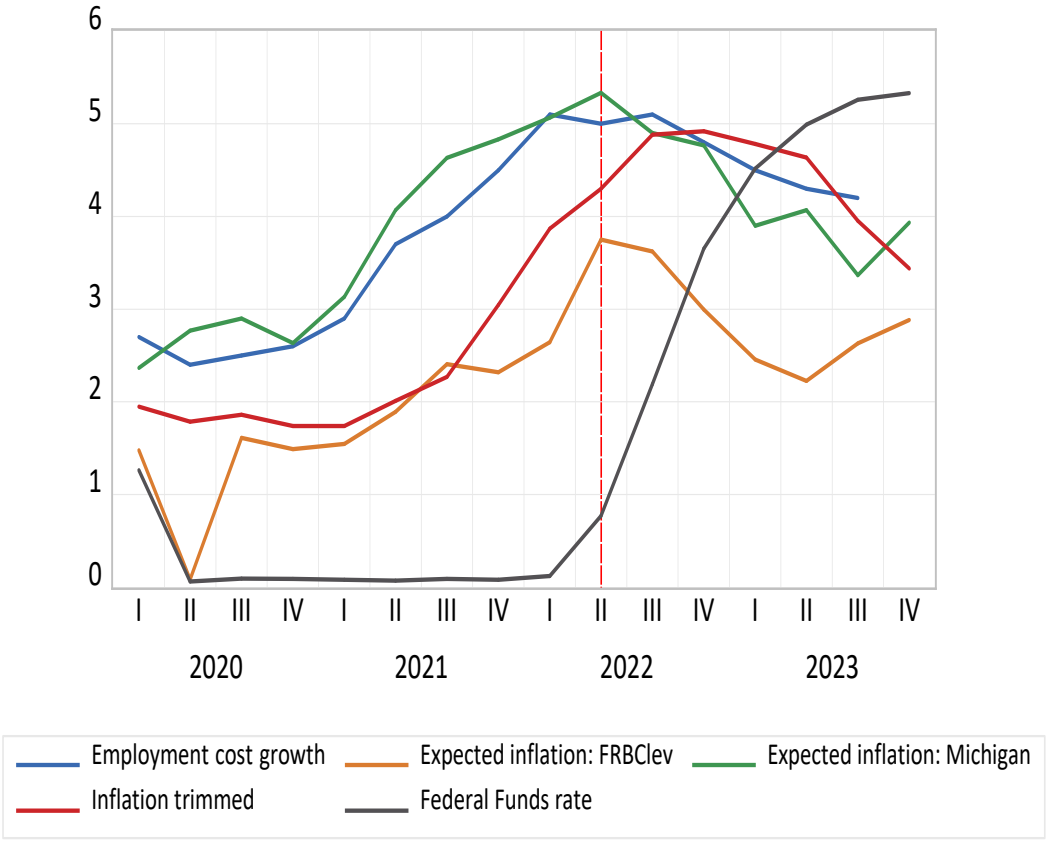
7 Model mechanism and the recent inflation surge

Our model has the implication that it is optimal for policymakers to initially look through supply shocks, but they should be ready to pivot sharply if shocks cumulate beyond some threshold level. While this pattern broadly squares with many countries' experiences during the global inflation surge that began in 2021, we use this section to more fully explore the extent to which the data

support the mechanisms that the model emphasizes.[18]

We start by noting that the model's primary mechanism works through expected inflation. With an unchanging policy stance, inflation expectations should co-move with a lagged overheating measure that puts significant weight on past inflation outcomes. On the other hand, pivoting from a loose policy stance to a tight one should quickly lead to a moderation in expected inflation, even if supply shocks help to keep contemporaneous realized outturns relatively high. The model also predicts that wage growth should drift up and down with expected inflation.

Figure 10: Evolution of key series around the recent spike in US inflation



Notes. The vertical dashed red line marks 2022Q2, the first full quarter of the Federal Reserve's tightening cycle.

Are such relationships evident in the data? To help answer this question, Figure 10 shows the

[18]From the outset, we stress that our empirical tests neither assume nor aim to confirm that central banks' real-world pivots were optimally timed or sized. Put differently, the focus in this section is only on the positive consequences of policy pivots (rather than normative consequences).

movements over 2020Q1-2023Q4 in five key US series: two measures of one year-ahead inflation expectations (one based on the University of Michigan Surveys of Consumers and another produced by the Federal Reserve Bank of Cleveland); wage growth, proxied by changes in an employment cost index; a trimmed measure of inflation; and the federal funds rate. The figure shows that both measures of expected inflation, as well as wage growth, kept trending up until 2022Q2, which marked the first full quarter of an aggressive tightening cycle initiated by the Federal Reserve. Following this critical shift away from the "low for long" policies previously in place, expected inflation and wage growth both immediately turned around, though inflation continued trending up for a couple more quarters. These patterns are all consistent with the model's main mechanism.

More formally, the model's key implied relationships between inflation, wage growth, expected inflation, and other variables can be summarized by the following three equations, where $\hat{W}_t \equiv \ln W_t - \ln W_{t-1} - \pi^*$ can be interpreted as a wage growth gap:

$$\hat{\pi}_t = \hat{W}_t - \hat{\theta}_t \tag{7.33}$$

$$\hat{W}_t = (1 - \lambda\phi_t)\, \mathbb{E}_{t-1}^{WS}(\hat{\pi}_t) \tag{7.34}$$

$$\mathbb{E}_{t-1}^{WS}(\hat{\pi}_t) = (1 - \lambda\phi_t)^{k-1} (\hat{\pi}_{t-1} + \lambda\hat{N}_{t-1}) \tag{7.35}$$

If we then let $\bar{\phi}$ denote a long-run average of the policy stance, equation (7.35) can usefully be rewritten as

$$\mathbb{E}_{t-1}^{WS}(\hat{\pi}_t) = (1 - \lambda\bar{\phi})^{k-1}(\hat{\pi}_{t-1} + \lambda\hat{N}_{t-1}) + u_t, \tag{7.36}$$

with a residual $u_t \equiv [(1-\lambda\phi_t)^{k-1} - (1-\lambda\bar{\phi})^{k-1}](\hat{\pi}_{t-1} + \lambda\hat{N}_{t-1})$. During periods when the economy is overheating but policymakers are still maintaining a "look-through" stance, we should then expect to have $\hat{\pi}_{t-1} + \lambda\hat{N}_{t-1} > 0$ with $\phi_t < \bar{\phi}$, leading to a situation where the residual in equation (7.36) is positive and likely rising. On the other hand, when policymakers eventually pivot to a tighter stance, we should expect the same residual to revert back toward zero and even turn negative.

To test these predictions, we take equation (7.36) to the data by fitting expected inflation on lagged measures of realized inflation and labour market tightness. Our baseline specification uses the above-noted Michigan series for inflation expectations, along with a trimmed measure of realized inflation. Following Benigno and Eggertsson (2023) and others, we also use the vacancy-to-unemployment ratio (V/U) as a proxy for labour market tightness.

Residuals from several different versions of this estimation are reported in Figure 11 and give us

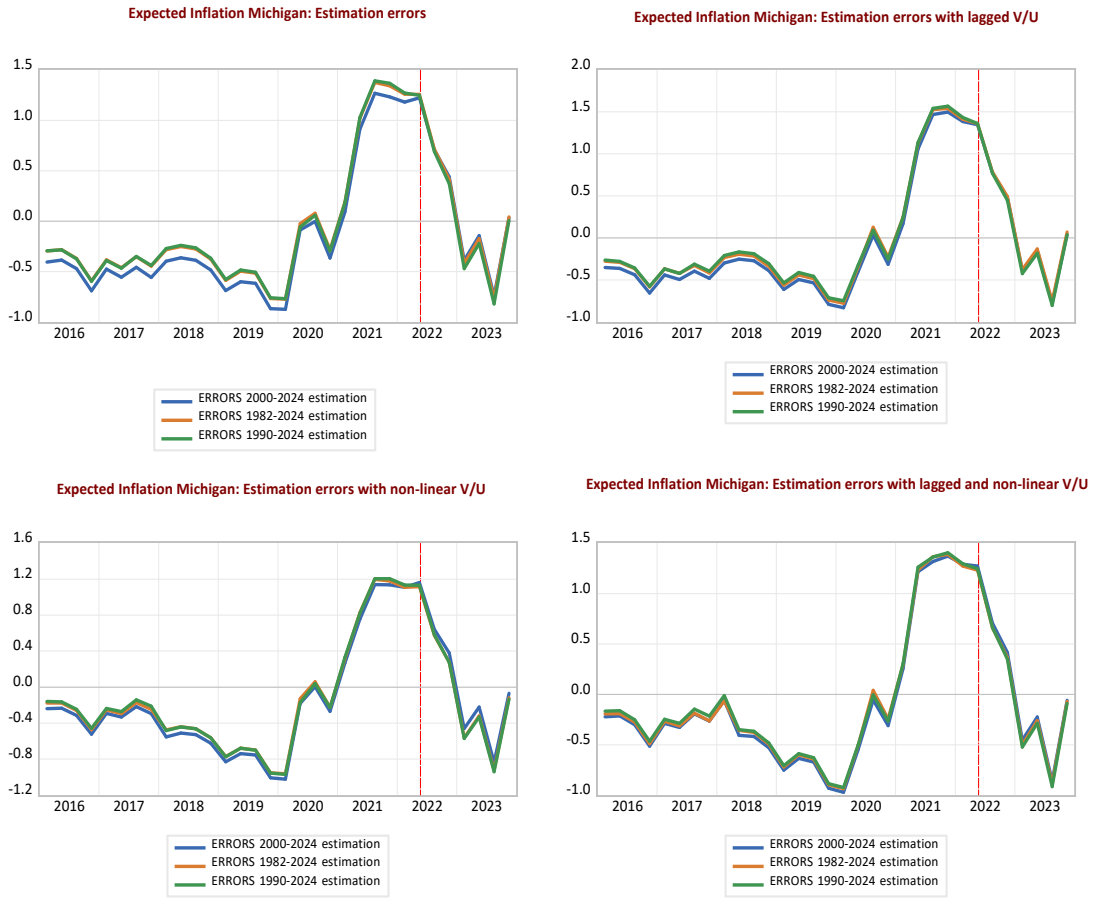
an empirical analogue for the residual in equation (7.36). As shown in the figure, we have estimated over three distinct sample periods: 1982Q1-2023Q4, 1990Q1-2023Q4, and 2000Q1-2023Q4. Some of our regressions also allow for non-linear effects in the labour market by including dummies for periods when $V/U > 1$, motivated by recent arguments from Benigno and Eggertsson (2023) that such non-linearities may have played an important role in driving post-pandemic inflation outcomes. To be more agnostic on the relevant timing assumptions, the figure also includes results for cases where we use contemporaneous values of the V/U ratio, rather than lagged values.

Across all of the specifications considered in Figure 11, a remarkably similar story emerges in terms of the way that the estimated residuals evolved over the course of the recent inflation surge. In particular, we consistently find that the residuals start rising around 2020Q2 and then keep rising until 2022Q2. At that point – i.e., the first full quarter of the US tightening cycle –, the residuals then begin declining back towards zero under all specifications and for all of the different sample periods considered. We view this as strong evidence in support of the model's main mechanism.[19]

More broadly, we view these findings as a complement to a growing body of empirical literature suggesting links between the timing of central banks' post-pandemic pivots and signs of re-anchoring in various indicators of inflation expectations. For example, Patzelt and Reis (2024) find that the first half of 2022 marked a peak in a dispersion-based measure of de-anchoring in European households' inflation expectations – with patterns then beginning to renormalize after the first ECB rate hike in July of that year. Using option price data to infer markets' perceived probability of a US "inflation disaster", Hilscher, Raviv, and Reis (2025) similarly note that the timing of a large drop in this probability "coincided almost to the month with the U-turn in monetary policy", arguing that "this provides support for a tight link between monetary policy and the inflation anchor". At the same time, Di Pace et al. (2025) document a significant impact of Bank of England rate hikes on UK firms' (own) price growth expectations, with evidence of non-linearities that appear to have amplified the impacts of larger hikes over the post-pandemic period. This last observation is especially notable, as it focuses on the expectations of players who are closer to real-world price-setting decisions – and also supports our model's prediction that larger pivots should have a disproprtionate impact on expectations (as stressed toward the end of Section 3).

[19] We have also run the expected inflation regressions with additional controls for supply conditions as measured by the Federal Reserve Bank of New York's Global Supply Chain Pressure Index (GSCPI). The key qualitative patterns depicted in Figure 11 remain unchanged with these additional global supply shock controls, independent of whether the GSCPI variable is included in levels or in changes. These results are available from the authors upon request.

Figure 11: Residuals in expected inflation regressions



Notes. The vertical dashed red lines mark 2022Q2, the first full quarter of the Federal Reserve's tightening cycle.

8 Conclusion

The appropriate monetary response to inflation driven by supply shocks is a question that has captured the attention of many policy-makers and analysts over the last few years. In this paper, we have aimed to contribute to this discussion by examining how best to stabilize inflation and employment under a range of assumptions on how agents form expectations, including cases where agents are sufficiently sophisticated to understand that policy decisions affect their environment but are not able to fully internalize all of the implications. The environment we examined was one where sticky wages can favour looking through supply shocks, while the sensitivity of inflation expectations to the policy stance and realized inflation makes that strategy risky.

We showed that under level-k thinking optimal policy generally takes the form of a threshold

rule. If inflation has been pushed only slightly above target, then policy should mainly look through supply shocks. However, if a series of shocks cumulate up to a specific threshold, we showed that policy needs to shift discontinuously to a more hawkish stance in order to control inflation expectations. While the hawkish policy stance involves reacting strongly to any subsequent inflation, the pivot itself – if executed optimally – can lead to a soft landing since its aim is to directly decrease inflation expectations and thereby inflation itself. To the extent that this allows policymakers to stabilize inflation without having to engineer significant slack, this indirect effect of tightening can largely counteract the negative direct effects, thereby creating potential for soft landings.

In 2022H1, many central banks changed their policy stances in manners that could be viewed as pivots. Through the lens of this model, this response represented a step in the right direction – but whether it was sufficiently aggressive is not clear. Especially in cases where policy-makers opted for less aggressive and/or more gradual pivots, our analysis suggests that a more decisive shift could potentially have led to a shorter tightening cycle and better odds of a soft landing.

References

Arad, A. and A. Rubinstein (2012). The 11-20 money request game: a level-k reasoning study. *American Economic Review 102*(7), 3561–73.

Benigno, P. and G. B. Eggertsson (2023, April). It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve. Working Paper 31197, NBER.

Bianchi Vimercati, R., M. S. Eichenbaum, and J. Guerriro (2021, August). Fiscal stimulus with imperfect expectations: Spending vs. tax policy. Working Paper 29134, NBER.

Blanchard, O. and N. Kiyotaki (1987). Monopolistic competition and the effects of aggregate demand. *American Economic Review 77*(4), 647–66.

Branch, W. and G. Evans (2006). A simple recursive forecasting model. *Economic Letters 91*(2), 158–166.

Calvo, G. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398.

Camerer, C. F., T. Ho, and J. Chong (2004, August). A cognitive hierarchy model of games. *Quarterly Journal of Economics 119*(3), 861–898.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2023, January). Anchored inflation expectations. *American Economic Journal: Macroeconomics 15*(1), 1–47.

Champagne, J., C. Hajzler, D. Matveev, H. Melinchuk, A. Poulin-Moore, K. Ozhan, Y. Park, and T. Taskin (2023, April). Potential output and the neutral rate in Canada: 2023 assessment. Working Paper 2023-6, Bank of Canada.

Costa-Gomes, M., V. P. Crawford, and N. Iriberri (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature 51*(1), 5–62.

di Giovanni, J., S. Kalemli-Özkan, A. Silva, and M. A. Yildirim (2022, July). Global supply chain pressures, international trade, and inflation. Working Paper 30240, NBER.

Di Pace, F., M. Giacomo, and R. Masolo (2025). Do firm expectations respond to monetary policy announcements? *Journal of Monetary Economics 149*(C).

Djeutem, E., M. He, A. Reza, and Y. Zhang (2022, March). Household heterogeneity and the performance of monetary policy frameworks. Working Paper 2022-12, Bank of Canada.

Farhi, E. and I. Werning (2019, November). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review 109*(11), 3887–3928.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications, Second edition*. Princeton University Press.

García-Schmidt, M. and M. Woodford (2019, January). Are low interest rates deflationary? A paradox of perfect-foresight analysis. *American Economic Review 109*(1), 86–120.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988, June). Investment, capacity utilization, and the real business cycle. *American Economic Review 78*(3), 402–417.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2022, August). The slope of the Phillips curve: evidence from U.S. states. *Quarterly Journal of Economics 137*(3), 1299–1344.

Hilscher, J., A. Raviv, and R. Reis (2025). How likely is an inflation disaster? *Review of Financial Studies*.

Iovino, L. and D. Sergeyev (2023). Central bank balance sheet policies without rational expectations. *Review of Economic Studies 90*(6), 3119–3152.

Kneeland, T. (2015). Identifying higher-order rationality. *Econometrica 83*(5), 2065–79.

Lorenzoni, G. and I. Werning (2023, January). Wage price sprials. Working paper, MIT.

Lucas, R. (1972). Expectations and the neutrality of money. *Journal of Economic Theory 4*(2), 103–124.

Macklem, T. (2025, March). Navigating tariff uncertainty. Speech at the University of Calgary.

Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics 54*(7), 2065–2082.

Milani, F. (2014). Learning and time-varying macroeconomic volatility. *Journal of Economic Dynamics and Control 47*, 94–114.

Patzelt, P. and R. Reis (2024, October). Estimating the rise in expected inflation from higher energy prices. Discussion Paper 18907, CEPR.

Preston, B. (2005, September). Learning about monetary policy rules when long-horizon expectations matter. *International Journal of Central Banking 1*(2), 81–126.

Reis, R. (2021, October). Losing the inflation anchor. Discussion Paper 16664, CEPR.

Reis, R. (2022). The burst of high inflation: How and why did we get here? Working paper, London School of Economics.

Sargent, T. J. (1982). The ends of four big inflations. In R. E. Hall (Ed.), *Inflation: Causes and Effects*, pp. 41–98. Chicago: Chicago University Press.

Wagner, J., T. Schlanger, and Y. Zhang (2022, Febraury). A horse race of alternative monetary policy regimes under bounded rationality. Working Paper 2022-4, Bank of Canada.

Woodford, M. (2004). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.