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

Paul Beaudry
Bank of Canada

Thomas J. Carter
Bank of Canada

Amartya Lahiri
University of British Columbia

October 24, 2022

Abstract

Central banks in most advanced economies have reacted similarly to the increase in inflation that started in 2021. They initially looked through the rising inflation by leaving monetary policy relatively unchanged. Then, after inflation continued to rise, central banks pivoted by quickly tightening policy. The pivot was explained, at least in part, as aiming to anchor drifting inflation expectations. Why might central banks want to look through supply-driven inflation sometimes and pivot away at other times? When does a tighter monetary policy stance help anchor expectations? Can aggressive monetary tightening be compatible with a soft landing? The paper develops a simple model to clarify these issues. The model has two key features: (a) risk of a wage-price spiral; and (b) the central bank views wage- and price-setters as having bounded rationality. We show how optimal policy in this environment mimics qualitative features of recent central bank behaviour, thereby offering a positive theory for the observed choices.

Keywords: Inflation expectations, supply shocks, monetary policy

JEL Classification: E3, E5

*The views expressed are those of the authors and do not necessarily reflect the views of the Bank of Canada. We would like to thank Roc Armenter, Jose Dorich, Ludwig Straub and seminar participants at UBC for helpful comments and discussions.

1 Introduction

Starting in mid-2021, as inflation started to rise globally, many central banks went through similar sequences of responses. First, they looked through the shocks in the sense of not reacting to rising inflation. This inert response was typically defended by pointing to the supply-side origins of the inflation as well as the likelihood that high inflation would prove to be temporary.[1] However, when inflation shocks kept materializing, central banks pivoted to a much more aggressive policy stance. Policy-makers then spent considerable effort defending the pivot as being necessary to anchor expectations in order to avoid igniting a wage-price spiral.

The objective of this paper is to argue that a large part of recent central bank behaviour may be better understood by adopting a bounded rationality perspective in an environment where wage-price spirals are possible. In particular, we argue that the observed monetary policy responses are hard to rationalize in environments where central banks view the public as being either fully rational (i.e., having rational expectations) or as simply forming adaptive expectations. Instead, we show that pivoting behaviour becomes optimal if central banks view agents as having bounded rationality and, more precisely, as using **level-k thinking**.

To this end, we present a simple model of price and wage determination that allows us to focus on the different ways in which agents can form expectations about inflation and the pricing decisions of others, and how this influences the design of optimal monetary policy. The model builds on the tractable framework first developed by Blanchard and Kiyotaki (1987), which we extend to capture the possibility of wage-price spirals.

The model aims to capture the interplay between prices and wages by introducing a timing structure that allows prices to adjust faster than wages in response to supply shocks. This feature has two consequences. First, it allows supply shocks to potentially increase inflation even when expectations are anchored and the economy is operating at its natural level.[2] Second, it introduces a potential wage-price spiral in the model, since wages are set based on expected inflation. Our set-up purposely departs from the canonical New Keynesian model by not exhibiting a "divine coincidence" whereby central banks can simultaneously keep inflation at target and output at its natural level when responding to shifts in the prices of key imported inputs like oil, productivity shocks, or other such supply shocks.

[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] When referring to supply shocks we are not including mark-up shocks.

Our model of expectation formation accommodates both rational and adaptive expectations as special cases. We use the model to first show that under rational expectations, it is optimal for central banks to always fully look through supply-driven inflation shocks when wages adjust more slowly than prices. In contrast, under adaptive expectations, we show that a policy-maker should never entirely look through inflation shocks. However, the optimal policy in this case does not involve a pivot. Rather, it is optimal for policy-makers to react to inflation in a way that involves a constant degree of look-through.

We then show that in the general case of boundedly rational agents forming expectations using level-k thinking, it becomes optimal for central banks to initially look through supply-driven inflationary shocks but then pivot to an aggressive monetary policy response if inflation shocks cumulate above a certain threshold. This prediction is in sharp contrast to those emerging under rational and adaptive expectations and reproduces the recent behaviour of many central banks. This is a key result of the paper.

Intuitively, level-k thinking makes inflation expectations a function of both past inflation and how one thinks others will set wages and prices. Since inflation depends on, amongst other factors, expected inflation, policy-makers face a trade-off: a "low look-through" policy that involves raising the policy rate to fight inflation helps to stabilize inflation expectations but simultaneously extracts an output and employment cost. In general, the solution to this trade-off depends on how far inflation has deviated from the central bank's target. For inflation deviations below a threshold, we show that it can be optimal for monetary policy to mainly look through supply shocks. However, beyond the threshold it becomes optimal to pivot to a strong anti-inflation stance.

The policy pivot described above occurs due to a non-convexity in the payoff function. Policy-makers in our setup minimize a loss function that depends on the squared deviations of inflation and employment from their targets. Monetary policy affects employment in our model through two channels: (a) a direct effect whereby economic activity is reduced by tighter policies for any given level of expected inflation; and (b) an indirect effect whereby tighter policy reduces inflation expectations. These two effects go in opposite directions: the direct effect decreases employment while the indirect effect raises it. As we show, the effect of a tight monetary stance on inflation expectations increases with the tightness of policy under level-k thinking, which is the source of the non-convexity. The increasing response of inflation expectations allows the indirect effect of policy tightening to potentially offset entirely the direct effect if policy is set to be sufficiently tight and the inflation deviation is high enough. Consequently, when deciding to tighten policy, there is a benefit to tightening by a sufficiently large amount once inflation deviations reach a threshold level since that generates a large decrease in inflation expectations. This indirect effect reduces the fall in employment induced directly by the aggressive policy stance. As a result, these counterbalancing forces open the possibility of a soft-landing for employment, despite a strong policy pivot, when policy is designed optimally.

In contrast to much of the New Keynesian literature, in this model, we downplay the role of multi-period nominal rigidities by assuming that prices and wages are pre-set for only one period. This allows us to concentrate on the cross-sectional interdependence of wage and price decisions, which is at the centre of wage-price spirals. Crucially, in our framework, wage-setters have to form expectations of prices, which are then set by firms following a standard mark-up pricing rule. Hence, forming expectations of prices amounts to forming expectations of the average wage which, in turn, requires forming expectations of the wages set by others. For this reason, the set-up embeds features of standard coordination games and thereby lends itself easily to an exploration of the potential implications of bounded rationality.

Throughout most of the paper, we disregard the role of demand shocks to focus on the dilemma associated with supply shocks. We adopt this focus despite the likelihood that demand forces played an important role in the recent inflation episode in many countries (as shown in di Giovanni et al., 2022). This choice reflects the fact that our approach has nothing novel to contribute regarding the demand side of the economy, as our framework shares the common property that optimal monetary policy should fully offset demand shocks. Accordingly, one should view our results regarding the monetary policy response to supply shocks as describing the central bank's preferred monetary stance over and above any policy adjustment that is required to offset demand shocks. We discuss this in Section 5.

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 and longer-standing 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 germane in emerging economies.[3]

[3]In the model, we subsume different types of supply shocks under the common rubric of productivity shocks. In

Our work is related to three distinct yet related strands of the macroeconomic literature. The first is, of course, the voluminous literature on New Keynesian (NK) macroeconomic models with nominal stickiness. In-depth and comprehensive summaries of the key implications of the NK approach to monetary policy can be found in Galí (2015) and Woodford (2004), amongst many others.

The second literature focuses on bounded rationality and its implications for macroeconomics. The problem of forecasting the forecasts of others and associated limits of deductive reasoning were famously discussed in Keynes (1936) in the context of newspaper competitions to judge beauty contests. Amongst others in the experimental game theory literature, Nagel (1995) and Costa-Gomes et al. (2001) offer evidence of bounded rationality in the form of level-k thinking. In addition, the implications of level-k thinking in terms of its ability to dampen the effects of monetary policy have been shown in Farhi and Werning (2019). Relative to previous work, we believe our paper is distinct in its focus on the implications of level-k thinking for optimal policy design in response to supply shocks.[4]

Our paper also contributes to a third, nascent literature on the origins of the current global inflationary episode, along with its consequences and implications for policy. Reis (2022) provides an excellent overview of the various hypotheses regarding the burst of inflation, the role of the inflation anchor, and the associated policy challenges for central bankers. Evidence of the importance of managing inflation expectations can be found in Reis (2021), which examines the role of unanchored inflation expectations in driving the Great Inflation of the 1970s in the United States.

The rest of the paper is organized as follows: the next section presents the model; Section 3 formalizes the optimal policy problem and presents a few special cases of the problem which can be solved analytically; Section 4 presents the fully dynamic policy problem and numerical simulations of the optimal policy; Section 5 discusses demand shocks; the last section concludes. All proofs are contained in the Appendix.

2 Environment

We consider an economy where wages are set before prices, with prices being more flexible in the sense that they are allowed to react to current productivity developments while wages cannot. This

Appendix A.4, we formally show an equivalence between productivity shocks and oil price shocks, which are one of the more volatile and frequent shocks that hit both emerging and advanced economies.

[4]Recent work by García-Schmidt and Woodford (2019) on reflective equilibria also explores the effects of relaxing rational expectations in favour of bounded rationality.

feature will be key in delivering inflation dynamics that are different from those in a canonical New Keynesian model (see, for example, Galí, 2015 or Woodford, 2004) and will help rationalize why looking through the effects of supply shocks on inflation can sometimes be warranted. As will be discussed in the Appendix, negative productivity changes in the model can be interpreted as reflecting positive changes in the price of an imported input such as oil, which is the more relevant interpretation in relation to recent events.

The model economy consists of a set of n infinitely lived private agents facing a demand for their labour that is decreasing in the wage they set, along with a set of m monopolistically competitive firms facing a demand for their goods that is decreasing in the price that they set. The agents and firms both consider themselves too small to affect aggregate outcomes. The central bank chooses monetary policy to minimize a weighted sum of squared deviations of output and inflation from targets. We will refer to the outcome of this minimization as optimal monetary policy, even though we are not deriving the central bank's objective from the fundamentals of the model.

The set-up differs from the canonical New Keynesian model in that price and wage rigidities last only one period. This allows the model to focus on the cross-sectional coordination problem that arises among price- and wage-setters at each point in time, as opposed to the inter-temporal trade-offs in price setting that arise with multi-period price rigidities. In this sense, the model can be seen as being closer in spirit to the environment first formalized in Blanchard and Kiyotaki (1987).

An important element of our analysis will be to consider different processes for how central banks may think private agents form their expectations, and how this affects optimal monetary policy decisions. Accordingly, when using the expectation operator to write $\mathbb{E}_{t-1}X_t$ as the expectation of an endogenous variable X_t based on $t-1$ information, we will not necessarily be referring to the rational expectation of X_t.

2.1 Individuals

There are n individuals, each of whom maximizes their expected lifetime utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln C_{it} - \eta N_{it}\right],$$

where C_{it} is the number of final goods consumed by individual i at date t, and N_{it} denotes labour supplied by individual i. Subscript t denotes time throughout the paper but will sometimes be suppressed when it should be clear from context.

The final good is produced by combining a continuum of intermediate goods according to

$$C_{it} = \left(\sum_{j=1}^{m} C_{ijt}^{\frac{\gamma-1}{\gamma}} \right)^{\frac{\gamma}{\gamma-1}}, \tag{2.1}$$

where C_{ijt} indicates consumption of intermediate good j by individual i. There is a constant set of m intermediate goods at every date.

The individual faces the periodic budget constraint

$$\sum P_{jt} C_{ijt} + B_{it+1} = W_{it} N_{it} + D_{it} + \tau_{it} + B_{it}(1 + \iota_t),$$

where B_{it+1} is the purchases of nominal bonds by individual i at date t, P_{jt} is the price of good j, W_{it} is the nominal wage of individual i, D_{it} are dividends received by i from ownership of firms, ι_t is the nominal interest rate and τ_{it} are transfers received from the government by i. Prices, dividends, interest rates and government transfers are taken as exogenous by the individual. In the following, it will be convenient to use the definition

$$I_{it} \equiv W_{it} N_{it} + D_{it} + \tau_{it} + B_{it}(1 + \iota_t). \tag{2.2}$$

Individuals in this model make three decisions in each period t, choosing their allocation of consumption spending across the various intermediate goods, the bond holdings that they plan to carry into the next period, and the wage they request for employment in the next period. This decision-making takes place in two stages. In the first stage, agents choose their optimal allocation of total consumption spending $I_{it} - B_{it+1}$ across the different intermediate goods. In the second stage, households make their consumption-savings decision and set their next-period wage to maximize lifetime utility, taking into account the effect of these choices on C_{it} through their impact on overall spending $I_{it} - B_{it+1}$.

The first-stage problem involves maximizing the expression on line 2.1, subject to the budget constraint $\sum_j P_{jt} C_{ijt} \leq I_{it} - B_{it+1}$. Note that the agent chooses consumption after having observed all relevant shocks for the period. This places optimal consumption demand for good j at

$$C_{ijt} = \left(\frac{P_{jt}}{P_t} \right)^{-\gamma} \left(\frac{I_{it} - B_{it+1}}{m P_t} \right), \tag{2.3}$$

where we have defined the price level as

$$P_t = \left(\frac{1}{m} \sum_{j=1}^{m} P_{jt}^{1-\gamma} \right)^{\frac{1}{1-\gamma}}. \tag{2.4}$$

Equation 2.3 can be used to derive the optimal demand for the final composite good C_{it} as

$$C_{it} = m^{\frac{1}{\gamma-1}} \left(\frac{I_{it} - B_{it+1}}{P_t} \right). \tag{2.5}$$

The second stage problem for the agent is to choose B_{it+1} and W_{it+1} to solve the dynamic problem represented by the value function

$$V\left(\frac{B_{it}}{P_t}, \frac{W_{it}}{P_t}\right) = max \ \left\{ \ln C_{it} - \eta N_{it} + \beta \mathbb{E}_t V\left(\frac{B_{it+1}}{P_{t+1}}, \frac{W_{it+1}}{P_{t+1}}\right) \right\}, \tag{2.6}$$

subject to equations 2.2 and 2.5 and the labour demand function facing the agent. We will solve the wage-setting problem after deriving the labour demand from the firm's problem below.

The optimal consumption-savings decision leads to the standard Euler equation

$$\frac{1}{P_t C_{it}} = \beta \mathbb{E}_t \left(\frac{1 + \iota_{t+1}}{P_{t+1} C_{it+1}} \right). \tag{2.7}$$

2.2 Firms

Firms in this economy produce intermediate goods using labour according to the production function

$$Y_{jt} = \theta_{jt} \left(\sum_{i=1}^{n} N_{ijt}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}},$$

where θ_{jt} is the firm's productivity. Productivity is stochastic and follows a process that is common knowledge. In the following, we shall use the definition

$$N_{jt} = \left(\sum_{i=1}^{n} N_{ijt}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}.$$

Firms maximize profits by choosing labour and the price of their product. All firm decisions are made after the firm observes wages and its own productivity for the period.

2.2.1 Labour demand

Firm j chooses its labour inputs to minimize its wage bill, subject to the production function. The firm's demand for labour is

$$N_{ijt} = n^{\frac{\rho}{1-\rho}} \left(\frac{W_{it}}{W_t} \right)^{-\rho} \frac{Y_{jt}}{\theta_{jt}}, \tag{2.8}$$

where

$$W_t = \left(\frac{1}{n} \sum_i W_{it}^{1-\rho} \right)^{\frac{1}{1-\rho}}$$

denotes the aggregate wage index.

2.2.2 Price-setting rule

Firms are monopolistically competitive in output markets. They set prices to maximize profits. Thus, firm j maximizes

$$D_{jt} = P_{jt} Y_{jt} - \sum_{i=1}^{n} W_{it} N_{ijt}.$$

This optimization is done subject to the following two constraints, along with the demand curve given by equation 2.3:

$$Y_{jt} = \sum_i C_{ijt}$$

$$\sum_{i=1}^{n} W_{it} N_{ijt} = n^{\frac{1}{1-\rho}} W_t \frac{Y_{jt}}{\theta_{jt}}.$$

Note that the total wage bill for firm j follows directly from equation 2.8.

This problem gives the optimal price as

$$P_{jt} = \left(\frac{\gamma}{\gamma - 1} \right) n^{\frac{1}{1-\rho}} \frac{W_t}{\theta_{jt}}. \tag{2.9}$$

This is a standard pricing rule involving a fixed mark-up over marginal cost. The expression also shows the usual inverse relationship between a firm's productivity and its price. As we shall show below, this inverse relationship between firm productivity and firm prices induces a similar negative relationship between aggregate inflation and aggregate productivity shocks.

2.3 Wage demands by households

Wages are set before the realization of productivity, prices and all aggregate outcomes. Households set wages based on the expected values of these variables. Formally, individuals solve the problem by choosing W_{it} to maximize the dynamic objective given by equation 2.6, subject to the constraints given by equations 2.2 and 2.8. Note that from equation 2.8, the total demand for household i's labour is

$$N_{it} = \left(\frac{W_{it}}{W_t}\right)^{-\rho} \frac{N_t}{n}, \tag{2.10}$$

where we have used the definition

$$N_t \equiv n^{\frac{1}{1-\rho}} \sum_j \frac{Y_{jt}}{\theta_{jt}}.$$

N_t is the aggregate employment index.

This problem gives the optimal wage as

$$W_{it+1} = \left(\frac{\rho\eta}{(\rho-1)m^{\frac{1}{\gamma-1}}}\right)\left(\frac{\mathbb{E}_t N_{t+1}}{\mathbb{E}_t(N_{t+1}/P_{t+1}C_{it+1})}\right). \tag{2.11}$$

Note that equation 2.11 implies that the current period wage W_{it} is based on the information set at date $t-1$. This is a consequence of our modelling of nominal rigidities as taking the form of wages being set a period in advance. This nominal stickiness is not persistent as wages are set for just one period at a time.

If we take the log of equation 2.11, and disregard Jensen's inequality by assuming that $\ln EX = E\ln X$, we see that wages are increasing in expected prices and expected consumption:

$$\ln W_{it} = \ln\left[\frac{\rho\eta}{(\rho-1)m^{\frac{1}{\gamma-1}}}\right] + \mathbb{E}_{t-1}\ln P_t + \mathbb{E}_{t-1}\ln C_{it},$$

where we have lagged the optimal wage expression in equation 2.11 by one period. This expression illustrates the coordination problem faced by wage-setters. To set wages for the next period, individuals have to form expectations of the average price level for the next period. Since price-setters follow a mark-up pricing rule, this implies that individual wage-setters have to form expectations about the average wage, which requires forming expectations regarding the wage-setting behaviour of others.

2.4 Aggregates

The only source of uncertainty in the model is firm productivity θ_{jt}. Given the cross-sectional variation of productivity across firms, we define aggregate productivity as

$$\theta_t = \left(\frac{1}{m} \sum_j \theta_{jt}^{\gamma-1} \right)^{\frac{1}{\gamma-1}}.$$

In the following, we shall make two simplifying[5] assumptions about the productivity process:

Assumption 2.1. $\theta_{jt} = \theta_t$ *for all* j

Assumption 2.2. $\ln \theta_t = \ln \theta_{t-1} + \epsilon_t$*, where* ϵ_t *is an i.i.d. shock with mean zero and variance* σ_θ^2

These assumptions imply that there is only aggregate uncertainty regarding firm productivity, with all firms receiving the aggregate productivity draw in each period. This is reasonable if we want to allow productivity shocks to be interpreted as the negative of oil price shocks.[6] Though individuals are unaware of the aggregate productivity realization for the next period at the time they set their wages, they know both the mean and variance of the innovation ϵ_t.

Using the solution for labour demand N_i from equation 2.10 gives us $\sum_i W_{it} N_{it} = W_t N_t$. Individuals in this economy own the firms, hence they receive all the firm dividends. This implies that

$$\sum_i I_{it} = P_t Y_t,$$

where $P_t Y_t = \sum_j \sum_i P_{jt} C_{ijt}$ is aggregate demand in the economy.

Since $\theta_{jt} = \theta_t$ for all j, it follows from equation 2.9 that $P_{jt} = P_t$ for all j. Using this in the definition $P_t Y_t = \sum_j P_{jt} Y_{jt}$ gives

$$Y_t = n^{\frac{1}{\rho-1}} \theta_t N_t,$$

where we have used the definition $N_t \equiv n^{\frac{1}{1-\rho}} \sum_j N_{jt}$.

2.4.1 Natural level of employment

In the following, we derive an expression for the natural rate of employment in the model economy, $\bar{N}_t$. To solve for it, note that the optimal wage of individual i when wages and prices are flexible

[5] Most of our results regarding optimal policy do not rely on the particular process for aggregate productivity.

[6] The equivalence between negative productivity shocks and positive oil price shocks is formally demonstrated in Appendix A.4

and determined simultaneously is given by

$$W_{it} = \left(\frac{\rho\eta}{\rho - 1}\right)\frac{P_t C_{it}}{m^{\frac{1}{\gamma-1}}}.$$

Summing this expression over i and noting that $W_{it} = W_t = \bar{W}_t$ for all i in the symmetric case gives

$$\bar{W}_t = \left(\frac{\rho\eta}{\rho - 1}\right) P_t Y_t.$$

Since $P_t Y_t = \frac{\gamma}{\gamma-1}\bar{W}_t \bar{N}_t$, this expression for $\bar{W}_t$ reduces to

$$\bar{N}_t = \left(\frac{\gamma - 1}{\gamma}\right)\left(\frac{\rho - 1}{\rho\eta}\right), \tag{2.12}$$

making $\bar{N}_t$ constant and, crucially, independent of θ_t. Denoting it $\bar{N}$ herein, the natural level of output at time t is given by

$$\bar{Y}_t = n^{\frac{1}{\rho-1}}\theta_t \bar{N}.$$

It follows that the output gap is given by $\ln\left(\frac{N_t}{\bar{N}}\right)$, since $\frac{Y_t}{\bar{Y}_t} = \frac{n^{\frac{1}{\rho-1}}\theta_t N_t}{n^{\frac{1}{\rho-1}}\theta_t \bar{N}}$.

2.4.2 Aggregate wages and inflation

To characterize the evolution of aggregate wages, it is useful to first note that in a symmetric equilibrium we have $W_{it} = W_t$ and $C_{it} = C_t$. Moreover, summing equation 2.5 over all i gives

$$\frac{P_t C_t}{m^{\frac{1}{\gamma-1}}} = I_t = P_t Y_t.$$

Substituting this expression into the solution for the optimal wage given by equation 2.11 and taking logs gives

$$\ln W_t = \ln\left(\frac{\rho\eta}{\rho - 1}\right) + \left(\frac{1}{\rho - 1}\right)\ln n + \mathbb{E}_{t-1}\ln P_t + \mathbb{E}_{t-1}\ln\theta_t + \mathbb{E}_{t-1}\ln N_t, \tag{2.13}$$

where we have used $Y_t = n^{\frac{1}{\rho-1}}\theta_t N_t$. Note that in deriving equation 2.13 we have approximated $\ln\mathbb{E}X$ by $\mathbb{E}\ln X$, thereby ignoring Jensen's inequality.[7] We will repeatedly use this approximation to get simple linear expressions. The subscript $t-1$ on the expectation operators in the equation indicate that the expectations are based on information available at the end of date $t-1$ when the

[7] This linear representation can alternatively be achieved by taking log-linear approximations.

wage decision for period t was taken. This timing of expectations arises because wages for period t were set at the end of period $t-1$, before the realization of θ_t.

Substituting equation 2.12 into the expression for $\ln W_t$ above gives aggregate wages as

$$\ln W_t = \mathbb{E}_{t-1} \ln P_t + \mathbb{E}_{t-1} \ln \theta_t + \mathbb{E}_{t-1} \ln \left(\frac{N_t}{\bar{N}} \right) + \ln \left(\frac{\gamma - 1}{\gamma} \right) + \left(\frac{1}{\rho - 1} \right) \ln n. \tag{2.14}$$

To derive the equilibrium aggregate price level, we can use the solution for the optimal price and the expression for the aggregate price level (equations 2.9 and 2.4, respectively) to get

$$\ln P_t = \ln \left(\frac{\gamma}{\gamma - 1} n^{\frac{1}{1-\rho}} \right) + \ln W_t - \ln \theta_t. \tag{2.15}$$

Equation 2.15 gives the equilibrium expression for the aggregate price level. Define

$$\omega_t \equiv \ln W_t - \ln P_{t-1} - \ln \left(\frac{\gamma - 1}{\gamma} \right) - \left(\frac{1}{\rho - 1} \right) \ln n$$

$$\pi_t \equiv \ln P_t - \ln P_{t-1}.$$

Substituting these definitions into equations 2.14 and 2.15 and combining the resulting expressions yields

$$\pi_t - \pi^* = \mathbb{E}_{t-1}(\pi_t - \pi^*) + \mathbb{E}_{t-1} \left(\ln N_t - \ln \bar{N} \right) - (\ln \theta_t - \mathbb{E}_{t-1} \ln \theta_t). \tag{2.16}$$

Equation 2.16 is the Phillips curve that emerges from the model. It says that inflation depends positively on inflation expectations and the expected output gap, and negatively on the productivity shock. This inflationary effect of negative productivity shocks is the mapping in the model to the inflationary effects of the negative supply shocks emphasized in the introduction.

This Phillips curve departs from the canonical Phillips curve of the New Keynesian literature in a few important ways. First, in our set-up, inflation at date t is driven by expectations of inflation formed at date $t-1$, as opposed to expectations of future inflation. Second, this Phillips curve does not imply a long-run trade-off between inflation and activity. Third, it is the expected output gap that drives inflation, since it is these expectations that drive wage-setting decisions. Fourth, productivity shocks have a direct effect on inflation even when expected inflation is on target and there is no expected output gap.[8]

[8] The goal of this paper is not to show that this Phillips curve fits the data better than a more standard Phillips

2.5 The Euler equation

By exploiting the symmetry $C_{it} = C_t$ for all i, we can write the individual Euler equation in equation 2.7 as

$$\ln C_t = \mathbb{E}_t \ln C_{t+1} - [(\iota_{t+1} - \bar{\iota}) - \mathbb{E}_t(\pi_{t+1} - \pi^*)],$$

where $\bar{\iota} \equiv \pi^* - \ln \beta$. Substituting the market clearing condition $C_t = m^{\frac{1}{\gamma-1}} Y_t = m^{\frac{1}{\gamma-1}} n^{\frac{1}{\rho-1}} \theta_t N_t$ into this aggregate Euler equation then gives

$$\iota_{t+1} - \bar{\iota} = \mathbb{E}_t(\ln N_{t+1} - \ln N_t) + \mathbb{E}_t(\pi_{t+1} - \pi^*), \tag{2.17}$$

where we have used the random walk property of θ_t.

2.6 Monetary policy rule

The aim of monetary policy is to minimize deviations of inflation and employment from pre-specified targets. Since the monetary authority can control the interest rate ι_t, we could immediately specify monetary policy in terms of a rate-setting rule and later optimize the parameters of that rule. However, given our interest in the policy implications of different inflation expectation formation processes, it is more convenient to think of monetary policy as directly aiming to control short-term employment N_t. We can then find the potentially time-varying interest rate rule that implements policy-makers' preferred employment outcomes and explore how this rule depends on the expectation formation process assumed in the private sector.

In light of the above, we will think of monetary policy as choosing how best to set N_t as a function of inflation, allowing the strength of this feedback to potentially vary over time. Allowing the feedback rule to change over time will enable us to ask under what conditions (if any) policy-makers would find it optimal to pivot in the sense of first not responding much to inflation pressures and then switching to a much stronger reaction if inflation picks up substantially.

To this end, we will consider monetary policy as being set in order to engineer employment outcomes in line with a feedback rule of the form

$$N_t = \bar{N} \left(\frac{1 + \pi_t}{1 + \pi^*} \right)^{-\phi_t},$$

curve. Instead, our claim is that this Phillips curve captures important aspects of how many central banks perceive the inflation process that are not reflected in the canonical New Keynesian Phillips curve.

where $\pi_t = \frac{P_t}{P_{t-1}}$ is the inflation rate, π^* is the central bank's target level of inflation, and $\bar{N}$ is the natural rate of employment. The parameter ϕ_t will govern the extent to which monetary policy is contractionary in response to inflation. This rule can be written in log-linear form as

$$\ln N_t - \ln \bar{N} = -\phi_t(\pi_t - \pi^*). \tag{2.18}$$

Once we solve the equilibrium under this assumed rule for employment, we can then use the Euler equation to find the interest rate rule that implements it. This rule will have to satisfy equation 2.17. As we shall show, under rational expectations, the optimal policy can be implemented by a very simple time-invariant interest rate rule.[9]

In Appendix A.3 we are more explicit about how to generate the feedback rule for monetary policy given in equation 2.18. In particular, we solve the Euler equation forward to express the employment deviation $\hat{N}_t$ as a negative function of the cumulated future deviations of the real interest rate from the natural rate of interest. We define this cumulated path of all future interest deviations as the stance of monetary policy and allow monetary authorities to set this quantity as a function of inflation.

To close the model, we need to specify the supply of bonds and any government transfers. We will assume that the net supply of bonds is zero and therefore the government does not need to raise any taxes to pay interest on bonds.

2.7 The equilibrium system

The equilibrium can be computed in recursive fashion by solving for equilibrium $\hat{\pi}$ and $\hat{N}$ from the Phillips curve and monetary policy rule given by equations 2.16 and 2.18, respectively. The interest rate rule that implements the equilibrium path for employment and inflation will then need to satisfy equation 2.17.

3 Optimal Policy

We now consider the design of optimal monetary policy in this model economy. Specifically, we want to address the following question: Given the structure of the economy, how aggressively should

[9]This formulation of monetary policy amounts to directly modelling policy-makers as recognizing that their policy response to deviations of inflation from target affects economic activity. The more aggressively they respond to inflation deviations, the higher is ϕ_t. The higher ϕ_t is, the greater the fall in employment.

policy-makers respond to deviations of inflation from the central bank's target? In particular, how does this response depend on the central bank's perceptions about the process of expectation formation?

Before going to optimal policy, it is helpful to first look at outcomes under different theories of expectation formation for arbitrary sequences of values for the policy stance parameter ϕ_t. In the following, we use the notation $\hat{\pi}_t = \pi_t - \pi^*, \mathbb{E}_{t-1}\hat{\pi}_t = \mathbb{E}_{t-1}(\pi_t - \pi^*)$, $\hat{\theta}_t = \ln\theta_t - \mathbb{E}_{t-1}\ln\theta_t$ and $\hat{N}_t = \ln N_t - \ln\bar{N}$.

1. **Rational expectations (RE)**: To determine the rational expectations outcome, we can take the date $t-1$ expectation of equation 2.16 to get $\mathbb{E}_{t-1}\hat{N}_t = 0$. Combining this with equation 2.18 gives $\mathbb{E}_{t-1}\hat{\pi}_t = 0$, assuming $\phi_t > 0$. Equations 2.16 and 2.18 then give

$$\hat{\pi}_t^{RE} = -\left(\ln\theta_t - \mathbb{E}_{t-1}\ln\theta_t\right) = -\hat{\theta}_t$$
$$\hat{N}_t^{RE} = \phi_t\left(\ln\theta_t - \mathbb{E}_{t-1}\ln\theta_t\right) = \phi_t\hat{\theta}_t,$$

   where $\hat{\pi}_t^{RE}$ and $\hat{N}_t^{RE}$ represent realizations of $\hat{\pi}$ and $\hat{N}$ under rational expectations. In this case, inflation is above (below) the policy target when the productivity shock is negative (positive). Correspondingly, the expected inflation rate in the model equals the inflation target under rational expectations for any $\phi_t > 0$. In the limit where $\phi_t \to 0$, employment would be stabilized and inflation would become an *i.i.d.* process.

2. **Adaptive expectations (AE)**: Consider the simplest version of adaptive expectations, wherein $\mathbb{E}_{t-1}\pi_t = \pi_{t-1}$ and $\mathbb{E}_{t-1}N_t = N_{t-1}$. Under these expectations, it is straightforward to verify that

$$\hat{\pi}_t^{AE} = \hat{\pi}_{t-1} + \hat{N}_{t-1} - \hat{\theta}_t$$
$$\hat{N}_t^{AE} = -\phi_t(\hat{\pi}_{t-1} + \hat{N}_{t-1} - \hat{\theta}_t),$$

   where $\hat{\pi}_t^{AE}$ and $\hat{N}_t^{AE}$ represent realizations of $\hat{\pi}$ and $\hat{N}$ under adaptive expectations. In this case, if monetary policy-makers decided to look through inflation pressures by setting $\phi_t = 0$, then inflation would become a random walk. This illustrates how the mapping between policy and inflation outcomes can change drastically depending on one's view of the inflation expectation process.

3.1 Level-k thinking

Our goal in this paper is to examine the policy implications of a more general framework for expectation formation that can accommodate rational and adaptive expectations within a broader range of possibilities. In our departure from rational expectations, we choose to exploit the concept of *level-k thinking* developed in the game theory literature. In the macroeconomics literature, this concept was recently used by Farhi and Werning (2019) to study monetary policy.

Under level-k thinking, individuals respond to any deviation from a rational expectations equilibrium by starting with an initial guess about the macroeconomic expectations of other agents and computing the aggregate outcome under that guess. The guess about others' expectations is then updated to reflect the aggregate outcome under the previous guess and used as the initial guess of a new iteration. This process is repeated recursively k times. The restriction to a finite number $k < \infty$ of iterations reflects some bounded computing power on the part of individual agents on account of limited resources or capacity for forecasting the future.

As we shall show, this expectation formation process converges to the rational expectations outcome when the number of iterations k goes to infinity. On the other hand, when $k = 0$ and agents use last period's outcomes as their initial guesses on the expectations of others, the model reduces to one with adaptive expectations.

To illustrate the impact of level-k thinking on expectation formation in the context of our model, recall that aggregate inflation and employment are given by

$$\hat{\pi}_t = \mathbb{E}_{t-1}\hat{\pi}_t + \mathbb{E}_{t-1}\hat{N}_t - \hat{\theta}_t$$
$$\hat{N}_t = -\phi_t[\mathbb{E}_{t-1}\hat{\pi}_t + \mathbb{E}_{t-1}\hat{N}_t - \hat{\theta}_t].$$

Level-k thinking starts with some initial (level-0) expectations of $\hat{\pi}_t$ and $\hat{N}_t$ that are used to generate the level-1 expectation of these variables. This process then continues recursively up to some finite k. Using $\mathbb{E}_{t-1}\hat{x}_t^k$ to denote the expectation of variable $\hat{x}$ that holds for iteration number k, we can write

$$\mathbb{E}_{t-1}\hat{\pi}_t^1 = \mathbb{E}_{t-1}\hat{\pi}_t^0 + \mathbb{E}_{t-1}\hat{N}_t^0$$
$$\mathbb{E}_{t-1}\hat{N}_t^1 = -\phi_t[\mathbb{E}_{t-1}\hat{\pi}_t^0 + \mathbb{E}_{t-1}\hat{N}_t^0],$$

where we have assumed that level-k thinking does not impact agents' expectations on exogenous variables, so $\mathbb{E}_{t-1}\hat{\theta}_t^k = 0$. Repeating the recursion above k times and substituting the level-k expectations into equations 2.16 and 2.18 gives

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

We shall assume throughout the following analysis that $\mathbb{E}_{t-1}\hat{\pi}_t^0 = \hat{\pi}_{t-1}$ and $\mathbb{E}_{t-1}\hat{N}_t^0 = \hat{N}_{t-1}$, i.e., the initial seeds for the level-k iteration on the private agent's inflation and employment expectations are the previous period values of these two variables. Under this assumption, the equilibrium system becomes

$$\hat{\pi}_t^{KLT} = (1-\phi_t)^k\left[\hat{\pi}_{t-1} + \hat{N}_{t-1}\right] - \hat{\theta}_t \tag{3.19}$$
$$\hat{N}_t^{KLT} = -\phi_t\left[(1-\phi_t)^k\left\{\hat{\pi}_{t-1} + \hat{N}_{t-1}\right\} - \hat{\theta}_t\right], \tag{3.20}$$

where $\hat{\pi}_t^{KLT}$ and $\hat{N}_t^{KLT}$ represent realizations of $\hat{\pi}$ and $\hat{N}$ under level-k thinking. It is easy to check that when $k = 0$, equations 3.19 and 3.20 reduce to the adaptive expectation case outlined above. Moreover, when k goes to infinity, $\hat{\pi}_t = -\hat{\theta}_t$ and $\hat{N}_t = \phi_t\hat{\theta}_t$, which are the rational expectations solutions we described earlier. Thus, level-k thinking nests these two cases as two ends of the iterative spectrum.

3.2 Optimal policy under three different theories of expectations

We now consider the optimal policy problem facing the policy-maker. We assume that a policy-maker with a discount factor β_G faces a dynamic problem of choosing ϕ_t to minimize the discounted value of periodic losses that are quadratic in inflation and employment,

$$\min_{\phi_t}\sum_{t=0}^{\infty}\beta_G^t\mathbb{E}_{t-1}\left(\hat{\pi}_t{}^2 + \mu\hat{N}_t^2\right),$$

subject to equations 3.19 and 3.20. Importantly, the policy-maker chooses ϕ_t at date t before observing the productivity shock θ_t for that period. Implicit in this formulation is a commitment by the monetary authority to carry out the policy prescribed by ϕ_t even if it may not be optimal to follow through once θ_t is realized.

The policy problem can be simplified by defining $x_t \equiv \hat{\pi}_t + \hat{N}_t$. This allows the problem to be restated in terms of only one state variable. This simplified problem reads as follows, where σ_θ^2 is the variance of $\hat{\theta}$:

$$\min_{\phi_t} \sum_{t=0}^{\infty} \beta_G^t \mathbb{E}\left[(1+\mu\phi_t^2)\left((1-\phi_t)^{2k}x_{t-1}^2+\sigma_\theta^2\right)\right],$$

subject to

$$x_t = (1-\phi_t)^{k+1}x_{t-1} - (1-\phi_t)\hat{\theta}_t.$$

We can now look at the properties of ϕ_t under three cases: RE ($k=\infty$), adaptive expectations ($k=0$) and level-k thinking ($0<k<\infty$).

3.2.1 Optimal policy under rational expectations ($k=\infty$)

Our model coincides with rational expectations when $k \to \infty$. In this case, the objective function becomes

$$\min_{\phi_t} \sum_{t=0}^{\infty} \beta_G^t \mathbb{E}\left[(1+\mu\phi_t^2)\sigma_\theta^2\right],$$

with $x_t = -(1-\phi_t)\hat{\theta}_t$.

The solution in this case is to set $\phi_t = 0$, since the loss to the policy-maker is increasing in ϕ_t^2.[10] Importantly, the rational expectations solution gives $\hat{N}_t = 0$.

The main takeaway from this result is that when expectations are rational, it is optimal for the policy-maker to fully look through any deviations of inflation from target while setting monetary policy. Intuitively, under rational expectations, private agents fully understand that all inflationary shocks are temporary, so their inflation expectations are always anchored to the inflation target π^*. This is true even if the economy is hit by a long sequence of negative supply shocks. Hence, just knowing that the central bank is committed to keeping inflation at target is enough to keep expected inflation at target.[11]

[10] Given the points made in our earlier discussion of the fully rational case, we interpret this solution as one where the central bank selects a vanishingly small value for ϕ_t.

[11] In terms of interest rates, this rational expectation solution implies that the central bank would need to keep the nominal interest rate fixed in response to supply shocks ($\iota_{t+1} = \bar{\iota}$). This is implied by the Euler equation. However, if the central bank simply stated that it was fixing the nominal interest rate, the system would admit multiple solutions. On the other hand, if, in addition to communicating a fixed interest rate rule, the central bank also communicated that it was aiming to have $\hat{N}_t = -\phi\hat{\pi}_t$ with $\phi = 0$, then this would remove the indeterminacy and leave only the one solution where $\hat{\pi}_t = -\hat{\theta}_t$ and $\hat{N}_t = 0$. In effect, by stating the central bank's policy in terms of desired real outcomes, such a formulation eliminates many multiple equilibrium problems that can arise with other rules.

3.2.2 Optimal policy under adaptive expectations ($k = 0$)

The adaptive expectations case coincides with $k = 0$. In this case, the policy-maker's problem can be written as a dynamic optimization problem where the policy-maker's value function is

$$V(x_{t-1}) = \min_{\phi_t} \left\{ (1 + \mu\phi_t^2)(x_{t-1}^2 + \sigma_\theta^2) + \beta_G \mathbb{E} V((1 - \phi_t)(x_{t-1} - \hat{\theta}_t)) \right\},$$

with $x_t = (1 - \phi_t)(x_{t-1} - \hat{\theta}_t)$.

We solve this problem using a guess-and-verify approach. Specifically, we conjecture that $V(x) = a_1 x^2 + a_2 \sigma_\theta^2$ where a_1 and a_2 are constants. The solution to this problem takes the form

$$\phi_t^{AE} = \frac{\beta_G a_1}{\mu + \beta_G a_1}$$
$$a_2 = \frac{1 + \mu\phi_t^2 + \beta_G a_1 (1 - \phi_t)^2}{1 - \beta_G},$$

where the positive constant a_1 solves

$$a_1 = 1 + \frac{\mu \beta_G a_1}{\mu + \beta_G a_1}.$$

Three features of the solution under adaptive expectations are noteworthy.[12] First, the optimal ϕ_t^{AE} is constant over time. Second, ϕ_t^{AE} lies between 0 and 1 for any $\beta_G > 0$. This implies that monetary policy would be viewed as more hawkish if monetary authorities believe agents have adaptive expectations, relative to the fully rational case discussed above. Third, in the special case of a completely myopic policy-maker with $\beta_G = 0$, the optimal policy reduces to $\phi_t^{AE} = 0$, which implies that the myopic policy-maker would completely look through any inflationary shocks and would thus, as we have seen, cause inflation to become a random walk.

[12] To derive the nominal interest rate policy that implements this outcome, it is not sufficient to only state how people form expectations of next-period endogenous variables $\hat{\pi}_{t+1}$ and $\hat{N}_{t+1}$, as we have been doing until now. It is necessary to specify how households form expectations of the entire future path for ι, $\hat{\pi}$ and $\hat{N}$. Under rational expectations, this is straightforward. However, once one departs from rational expectations, different ways can be specified. To keep the focus of the paper, we do not pursue this implementation issue further here.

3.2.3 Optimal policy under level-k thinking ($0 < k < \infty$)

We now turn to the general case of level-k thinking with $0 < k < \infty$. The policy-maker's value function in this general case is

$$V(x_{t-1}) = \max \left\{ (1 + \mu\phi_t^2)\left((1-\phi_t)^{2k} x_{t-1}^2 + \sigma_\theta^2\right) + \beta_G \mathbb{E} V(x_t) \right\},$$

with the associated transition equation for x being

$$x_t = (1-\phi_t)^{k+1} x_{t-1} - (1-\phi_t)\hat{\theta}_t.$$

This problem in general does not admit closed-form solutions, nor easy qualitative characterizations. As a result, we proceed in two steps. First, we examine the problem when $\beta_G = 0$ and $k = 1$. This corresponds to a myopic policy-maker facing private agents who form expectations using just one iteration to update their initial guesses. Although this is a very special case, it allows us to derive important features of the optimal ϕ_t when viewed as a function of key parameters of the model. As a second step, we then show in our next section that key properties of the optimal solution when $(\beta_G, k) = (0, 1)$ carry over to cases where $\beta_G > 0$ and $k > 1$. This involves solving the full dynamic problem using numerical methods and allows us to show the robustness of the qualitative results derived for the special case where $\beta_G = 0$ and $k = 1$.

Special case: $\beta_G = 0, k = 1$

In this special case, the first-order condition for the optimal choice of ϕ_t is

$$\mu\phi_t \left[(1-\phi_t)^2 \tilde{x}_{t-1} + 1\right] = (1 + \mu\phi_t^2)(1-\phi_t)\tilde{x}_{t-1}, \tag{3.21}$$

where we have defined $\tilde{x}_{t-1} \equiv \frac{x_{t-1}^2}{\sigma_\theta^2}$. Clearly, the solution for ϕ_t will be a function of $\tilde{x}_{t-1}$. However, the first-order condition illustrates the richness that is introduced to the policy problem under level-k thinking. Specifically, even with $k = 1$, equation 3.21 represents a cubic function of ϕ_t. So, in general, there will not be a unique solution to the equation. Rather, the solution can take the form of a correspondence from the state variable $\tilde{x}$ to ϕ.

Let $\{\hat{\phi}_t^j\}$ denote the set of all permissible solutions to equation 3.21 and $L_t(\hat{\phi}_t^j)$ denote the policy-maker's loss when $\phi_t = \hat{\phi}_t^j$. All permissible solutions have to be feasible and satisfy both the first- and second-order conditions. The optimal solution $\hat{\phi}_t$ will solve

$$\min_{\hat{\phi}_t^j} \left\{ L_t(\hat{\phi}_t^j) \right\}.$$

In other words, the solution to the optimal policy problem would involve first eliminating solutions to equation 3.21 that do not satisfy the second-order condition. Then, from the remaining permissible solutions, the optimal ϕ_t would be the one that generates the global minimum.

To understand the challenge associated with solving this problem, it is helpful to consider the case where $\hat{\pi}_{t-1} \neq 0$ and $\hat{N}_{t-1} = 0$. In this case, the expected squared deviation of inflation from target is given by $(1-\phi)^2\hat{\pi}_{t-1}^2 + \sigma_\theta^2$, while the expected squared deviation of employment from target is given by $\phi^2(1-\phi)^2\hat{\pi}_{t-1}^2 + \phi^2\sigma_\theta^2$. This problem is non-convex, and that is why the first-order condition is not sufficient to describe the solution. The source of non-convexity that favours a potentially discontinuous response comes from the term $\phi^2(1-\phi)^2\hat{\pi}_{t-1}^2$; that is, it comes from the role of policy in affecting employment through managing expectations. This term is minimized at either $\phi = 0$ or $\phi = 1$, as opposed to having a unique minimizer. In contrast, $(1-\phi)^2\hat{\pi}_{t-1}^2 + \sigma_\theta^2$ is minimized with $\phi = 1$, while $\phi^2\sigma_\theta^2$ is minimized at $\phi = 0$. As we shall show below, the term $\phi^2(1-\phi)^2\hat{\pi}_{t-1}^2$ favours a jump from a low value of ϕ when $\hat{\pi}_{t-1}$ is small to a high value of ϕ when $\hat{\pi}_{t-1}$ becomes sufficiently large.

We can now turn to looking at the optimal policy in steps. In what follows it will be helpful to define $dd \equiv (1-\phi_t)(1+\mu\phi_t^2) - \mu\phi_t(1-\phi_t)^2$ and note that the inverse relationship from ϕ to $\tilde{x}$ implied by 3.21 is given by

$$\tilde{x}_{t-1} = \frac{\mu\phi_t}{(1-\phi_t)(1+\mu\phi_t^2) - \mu\phi_t(1-\phi_t)^2} = \frac{\mu\phi_t}{dd}. \tag{3.22}$$

To characterize the optimal solution $\hat{\phi}_t = \hat{\phi}(x_{t-1})$, it is useful to exploit equation 3.22, as it defines a function (as opposed to a correspondence).

The derivative of the inverse function 3.22 with respect to ϕ_t is given by

$$\frac{\partial \tilde{x}_{t-1}}{\partial \phi_t} = \frac{\mu - 3\mu^2\phi_t^2 + 4\mu^2\phi_t^3}{[(1-\phi_t)(1+\mu\phi_t^2) - \mu\phi_t(1-\phi_t)^2]^2} = \frac{nn}{dd^2}. \tag{3.23}$$

For convenience in notation, we will refer to the numerator of this derivative as nn, and the denominator as dd^2.

We characterize the optimal solution using a sequence of lemmas and propositions below. All proofs are provided in the Appendix.

Lemma 3.1. $0 \leq \hat{\phi}(\tilde{x}_{t-1}) \leq 1$ *for all* $\tilde{x}_{t-1} \geq 0$, *with* $\hat{\phi}(0) = 0$ *and* $\hat{\phi}(\infty) = 1$.

Lemma 3.1 says that the optimal solution is always bounded between 0 and 1. Moreover, the optimal ϕ goes to 0 when $\tilde{x}$ goes to zero, while it goes to 1 when $\tilde{x}$ goes to infinity.

Lemma 3.2. *For* $\mu < 8$, *dd is never zero on the interval [0,1[. For* $\mu \geq 8$, *there exists* ϕ_a *and* ϕ_b *between 0 and 1, such that*

1. $\phi_a = \frac{1}{4} - (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}; \phi_b = \frac{1}{4} + (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}$

2. $dd = 0$ *when* ϕ_t *equals* $\phi_a \in [0, 1[$ *or* $\phi_b \in [0, 1[$

3. $dd < 0$ *for* $\phi_t \in]\phi_a, \phi_b[$

Lemma 3.2 allows us to eliminate some candidate solutions. From equation 3.22 we know that $dd < 0$ implies $\tilde{x} < 0$. But this contradicts $\tilde{x} \geq 0$. It follows that any $\phi \in]\phi_a, \phi_b[$ cannot be optimal since $dd < 0$ in this range.

The requirement that $dd > 0$ has important implications for the relationship between any candidate optimal ϕ_t and the state variable $\tilde{x}_{t-1}$. Totally differentiating the first-order condition for ϕ_t and evaluating it around an optimum gives

$$\frac{d\phi_t}{d\tilde{x}_{t-1}} = \frac{dd}{SOC}, \tag{3.24}$$

where SOC is the derivative of the first-order condition with respect to ϕ_t. Since SOC needs to be positive, while $dd > 0$ from Lemma 3.2, it follows that any optimal function $\hat{\phi}(\tilde{x})$ must be increasing in $\tilde{x}$.

Lemma 3.3. *If* $\mu \leq 4$, *then* $nn \geq 0$ *for* ϕ *in* $[0, 1]$. *If* $\mu > 4$, *then there exists* ϕ_c *and* ϕ_d *between 0 and 1, such that* $nn \geq 0$ *for* $\phi \in [0, \phi_c]$, $nn < 0$ *for* $\phi \in]\phi_c, \phi_d[$, *and* $nn \geq 0$ *for* $\phi \in [\phi_d, 1]$. *Moreover, when* $\mu \geq 8$, $nn > 0$ *when* $\phi = \phi_a$ *and* $nn < 0$ *when* $\phi = \phi_b$.

Lemma 3.3 characterizes the behaviour of the numerator of equation 3.23 for the entire permissible range of ϕ. The key aspect to note is that the sign of the derivative $\frac{\partial \tilde{x}_t}{\partial \phi_{t-1}}$ in equation 3.23 depends on the sign of nn since $dd > 0$ for permissible values of ϕ. Lemma 3.3 shows that the sign of nn depends on μ. For $\mu > 4$, nn switches sign twice in the interior of the range $]0, 1[$ for ϕ, once at ϕ_c and again at ϕ_d.

The implication of the change of sign of nn is that $\frac{\partial \tilde{x}_t}{\partial \phi_{t-1}} > 0$ for all $\phi < \phi_c$ and $\phi > \phi_d$. Correspondingly, $\frac{\partial \tilde{x}_t}{\partial \phi_{t-1}} < 0$ for all $\phi \in]\phi_c, \phi_d[$. But $\frac{\partial \tilde{x}_t}{\partial \phi_{t-1}} < 0$ implies a negative relationship between ϕ_t and $\tilde{x}_{t-1}$, which violates the requirement from equation 3.24 that any optimal ϕ_t must co-move positively with $\tilde{x}_{t-1}$. Hence, when $\mu > 4$, any $\phi \in]\phi_c, \phi_d[$ cannot be optimal.

Lemmas 3.2 and 3.3 in conjunction with the requirement that all optimal solutions must reflect a positive relationship between ϕ_t and $\tilde{x}_{t-1}$ impose restrictions on the permissible values for ϕ. In particular, any candidate optimal solution for ϕ must be consistent with $dd > 0$ and $nn > 0$.

Lemma 3.4. *If $\mu \leq 4$, the first-order condition given by equation 3.21 implicitly defines a unique $\phi(\tilde{x})$ function that is increasing in $\tilde{x}$. $\hat{\phi}(\tilde{x})$ is therefore given by this implicitly defined function. If $\mu > 4$, then the first- and second-order conditions define two monotonically increasing and continuous functions $\phi_1(\tilde{x})$ and $\phi_2(\tilde{x})$ that represent local optima, where $\phi_1(\cdot)$ is defined over $\tilde{x} \in (0, z_1)$ and $\phi_2(\cdot)$ is defined over $\tilde{x} \in (z_2, \infty)$ with $0 < z_2 < z_1$ (z_1 can be ∞), and $\phi_2(z_2) > \phi_1(z_1)$.*

Intuitively, Lemma 3.4 uses the unique mapping from ϕ to $\tilde{x}$ in the inverse function 3.22 implied by Lemmas 3.2 and 3.3 to characterize the mapping from $\tilde{x}$ to all permissible values of ϕ. For $\mu < 4$, the inverse function 3.22 is monotone and continuous. Consequently, there is a unique solution for ϕ_t when $\mu < 4$, and this solution is monotone in $\tilde{x}_{t-1}$.

When $4 < \mu < 8$, the inverse function 3.22 is continuous but non-monotone. Since negative relationships between $\tilde{x}$ and ϕ are not permissible, the range of feasible ϕ is discontinuous. Consequently, the mapping from $\tilde{x}$ to the optimal ϕ defines two strictly increasing functions with an overlapping range of $\tilde{x}$ values, but with one function lying strictly above the other.

We show the case $4 < \mu < 8$ graphically in Figure 1. The left-hand panel of the figure uses the inverse function in equation 3.22 to express $\tilde{x}$ as a function of ϕ, while the right-hand panel shows the associated ϕ_1 and ϕ_2 functions that represent local optima.

The case $\mu > 8$ is similar to the case $4 < \mu < 8$ except that here, the inverse function is neither continuous nor monotone. The inverse function from ϕ to $\tilde{x}$ and the associated local optima functions are depicted graphically in Figure 2.

Lemma 3.4 directly leads to the following proposition:

Proposition 3.1. *The optimal policy is a function $\hat{\phi}\left(\tilde{x}_{t-1}\right)$ that is monotonically increasing in $\tilde{x}$, with $\phi(0) = 0$ and $\phi(\infty) = 1$.*

The fact that for $\mu > 4$ there are two strictly increasing functions $\phi_1(\tilde{x}_{t-1})$ and $\phi_2(\tilde{x}_{t-1})$ representing local optima leaves the determination of the global optimum still unresolved. Moreover, since the two functions have an overlapping range over which $\phi_2(\cdot)$ is consistently strictly greater than $\phi_1(\cdot)$, it is possible for the globally optimal choice on ϕ to jump discretely. We address this issue in the following lemma:

Figure 1: Case $4 < \mu < 8$

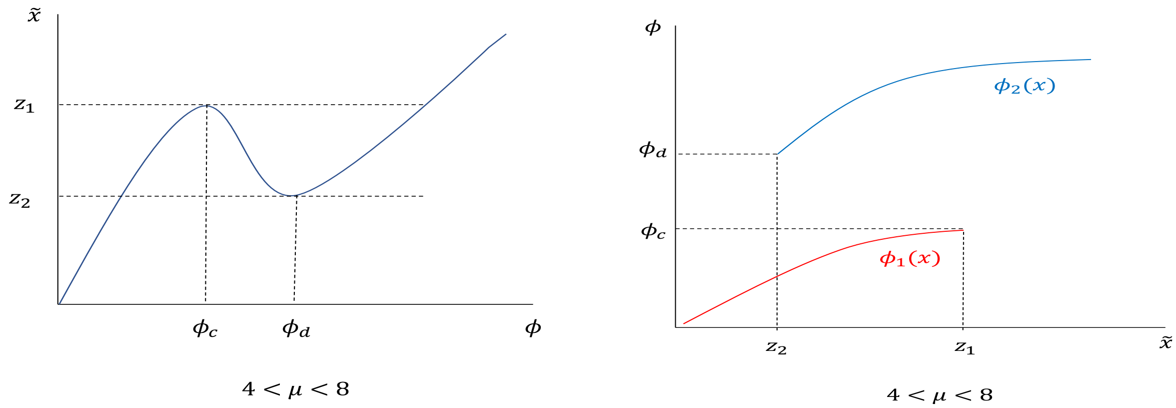

(a) (b)

Notes: Panel (a) of this figure depicts the inverse function in equation 3.22 when $4 < \mu < 8$. Panel (b) shows the associated ϕ_1 and ϕ_2 functions that represent local optima.

Figure 2: Case $\mu > 8$

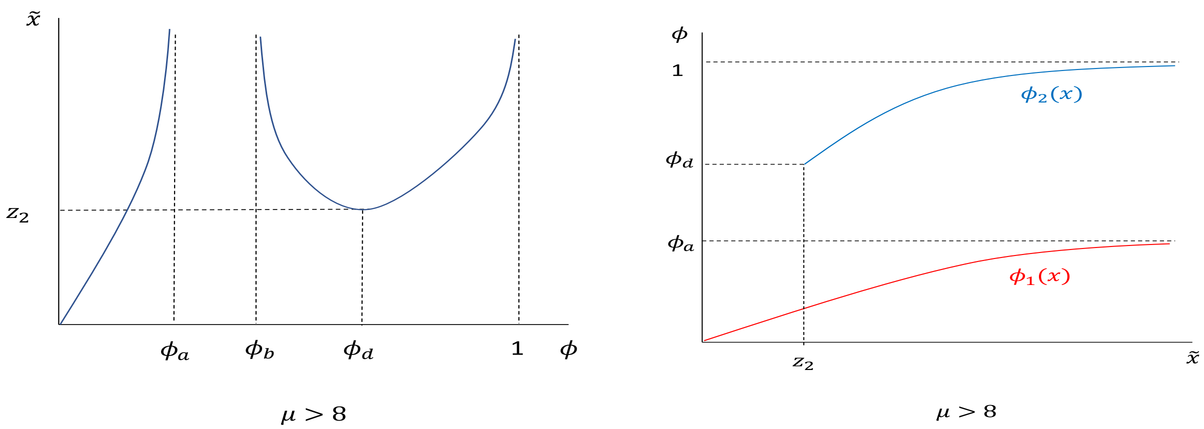

(a) (b)

Notes: Panel (a) of this figure depicts the inverse function in equation 3.22 when $\mu > 8$. Panel (b) shows the associated ϕ_1 and ϕ_2 functions that represent local optima.

Lemma 3.5. *Conditional on $\mu > 4$, there exists a $z_3 \in (z_2, z_1)$ such that the $\hat{\phi}(\tilde{x}_{t-1})$ function will correspond to $\phi_1(\tilde{x}_{t-1})$ for $\tilde{x}_{t-1} \in (0, z_3)$ and correspond to $\phi_2(\tilde{x}_{t-1})$ for $x \in (z_3, \infty)$.*

The key result in Lemma 3.5 is that the optimal function jumps exactly once from $\phi_1(\tilde{x}_{t-1})$ to $\phi_2(\tilde{x}_{t-1})$ at point z_3. One can use this to state our second key result:

Proposition 3.2. *If μ is sufficiently big, there exists a unique cutoff for $\tilde{x}_{t-1}$, such that at this cutoff, $\hat{\phi}\left(\tilde{x}_{t-1}\right)$ jumps up discontinuously.*

Proposition 3.2 shows that when the central bank cares enough about employment, the framework generates a discontinuous shift in the aggressiveness with which the optimal policy responds to inflation shocks. Specifically, for a range of realizations of the state variable $\tilde{x}_{t-1}$, the optimal policy is to look through shocks in the sense of keeping ϕ_t relatively low. However, once the shocks cumulate up to a threshold level, the optimal policy pivots to an aggressive response encapsulated in a discontinuously higher ϕ_t.

Figure 3 shows the optimal $\hat{\phi}$ representing the global optimum for the two cases $4 < \mu < 8$ and $\mu > 8$ where $\hat{\phi}$ jumps.

Figure 3: Policy pivot – jumps in $\hat{\phi}$

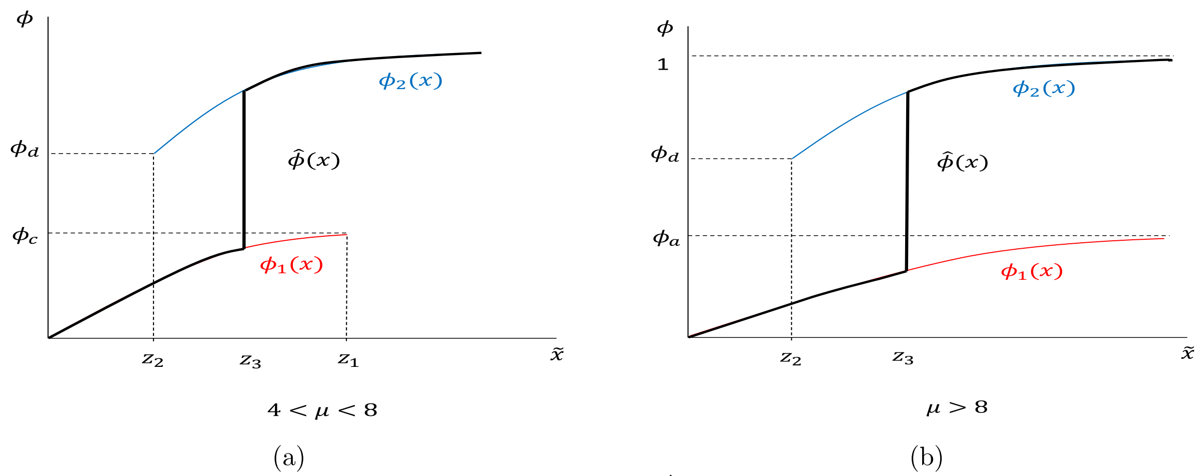

Notes: Panel (a) of this figure depicts the global optimum $\hat{\phi}$ when $4 < \mu < 8$. Panel (b) shows the global optimum $\hat{\phi}$ when $\mu > 8$.

To gain some intuition for Proposition 3.2, recall that $x_{t-1} = \hat{\pi}_{t-1} + \hat{N}_{t-1} = (1-\phi_t)\hat{\pi}_{t-1}$. Hence, movements in x are largely driven by inflation outcomes. The proposition essentially reflects the importance of managing inflation expectations. When inflation is close to target, it is optimal for monetary policy not to respond aggressively to inflation shocks, as the employment cost of doing so is relatively high while inflation expectations and inflationary pressures remain relatively modest. However, once inflation goes above a threshold level, its persistent impact on inflation expectations makes the cost of *not* fighting inflation too high. At this point, the optimal policy pivots to a much more aggressive response to inflation shocks.

The pivot in policy arises due to a non-convexity in the model. To see this more clearly, recall that the policymaker's loss function in the myopic case is

$$L_t = (1 + \mu\phi_t^2)\left((1 - \phi_t)^{2k} x_{t-1}^2 + \sigma_\theta^2\right).$$

Differentiating the loss function with respect to ϕ_t and rearranging terms gives

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

Define $A \equiv \frac{\phi_t}{2L_t}\frac{dL_t}{d\phi_t}$; $B \equiv \frac{\mu\phi_t^2}{1+\mu\phi_t^2}$; $\eta \equiv \frac{k\phi_t}{1-\phi_t}$; and $D \equiv \frac{(1-\phi_t)^k x_{t-1}}{(1-\phi_t)^k x_{t-1}+\sigma_\theta^2}$. Using these definitions we can rewrite the last equation above as $A = B - \eta D$. At an optimum we must have $A = 0$.

It is easy to check that B and η are both rising in ϕ while D is decreasing in ϕ. Moreover, since $\mathbb{E}_{t-1}\hat{\pi}_t = (1-\phi_t)^k(\hat{\pi}_{t-1} + \hat{N}_{t-1})$ under level-k thinking (see equation 3.19 above), the term $\eta = \frac{k\phi_t}{1-\phi_t}$ is proportional to the elasticity of $\mathbb{E}_{t-1}\hat{\pi}_t$ with respect to ϕ_t. Hence, the policy elasticity of inflation expectations is rising in ϕ

To better understand the results, first note that the employment deviation $\hat{N}$ impacts the loss function L both directly through the term $\mu\phi^2$ as well as indirectly through its effect on the expected inflation deviation $\mathbb{E}_{t-1}\hat{\pi}$. The term B captures the direct effect of ϕ on L through its effect on employment while the term ηD captures the indirect effect going through the effect of ϕ on the expected inflation deviation.

The last term ηD is the source of the non-convexity that cause the pivot. Since η is increasing in ϕ, $\mathbb{E}_{t-1}\hat{\pi}_t$ declines proportionately more in response to a policy tightening when ϕ is high. The fall in $\mathbb{E}_{t-1}\hat{\pi}_t$ induces a gain in the loss function for the policymaker as it reduces the employment contraction associated with the policy tightening. The magnitude of this gain depends on the magnitude of D. Since D depends on $\hat{\pi}_{t-1}$, this implies that there is a much greater payoff from tightening policy if the tightening is sharper and delayed till the inflation deviation is high. This latter effect causes a switch in the global optimum from a lower to a higher ϕ only after a threshold level of $\hat{\pi}$ is reached.

Related to the above, it is important to note that even though the pivot in policy stance may be drastic, the actual effect on employment is unclear as the adjustment in inflation expectations induced by the change in policy stance directly reduces inflation pressures and therefore reduces the amount of actual tightening needed. This is reflected in the fact that the effect of the pivot on expected employment is driven by the product $-\phi(1-\phi)$. As shown in Proposition 3.3, a property of the optimal policy is for $-\phi(1-\phi)$ to remain constant at the pivot point. In other words, the pivot is associated with a change in policy stance that helps reduce inflation with no expected effect on employment. This can be referred to as a "soft landing." The pivot is in fact about talking tough in order to modify expectations and thereby limit the actual amount of contraction in employment needed to reduce inflation.

Now, if the optimally timed pivot is expected to allow for a soft landing, why wait to implement it? The answer relates to the effect of the pivot on the variance of employment. As also indicated in Proposition 3.3, when the policy stance pivots from less aggressive to more aggressive, the variance of employment increases discontinuously, as the variance of employment is proportional to ϕ^2, which is increasing at the point of pivot. The cost of taking the aggressive policy stance relates to the commitment to respond strongly to inflation, whether it be driven by higher expectations or shocks. The change in stance implies that policy will stop looking through shocks and instead will react to them strongly. This causes the variance of employment to increase. Hence, the change in policy stance – when optimally timed – is effectively aimed at creating a "risky soft landing," or a "narrow soft landing," in the sense of reducing inflation at the cost of more uncertain outcomes in terms of employment.

Proposition 3.3. *If ϕ and ϕ' ($\phi' > \phi$) represent the two stances of monetary policy at the point of discontinuity of optimal policy (at the pivot point), then $\phi(1-\phi) = \phi'(1-\phi')$. This implies that at the optimal point of pivot, inflation is expected to fall but employment is expected to stay constant. However, the variance of employment increases discontinuously at the pivot point.*

4 Dynamic Model: Numerical Simulations

In the previous section, we provided conditions under which the model exhibits discontinuous jumps or pivots in the optimal policy response to inflation shocks under level-k thinking. However, we were able to derive these results only in the special case of $\beta_G = 0$ and $k = 1$.

How robust is the policy pivot result to extending the model to a non-myopic policy-maker with $\beta_G > 0$ and/or to private agents who iterate more than once when forming expectations, i.e., when $k \in]1, \infty[$? This section explores this question using numerical simulations of the model. Before doing so, however, we provide a generalization of the Phillips curve relationship in the model to allow its slope to be an arbitrary number that is not necessarily unity as in the baseline model. This generalization is useful for both the interpretation of the numerical results as well as illustrating its generality.

4.1 Generalization to case with an arbitrary slope of the Phillips curve

Recall that our assumed preference structure induced a Phillips curve given by

$$\hat{\pi}_t = \mathbb{E}_{t-1}\hat{\pi}_t + \mathbb{E}_{t-1}\hat{N}_t - \hat{\theta}_t,$$

where there is a coefficient of exactly unity on $\mathbb{E}_{t-1}\hat{N}_t$. This implies a steep Phillips curve, which can be seen as quite restrictive. In Section A.2 of the Appendix, we show that it is feasible to change preferences to generate a more general Phillips curve, for example by using GHH preferences. Specifically, we show that with a lifetime utility function of the GHH form

$$\mathbb{E}\sum_{t=0}^{\infty}\beta^t \ln\left(C_{it} - \eta\theta_t N_{it}^{1+\lambda}\right),$$

the associated Phillips curve becomes

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

where $\lambda \in]0, \infty[$.

It turns out that this extension of the model is in fact already embedded in the baseline case as long as one reinterprets μ and ϕ_t appropriately for the case where $\lambda \neq 1$. To show this, first note that when $\lambda \neq 1$, the system describing equilibrium outcomes becomes

$$\begin{aligned}
\hat{\pi}_t &= (1-\lambda\phi_t)^k\left[\hat{\pi}_{t-1} + \lambda\hat{N}_{t-1}\right] - \hat{\theta}_t \\
\hat{N}_t &= -\phi_t\left[(1-\lambda\phi_t)^k\left\{\hat{\pi}_{t-1} + \lambda\hat{N}_{t-1}\right\} - \hat{\theta}_t\right].
\end{aligned}$$

Defining the new state variable $x_t = \hat{\pi}_t + \lambda\hat{N}_t$ then gives the policy-maker's problem in the general case as

$$\min_{\phi_t}\sum_{t=0}^{\infty}\beta_G^t\mathbb{E}\left[(1+\mu\phi_t^2)\left((1-\lambda\phi_t)^{2k}x_{t-1}^2 + \sigma_\theta^2\right)\right],$$

subject to

$$x_t = (1-\lambda\phi_t)^{k+1}x_{t-1} - (1-\lambda\phi_t)\hat{\theta}_t.$$

Defining $\tilde{\mu} \equiv \frac{\mu}{\lambda^2}$ and $\tilde{\phi}_t \equiv \lambda \phi_t$ allows us to rewrite this problem as

$$\min_{\tilde{\phi}_t} \sum_{t=0}^{\infty} \beta_G^t \mathbb{E} \left[(1 + \tilde{\mu} \tilde{\phi}_t^{\,2}) \left((1 - \tilde{\phi}_t)^{2k} x_{t-1}^2 + \sigma_\theta^2 \right) \right],$$

subject to

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

As can be seen, this is the exact same problem as before but with $\tilde{\phi}$ and $\tilde{\mu}$ replacing ϕ and μ. This makes solving the general case of $\lambda \neq 1$ equivalent to solving the original problem with a modified weight on employment in the central bank's loss function, though one must be mindful to infer ϕ_t via $\tilde{\phi}_t = \lambda \phi_t$.

This generalization is important for a few reasons. First, it shows that the qualitative properties of the solution we derived above are not dependent on having $\lambda = 1$. Second, since the literature generally favours parameter values for λ much less than one, it implies a significant widening in the range of values for the "true weight" μ for which the pivoting behaviour described in Proposition 3.2 emerges in the special case $(\beta_G, k) = (0, 1)$. Third, it also associates pivots with wider swings in the monetary stance parameter ϕ_t since $\phi_t = \tilde{\phi}_t / \lambda$.

With these observations in mind, we now use the remainder of this section to explore the model's properties and implied potential for pivots outside the special case $(\beta_G, k) = (0, 1)$.

4.2 Baseline parameterization and solution method

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 identify model periods with years. We therefore set policy-makers' discount factor to $\beta_G = 1/1.005$, aligning with the midpoint of Bank of Canada staff's current assessed range of 0 to 1% for the real neutral rate in Canada (Faucher et al., 2022).

We also 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 employment deviations and deviations of inflation from target. In addition, we set the Phillips curve slope to $\lambda = 0.092$, a value broadly in line with Djeutem et al. (2022) and Wagner et al. (2022), and assume $k = 2$ levels of thinking in the private sector, placing us roughly in the middle of the range of $k = 1$ to 4 considered in Farhi and Werning (2019). The only remaining parameter is then the productivity-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. Briefly, this involves the following five steps:

1. Fix some grid X_G of values for the state variable x_{t-1};

2. Make some guess on the value that the central bank's value function takes at each $x_{t-1} \in X_G$;

3. Use a fine grid search to find the central bank's implied optimal choice on $\tilde{\phi}_t$ at each $x_{t-1} \in X_G$, evaluating any relevant expectations using a combination of Gauss-Hermite quadrature and linear interpolation – i.e.,

$$\tilde{\phi}_t^*(x_{t-1}) \in \arg\min_{\tilde{\phi}_t \in \tilde{\Phi}} \left\{ \begin{array}{l} \left[1 + \left(\mu/\lambda^2\right)\tilde{\phi}_t^2\right]\left[\left(1-\tilde{\phi}_t\right)^{2k} x_{t-1}^2 + \sigma_\theta^2\right] \\ \\ +\beta_G \sum_{n=1}^N \omega_n V\left[\left(1-\tilde{\phi}_t\right)^{k+1} x_{t-1} - \left(1-\tilde{\phi}_t\right)\hat{\theta}_n\right] \end{array} \right\},$$

   where $\tilde{\Phi}$ is a fine grid; $\{\omega_n\}_{n=1}^N$ and $\{\hat{\theta}_n\}_{n=1}^N$ respectively denote the weights and nodes of an N-point quadrature rule; and $V(\cdot)$ denotes a linear interpolant constructed from our current guesses on the central bank's value function;

4. For each $x_{t-1} \in X_G$, use $\tilde{\phi}_t^*(x_{t-1})$ from our previous step to update our guess on the value that the central bank's value function takes at the grid point in question;

5. Repeat steps three and four until successive guesses fall within some small tolerance of one another.

Details on this algorithm and the accuracy of the solution it delivers are available on request.

4.3 Results

Results for the baseline parameterization. Motivated by the analytical results presented earlier, our main numeric results concern the relationship between (i) the normalized monetary stance parameter $\tilde{\phi}_t \equiv \lambda\phi_t$, for which higher values indicate a lower willingness among policy-makers to look through periods of high inflation; and (ii) the state variable $x_{t-1} \equiv \hat{\pi}_{t-1} + \lambda\hat{N}_{t-1}$, which we can interpret as a measure of the degree of overheating that occurred in the previous period.

The top panel of Figure 4 plots the relationship between these two variables, holding all parameters at the baseline values described above. Results indicate that pivoting remains a key feature of the central bank's behaviour, with realizations of x_{t-1} above roughly 4% in absolute value triggering a sudden increase in $\tilde{\phi}_t$. This pivoting behaviour is further illustrated in the lower panel of the figure, which focuses on the relationship between x_{t-1} and the central bank's inflation expectations at the time it sets $\tilde{\phi}_t$ – i.e., $\mathbb{E}^{RE}_{t-1}\hat{\pi}_t$. As the panel clearly shows, pivoting is associated with a sharp re-normalization of these expectations.

In Figure 5, we zoom in on the model's behaviour around its positive pivot point, focusing on inflation deviations (left-hand panel) and deviations of employment from its natural level (right-hand panel). In each panel, a solid line reports expected outcomes before observing the productivity shock $\hat{\theta}_t$, while the surrounding bands report realized outcomes for a range of potential values of $\hat{\theta}_t$. Results indicate that many of the "risky soft landing" properties emphasized in Section 3 continue to hold qualitatively outside the special case $(\beta_G, k) = (0, 1)$. For example, the change in expected employment outcomes around the pivot point is relatively modest. As we explained in Section 3, this reflects offsetting forces associated with the fact that shifting to a tighter policy stance induces changes in inflation expectations that reduce the degree of actual economic slack needed to keep inflation close to target.

However, since a tighter stance limits policy-makers' ability to stabilize employment in response to shocks, pivoting also leads to an increase in the level of risk surrounding the employment outlook. This pattern is illustrated by the sharp widening of the band in the right-hand panel of Figure 5. Taken all together, the results in the figure thus suggest that pivoting in the numeric model is compatible with a soft landing *in expectation* but also entails a significant risk of harder landings, much as was the case when $(\beta_G, k) = (0, 1)$.

It should nevertheless be noted that even if the pivot does not cause much expected change in employment, prior to the pivot, employment is depressed in our baseline parameterization. This can be seen in Figure 5, where expected employment is more that 2% below target before and after the pivot. This reflects the fact that when $\beta_G > 0$ and $k > 1$, even during the "looking through" phase, policy can become sufficiently restrictive to cause employment to be materially below target. However, this pre-pivot policy stance is not sufficient to bring inflation back to target, and this is what eventually requires a pivot. This is why we think it is appropriate to refer to the pre-pivot monetary stance as a "looking-through" stance.

Impact of changing the central bank's discount factor. With the points above in mind, we now use the remainder of this section to explore the roles that key parameters play in driving and shaping the central bank's pivoting behaviour. Figure 6 focuses on the role played by policy-makers' discount factor β_G, comparing our baseline parameterization against an otherwise comparable "myopic benchmark" under which we set $\beta_G = 0$. For both these parameterizations, the left-hand panel in the figure plots the relationship between $\tilde{\phi}_t$ and the absolute value $|x_{t-1}|$, making use of the fact that the problem facing policy-makers is symmetric across positive and negative values of x_{t-1}. The right-hand panel then repeats, focusing on the relationship between $|x_{t-1}|$ and the absolute expected inflation deviation $|\mathbb{E}^{RE}_{t-1}\hat{\pi}_t|$.

The results in Figure 6 clearly associate the myopic benchmark with a situation where the central bank is willing to look through periods of off-target inflation for a much wider range of values for x_{t-1}. Put differently, shifting from myopia to a forward-looking policy approach "pulls forward" the threshold around which the central bank is prepared to pivot in response to overheating in the previous period. This 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 by helping to keep future inflation expectations close to target.

Impact of changing the sophistication of private-sector expectations. In Figure 7, we turn our attention to the parameter k, which gives the levels of thinking in which the private sector engages when forming expectations. The figure reports policy functions for values of this parameter in the range of $k = 1$ to 3, holding all other parameters constant at their baseline values. Results indicate that pivots occur for all choices on k in the aforementioned range. However, higher k values are associated with smoother and generally lower profiles for $\tilde{\phi}_t$, with smaller jumps around pivot points. These patterns are consistent with our earlier observations that expectations become rational as $k \to \infty$, and in this case the central bank's optimal policy is to fully look through periods of off-target inflation in the sense of keeping $\tilde{\phi}_t$ constant at zero. To further emphasize these points, the figure also reports results for a case under which expectations are relatively close to rational in the sense that k takes a very high value of 40, and in this case no pivoting occurs even when x_{t-1} rises as high as 10%.

A useful way to understand the patterns in Figure 7 is to recognize that higher values for k make inflation expectations more sensitive to a given change in the central bank's policy stance. As

Figure 4: Pivoting under the baseline parameterization

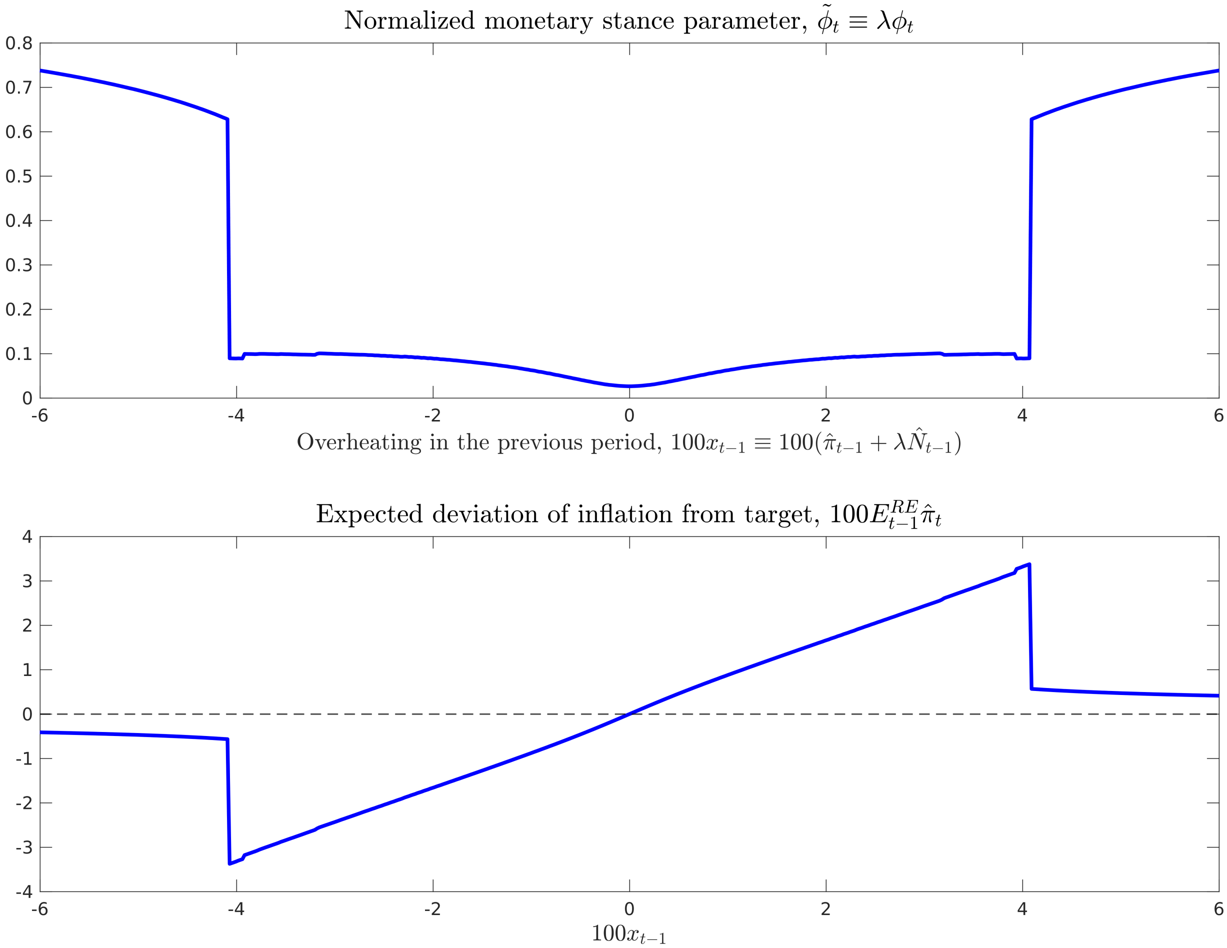

Figure 5: Risky soft landing under the baseline parameterization

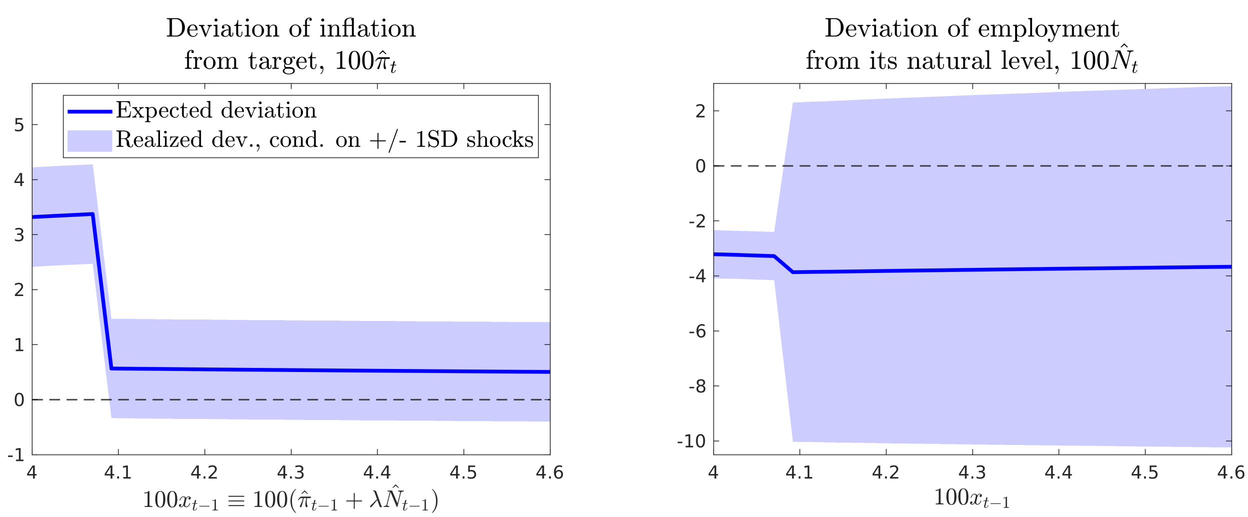

Figure 6: Impact of changing policy-makers' discount factor, β_G

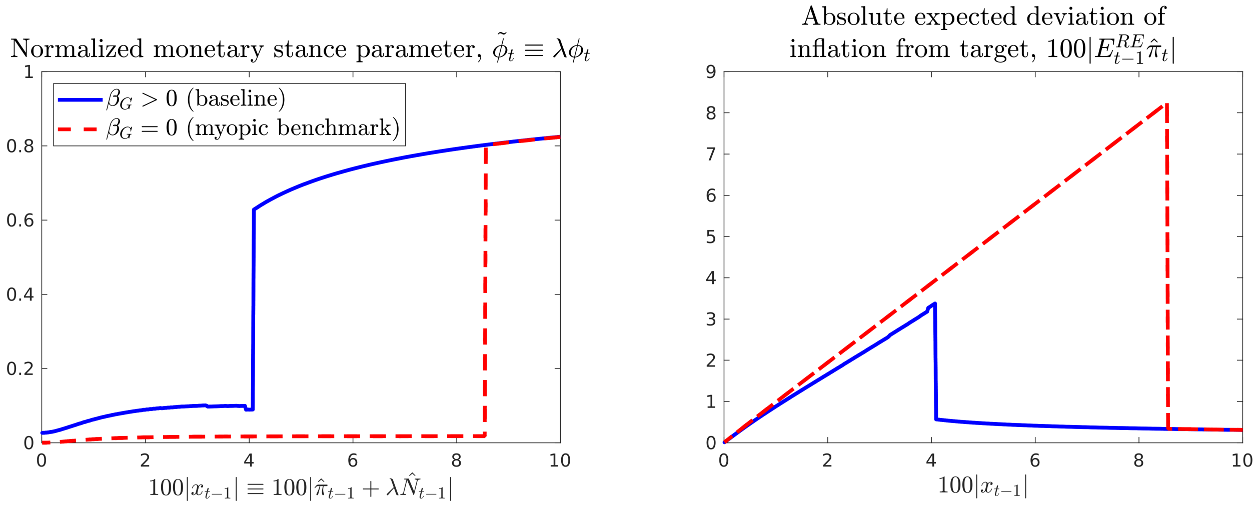

a result, central banks facing higher values for k can make do by increasing the stance parameter $\tilde{\phi}_t$ by smaller margins around pivot points. This is important because smaller jumps in $\tilde{\phi}_t$ leave policy-makers with more room to stabilize employment in response to shocks, suggesting that the "risky soft landing" problem discussed earlier should be less of an issue when k is high.

To confirm this intuition, Figure 8 zooms in on the economy's behaviour around its positive pivot point, assuming the same $k = 1$ through $k = 3$ parameterizations described above. Both panels of the figure were constructed along lines similar to Figure 5, with the exception that the quantity on the horizontal axis has now been normalized so that pivoting occurs at a value of zero for all three of the parameterizations in question. This allows us to compare parameterizations in terms of the widening in the range of potential employment outcomes that occurs around pivot points. The fact that this widening is much less pronounced when k is high thus confirms our intuition that higher values for k should help policy-makers mitigate the risk of a hard landing.

We interpret this finding as one with strong implications for central banks' communication strategies. Though k is a fixed parameter in the model, the cognitive frictions that it aims to capture are likely to vary over time, rising in novel economic environments where the private sector finds it more challenging to think through the full implications of different policy stances. To the extent that clear, effective communication from policy-makers can aid and inform this reasoning process, it should be associated with higher effective values for k, leading in turn to a lower risk of hard landings.

Impact of changing policy-makers' effective weight on employment deviations. In Fig-

Figure 7: Impact of changing the levels of private-sector thinking, k

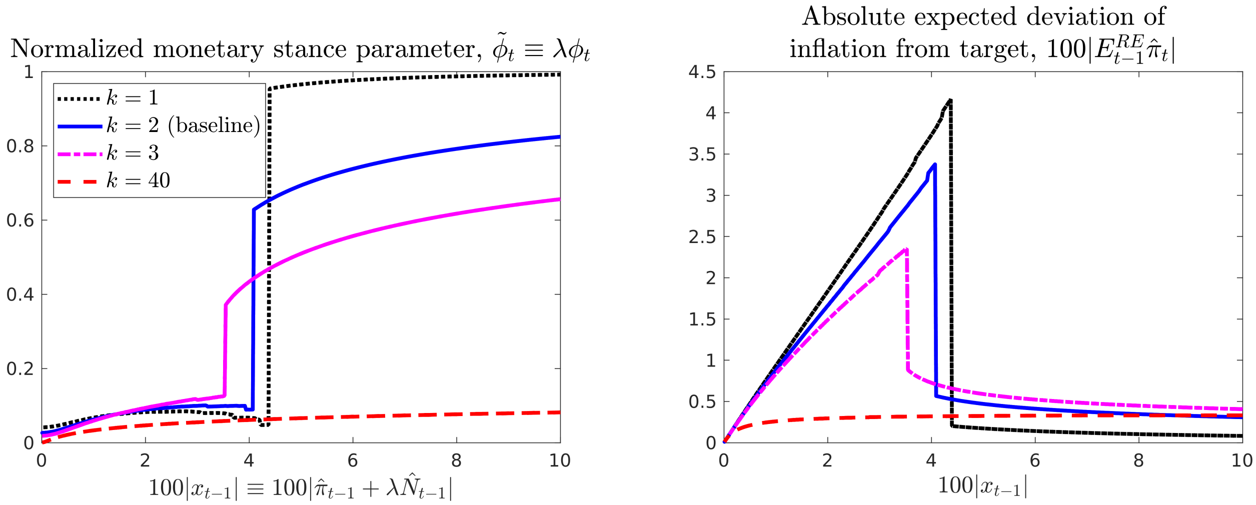

Figure 8: Risky soft landing under different levels of private-sector thinking (k)

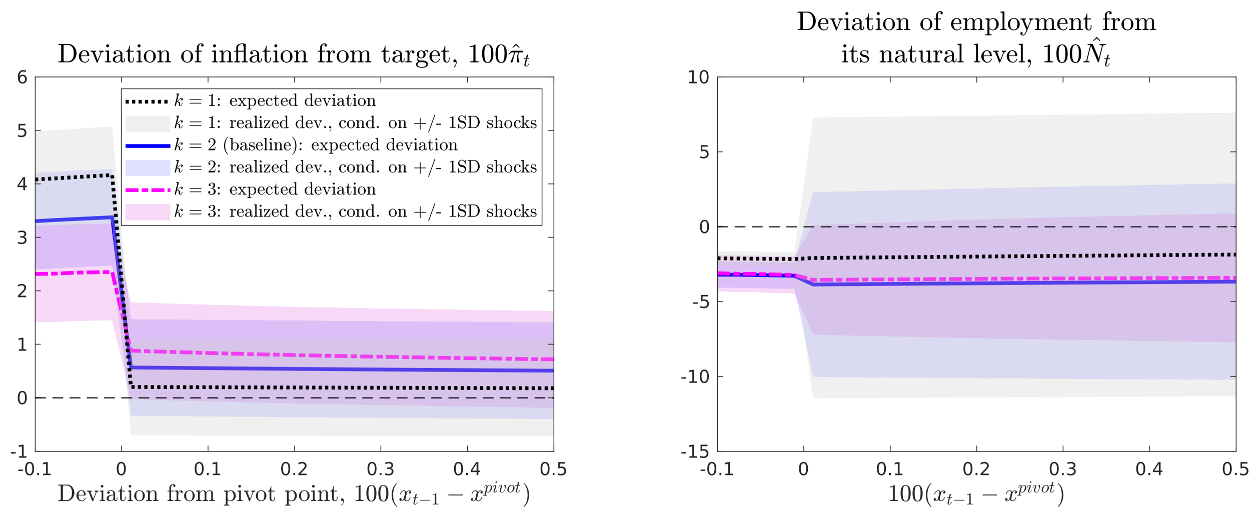

ure 9, we finally turn our attention to $\tilde{\mu} \equiv \mu/\lambda^2$, the effective weight that policy-makers place on minimizing deviations of employment from its natural level. The figure reports policy functions for choices on this weight spanning three orders of magnitude, holding all other parameters in line with our baseline calibration. Consistent with some of the insights emerging from our earlier analytical work, the policy functions that obtain when $\tilde{\mu}$ is very small are smooth, 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 off-target inflation among central banks placing greater weight on employment outcomes.

A related and important feature of Figure 9 is that when x_{t-1} reaches levels high enough to trigger a pivot for a central bank with a very large $\tilde{\mu}$, the margins by which both $\tilde{\phi}_t$ and $\mathbb{E}^{RE}_{t-1}\hat{\pi}_t$ adjust are larger than is the case for smaller values of $\tilde{\mu}$ – i.e., the pivot may come late, but it's also more dramatic. Mindful of this pattern, coupled with our previous finding that higher values for $\tilde{\mu}$ can equivalently be associated with higher values for the true weight μ or lower values for the Phillips curve slope λ, we note that the recent pivots observed in advanced economies occurred specifically at a time when (i) some central banks had recently announced framework changes placing greater emphasis on employment outcomes; and when (ii) a wide range of studies suggested that Phillips curves in advanced economies had flattened significantly. To the extent that both factors contributed to a higher value for $\tilde{\mu}$, they would have set the scene for a later but more pronounced pivot when viewed through the lens of this modelling framework.

Figure 9: Impact of changing policy-makers' effective weight on employment deviations, $\tilde{\mu} \equiv \mu/\lambda^2$

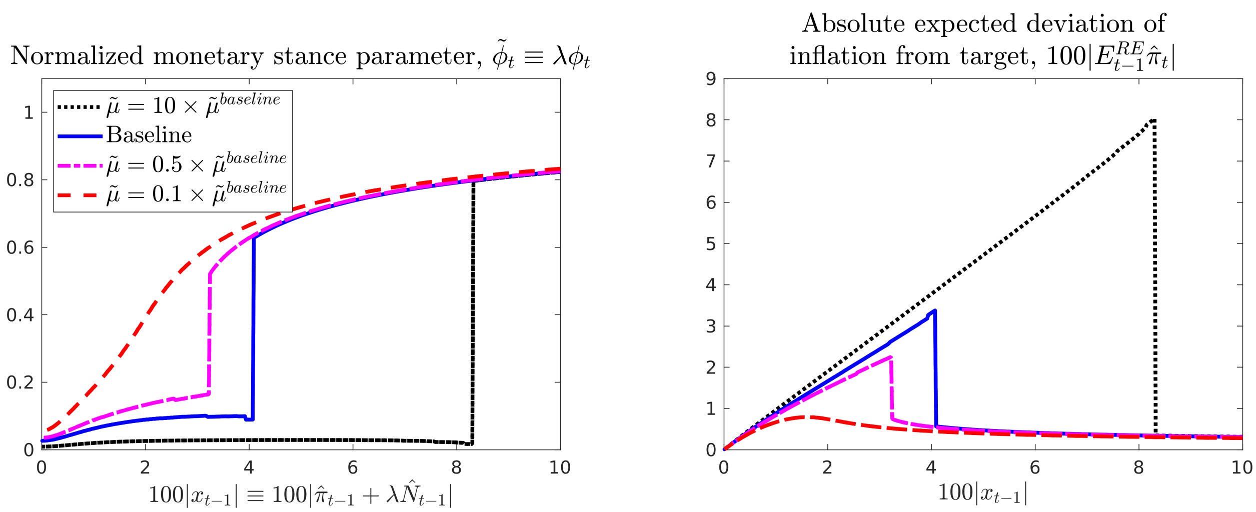

5 What about demand shocks?

Up to now, we have ignored demand shocks, despite the fact that they certainly played a role in the recent inflation episode. It is simple to extend our framework to incorporate demand shocks. A common way to think about demand shocks is to view them as shocks to the household's discount rate. The addition of this form of demand shocks simply adds a stochastic element in the Euler equation. In this case, if we denote the demand shock by d_t (which can follow an arbitrary process), the forward expansion of the Euler equation becomes:

$$\ln N_t - \ln \bar{N} = -\sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)] + \sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1} d_{t+h-1}. \tag{5.25}$$

In the absence of supply shocks, optimal policy would prescribe generating an expected path for real interest rates such that

$$\sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)] = \sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1} d_{t+h-1}.$$

If effectively implemented, this policy stance would stabilize both employment and inflation and therefore would be optimal. In this sense, our framework exhibits a common property whereby optimal policy should aim to fully offset demand shocks.

Accordingly, in the presence of demand shocks, the effective measure of monetary stimulus can be thought of as the expected real interest rate path after netting out the part that is needed to offset demand shocks. This would be given by

$$\sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)] - \sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1} d_{t+h-1}.$$

Thus, in the model with demand shocks, we can still view the policy dilemma facing central banks as an issue of determining how much stimulus to inject as a function of inflation, though stimulus must be properly redefined. In particular, the only required change involves specifying policy as

$$\phi_t(\hat{\pi}_t) = \sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)] - \sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1} d_{t+h-1}.$$

As can be easily seen, this reformulation does not change our optimal policy problem, as we still have $\ln N_t - \ln \bar{N} = -\phi_t(\pi_t - \pi^*)$; it changes only how one would need to set interest rates to implement the stimuli prescribed by $\phi_t(\hat{\pi}_t)$. In particular, the required nominal interest rates would now depend not only on inflation but also on the path of d_t.

Optimal monetary policy in response to both demand and supply shocks can therefore be thought of as a two-stage process, first adjusting monetary policy to fully offset demand shocks, and then deciding how long to look through supply shocks and when to pivot in favour of controlling inflation expectations. However, if a central bank mistakenly interpreted some demand shocks as supply shocks, or misjudged the strength of demand, this could cause it to adopt an overly stimulative stance because it would be inappropriately looking through these shocks. This type of error narrative behind the recent inflation episode could potentially be incorporated into our analysis. However, this would require expanding the setting to introduce a signal extraction problem, which we leave for future work.

6 Conclusion

The appropriate response to supply-driven inflation shocks is a question that has captured the attention of many policy-makers and analysts around the world over the past several months. Most monetary authorities initially resisted responding to the unfolding rise in inflation before pivoting sharply into an accelerated tightening of monetary policy. The initial approach of looking through the rising inflation numbers was justified by pointing to their predominantly supply-side origins, while the subsequent pivot was explained as an attempt at anchoring inflation expectations. In this paper, we have provided a framework to study supply-driven inflation. In doing so, our goals are to both explain central bank behaviour as well as provide a rubric for studying the design of optimal policy.

An important element in our modelling of the inflation process is allowing prices to be more flexible than wages. This feature is key to understanding why looking through supply shocks may at times be desirable. It also gives rise to a Phillips curve that departs slightly from a canonical New Keynesian Phillips curve in that the residual term becomes a supply shock instead of a mark-up shock. In particular, the supply shock reflects the forecast errors made by wage-setters when setting wages before seeing the prices set by firms.

Within our framework we have shown two main results. First, the type of pivoting behaviour recently exhibited by many central banks should not arise if policy-makers view agents as either fully rational or having simple adaptive inflation expectations. For example, under rational expectations, the role of policy in anchoring expectations becomes easy, because private agents are able to see through the supply-side origins of inflation, obviating the need to tighten in response to supply shocks. In contrast, with adaptive expectations, central banks should tighten policy in response to supply shocks, but the strength of this response should not change over time.

Our second main result is that when agents are not fully rational but instead form expectations under bounded rationality in the form of **level-k thinking**, optimal policy can involve an initial period of looking through inflation shocks, followed by a discontinuous jump into aggressive tightening if shocks accumulate above a threshold. This prediction matches recently observed policy behaviour in many economies.

A competing explanation for the abrupt policy pivots performed by many central banks is that there was a belated recognition that some of the observed increase in inflation was due to demand pressures rather than temporary supply shocks. While we do not include demand shocks in our baseline model, we have shown that the optimal response to demand shocks in the model is to tighten monetary policy, which is the standard prediction of most versions of the New Keynesian model. More crucially, we also showed that the optimal monetary policy response to supply shocks in our model can be easily reinterpreted as the policy stance net of the response to demand shocks. Consequently, the existence of supply shocks would induce a policy pivot even after accounting for a belated recognition of demand shocks. In this sense, our mechanism is complementary to the "policy error" mechanism.

We believe our results are important not only in terms of explaining the recent experience in industrial economies but also in terms of characterizing long-term policy challenges in emerging economies. In these economies, policy-makers often have to manage inflation in the context of volatile food and fuel prices driven by global factors and consumer expenditure baskets with very high shares of food and fuel expenditures. Our results suggest that it is optimal for policy-makers in such economies to look through moderate inflation shocks, but then respond aggressively if the shocks become very large. Indeed, not pivoting in such cases could induce long periods of high inflation driven solely by persistent deviations of inflation expectations.

References

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

Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001, September). Cognition and behavior in normal-form games: An experimental study. *Econometrica 69*(5), 1193–1235.

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, National Bureau of Economic Research.

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.

Faucher, G., C. Hajzler, D. Matveev, Y. Park, and T. Taskin (2022, April). Potential output and the neutral rate in Canada: 2022 assessment. Working Paper 2022-3, Bank of Canada.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications, Second edition*. Number 10495 in Economics Books. 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.

Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money*. Harcourt Brace.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review 85*(5), 1313–26.

Reis, R. (2021, October). Losing the inflation anchor. Discussion Paper 16664, Centre for Economic Policy Research.

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

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.

A Appendix

A.1 Proofs

Lemma 3.1: $0 \leq \hat{\phi}(\tilde{x}_{t-1}) \leq 1$ for all $\tilde{x}_{t-1} \geq 0$, with $\hat{\phi}(0) = 0$ and $\hat{\phi}(\infty) = 1$.

Proof: First part: For any candidate $\phi > 1$, $\phi = 1$ dominates, since $(1 + \mu) < (1 + \mu\phi^2)(1 - \phi)^2\tilde{x} + (1 + \mu\phi^2)$ for all $\phi > 1$. For any candidate $\phi < 0$, $\phi = 0$ dominates, since $1 + \tilde{x} < (1+\mu\phi^2)(1-\phi)^2\tilde{x}+(1+\mu\phi^2)$. Second part: If $\tilde{x} = 0$, the minimization becomes $\min(1+\mu\phi^2)$, which implies that $\phi = 0$ is optimal, while for $\tilde{x} = \infty$, the minimization becomes $\min(1 + \mu\phi^2)(1 - \phi)^2$, which implies that $\phi = 1$ is optimal.

Lemma 3.2: For $\mu < 8$, dd is never zero on the interval [0,1[. For $\mu \geq 8$, there exists ϕ_a and ϕ_b between 0 and 1, such that

1. $\phi_a = \frac{1}{4} - (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}; \phi_b = \frac{1}{4} + (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}$
2. $dd = 0$ when ϕ_t equals $\phi_a \in [0, 1[$ or $\phi_b \in [0, 1[$
3. $dd < 0$ for $\phi_t \in]\phi_a, \phi_b[$

Proof: Since $dd = (1 - \phi)[(1 + \mu\phi^2) - \mu\phi(1 - \phi)]$, it has three roots, with $\phi = 1$ always being a real root. If $\mu \geq 8$, the two other roots are real and between 0 and 1, as they are given by $\frac{1}{4} - (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}$ and $\frac{1}{4} + (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}$. If $\mu < 8$, the other two roots are imaginary and hence dd is always positive on [0,1[.

Lemma 3.3: If $\mu \leq 4$, then $nn \geq 0$ for ϕ in $[0, 1]$. If $\mu > 4$, then there exists ϕ_c and ϕ_d between 0 and 1, such that $nn \geq 0$ for $\phi \in [0, \phi_c]$, $nn < 0$ for $\phi \in]\phi_c, \phi_d[$, and $nn \geq 0$ for $\phi \in [\phi_d, 1]$. Moreover, when $\mu \geq 8$, $nn > 0$ when $\phi = \phi_a$ and $nn < 0$ when $\phi = \phi_b$.

Proof: This proof relies on understanding the shape of the cubic function $nn = \mu - 3\mu^2\phi^2 + 4\mu^2\phi^3$ as a function of ϕ. First note that at $\phi = 0$, $nn = \mu$, and the derivative of the cubic at this point is 0. On the other hand, at $\phi = 1$, nn is again positive but now with a positive slope. Hence, this cubic will have one real root that is negative, and will have two other real roots that are between 0 and 1 if the discriminant of this cubic is positive (note: the discriminant of a cubic of the form $az^3 + bz^2 + d$ is $-4b^3 - 27a^2$). Otherwise, it will have only one real root. For the discriminant to be positive, we need $\mu \geq 4$. Accordingly, for $\mu \leq 4$, nn is always positive on the interval $[0, 1]$. If $\mu > 4$, then $nn \geq 0$ for $\phi \in [0, \phi_c]$, $nn < 0$ for $\phi \in]\phi_c, \phi_d[$, and $nn \geq 0$ for $\phi \in [\phi_d, 1]$ where

ϕ_c and ϕ_d are the two positive real roots of the cubic. The last part follows from replacing ϕ in $\mu - 3\mu^2\phi^2 + 4\mu^2\phi^3$ by either $\phi = \frac{1}{4} - (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}$ or $\phi = \frac{1}{4} + (\frac{1}{16} - \frac{1}{2\mu})^{\frac{1}{2}}$.

Lemma 3.4: If $\mu \leq 4$, the first-order condition given by equation 3.21 implicitly defines a unique $\phi(\tilde{x})$ function that is increasing in $\tilde{x}$. $\hat{\phi}(\tilde{x})$ is therefore given by this implicitly defined function. If $\mu > 4$, then the first- and second-order conditions define two monotonically increasing and continuous functions $\phi_1(\tilde{x})$ and $\phi_2(\tilde{x})$ that represent local optima, where $\phi_1(\cdot)$ is defined over $\tilde{x} \in (0, z_1)$ and $\phi_2(\cdot)$ is defined over $\tilde{x} \in (z_2, \infty)$ with $0 < z_2 < z_1$ (z_1 can be ∞), and $\phi_2(z_2) > \phi_1(z_1)$.

Proof: Lemmas 3.2 and 3.3 provide the key elements for a characterization of the inverse function $\tilde{x}_{t-1} = \frac{\mu\phi_t}{(1-\phi_t)(1+\mu\phi_t^2)-\mu\phi_t(1-\phi_t)^2}$. On the intervals where this function is positive and increasing, its inverse offers candidate functions for $\hat{\phi}(\tilde{x}_{t-1})$, since they will satisfy both the first-order condition (FOC) and the second-order condition (SOC) on the relevant ranges.

If $\mu < 4$, Lemmas 3.2 and 3.3 imply that the FOC defines a function $\phi(\tilde{x}_{t-1})$ that is continuous and monotonically increasing on $\tilde{x}_{t-1} \geq 0$, since the inverse function is monotonic and continuous on the interval [0,1]. When $\mu \geq 4$, the inverse function $\tilde{x}_{t-1} = \frac{\mu\phi_t}{(1-\phi_t)(1+\mu\phi_t^2)-\mu\phi_t(1-\phi_t)^2}$ provides two candidate functions $\phi_1(\tilde{x}_{t-1})$ and $\phi_2(\tilde{x}_{t-1})$ and their ranges, and these satisfy both the FOC and SOC.

In the case where $4 < \mu < 8$, foregoing lemmas imply that the inverse function $\tilde{x}_{t-1} = \frac{\mu\phi_t}{(1-\phi_t)(1+\mu\phi_t^2)-\mu\phi_t(1-\phi_t)^2}$ is continuous but not monotonic. In particular, on the interval [0,1], the inverse function will first be increasing, then decreasing, and then increasing again. Hence the FOC implies a correspondence between ϕ_t and $\tilde{x}_{t-1}$ on a range for $\tilde{x}_{t-1}$. This implies that the FOC and the SOC together define two positively sloped relationships between $\tilde{x}_{t-1}$ and ϕ_t that share a range but do not cross. Using ϕ_c and ϕ_d defined in Lemma 3.3, we need to set $z_1 = \frac{\mu\phi_c}{(1-\phi_c)(1+\mu\phi_c^2)-\mu\phi_c(1-\phi_c)^2}$ and $z_2 = \frac{\mu\phi_d}{(1-\phi_d)(1+\mu\phi_d^2)-\mu\phi_d(1-\phi_d)^2}$ to define the intervals over which $\phi_1(\tilde{x}_{t-1})$ and $\phi_2(\tilde{x}_{t-1})$ are defined.

Finally, if $\mu \geq 8$, then the inverse function is neither continuous nor monotonic. Instead, it will have two points of discontinuity, and the sign of its slope (defined everywhere except at these two points of discontinuity) will change (as in the case of $4 < \mu < 8$) exactly twice on the interval [0,1] (again: first positive, then negative and then positive). In particular, this implies that the FOC and SOC again define two monotonically increasing functions (for ϕ as a function of $\tilde{x}_{t-1}$) that share a range but do not cross. Moreover, one branch must start at $(0, 0)$, while the other branch must end at $(\infty, 1)$. Using previous lemmas, the lower branch will be defined over $[0, \infty[$ (i.e., $z_1 = \infty$), while the higher branch will be defined over $[z_2, \infty[$ where $z_2 = \frac{\mu\phi_d}{(1-\phi_d)(1+\mu\phi_d^2)-\mu\phi_d(1-\phi_d)^2}$.

Lemma 3.5: Conditional on $\mu > 4$, there exists a $z_3 \in (z_2, z_1)$ such that the $\hat{\phi}(\tilde{x}_{t-1})$ function will correspond to $\phi_1(\tilde{x}_{t-1})$ for $\tilde{x}_{t-1} \in (0, z_3)$ and correspond to $\phi_2(\tilde{x}_{t-1})$ for $x \in (z_3, \infty)$.

Proof: The optimal function $\hat{\phi}(\tilde{x}_{t-1})$ is composed of segments of $\phi_1(\tilde{x}_{t-1})$ and $\phi_2(\tilde{x}_{t-1})$ as defined in Lemma 3.4. It is clear that the optimal function has to jump at least once between $\phi_1(\tilde{x}_{t-1})$ and $\phi_2(\tilde{x}_{t-1})$, since we know that $\hat{\phi}(\tilde{x}_{t-1})$ needs to start at zero and goes to 1 as $\tilde{x}_{t-1}$ goes from 0 to ∞. To prove that this is the only point at which jumping between branches can occur, define for $i = 1, 2$,

$$U_i(\tilde{x}_{t-1}) = (1 + \mu\phi_i(\tilde{x}_{t-1})^2)(1 - \phi_i(\tilde{x}_{t-1}))^2\tilde{x}_{t-1} + (1 + \mu\phi_i(\tilde{x}_{t-1})^2).$$

This function gives the social cost of inflation and employment deviations that arises when following the policy rule defined by $\phi_i(\tilde{x}_{t-1})$ discussed in Lemma 3.4. Since previous lemmas imply $\phi_2(\tilde{x}_{t-1}) > \phi_1(\tilde{x}_{t-1})$, this means that at any positive value for $\tilde{x}_{t-1}$ at which $U_1(\tilde{x}_{t-1}) = U_2(\tilde{x}_{t-1})$ must satisfy $(1 + \mu\phi_1(\tilde{x}_{t-1})^2)(1 - \phi_1(\tilde{x}_{t-1}))^2 > (1 + \mu\phi_2(\tilde{x}_{t-1})^2)(1 - \phi_2(\tilde{x}_{t-1}))^2$. However, the envelope theorem makes this last inequality equivalent to $\frac{\partial U_1(\tilde{x}_{t-1})}{\partial \tilde{x}_{t-1}} > \frac{\partial U_2(\tilde{x}_{t-1})}{\partial \tilde{x}_{t-1}}$. Conclude that any point of intersection between the two functions must have the property that $U_1(\tilde{x}_{t-1})$ approaches $U_2(\tilde{x}_{t-1})$ strictly from below, thus precluding multiple intersections.

Propositions 3.1 and 3.2 follow from Lemmas 3.4 and 3.5.

Proposition 3.3: If ϕ and ϕ' represent the two stances of monetary policy at the pivot point, then $\phi(1 - \phi) = \phi'(1 - \phi')$. This implies that at the optimal point of pivot, inflation is expected to fall but employment is expected to stay constant. However, the variance of employment increases discontinuously at the pivot point.

Proof: If there is a point of pivot at the value of $\tilde{x} = \tilde{x}'$, which arises only if $\mu > 4$, then the following three equations must be satisfied:

1) The payoff must be the same at $\tilde{x}'$ whether the monetary stance is given by ϕ and ϕ', that is

$$(1 + \mu\phi^2)(1 - \phi)^2\tilde{x}' + (1 + \mu\phi^2) = (1 + \mu\phi'^2)(1 - \phi')^2\tilde{x}' + (1 + \mu\phi'^2).$$

2) Both ϕ and ϕ' must satisfy the first-order condition (FOC), that is

$$[\mu\phi(1 - \phi)^2 - (1 - \phi)(1 + \mu\phi^2)]\tilde{x}' + \mu\phi = 0$$

$$[\mu\phi'(1 - \phi')^2 - (1 - \phi')(1 + \mu\phi'^2)]\tilde{x}' + \mu\phi' = 0.$$

Using the FOC for ϕ, we know that

$$\tilde{x}' = \frac{\mu\phi}{(1-\phi)(1+\mu\phi^2) - \mu\phi(1-\phi)^2}.$$

This allows us to reduce these three equations to the following two equations in ϕ and ϕ':

$$\frac{\mu\phi(1+\mu\phi^2)(1-\phi)^2}{(1-\phi)(1+\mu\phi^2) - \mu\phi(1-\phi)^2} + (1+\mu\phi^2) = \frac{\mu\phi(1+\mu\phi'^2)(1-\phi')^2}{(1-\phi)(1+\mu\phi^2) - \mu\phi(1-\phi)^2} + (1+\mu\phi'^2)$$

$$\frac{[\mu\phi'(1-\phi')^2 - (1-\phi')(1+\mu\phi'^2)]\mu\phi}{(1-\phi)(1+\mu\phi^2) - \mu\phi(1-\phi)^2} + \mu\phi' = 0$$

It can be verified that $\phi = \frac{\mu-(\mu^2-4\mu)^{.5}}{2\mu}$ and $\phi' = \frac{\mu+(\mu^2-4\mu)^{.5}}{2\mu}$ comprise a solution to these two equations when $\mu > 4$. Note that these values for ϕ and ϕ' satisfy the relationship $\phi' = 1 - \phi$.[13] Moreover, since we previously proved that there can be at most one pivot, these values for ϕ and ϕ' characterize the unique solution.

Given that the change in expected employment at the point of pivot is proportional to $-\phi'(1-\phi') + \phi(1-\phi)$, and given that we have shown that $\phi' = 1 - \phi$, the expected change in employment at the pivot is zero. In contrast, the change in the variance of employment at the pivot point is given by $(\phi'^2 - \phi^2)\sigma_\theta^2$, which experiences a jump at the pivot value $\tilde{x}'$ since $\phi' > \phi$.

A.2 A more general Phillips curve

Under the log-linear preferences assumed in the main text, the Phillips curve generated by the model is

$$\hat{\pi}_t = \mathbb{E}_{t-1}\hat{\pi}_t + \mathbb{E}_{t-1}\hat{N}_t - \hat{\theta}_t.$$

As we noted in Section 4.1, this expression gives a coefficient of unity on $\mathbb{E}_{t-1}\hat{N}_t$, which implies a very steep Phillips curve. This feature can be relaxed quite easily by changing preferences to GHH preferences. Assume that expected lifetime utility of households is given by

$$\mathbb{E}\sum_{t=0}^{\infty}\beta^t \ln\left(C_{it} - \eta\theta_t N_t^{1+\lambda}\right).$$

[13] One way of seeing that these values for ϕ and ϕ' satisfy the two equations is to begin by conjecturing that $\phi' = 1 - \phi$ and use this to replace ϕ'. Then each of these two equations individually becomes a function of only ϕ. Both of these two new equations in ϕ only then can be reduced to the same cubic function of ϕ given by $2\mu\phi^3 - 3\mu\phi^2 + (2+\mu)\phi - 1 = 0$. This cubic function can then be factored as $(\phi - .5)(2\mu\phi^2 - 2\mu\phi + 2) = 0$. The postulated ϕ and ϕ' are the roots of $(2\mu\phi^2 - 2\mu\phi + 2) = 0$, and these satisfy the conjecture.

This maximization is done subject to the same budget constraint as before:

$$\sum P_{jt}C_{ijt} + B_{it+1} = W_{it}N_{it} + D_{it} + \tau_{it} + B_{it}(1+\iota_t).$$

As in the text, W_{it} is predetermined at date t. At every date, individuals choose consumption and bond purchases as well as the nominal wage for the next period. The first-order conditions for this problem are

$$\frac{1}{P_t\left(C_{it} - \eta\theta_t N_{it}^{1+\lambda}\right)} = \beta\mathbb{E}_t\left[\frac{1+\iota_{t+1}}{P_{t+1}\left(C_{it+1} - \eta\theta_t N_{it+1}^{1+\lambda}\right)}\right] \tag{A.26}$$

$$W_{it+1} = \left(\frac{(1+\lambda)\rho\eta}{\rho-1}\right)\frac{\mathbb{E}_t\left[\frac{\theta_{t+1}N_{it+1}^{1+\lambda}}{C_{it+1}-\eta\theta_t N_{it+1}^{1+\lambda}}\right]}{\mathbb{E}_t\left[\frac{N_{it+1}}{P_{t+1}\left(C_{it+1}-\eta\theta_t N_{it+1}^{1+\lambda}\right)}\right]}. \tag{A.27}$$

Equation A.26 is the Euler equation governing the saving decision, while equation A.27 is the optimal wage set by individual i for period $t+1$.

Under a symmetric equilibrium, we have $C_{it} = C_t, N_{it} = \frac{N_t}{n}$ and $W_{it} = W_t$ for all i. Using this and approximating $\ln\mathbb{E}(x) \approx \mathbb{E}\ln(x)$, we can rewrite equation A.27 as

$$\ln W_t = \ln\left(\frac{\rho(1+\lambda)\eta}{\rho-1}\right) - \lambda\ln n + \mathbb{E}_{t-1}\ln P_t + \mathbb{E}_{t-1}\ln\theta_t + \lambda\mathbb{E}_{t-1}N_t, \tag{A.28}$$

where we have lagged the aggregate wage equation by one period. Equation A.28 is identical to the corresponding expression in the log-linear utility case given by equation 2.13, except for constants and the λ instead of one multiplying $\mathbb{E}_{t-1}N_t$.

To compute the natural level of employment in this case, suppose for a moment that wages and prices are both flexible and determined simultaneously. In this case, equation A.27 should be replaced with the following:

$$W_{it} = \left(\frac{(1+\lambda)\rho\eta}{\rho-1}\right)P_t\theta_t N_{it}^{\lambda}.$$

Since $PY = \frac{\gamma}{\gamma-1}WN$ and $Y = n^{\frac{1}{\rho-1}}\theta N$, the above can be rewritten as

$$\frac{\gamma-1}{\gamma}n^{\frac{1}{\rho-1}} = \left(\frac{(1+\lambda)\rho\eta}{\rho-1}\right)P\left(\frac{N}{n}\right)^{\lambda}.$$

This gives the natural level of employment under GHH preferences as

$$\bar{N} = \left[\left(\frac{\gamma-1}{\gamma}\right)\left(\frac{\rho-1}{\rho\eta(1+\lambda)}\right)n^{\frac{1}{\rho-1}+\lambda}\right]^{\frac{1}{\lambda}}. \tag{A.29}$$

Using the expression for $\bar{N}$ in equation A.29 to add and subtract $\lambda\ln\bar{N}$ from equation A.28 gives

$$\ln W_t = \ln\left(\frac{\gamma-1}{\gamma}\right) + \left(\frac{1}{\rho-1}\right)\ln n + \mathbb{E}_{t-1}\ln P_t + \lambda\mathbb{E}_{t-1}\hat{N}_t + \mathbb{E}_{t-1}\ln\theta_t. \tag{A.30}$$

From the price-setting equation, we have

$$\ln P_t = \ln\left(\frac{\gamma}{\gamma-1}\right) + \left(\frac{1}{1-\rho}\right)\ln n + \ln W_t - \ln\theta_t.$$

Subtracting $\ln P_{t-1}$ from both sides and using equation A.30 to get $\ln W_t - \ln P_{t-1}$, we get

$$\hat{\pi}_t = \mathbb{E}_{t-1}\hat{\pi}_t + \lambda\mathbb{E}_{t-1}\hat{N}_t - \hat{\theta}_t. \tag{A.31}$$

This is the same Phillips curve as before but with slope coefficient λ.

A.3 The Euler equation and the monetary policy rule

By exploiting the symmetry $C_{it} = C_t$ for all i, we can write the individual Euler equation in equation 2.7, after taking logs and making the usual approximations, as

$$\ln C_t = \mathbb{E}_t\ln C_{t+1} - [(\iota_{t+1} - \bar{\iota}) - \mathbb{E}_t(\pi_{t+1} - \pi^*)],$$

where $\bar{\iota} \equiv \pi^* - \ln\beta$. Substituting the market clearing condition $C_t = m^{\frac{1}{\gamma-1}}Y_t = m^{\frac{1}{\gamma-1}}n^{\frac{1}{\rho-1}}\theta_t N_t$ into this aggregate Euler equation then gives

$$\ln N_t - \ln\bar{N} = \mathbb{E}_t(\ln N_{t+1} - \ln\bar{N}) - [\iota_{t+1} - \bar{\iota} - \mathbb{E}_t(\pi_{t+1} - \pi^*)],$$

where $\pi_t = \frac{P_t}{P_{t-1}}$ is the inflation rate, π^* is the central bank's target level of inflation and $\bar{N}$ is the natural rate of employment, and where we have used the random walk property of θ_t.

We can then iterate forward on this equation to get[14]

[14] Given that we consider departures from rational expectations, we are not imposing the law of iterated expectations in general. When we look at the case with rational expectations, this iteration reduces to $\mathbb{E}_t\sum_{h=1}^{\infty}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)]$.

$$\ln N_t - \ln \bar{N} = -\sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)], \tag{A.32}$$

where employment is now expressed as deviations of future expected real interest rates from the natural real rate of interest. Equations 2.16 and A.32 represent the constraints on the central bank regarding how monetary policy can affect inflation and employment. In fact, $\sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)]$ can be viewed as the amount of monetary stimulus injected into the system at any point in time, and equation A.32 represents how that monetary stimulus affects real activity.

We could complete this environment by specifying a class of policy rules for setting ι_{t+1}. However, it will be more convenient to think of the central bank as setting policy with the aim of controlling the stimulus path $\sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)]$. Given the importance that central banks associate with communicating a stance of monetary policy that goes beyond only the current policy rate setting, this approach can be viewed as both more realistic and encompassing. Accordingly, instead of thinking only of ι_{t+1} as being set as a function of inflation, as would be typical with a Taylor rule, here we think of the central bank as setting an expected stimulus path as a function of inflation according to

$$\sum_{h=1}^{\infty} \mathbb{E}_t \cdot \cdot \mathbb{E}_{t+h-1}[\iota_{t+h} - \bar{\iota} - (\pi_{t+h} - \pi^*)] = \phi_t(\pi_t - \pi^*), \tag{A.33}$$

where ϕ_t captures the extent to which the central bank adds or removes stimulus depending on current inflation. Combining equations A.32 and A.33 then gives

$$\ln N_t - \ln \bar{N} = -\phi_t(\pi_t - \pi^*), \tag{A.34}$$

which is the same as equation 2.18 that we derived in the main text.

With this perspective, the economy inherits a simple recursive structure whereby, for a given policy stance captured by $\{\phi_t\}$, equations 2.16 and A.34 determine inflation and employment. Once the paths for inflation and employment are determined, one can use the Euler equation to find the expected nominal interest rate process that would implement the outcome.

A.4 Oil price shocks as productivity shocks

One of the objectives of this paper has been to examine optimal monetary policy when an economy faces supply-side shocks, with oil price shocks being one especially relevant example. However, in the model, the only source of aggregate uncertainty is a productivity shock. We now show the mapping between oil price shocks and productivity shocks to the value-added production function of an economy.

Let the gross output of the economy be given by

$$Q_t = F(Z_t N_t, O_t), \tag{A.35}$$

where $F(\cdot,\cdot)$ is a constant returns-to-scale (CRS) function in its two arguments, Z_t is a labour-augmenting productivity factor, N_t is labour and O_t is oil. The corresponding value-added function is

$$\frac{Y_t}{N_t} = F\left(Z_t, \frac{O_t}{N_t}\right) - P_t^o \frac{O_t}{N_t},$$

where P_t^o denotes the price of a unit of oil. Note that we have used the CRS property of F to write the value-added function in per-worker terms.

The first-order condition governing optimal oil usage is

$$F_2\left(Z_t, \frac{O_t}{N_t}\right) = P_t^o,$$

where we use $F_2(\cdot,\cdot)$ to denote the derivative of F with respect to its second argument. Since F is CRS in its two arguments,

$$F_2\left(Z_t, \frac{O_t}{N_t}\right) = F_2\left(1, \frac{O_t}{Z_t N_t}\right) = f'\left(\frac{O_t}{Z_t N_t}\right),$$

where $f(\cdot) \equiv F(1,\cdot)$, and thus $f''(\cdot) < 0$. We can use this expression along with the first-order condition to derive optimal oil demand as

$$\frac{O_t}{N_t} = Z_t (f')^{-1}(P_t^o). \tag{A.36}$$

It is straightforward to verify that the demand for oil is decreasing in its price.

We can use the optimal demand for oil from equation A.36 in the value-added production function to get

$$Y_t = N_t Z_t h(P_t^o),$$

where $h(P_t^o) \equiv f[(f')^{-1}(P_t^o)] - P_t^o (f')^{-1}(P_t^o)$. Differentiating h with respect to P_t^o and evaluating it around the optimum gives

$$h'(P_t^o) = -(f')^{-1}(P_t^o) < 0.$$

Defining $\theta_t \equiv Z_t h(P_t^o)$, we get the value-added function to be

$$Y_t = N_t \theta_t, \tag{A.37}$$

where labour productivity θ_t is a decreasing function of the oil price P_t^o.