Staff Discussion Paper/Document d’analyse du personnel—2024-7

Last updated: July 2, 2024

The Output-Inflation Trade-off in Canada

by Stefano Gnocchi, Fanny McKellips, Rodrigo Sekkel, Laure Simon, Yinxi Xie and Yang Zhang

Canadian Economic Analysis Department
Bank of Canada
SGnocchi@bankofcanada.ca, FMckellips@bankofcanada.ca, RSekkel@bankofcanada.ca, LSimon@bankofcanada.ca, YinxiXie@bankofcanada.ca, YangZhang@bankofcanada.ca

DOI: https://doi.org/10.34989/sdp-2024-7 | **ISSN** 1701-9397

Acknowledgements

We thank Don Colleti, Olivier Gervais, Marc-André Gosselin, Martín Harding, Sharon Kozicki, Dmitry Matveev, Tolga Özden and seminar participants at the Bank of Canada. All errors are our own. The views in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

Abstract

We explain how the Bank of Canada’s policy models capture the trade-off between output and inflation in Canada. We start by briefly revisiting the determinants of the New Keynesian Phillips curve. Next, we provide an overview of the Phillips curves that are currently embedded in the two main policy models the Bank uses for macroeconomic projections and analysis, known for short as ToTEM and LENS. We then discuss the challenges in identifying the trade-off between output and inflation and provide new estimates of the trade-off using recently proposed methods. Finally, we contrast these estimates with the ones in the Bank’s policy models.

Topics: Business fluctuations and cycles; Econometric and statistical methods; Inflation and prices; Monetary policy transmission
JEL codes: E31, E52

Résumé

Nous expliquons comment les modèles économiques de la Banque du Canada rendent compte de l’arbitrage entre la production et l’inflation au Canada. Tout d’abord, nous réexaminons brièvement les déterminants de la courbe de Phillips néo-keynésienne. Ensuite, nous présentons les courbes de Philips actuellement intégrées aux deux principaux modèles utilisés par la Banque pour ses projections et analyses macroéconomiques, c’est-à-dire les modèles TOTEM et LENS. Nous abordons également les difficultés à identifier l’arbitrage entre la production et l’inflation, et nous fournissons de nouvelles estimations de cet arbitrage réalisées à l’aide de méthodes récemment proposées. Enfin, nous comparons ces estimations à celles issues des modèles de la Banque.

Sujets : Cycles et fluctuations économiques; Inflation et prix; Méthodes économétriques et statistiques; Transmission de la politique monétaire
Codes JEL : E31, E52

1 Introduction

The trade off between output and inflation represents one of the key issues in macroeconomics and monetary policy. In this paper, we quantify such a trade-off based on the Bank of Canada's estimated policy models, which we complement with new reduced-form empirical evidence for Canada.

We learn several lessons.

First, inferring the output-inflation trade-off from a Phillips curve alone suffers from several limitations, both conceptual and empirical. From a conceptual standpoint, the Phillips curve, which provides the theoretical underpinning of the very existence of an output-inflation trade-off, is not sufficient by itself to quantify the output cost of stabilizing inflation. From an empirical standpoint, specification uncertainty and identification issues make it very challenging to rely exclusively on the estimation of a single equation to reach reliable conclusions. Overall, our view is that, parallel to recent developments in the literature on fiscal policy, a better strategy is to resort jointly to two complementary approaches: (i) measures for the output-inflation trade-off based on estimated, fully fledged models of the whole economy; and (ii) its direct estimation, without needing to specify a Phillips curve model. These two routes, combined, reinforce each other's strengths and balance out some of the risks.

Second, after applying state-of-the-art techniques, we quantify the output-inflation trade-off in Canada. We conclude that, overall, our two main policy models are not too far from each other and not too far off from the data.

Last but not least, we cannot avoid noticing that identifying the Canadian trade-off seems harder than the US trade-off. Nevertheless, our findings, together with previous work by Champagne and Sekkel (2017), seem to point in the direction of a far stronger trade-off in Canada than in the United States.

We proceed as follows. To fix ideas, we start by laying out the textbook New Keynesian Phillips curve (NKPC) as in Galí (2015). We discuss the components that drive inflation dynamics as implied by the NKPC, namely economic slack, inflation expectations and trade-off shocks. We then discuss the variants of the Phillips curve embedded in the Bank's two main in-house models, known best by their acronyms ToTEM and LENS.[1] ToTEM is a dynamic stochastic general equilibrium (DSGE) model that imposes restrictions across equations, whereas LENS imposes far less structure on the data. For each model, we compute measures of the output-inflation trade-off. First, we compute the so-called sacrifice ratio, which measures the cumulative loss of gross domestic product (GDP) after a one-period contractionary monetary policy shock that lowers inflation by 1% at its peak. Effectively, this ratio measures the cost, in terms of a drop in output, that a central bank faces when fighting inflation. Next, following Barnichon and Mesters (2021), we also compute the Phillips multiplier for both models, that is, by how much inflation decreases following a monetary policy shock that raises the unemployment rate, or decreases GDP, by 1%. The smaller the Phillips multiplier, the more severe is the trade-off between inflation and output. We show that the sacrifice ratio in ToTEM is 18.1 at a 12-quarter horizon, whereas in LENS it is almost double at 35.5. This implies a worse output-inflation trade-off in the latter compared with the former. Similarly, the Phillips multiplier, also at a 12-quarter horizon, is higher in ToTEM (0.44) than in LENS (0.23), again reflecting a worse output-inflation trade-off in the latter model.

ToTEM and LENS both rely on economic theory to varying degrees, imposing several restrictions on the data that could bias their parameters due to misspecifications. We next discuss the possibility and challenges associated with directly estimating the sacrifice ratio and Phillips multipliers without specifying a full macroeconomic model. We review different techniques proposed in the literature and apply them to Canadian data to complement the findings of our policy models. First, we follow Barnichon and Mesters (2021) and estimate the Phillips multipliers for Canada using exogenously identified monetary policy shocks. We find a Phillips multiplier of approximately 0.27 after 12 quarters. This figure is slightly higher than the multiplier for LENS at the same horizon (0.23), but smaller than the one for ToTEM (0.44).

1. See Corrigan et al. (2021) for the details of ToTEM and Gervais and Gosselin (2014) for LENS.

Lastly, we estimate a Phillips curve for Canada using regional variations in economic activities, following the approach in McLeay and Tenreyro (2020) and Hazell et al. (2022). Since central banks do not respond to regional differences in demand conditions, these differences could potentially provide information about the slope of the Phillips curve. While not speaking directly to our policy models, this exercise is useful to place Canadian evidence into an international context.

It is important to acknowledge that there are limits to how much we can learn from the data. Small monetary policy shocks since the Bank implemented an inflation-targeting regime imply larger confidence intervals in the response of GDP and inflation to those shocks. Similarly, strong co-movement of economic activity between provinces might imply little additional information in provincial data. Still, our reading of the literature suggests that the numbers we provide are the best available for Canada.

2 Theoretical background

In this section, we flesh out the standard NKPC and the theoretical underpinning of the trade-off between inflation and real economic activity. We also illustrate the main drivers of the real cost of fighting inflation.

For the sake of clarity, we focus on the simplest version of the NKPC, which is insightful for understanding the main mechanisms at play. In the discussion, however, we also highlight some of the most recent advances in the field that propose more sophisticated versions of the NKPC. The intent is to motivate the adoption of the richer models that central banks use as quantitative tools, and thus set the stage for the rest of the paper.

As well, our selection of recent literature provides additional context for the ongoing development of the fourth generation of the Bank's policy models, which complements the recent discussion made by Coletti (2023).

2.1 The New Keynesian Phillips curve

Consider a textbook New Keynesian model where the economy is subject to some form of price rigidity: firms infrequently adjust their prices to changes in market conditions, such as consumers' demand or production costs.[2] The solution to the firms' pricing problem in log-linearized form yields the following expression,

$$\pi_t = \beta \, \mathbb{E}_t \, \pi_{t+1} + \lambda mc_t, \tag{1}$$

where π_t and and mc_t are aggregate inflation and the aggregate real marginal cost, respectively, and $\mathbb{E}_t \, \pi_{t+1}$ is the inflation rate that firms expect will prevail in the future, conditional on the information they have available at time t. All variables are expressed as percentage differences from their long-run level, parameter $\lambda > 0$, and $0 < \beta < 1$.[3]

Intuitively, a positive relation exists between inflation and real marginal cost, as in the simplest case where all firms reset their prices each time a shock occurs. This is commonly referred to as the flexible-price equilibrium. Most importantly, the additional term $\mathbb{E}_t \, \pi_{t+1}$ arises. In fact, firms that reset prices rationally foresee that they may not have a chance of revising them for a while. Hence, their current decisions factor in expected future economic conditions, of which expected inflation is a sufficient statistic.

In light of the goal of our analysis, we need to rewrite equation (1) and relate inflation to some notion of economic slack. This can be done by relying on the equilibrium relation between the real marginal cost and the gap between actual output and the flexible-price level of output y_t, and so obtain:

$$\pi_t = \beta \, \mathbb{E}_t \, \pi_{t+1} + \kappa y_t, \tag{2}$$

2. See Galí (2015, chap. 3) for a textbook treatment and the algebraic derivations for this section, where firms adjusting prices are selected at random. The degree of price stickiness, captured by the probability of keeping the current price unchanged, is a stand-in for anything that makes a price change costly, above and beyond its immediate implications for the demand for goods. The cost of reprinting menus or of acquiring and processing information are just two common examples in the literature.

3. In what follows, inflation is assumed to be nil in the long run. All of our conclusions can be generalized to the case of positive steady-state inflation.

where $\kappa > 0$. Equation (2) implies that stabilizing both inflation and real activity around the flexible-price equilibrium is feasible. The economic intuition is straightforward: if $\pi_t = 0$ at all times, no firm ever has an incentive to change its price, even if it could. Hence, price rigidity cannot be binding. As a result, the economy achieves the same level of output that it would under flexible prices so that the gap y_t must remain closed.

However, the gap y_t may not be a policy-relevant measure of slack. For reasons that our stylized model does not capture, various temporary forces can indeed prevent the economy from reaching its potential, meaning the level of output at which all the factors of production are fully employed even if inflation is nil.[4] Any of these forces, which will be discussed below, would open a wedge between the flexible-price equilibrium and potential output. Hence, it is more appropriate to rewrite the NKPC,

$$\pi_t = \beta \, \mathbb{E}_t \, \pi_{t+1} + \kappa x_t + \mu_t, \tag{3}$$

where x_t finally represents deviations of output from its potential and μ_t, the supply shock, stands for any wedge between potential output and its flexible-price counterpart that continues to persist when inflation is nil.[5] Henceforth, we will refer to x_t as the output gap. Alternatively, an equivalent well-known formulation of the NKPC is written in terms of the unemployment gap, measured by the difference between the unemployment rate, u_t, and its natural rate, u_t^n.[6] The latter represents relatively small frictional unemployment, which, in an economy that works at full capacity, is due to unavoidable workers' churn between

4. We deliberately leave undefined the concept of potential because it is outside the scope of this paper. In the simplest case, a natural concept of potential could be the Pareto efficient equilibrium. However, the Pareto efficiency is not a good benchmark in most complex cases, so definitions other than efficiency may also be appealing. See Galí (2015) and Woodford (2003) for a detailed discussion and various references on these issues.

5. In this context, the use of the label "supply shock" is widespread among central bankers because the NKPC is typically conceptualized as an aggregate supply curve. Cost-push or trade-off shocks are alternative terms that are more common in the academic literature because some supply shocks, such as shocks to total factor productivity, do not necessarily generate a trade-off and hence do not give rise to any residual in (3).

6. Among many others, see Blanchard and Galí (2010) and Hazell et al. (2022).

different jobs and occupations. Hence,

$$\pi_t = \beta\,\mathbb{E}_t\,\pi_{t+1} - \kappa(u_t - u_t^n) + \mu_t. \tag{4}$$

Both formulations of the NKPC—equations (3) and (4)—include three terms: economic slack, supply shocks and inflation expectations. We next discuss each of them.

Economic slack. A trade-off between inflation and economic slack is built into equations (3) and (4). In other words, output has to fall short of its potential or unemployment has to rise above its natural level to tame inflation. The slope κ captures the sensitivity of inflation to real economic activity, other things being equal, and plays a key role in gauging the trade-off between output and inflation. Anything that impacts the pass-through from the real marginal cost to final goods prices also affects the slope of the Phillips curve. Broadly, there are two main factors for the pass-through:

- *Price stickiness.* The stickier the prices, the smaller is the value of κ. For greater nominal rigidity, the same change in the inflation rate calls for a larger variation in production relative to potential. As a result, the NKPC flattens and makes lowering inflation more costly in real terms. Evidence suggests that the frequency of price adjustments has been falling over the last decade. For example, Cotton and Garga (2022) document this long-term trend in the US economy due to a shift in the industrial composition from manufacturing toward services, which have stickier prices. The recent surge in inflation may have contributed to countering this trend (see section 2.3 and related contributions by Ascari and Haber (2022) and Costain, Nakov, and Petit (2022)).

- *Production structure.* The output-inflation trade-off is also shaped by the degree of real rigidity, which is in turn affected by the structure of production costs. Real rigidity is a term introduced by Ball and D. Romer (1990) to describe an economy with low elasticity of aggregate supply, meaning an economy where each industry's desired relative price is not very sensitive to aggregate economic conditions. Under such circumstances, the producers that adjust prices in response to a surge in demand do so to a very limited extent, reinforcing the aggregate effects of nominal rigidity. As a consequence, real rigidity reduces κ for a given degree of price stickiness. A prominent driver of real rigidity is the sectoral composition of production. Recent contributions extend the textbook version of the New Keynesian model to the case of multiple sectors and discuss their relevance for the output-inflation trade-off. For example, Rubbo (2020) shows that in a multi-sector New Keynesian model, several sector-specific Phillips curves can be combined to obtain an aggregate NKPC with the same form as equation (3), where κ depends on the share of intermediate inputs in the production of final goods. This framework predicts that the NKPC has flattened by about 30% over the last 70 years. Huang and Z. Liu (2005) and Wei and Xie (2020) also show that nominal rigidity compounds along the production chain, so that output prices display more inertia downstream.

Supply shocks. The output-inflation trade-off has no bite if $\mu_t = 0$ because equations (3) and (4) would reduce to equation (2). As already discussed, this implies the possibility of simultaneously and fully stabilizing inflation and the output gap. The feasibility of this case—the so-called divine coincidence—can however be undermined by various factors. For example, there may be a rise in price markups above and beyond the one implied by the nominal rigidity if firms become better able to retain their customers following a price increase (Galí 2015, chap. 5). Additional sources of nominal rigidity, such as wage stickiness, can also open a wedge between the flexible-price level of output and the one that a central bank would want to stabilize (chap. 6). Changes to value-added or labour-income taxes are another type of supply shock that requires inflation to adjust if one seeks to maintain the economy at its potential. As well, Rubbo (2020) shows that equation (3) generalizes to a multi-sector economy, where changes in productivity alone make it impossible to simultaneously stabilize inflation and the output gap. Finally, and most importantly for Canada, supply shocks in the Phillips curve in commodity-exporting small open economies may arise due to the presence of the commodity sector (see, for example, Ferrero and Seneca 2019). In general, the divine coincidence may not hold in small open economies, depending on the elasticity of substitution between imported and domestically produced goods (De Paoli 2009).

Inflation expectations. The Phillips curve (equation 3) suggests that a tight link exists between current and expected inflation. This forward-looking relation reflects the prevailing modeling approach in macroeconomics, which endows economic agents with full information (FI) about the underlying economy. It also assumes that they form rational expectations (RE) based on this comprehensive knowledge.[7] The Full Information Rational Expectations (FIRE) model of the economy serves as a useful benchmark where economic agents do not make systematic mistakes. However, the FIRE approach has come under scrutiny with the increasing availability of survey data. Recent empirical studies typically find that the approach fails to accurately capture the empirical regularities in the behaviour of inflation expectations observed in the data. A non-exhaustive list includes Coibion and Gorodnichenko (2015), Bordalo et al. (2020), Kohlhas and Walther (2021) and Adam, Matveev, and Nagel (2021).

This growing evidence has inspired a number of alternative approaches to modelling frictions in information and the formation of expectations. Some of these include:

- adaptive learning, where agents observe the data and learn from their mistakes (Evans and Honkapohja 2001; Carvalho et al. 2023)
- behavioural inattention, where agents are allowed to be myopic and discount the future disproportionately (Gabaix 2020), or rational inattention where agents choose what

7. In FIRE models, economic agents are still subject to a forecast error, which is, however, zero on average and uncorrelated with the state of the economy. In other words, mistakes are small and random.

kind of information they want to absorb and how much of it they want to absorb based on a cost-benefit analysis (Sims 2010)

- level-k thinking, where agents have a limited understanding of second-round effects and general equilibrium outcomes (Farhi and Werning 2019)
- finite planning horizons, where decision makers have limited foresight and can only plan into a finite distance into future (Woodford 2019; Woodford and Xie 2022)

In more recent studies, alternative approaches to forming expectations have been incorporated into otherwise standard quantitative macroeconomic models, which have been estimated and used as forecasting tools. These contributions find evidence against rational expectations in favour of the alternatives (Meggiorini 2023 and Özden, Forthcoming). Importantly, departures from FIRE not only improve the data fit of the model, but they are also capable of enhancing the reliability of policy counterfactuals (Angeletos and Lian 2022).

Given the importance of data availability on inflation expectations, many researchers and central banks have continuously put efforts into developing survey-based measures, including the Survey of Professional Forecasts available at the Federal Reserve Bank of Philadelphia, the Blue Chip Economic Indicators from Wolters Kluwer, and the Survey of Consumer Expectations published by the Federal Reserve Bank of New York. Since 2014, the Bank of Canada has conducted quarterly the Canadian Survey of Consumer Expectations (CSCE), which includes short- and long-term expectations of both inflation and labour market conditions. Coibion, Gorodnichenko, and Kamdar (2018) provide a comprehensive review on using survey expectations to test the validity of the NKPC.

2.2 Model-implied trade-off between output and inflation

It is crucial to recognize that the NKPC alone is not sufficient to determine the output cost of stabilizing inflation, and even less so its slope κ.

To begin with, the forward-looking nature of the NKPC makes both the current and future output gaps relevant for current inflation. To see this, one can iterate the NKPC forward and, for the case of equation (3), write,

$$\pi_t = \sum_{j=0}^{\infty} \mathbb{E}_t\, \beta^j \left[\kappa x_{t+j} + \mu_{t+j}\right]. \tag{5}$$

This intertemporal version of the curve states that for a given choice of the current inflation rate, π_t, there is a relation between current and future output losses, which cannot thus be chosen independently. In other words, the trade-off between output and inflation is inherently dynamic.

Two main implications arise from this dynamic relationship. On one hand, one needs to solve for the whole *path* of the output gap that is consistent with the inflation rate π_t that a central bank wants to bring about. Hence, a fully fledged DSGE model has to be specified. On the other hand, the output cost of changing the inflation rate by a given amount has to be defined over a specific time horizon. We discuss these two implications in turn.

A natural model for this context is the textbook three-equation New Keynesian model (Galí 2015, chap. 3), which complements the NKPC with an investment-savings (IS) curve and a monetary policy rule,

$$x_t = E x_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1}) + \varepsilon_t, \tag{6}$$

$$i_t = \Phi_\pi \pi_t + \Phi_x x_t + \delta_t, \tag{7}$$

respectively, where i_t denotes the nominal interest rate, terms ε_t and δ_t are exogenous disturbances, and σ, Φ_π, and Φ_x are constant parameters. The IS equation (6) captures aggregate demand as a function of the real interest rate, whereas the Taylor type rule (equation 7) describes the conduct of monetary policy, which systematically responds to inflation and the output gap. As already discussed extensively, the NKPC (3) is the aggregate supply relation that closes the model. One can then find the sequence $\{x_t, \pi_t, i_t\}$ that solves this system of equations when the realization of exogenous shocks is observed and their stochastic processes are specified. The importance of considering all of the model equations, in addition to the NKPC, is understated in the context of this simple example. It becomes evident, however, in larger models that depart from rational expectations or that feature a more complex production structure, for example. This is the case for models of the Canadian economy examined in the next section.

With a solution to the general equilibrium, one can engineer an exogenous fall in inflation of 1 percentage point by suitably choosing a realization for the monetary policy shock, and then measuring the associated output loss predicted by the model. In light of the theoretical model, the policy-relevant measure for the output loss should be the output gap. However, distinguishing between actual output and the output gap is immaterial under the commonly maintained assumption that a monetary policy shock does not affect potential output.[8] Hence, we focus on the former for the rest of the paper. The trade-off between output and inflation can then be measured by the sacrifice ratio, defined as the cumulative sum of foregone output that follows the exogenous monetary policy tightening. The literature has so far focused on a permanent reduction in inflation to quantify transition costs due to a change in the inflation target (e.g., Cecchetti and Rich 2001; Mankiw 2001; Tetlow 2022). We depart from previous contributions by restricting our attention to an exogenous and *transient* reduction in the inflation rate. In fact, our aim is to quantify the output cost of stabilizing inflation around its target when shocks hit the economy. As an example, a sacrifice ratio of two means that bringing down inflation by 1 percentage point implies a loss of 2 percentage points of output.[9] The sacrifice ratio depends on the persistence of the shock and on the specific horizon over which the output loss is computed. We discuss this in the next section where we quantify the output-inflation trade-off in two models of the Canadian economy.

8. In larger models, the change in business investment triggered by the shock may affect potential. Since the impact of this change would likely dissipate over the relevant time horizon, we choose to disregard it.

9. Another typical definition of this ratio, proposed by Ball (1994), is the cumulative change in unemployment from a one percentage point reduction in inflation.

The Phillips multiplier, recently introduced by Barnichon and Mesters (2021), is another measure for the output-inflation trade-off. The multiplier can be defined as the expected average change in inflation over a pre-specified horizon caused by a monetary policy shock that lowers expected average GDP or increases expected average unemployment by 1 percentage point. In other words, it quantifies a central bank's ability to use its interest rate policy to transform GDP and unemployment into inflation. This approach draws from the literature on fiscal shocks, where the government spending multiplier is defined as the ratio of the cumulative impulse responses of output and government spending (see, for example, Ramey and Zubairy 2018). Intuitively, the Phillips multiplier is smaller because the output-inflation trade-off is more severe and, as a result, is inversely related to the sacrifice ratio. In section 3, we compute the Phillips multipliers associated with our policy models and compare them with sacrifice ratios. In section 5, we estimate them with Canadian data.

2.3 Beyond the textbook model

Given its stylized nature, the standard New Keynesian model has been extended. Various studies propose more sophisticated versions of the NKPC that are more relevant to specific environments or better fit the data.

To begin with, Galí and Monacelli (2005) show that in the context of an economy that is open to international trade and capital flows, the NKPC takes the same form as in equation (2) under some conditions, where π_t stands for domestic inflation. Such isomorphism, however, ceases to hold under pricing protocols that look more empirically relevant and the NKPC becomes more complex. This is the case under local currency pricing, which is if firms set their prices in the currency of the country of export (LCP) or in US dollars (DCP) regardless of where goods are produced and consumed. Under the LCP, the NKPC indeed includes an extra term measuring deviations from purchasing power parity (Engel 2011), whereas the terms of trade become a driver of inflation together with the output gap under DCP (Gopinath et al. 2020; Gopinath and Itskhoki 2022).

The pivotal assumption of price and wage stickiness has also spurred substantial theoretical and empirical literature. Despite solid evidence in favour of nominal rigidity (Bils and Klenow 2004; Klenow and Kryvtsov 2008; Nakamura and Steinsson 2008; Eichenbaum, Jaimovich, and Rebelo 2011), the specific pricing protocol that is most appropriate is still an object of debate. As opposed to the textbook model, which assumes that the likelihood of a price change is constant across firms and over time, several studies that follow the seminal contribution by Golosov and Lucas (2007) show how the number of firms that decide to change prices, the frequency and size of price changes all depend on the state of the economy. Two implications are relevant to our discussion. First, large monetary policy shocks lead to proportionally larger adjustments of inflation but weaker effects on output, as if the NKPC was steeper. The result is explained by the fact that more firms react to the shock, and to a greater extent. Second, and for similar reasons, a monetary policy shock has a greater impact on inflation when inflation is persistently higher than usual. Both of these predictions are empirically validated by Ascari and Haber (2022), whereas Costain, Nakov, and Petit (2022) propose a model that rationalizes these findings.

More broadly, various studies depart from the linearity of the textbook model and feature an NKPC with a slope that depends on the state of the economy. As documented by Forbes, Gagnon, and Collins (2021), the Phillips curve becomes steeper when output is above potential and flatter when it is below. Harding, Lindé, and Trabandt (2022) propose a New Keynesian model consistent with this evidence, where the non-linearity of the NKPC is due to the fact that the elasticity of demand varies with economic activity as in Kimball (1995), instead of being constant as in the baseline model.

All of these recent advances highlight how the textbook model must be enriched to quantitatively account for the output cost of stabilizing inflation.[10] This is an important consideration that motivates central banks to use much more complex models for their policy analysis. A discussion of models used at the Bank of Canada is the subject of the next

10. Recent studies also suggest that fiscal policy plays a significant role in affecting the estimates of Phillips curve; see, for example, Jørgensen and Ravn (2022) and C. Liu and Xie (2022).

section.

3 Output-inflation trade-off in Bank of Canada models

We now turn to the analysis of the output-inflation trade-off embedded in the two major Bank of Canada policy models.

3.1 ToTEM, LENS and the Phillips curve

ToTEM and LENS are two large-scale open-economy models that capture salient features of the Canadian economy. ToTEM is a DSGE model described in Corrigan et al. (2021). LENS follows a methodology more similar to that of the Federal Reserve Board's FRB/US model. For an in-depth discussion, we refer to Gervais and Gosselin (2014). While its reduced-form equations are informed by theory, the specification of LENS, relative to ToTEM, is more heavily driven by empirical fit.

The two models complement each other in our policy analysis. ToTEM hinges on a detailed set of assumptions about markets and economic agents' behaviour. As a result, conditional on the assumptions being valid, the model provides some causal interpretation of observed economic dynamics, disentangling the role played by monetary policy from other economic fundamentals. In contrast, LENS is less prone to misspecification error owing to its greater flexibility. Relative to a fully fledged DSGE, it may also have an advantage as a forecasting tool. In sum, the joint use of the two models strikes a balance between the benefits and risks of imposing structure onto the data. As well, the two models, in light of their differences, naturally lend themselves to different estimation procedures. While both models are estimated with Bayesian methods, all equations are estimated at once in ToTEM whereas LENS is estimated by blocks of equations.

ToTEM. There are several differences between the NKPC presented in section 2 and the one implied by ToTEM. To begin with, ToTEM models various sectors, each corresponding to a component of gross domestic product (GDP), and features various sectoral Phillips curves. Each curve relates an inflation rate to its relevant marginal cost measure, and the curves are not combined into a single aggregate one. As a result, analytically, a structural relation between inflation and aggregate economic slack akin to equation (3) is not readily available. At the same time, the set of equations implies an equilibrium relation between real activity and the consumer price index because capacity utilization drives the real marginal cost in the sector that produces consumption goods. Furthermore, the pricing behaviour of firms is more elaborate than in the textbook New Keynesian model: a fraction of firms use a rule of thumb to set their price, using a linear combination of past inflation and the inflation target. Therefore, the resulting Phillips curve includes lagged inflation and the inflation target in addition to expected inflation.

LENS. Similar to equation (3), the Phillips curve in LENS features a measure of slack captured by potential output. Like in ToTEM, firms' pricing protocol is more complex than in the textbook version of the NKPC and is consistent with the presence of pure forward-looking and rule-of-thumb behaviour. Most importantly, the more recent version of the model distinguishes between short-run and long-run expectations. Finally, the curve is also augmented with import prices excluding commodities.

Differences in modelling assumptions as well as the greater complexity of both models, relative to the simple example discussed in the previous section, demonstrate that drawing comparisons by merely looking at the slope of one Phillips curve is at best opaque and, at worst, misleading.

3.2 Sacrifice ratios and Phillips multipliers

As anticipated in the previous section, the sacrifice ratio is defined as the cumulative sum of foregone output following an exogenous monetary policy tightening that reduces inflation by one percentage point at peak. We assume that in each of the two models the inflation reduction dissipates at the speed implied by the estimated persistence of the monetary policy shock.[11] In ToTEM, we compute inflation by using the Phillips curve of the consumption sector, whereas in LENS we calculate the relevant inflation rate from the readily available aggregate Phillips curve.

Figure 1 shows the impact of the contractionary monetary policy shock in LENS and in ToTEM. The first row displays the impulse responses of inflation, while the second row displays the impulse responses of output (GDP). The prolonged output contraction, which is far more pronounced in LENS, captures the real cost that a central bank has to tolerate for inflation to come off. In our baseline measure of the sacrifice ratio, we sum period output losses up until inflation peaks. In LENS, a reduction of 1 percentage point in inflation is achieved after seven quarters and the sacrifice ratio is 22.41. In ToTEM, the 1 percentage point drop in inflation is achieved after four quarters, and the sacrifice ratio is 8.91. As alternative measures, we sum output losses at 4, 8 and 12 quarters. We find corresponding sacrifice ratios of 11.39, 18.90 and 25.65 for LENS, and 8.91, 16.48 and 11.39 for ToTEM.[12] All of these figures are collected in **Table** 1. Similar to the baseline case, the output cost of curbing inflation is consistently larger for LENS than for ToTEM. Differences across models are driven by different assumptions about expectations and pricing behaviour, which imply a larger weight on past inflation realizations in the Phillips curve featured by LENS. As well, sacrifice ratios increase with the horizon in both models, being positive and bounded above, although their growth rates fall over time as output gradually returns to its steady-state level.

We also compute Phillips multipliers like in Barnichon and Mesters (2021). For the sake of comparison with sacrifice ratios, we consider 4-, 8- and 12-quarter horizons. As an example,

11. We let the size and persistence of the exogenous shock differ across models to make comparing sacrifice ratios easier. The size of the shock is adjusted to normalize the response of inflation, whereas the persistence of the shock is set such that the persistence of the nominal interest rate is the same. Hence, differences across models are entirely driven by their structure, not by differences in the process of the shock.

12. As a reference, Tetlow (2022), who computes the sacrifice ratio associated with a *permanent* reduction in inflation for 40 different structural models of the US economy, finds a mean (median) estimate of 16 (9) with a large standard deviation of 16.

the Phillips multiplier at four quarters for ToTEM is 0.27. Hence, sacrificing a percentage point of GDP brings about a 0.27% reduction in inflation over four quarters. **Table 1** shows that the Phillips multipliers are uniformly smaller in LENS than in ToTEM. After recalling that the Phillips multiplier and the sacrifice ratios are inversely related at a given horizon, one can conclude that these findings are in line with those for sacrifice ratios.

Although the two models present a very similar picture when either of the two trade-off measures is used, there are some unavoidable differences due to their different dynamics. It would thus be of interest to compare both models with the data and, with that comparison in mind, to reassess the economic and statistical significance of such differences. We focus on this topic next.

Table 1: Sacrifice ratios and Phillips multipliers in ToTEM and LENS

	Sacrifice ratios		Phillips multipliers	
Horizon	ToTEM	LENS	ToTEM	LENS
Baseline	8.91	22.41	n/a	n/a
4 quarters	8.91	11.39	0.27	0.12
8 quarters	16.48	25.65	0.35	0.20
12 quarters	18.13	35.49	0.44	0.23

Notes: ToTEM and LENS are two large-scale open-economy models that capture salient features of the Canadian economy. ToTEM is a DSGE model described in Corrigan et al. (2021). LENS is a model following the methodology similar to that of the Federal Reserve Board's FRB/US model; for details, see Gervais and Gosselin (2014).

4 Identification challenges

Since both ToTEM and LENS rely on economic theory to some extent, they identify the parameters of interest by imposing several restrictions on the data. As a consequence, they are exposed to the risk of misspecification, which would bias our estimates of the sacrifice ratio. The use of alternative estimation techniques that are more agnostic about

Figure 1: Impulse response to a one-period contractionary monetary policy shock

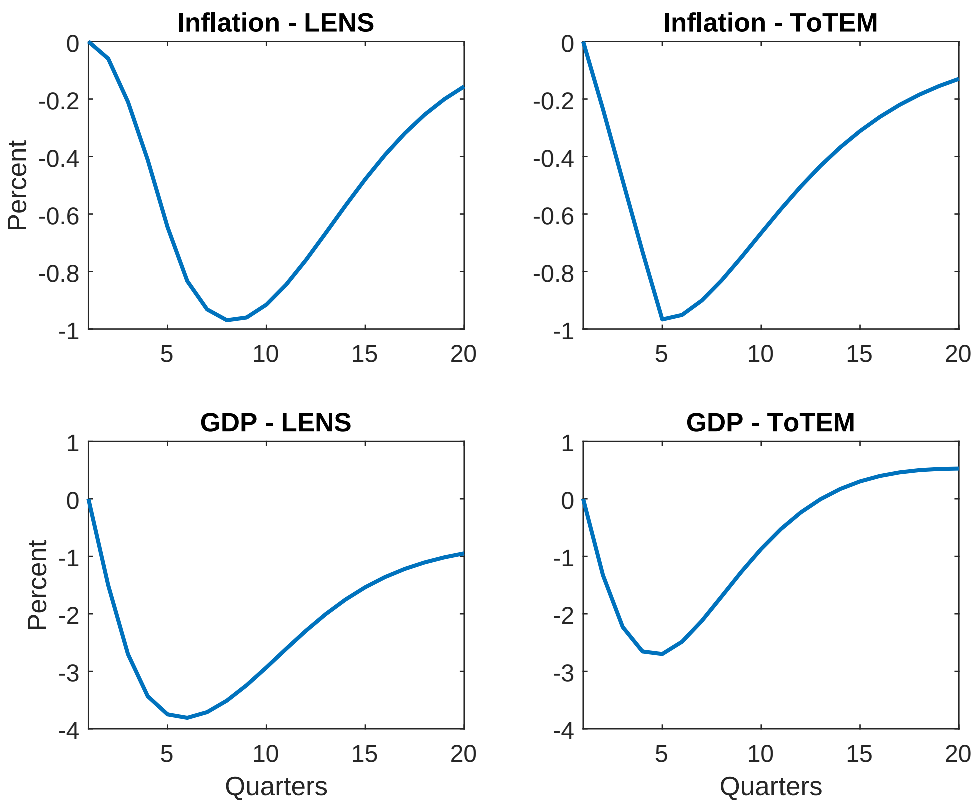

Notes: This figure shows the impulse responses of inflation and gross domestic product (GDP) to a one-period contractionary monetary policy shock in LENS and ToTEM that lowers inflation by 1% at peak. The horizontal axes report the number of quarters since the shock hits, and the vertical axes report inflation and GDP in percentage deviations from the steady state for both models. ToTEM is a DSGE model described in Corrigan et al. (2021) and LENS is a model following the methodology similar to that of the Federal Reserve Board's FRB/US model (see Gervais and Gosselin 2014).

the structure of the economy would be a good complement to cross-check the results derived from our estimated policy models. One possibility is to estimate the Phillips curve alone, rather than the full model. Another possibility is to directly estimate the sacrifice ratios or the Phillips multipliers. These routes, however, are not free of challenges. A key difficulty is the presence of supply shocks that simultaneously affect inflation and the slack variable, biasing the estimates of the parameters. Magnifying this simultaneity bias is the systematic response of monetary policy, which tends to offset demand shocks.

The goal of this section is to present state-of-the-art techniques designed to overcome those challenges. We conclude that conditioning on identified monetary policy shocks and directly estimating measures of the trade-off without even specifying a Phillips curve model are the most promising avenues, which becomes clearer in what follows.

4.1 Endogeneity of slack and monetary policy

To illustrate the first key issue an econometrician faces when estimating the Phillips curve, let us consider the basic three-equation New Keynesian model presented in section 2. For simplicity, we consider a simpler Taylor rule by assuming $\Phi_x = 0$ and exclude monetary policy shocks. After substituting i_t in the investment-saving (IS) equation (6) using the monetary policy rule (equation7), the model reduces to a system of two equations corresponding to aggregate demand (AD) (8) and supply (AS) (9) curves:

$$\pi_t = -\frac{\sigma}{\Phi_\pi}(x_t - E_t x_{t+1} - w_t), \tag{8}$$

$$\pi_t = \kappa x_t + v_t, \tag{9}$$

where $v_t \equiv \beta E_t \pi_{t+1} + \mu_t$ and $w_t \equiv \sigma^{-1} E_t \pi_{t+1} + \varepsilon_t$.[13]

13. ε_t and μ_t are assumed to follow exogenous AR(1) processes with persistence ρ_ε and ρ_μ, respectively. That is, $\varepsilon_t = \rho_\varepsilon \varepsilon_{t-1} + e_t$ and $\mu_t = \rho_\mu \mu_{t-1} + \eta_t$, where e_t and η_t are independent and identically distributed with mean zero and variances given by σ_ε^2 and σ_μ^2, respectively.

Figure 2 plots the AD and AS curves under different scenarios, maintaining the simplifying assumption that all structural shocks are distributed independently and identically. In equilibrium, only the intersection of these two curves is observed. Demand shocks shift the aggregate demand curve, leading to different equilibrium outcomes that trace out the Phillips curve (**Figure 2**, panel 2a). In this case, absent supply shocks, the slope of the Phillips curve can be consistently estimated.

However, in response to supply shocks, the Phillips curve itself shifts, leading to movements of the equilibrium along the aggregate demand curve (**Figure 2**, panel 2b). Since supply shocks yield a negative correlation between inflation and the output gap, the estimated slope of the Phillips curve is negatively biased.[14] As supply shocks become relatively more volatile, the correlation between inflation and economic activity further decreases. If only supply shocks were present, the estimated slope coefficient would be negative, even though the structural relation is positive, and the slope of the Phillips curve would not be identified.

As illustrated in panels 2c and 2d of **Figure 2**, the correlation between slack and inflation may also decrease due to:

- a flattening of the Phillips curve (i.e., a lower κ (slope hypothesis))
- a more aggressive stance of monetary policy to stabilize inflation (i.e., a higher Φ_π (policy hypothesis))

This configuration makes identifying the Phillips curve more complicated.[15]

To discriminate between these two underlying causes, one approach is to study the impulse response functions of inflation and output to demand shocks, which are represented in **Figure 3** for the more general case of serially correlated structural disturbances. With a flatter Phillips curve, a shift in demand has a larger impact on output but a smaller effect on inflation because the policy rate needs to move very little. In contrast, when monetary policy leans

14. See Appendix C for a derivation of the Ordinary Least Square estimator and the sources of biases.

15. In **Appendix A**, we further illustrate these challenges using simulated data from this theoretical model.

Figure 2: Graphical representation of aggregate demand and supply curves

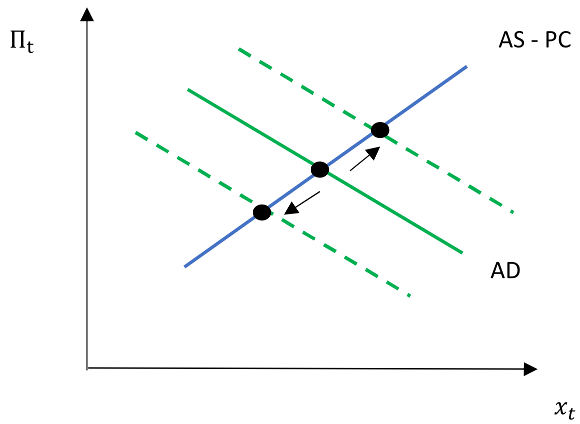

(a) Only demand shocks

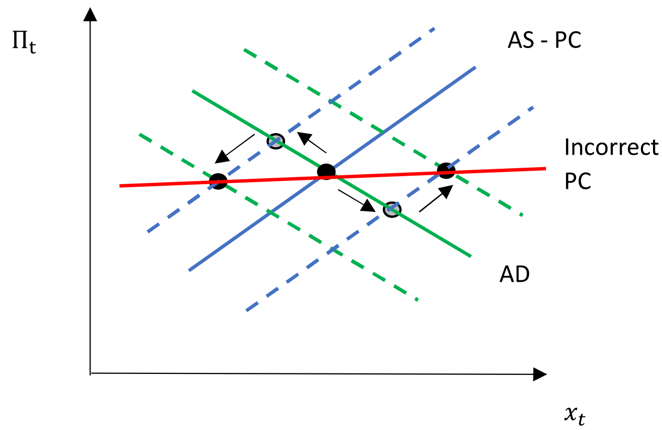

(b) Both demand and supply shocks

Note: These graphs plot the aggregate demand (AD) and aggregate supply (AS) curves, that is, the Phillips curve (AS - PC). Demand shocks shift the AD curve, while supply shocks shift the AS curve. Due to the presence of supply shocks, different equilibrium outcomes trace out an incorrect Phillips curve.

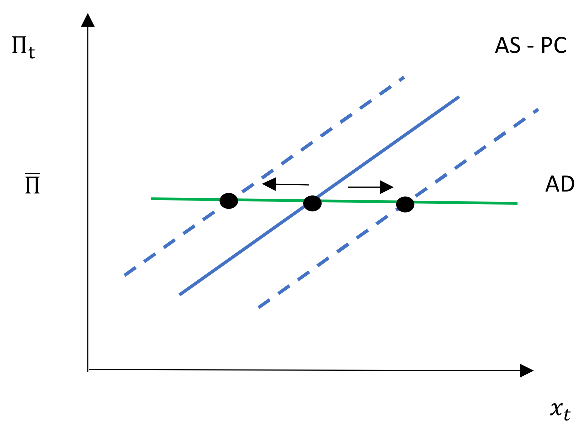

(c) Flat aggregate demand curve

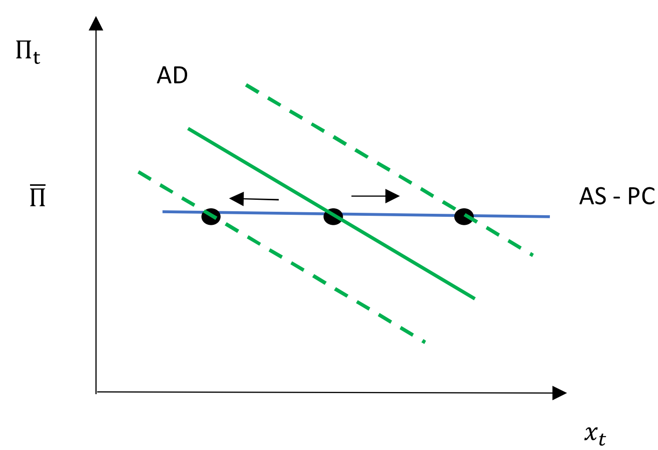

(d) Flat Phillips curve

Note: Either a flat aggregate demand curve or a flat Phillips curve may result in stable inflation.

heavily against inflation, a demand shock has a smaller effect on both output and inflation compared with the baseline.[16] Using this insight to study the sources of the disconnect between inflation and economic activity in the United States, Del Negro et al. (2020) find that the response of inflation to demand disturbances has become much weaker than the responses of real variables, which suggests that the Phillips curve has flattened.

Figure 3: Impulse response functions of inflation and output gap to demand (first row) and supply (second row) shocks

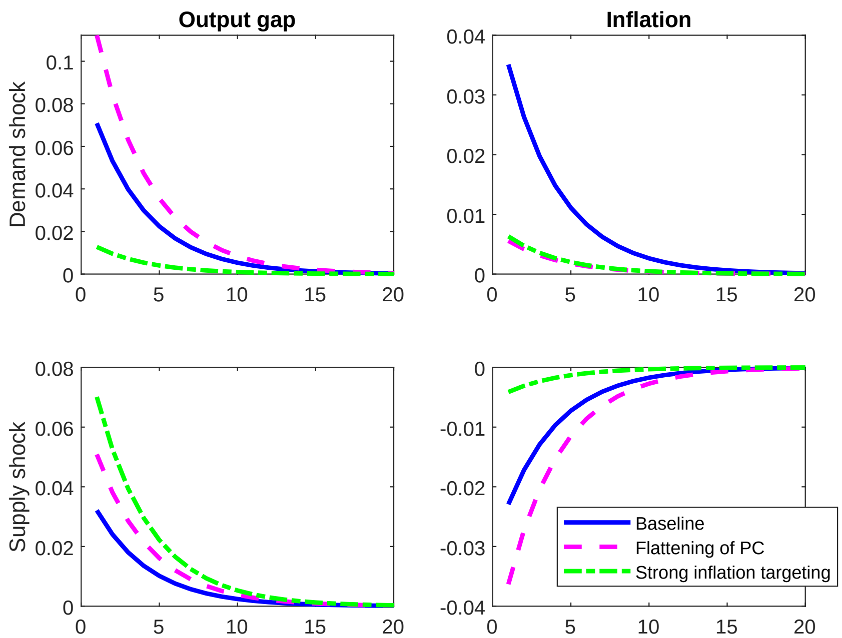

4.2 Reduced-form estimation strategies

The literature proposes several strategies to identify the output cost of stabilizing inflation. We discuss them in what follows.

Controlling for supply shocks. To isolate the demand variation in the slack variable that correctly identifies the Phillips curve and to limit the bias in the estimates, one approach

16. See **Appendix B** for more details.

is to control for the effects of supply shocks. Using the core measure of inflation rather than headline inflation when estimating the Phillips curve partially alleviates the influence of cost-push shocks because prices of food and energy are more impacted by such shocks. Gordon (2013) emphasizes the importance of supply shocks and estimates a triangle model developed to account for the shift in inflation dynamics due to the 1973–75 oil shock. Compared with the New Keynesian Phillips curve, the model also includes long lags on the dependent variable and a set of explicit variables to represent the supply shocks, namely food and energy price inflation, relative import price inflation, changes in trend labour productivity and dummies for the effect of the 1971–74 Nixon-era price controls.

However, this approach would require controlling for many other trade-off shocks. For instance, according to Gilchrist et al. (2017), the lack of disinflation during the Great Recession could be explained by the intensification of financial frictions that created an incentive for constrained firms to raise prices to avoid costly external financing. Therefore, financial frictions may also induce supply-side effects and would need to be controlled for. In addition, the relative importance of these trade-off-inducing shocks likely varies over time.

Using instrumental variables. Another approach to identify the causal impact of slack on inflation uses instrumental variables to identify movements along the Phillips curve. A valid instrument should be correlated with the demand variation in the output gap and be orthogonal to supply shocks. A first-stage regression of the slack variable on the instrument aims to isolate the variation in slack that is due to demand shocks. The Phillips curve is then estimated using the predicted series as the regressor. Notably, Barnichon and Mesters (2021) estimate an instrumental variable regression of cumulative inflation on cumulative unemployment, where monetary policy shocks are used as instruments.

Using cross-sectional regional variation. Another approach to identify the Phillips curve is to use regional data. The analysis is typically conducted using panel regressions of regional inflation on regional unemployment rates, with time-fixed effects that control for common trends—in particular due to aggregate shocks—and region-fixed effects that control for time-invariant cross-sectional heterogeneity.

The motivation behind this strategy is that deviations of regional variables from the national average should be independent from monetary policy. Indeed, monetary policy responds to national economic conditions but cannot offset idiosyncratic regional shocks. In addition, data at the regional level have more variation than aggregate data. Thus, this approach helps alleviate the endogeneity and lack of power issues that are associated with estimating aggregate time series. Several papers use city- or state-level data in the United States to estimate the slope of regional Phillips curves (see, for example, Fitzgerald and Nicolini 2014; Hooper, Mishkin, and Sufi 2020; McLeay and Tenreyro 2020; Hazell et al. 2022). Hazell et al. (2022) also show that this approach provides a solution to the problem of shifting values of long-term inflation expectations confounding the estimation of the slope of the Phillips curve. Indeed, these long-run changes can be absorbed by time-fixed effects when using a panel data specification since they are likely similar across regions.

5 Reduced-form evidence for Canada

In this section, we apply to Canadian data what we see as the two most promising empirical strategies proposed in the literature. The results complement the findings based on our policy models and add to a growing literature that applies similar methods to other countries.

5.1 Phillips multipliers

We follow Barnichon and Mesters (2021) and estimate Phillips multipliers for Canada. The Phillips multiplier is a statistic to non-parametrically characterize the output-inflation trade-off. At each horizon h, the Phillips multiplier is defined by

$$PM_h = IRF_h^{\bar{\pi}} / IRF_h^{\bar{y}}, h = 0, 1, 2, ..., \tag{10}$$

where $IRF_h^{\bar{\pi}}$ and $IRF_h^{\bar{y}}$ are the impulse-response functions of average inflation and output, respectively, to a one-unit monetary policy shock, ε_t. We use the vector autoregression (VAR) model in Champagne and Sekkel (2017) to estimate the impulse-response functions of GDP and inflation to a monetary policy shock that is identified through narrative methods, as in C. D. Romer and D. H. Romer 2004. The VAR contains data on GDP, total consumer price index (CPI) inflation, the Bank of Canada commodity price index (BCPI) and the Bank rate, which is instrumented by the narrative monetary policy shocks. The model contains 12 lags and is estimated with data from January 1992 to December 2019.

The Phillips multiplier statistic overcomes some of the challenges described in the previous section by using monetary policy shocks as an instrument to identify exogenous variation in real activity. Furthermore, the method does not rely on a model—not even of a Phillips curve—and it thus avoids specification uncertainty. In a sense, from a methodological perspective, it parallels the literature on fiscal multipliers and complements a model-based approach with an instrumental-variable based approach. In this respect, Ramey and Zubairy (2018) constitute a prominent example and Ramey (2016) provides an exhaustive review of existing studies in that field.

Results are displayed in the third column of **Table 2**, which also reports the Phillips multipliers based on our estimated policy models in the first two columns. Similar to Barnichon and Mesters (2021), multipliers are not significant at short horizons. However, they become significantly positive and reach about 0.27 after four quarters. This points to a stronger trade-off than in the United States if we compare with the more sizeable estimates of Phillips multipliers from Barnichon and Mesters (2021) using US data. Champagne and Sekkel (2017) also find a weaker price level response to monetary policy shocks in Canada than in the United States and the United Kingdom. A reason for this could be that the Bank of Canada's conduct of monetary policy has been more predictable, making true monetary policy shocks rarer and limiting the exogenous variation, thus weakening the instrument. Relative to the data, the models do not seem far off. The larger output-inflation trade-off embedded in LENS compared with ToTEM seems closer to the data. Overall, the story told by our models aligns well with the empirical estimates.

Table 2: Phillips multipliers - Bank of Canada models and empirical estimates

Horizon	ToTEM	LENS	Empirical
4 quarters	0.27	0.12	0.01
			(-0.26, 0.28)
8 quarters	0.35	0.20	0.18
			(0.01, 0.45)
12 quarters	0.44	0.23	0.27
			(0.09, 0.59)

Note: The numbers in parentheses for the empirical estimates are the 68% bootstrapped confidence intervals.

5.2 Regional variation

Another empirical approach used to identify the slope of the Phillips curve is by exploiting regional variation. This approach is described by McLeay and Tenreyro (2020), who apply it to US metro-area data. While monetary policy seeks to undo demand shocks at the national level, it does not respond to idiosyncratic regional shocks. Therefore, this identification strategy reduces the endogeneity issue arising from monetary policy and aggregate supply shocks. We apply the approach to Canada using provincial data. Although it does not speak directly to our models, this exercise is useful in placing the Canadian case into an international context.

Panel regressions are run using data from the 10 Canadian provinces. Quarter-over-quarter CPI inflation excluding food and energy, seasonally adjusted at annual rates, is regressed on the seasonally adjusted, quarterly unemployment rate and controls. The sample is from 1990 to 2019. Controls include a lag of inflation and a constant. Two specifications of the panel regressions are used: i) pooled, and ii) regional- and time-fixed effects. A simple regression using aggregate Canadian data only is also shown for comparison.

Results are shown in **Table 3**. As expected, using regional data leads to better identification of the slope than the aggregate ordinary least squares (OLS) equation. The slope of the aggregate OLS equation is not statistically significant and has a positive sign. In comparison, the panel analysis yields a negative slope, both under the pooled and fixed effects specifications. When including fixed effects, the slope is steeper and statistically significant at the 5% level of confidence.

Table 3: Phillips curves for Canadian provinces

Variables	OLS on aggregate (Canada)	Pooled	Regional
Inflation(-1)	0.48***	0.38***	0.17***
	(0.00)	(0.00)	(0.00)
Unemployment rate	0.021	-0.00032	-0.069**
	(0.78)	(0.98)	(0.04)
Constant	0.76	1.1***	2.1***
	(0.20)	(0.00)	(0.00)
Observations	120	1,200	1,200
Adjusted R-squared	0.23	0.15	0.62
Time-fixed effects	No	No	Yes

Note: The sample is 1990 to 2019 and includes data for the 10 provinces. *, ** and *** indicate significance levels at 10%, 5% and 1%, respectively. Inflation is quarterly consumer price index (CPI) excluding food and services seasonally adjusted at annual rates. The unemployment rate is quarterly seasonally adjusted unemployed persons over the labour force. Quarterly values are obtained from the average of monthly values.

Our results suggest a steeper coefficient on slack for Canada than in the euro area compared with Eser et al. (2020), who use a similar methodology. They estimate a coefficient of -0.01 (p-value = 0.001) when using the unemployment rate as the measure of slack with regional and time-fixed effects using data from 18 euro area countries from the first quarter of 1999 to the second quarter of 2019. Meanwhile, our results suggest a flatter coefficient on slack than for the United States compared with McLeay and Tenreyro (2020), who also use a similar methodology.[17] They use data from 28 US metropolitan areas from the first half of 1990 to the second half of 2017 and estimate a coefficient of -0.379 (p-value = 0.052) when including regional- and time-fixed effects. Finally, exploiting a similar regional analysis but using the old Keynesian framework, Fitzgerald and Nicolini (2014) also find a coefficient for the United States of around -0.3, which is relatively stable across models and periods.

The panel regression specifications address some identification challenges, although some limits remain. Inflation expectations and the natural rate of unemployment are omitted variables. Regional-fixed effects control for differences across the cross-section, while time-fixed effects control for time-varying change. In addition, the lag of inflation can also help control for inflation expectations.[18] [19] However, fixed effects will not capture any cross-sectional differences in inflation expectations or in the natural rate of unemployment that vary over time.

Hazell et al. (2022) raise some caveats to this approach. They show that the slope of regional Phillips curves for non-tradeables would have the same slope as the aggregate Phillips curve because tradeables are priced nationally rather than regionally. The slope of the regional Phillips curves with total CPI would hence be smaller. However, they also find that estimates of the regional Phillips curve slopes tend to be larger than estimates based on aggregate data. Their model shows that to estimate the true slope, one would have to control for one-period-ahead regional inflation expectations or use the truncated sum of unemployment rates. In their paper, the authors use the latter approach. This suggests the true slope for Canada could be even smaller than reported in **Table 3**.[20] Nonetheless,

17. Their specification also includes 12-month inflation expectations from the University of Michigan Consumer Survey for four regions.

18. The lag of inflation can also control for backward-looking pricing behaviour on the part of firms, as in the ToTEM and LENS Phillips curves described in section 2.

19. Hazell et al. (2022) argue that not controlling for one-period-ahead regional inflation expectations will bias the estimates. However, provincial-level inflation expectations are not available for Canada. McLeay and Tenreyro (2020) also do not control for one-period-ahead inflation expectations for each metropolitan area, although they include 12-month inflation expectations from the University of Michigan Consumer Survey for four regions and match these to the metropolitan areas. Eser et al. (2020) do not control for inflation expectations.

20. While the approach of Hazell et al. (2022) cannot be replicated for Canada because regional inflation excluding tradables is not available, a test using generalized method of moments (GMM) on the truncated

the specification using inflation excluding food and energy is comparable with McLeay and Tenreyro (2020) and Eser et al. (2020). Therefore, these results can be compared with the United States and euro area.

The panel regressions using data from 10 Canadian provinces lead to a slope of the Phillips curve that is steeper and more statistically significant—especially when including regional- and time-fixed effects—when compared with a simple OLS regression on aggregate Canadian data. Thus, the identification strategy described in this section is helpful to identify the slope of the Phillips curve for Canada and supports the finding in LENS and ToTEM that the slope of the Canadian Phillips curve is quite small.

6 Conclusion

This paper provides a comprehensive summary of existing academic research and the Bank of Canada's in-house policy models, and adds new empirical evidence for Canada regarding the trade-off between output and inflation.

We quantify the trade-off embedded in the Bank's two main in-house estimated policy models, ToTEM and LENS, and complement these results with reduced-form evidence that hinges on a far lighter battery of assumptions regarding the structure of the economy. Following state-of-the-art techniques, such as using monetary policy shocks as an instrument to identify exogenous variations in real economic activity or using regional variations, we find that the conclusions reached by our policy models line up well with the data.

Our analysis, however, also raises several caveats that lay the groundwork for future research. Recent events, including but not limited to the fast-evolving world economy and geopolitical developments since the COVID-19 pandemic, have prompted Bank staff to work on a new generation of models, following the call to action and strategy proposed in Coletti (2023). Some of the most pressing issues are:

sum of unemployment rates with the four-quarter-lag of unemployment as an instrument suggests an even flatter slope than in **Table 3**.

- the complex and evolving Canadian production network through which important structural shocks (such as oil price fluctuation) propagate and influence inflation

- unexplored drivers of markups that reflect state-dependent firms' pricing behaviour over the business cycle and the resulting non-linearity in the Phillips curve

- a rethinking of how to best model inflation expectations

- agents learning in complex economic environments

Another, even more radical, challenge ahead is reconsidering the commonly held assumption that monetary policy does not affect potential (for example, see Moran and Queralto 2018; Jordà, Singh, and Taylor 2020; Baqaee, Farhi, and Sangani 2021; Garga and Singh 2021; Meier and Reinelt 2022).

To conclude, while striving to give the most up-to-date picture of the output-inflation trade-off in Canada, we also hope that all the models described here will soon become obsolete and be replaced by a new and improved generation. Rather than representing a finish line, this analysis really has to be seen as a beginning.

Appendix

A Simulation of the basic New Keynesian model

To further illustrate the key identification challenges, we generate artificial data from the simple model described in section 3, subject to a sequence of supply and demand shocks over 500 periods. The parameters in the baseline scenario are calibrated as follows:

- the slope of the Phillips curve is set to $\kappa = 0.1275$
- the weight on inflation in the monetary policy rule is set to $\Phi_\pi = 1.1$
- the discount factor β is set to 0.99
- the risk aversion coefficient σ is set to 1
- the persistence of demand and supply shocks is set to $\rho_\varepsilon = \rho_\mu = 0.75$

In **Figure A.4**, we plot artificial data generated conditional on demand shocks only (left column) and on both demand and supply shocks (right column). In the former case, as explained in section 4.1, the data points trace out the true Phillips curve. For each plot, the baseline case is depicted in blue circles. It is compared with alternative scenarios in red triangles:

- an increased volatility of supply shocks (first row)
- a reduced slope of the Phillips curve (second row)
- a stronger weight on inflation stabilization in the monetary policy rule (third row)

In each case, the scatter plots are augmented with regression lines.

As illustrated in plots in the first row, the correlation between inflation and economic activity declines in the presence of supply shocks as they move the equilibrium output gap and inflation in opposite directions. Thus, the observed Phillips curve appears flatter and the estimate of the slope from a naive least-square regression of inflation on the output gap is negatively biased. This bias is amplified as the volatility of supply shocks increases.

Plots in the second row show that the correlation between inflation and the output gap decreases when the Phillips curve flattens, with more variability in the output gap than in inflation. Indeed, demand shocks have much smaller effects on inflation under a flat Phillips curve. However, because supply shocks have larger effects on both output and inflation, the distribution of both variables becomes more dispersed if the volatility of supply shocks increases. Fluctuations in the output gap are further magnified if the aggregate demand curve also flattens.

Plots in the third row of **Figure A.4** illustrate that under stricter inflation targeting, monetary policy offsets demand shocks even more. Therefore, supply shocks largely drive the remaining variation in the data. This reinforces the simultaneity problem and leads to a seemingly flattened Phillips curve, as explained in McLeay and Tenreyro (2020). An increased weight on stabilizing inflation contributes to a flattening of the aggregate demand curve, which leads to smaller fluctuations in inflation after aggregate shocks. In response to supply shocks, the monetary authority prefers to create output gap fluctuations to stabilize inflation, which results in larger volatility of output after these shocks.

B Solution to the static version of the model

Here we present the solution to the static version of the basic three-equation New Keynesian model described in section 3. To this end, we assume that the shocks ε_t and μ_t are independently and identically distributed, with variances σ_ε^2 and σ_μ^2, respectively. Thus $E_t\pi_{t+1} = E_t x_{t+1} = 0$. For simplicity, we assume $\Phi_x = 0$. Then, the model (6)-(7) can be rewritten:

Figure A.4: Simulated data

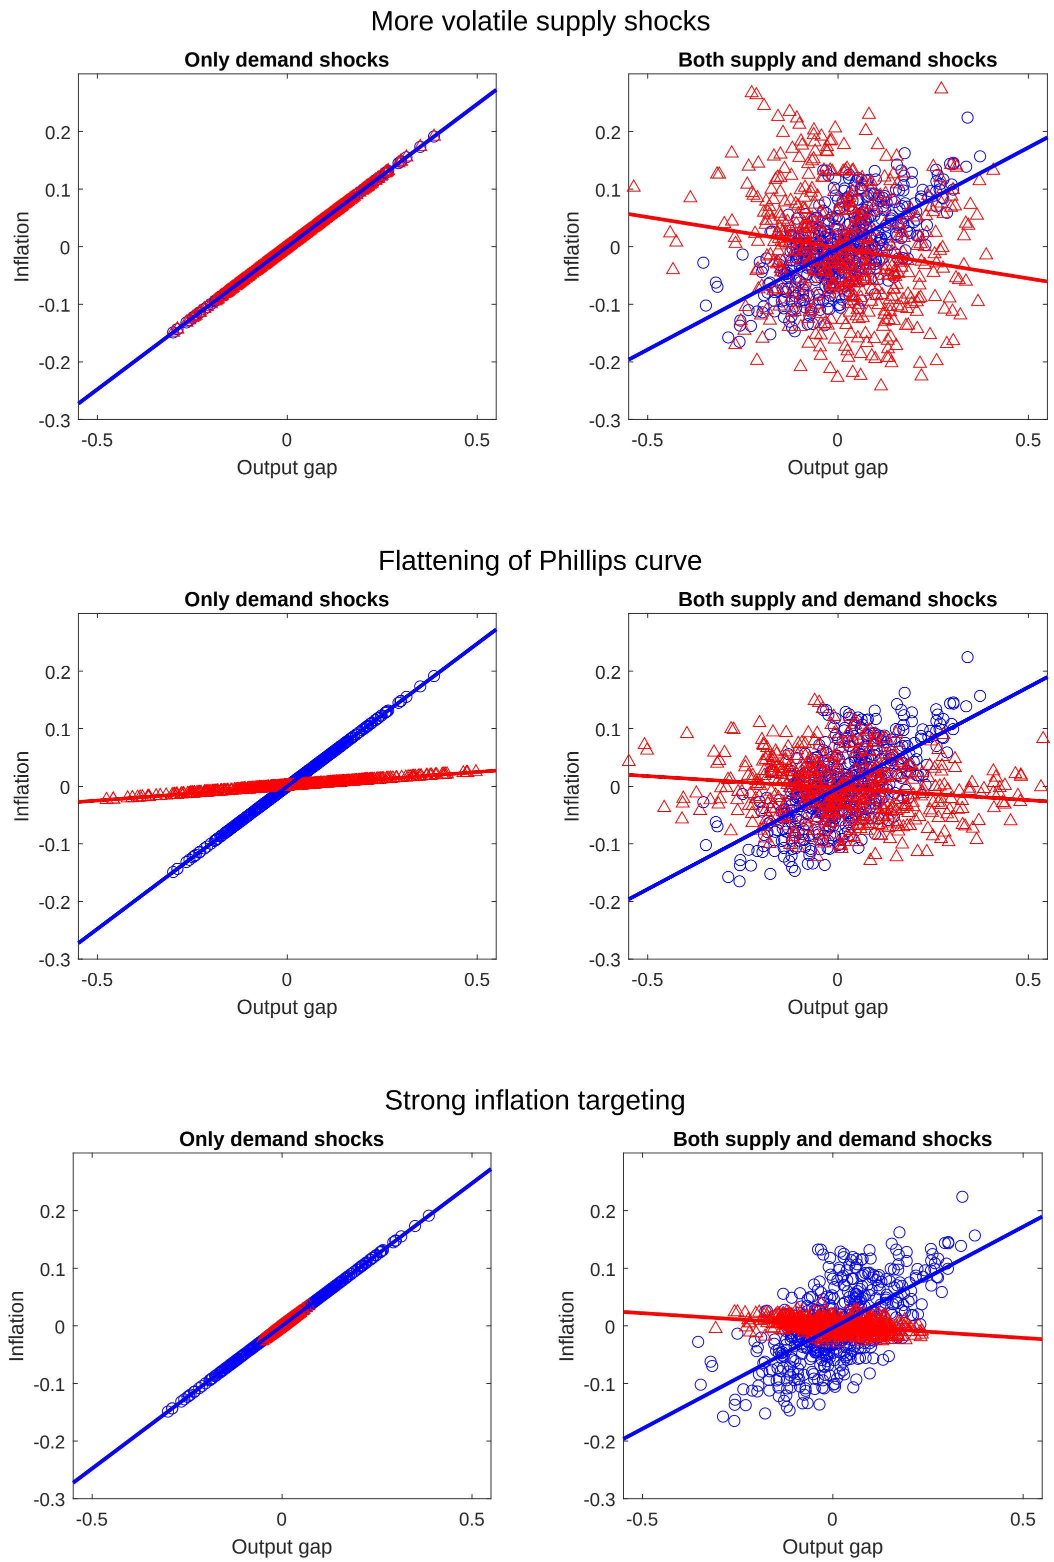

$$x_t = -\sigma^{-1}(\Phi_\pi \pi_t) + \varepsilon_t \tag{B.11}$$

$$\pi_t = \kappa x_t + \mu_t \tag{B.12}$$

Solving this system gives the equilibrium output gap and inflation rate, which are linear combinations of the supply and demand disturbances:

$$x_t = \frac{1}{\Omega}\varepsilon_t - \frac{\Phi_\pi}{\sigma\Omega}\mu_t \tag{B.13}$$

$$\pi_t = \frac{\kappa}{\Omega}\varepsilon_t + \frac{1}{\Omega}\mu_t \tag{B.14}$$

where $\Omega \equiv 1 + \sigma^{-1}\kappa\Phi_\pi$.

These expressions show that demand shocks, ε_t, lead to a positive correlation between inflation and the output gap, while supply shocks, μ_t, induce a negative correlation. Thus, the correlation between inflation and the output gap is not informative on any of the parameters if the shocks are not observed.

However, if the Phillips curve flattens—that is, κ decreases—then x_t becomes more responsive to both supply and demand shocks, while π_t becomes more responsive to supply shocks and less responsive to demand shocks. In contrast, if monetary policy leans more heavily against inflation—meaning Φ_π increases—then x_t becomes less responsive to demand shocks and more responsive to supply shocks, whereas π_t becomes less responsive to both supply and demand shocks. This provides a way to overcome the identification challenge.

C Ordinary least square estimator

Using the analytical solution derived in **Appendix B**, the ordinary least square estimator of the slope coefficient is given by:

$$\kappa^{OLS} = \frac{cov(\pi_t, x_t)}{var(x_t)} = \frac{\kappa\sigma^2 - \Phi_\pi \sigma \sigma_\mu^2 / \sigma_\varepsilon^2}{\sigma^2 + \Phi_\pi^2 \sigma_\mu^2 / \sigma_\varepsilon^2} \tag{C.15}$$

The estimator is therefore downwardly biased in the presence of supply shocks, that is if $\sigma_\mu^2 > 0$. The bias is amplified when monetary policy reacts more strongly to deviations of inflation from target (i.e., if Φ_π increases), and when supply shocks become more volatile relative to demand shocks (i.e., if $\sigma_\mu^2/\sigma_\varepsilon^2$ increases).

References

Adam, K., D. Matveev, and S. Nagel. 2021. "Do Survey Expectations of Stock Returns Reflect Risk Adjustments?" *Journal of Monetary Economics* 117:723–740.

Angeletos, G.-M., and C. Lian. 2022. "Dampening General Equilibrium: Incomplete Information and Bounded Rationality." In *Handbook of Economic Expectations,* edited by R. Bachmann, G. Topa, and W. van der Klaauw. London, UK: Elsevier Science.

Ascari, G., and T. Haber. 2022. "Non-Linearities, State-Dependent Prices and the Transmission Mechanism of Monetary Policy." *Economic Journal* 132 (641): 37–57.

Ball, L. 1994. "What Determines the Sacrifice Ratio?" In *Monetary Policy,* edited by N. G. Mankiw, 155–193. Chicago, IL: The University of Chicago Press.

Ball, L., and D. Romer. 1990. "Real Rigidities and the Non-Neutrality of Money." *Review of Economic Studies* 57 (2): 183–203.

Baqaee, D., E. Farhi, and K. Sangani. 2021. "The Supply-Side Effects of Monetary Policy." *National Bureau of Economic Research Working Paper,* no. 28345.

Barnichon, R., and G. Mesters. 2021. "The Phillips Multiplier." *Journal of Monetary Economics* 117:689–705.

Bils, M., and P. J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (5): 947–985.

Blanchard, O., and J. Galí. 2010. "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment." *American Economic Journal: Macroeconomics* 2 (2): 1–30.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review* 110 (9): 2748–2782.

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

Cecchetti, S. G., and R. W. Rich. 2001. "Structural Estimates of the US Sacrifice Ratio." *Journal of Business & Economic Statistics* 19 (4): 416–427.

Champagne, J., and R. Sekkel. 2017. "Changes in Monetary Regimes and the Identification of Monetary Policy Shocks: Narrative Evidence from Canada." *Bank of Canada Staff Working Paper,* no. 2017-39.

Coibion, O., and Y. Gorodnichenko. 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8): 2644–2678.

Coibion, O., Y. Gorodnichenko, and R. Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature* 56 (4): 1447–1491.

Coletti, D. 2023. "A Blueprint for the Fourth Generation of Bank of Canada Projection and Policy Analysis Models." *Bank of Canada Staff Discussion Paper,* no. 2023-23.

Corrigan, P., H. Desgagnés, J. Dorich, V. Lepetyuk, W. Miyamoto, and Y. Zhang. 2021. "ToTEM III: The Bank of Canada's Main DSGE Model for Projection and Policy Analysis." *Bank of Canada Technical Report,* no. 119.

Costain, J., A. Nakov, and B. Petit. 2022. "Flattening of the Phillips Curve with State-Dependent Prices and Wages." *Economic Journal* 132 (642): 546–581.

Cotton, C., and V. Garga. 2022. "The Role of Industrial Composition in Driving the Frequency of Price Change." *Federal Reserve Bank of Boston Research Department Working Paper,* no. 22-9.

De Paoli, B. 2009. "Monetary Policy and Welfare in a Small Open Economy." *Journal of International Economics* 77 (1): 11–22.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti. 2020. "What's Up with the Phillips Curve?" *Brookings Papers on Economic Activity* (Spring):301–373.

Eichenbaum, M., N. Jaimovich, and S. Rebelo. 2011. "Reference Prices, Costs, and Nominal Rigidities." *American Economic Review* 101 (1): 234–262.

Engel, C. 2011. "Currency Misalignments and Optimal Monetary Policy: A Reexamination." *American Economic Review* 101 (6): 2796–2822.

Eser, F., P. Karadi, P. R. Lane, L. Moretti, and C. Osbat. 2020. "The Phillips Curve at the ECB." *Manchester School* 88 (S1): 50–85.

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

Farhi, E., and I. Werning. 2019. "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review* 109 (11): 3887–3928.

Ferrero, A., and M. Seneca. 2019. "Notes on the Underground: Monetary Policy in Resource-Rich Economies." *Journal of Money, Credit and Banking* 51 (4): 953–976.

Fitzgerald, T. J., and J. P. Nicolini. 2014. "Is There a Stable Relationship Between Unemployment and Future Inflation? Evidence from U.S. Cities." *Federal Reserve Bank of Minneapolis Working Paper* 713.

Forbes, K., J. Gagnon, and C. G. Collins. 2021. "Low Inflation Bends the Phillips Curve around the World." *National Bureau of Economic Research Working Paper,* no. 29323.

Gabaix, X. 2020. "A Behavioral New Keynesian Model." *American Economic Review* 110 (8): 2271–2327.

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

Galí, J., and T. Monacelli. 2005. "Monetary Policy and Exchange Rate Volatility in a Small Open Economy." *Review of Economic Studies* 72 (3): 707–734.

Garga, V., and S. R. Singh. 2021. "Output Hysteresis and Optimal Monetary Policy." *Journal of Monetary Economics* 117:871–886.

Gervais, O., and M.-A. Gosselin. 2014. "Analyzing and Forecasting the Canadian Economy through the LENS Model." *Bank of Canada Technical Report,* no. 102.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek. 2017. "Inflation Dynamics during the Financial Crisis." *American Economic Review* 107 (3): 785–823.

Golosov, M., and R. Lucas. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (2): 171–199.

Gopinath, G., E. Boz, C. Casas, F. J. Díez, P.-O. Gourinchas, and M. Plagborg-Møller. 2020. "Dominant Currency Paradigm." *American Economic Review* 110 (3): 677–719.

Gopinath, G., and O. Itskhoki. 2022. "Dominant Currency Paradigm: A Review." In *Handbook of International Economics: International Macroeconomics,* edited by G. Gopinath, E. Helpman, and K. Rogoff, 6:45–90. Oxford, UK: Elsevier.

Gordon, R. J. 2013. "The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery." *National Bureau of Economic Research Working Paper,* no. 19390.

Harding, M., J. Lindé, and M. Trabandt. 2022. "Resolving the Missing Deflation Puzzle." *Journal of Monetary Economics* 126:15–34.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson. 2022. "The Slope of the Phillips Curve: Evidence from U.S. States." *Quarterly Journal of Economics* 137 (3): 1299–1344.

Hooper, P., F. S. Mishkin, and A. Sufi. 2020. "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?" *Research in Economics* 74 (1): 26–62.

Huang, K. X. D., and Z. Liu. 2005. "Inflation Targeting: What Inflation Rate to Target?" *Journal of Monetary Economics* 52 (8): 1435–1462.

Jordà, Ò., S. R. Singh, and A. M. Taylor. 2020. "The Long-Run Effects of Monetary Policy." *National Bureau of Economic Research Working Paper,* no. 26666.

Jørgensen, P. L., and S. H. Ravn. 2022. "The Inflation Response to Government Spending Shocks: A Fiscal Price Puzzle?" *European Economic Review* 141:103982.

Kimball, M. S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit, and Banking* 27 (4): 1241–1277.

Klenow, P. J., and O. Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123 (3): 863–904.

Kohlhas, A. N., and A. Walther. 2021. "Asymmetric Attention." *American Economic Review* 111 (9): 2879–2925.

Liu, C., and Y. Xie. 2022. "Understanding Inflation Dynamics: The Role of Government Expenditure." Available at SSRN: https://ssrn.com/abstract=4121796.

Mankiw, N. G. 2001. "The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment." *Economic Journal* 111 (471): 45–61.

McLeay, M., and S. Tenreyro. 2020. "Optimal Inflation and the Identification of the Phillips Curve." *NBER Macroeconomics Annual* 34 (1): 199–255.

Meggiorini, G. 2023. "Behavioral New Keynesian Models: An Empirical Assessment." *Journal of Macroeconomics* 77:103538.

Meier, M., and T. Reinelt. 2022. "Monetary Policy, Markup Dispersion, and Aggregate TFP." *Review of Economics and Statistics,* 1–45.

Moran, P., and A. Queralto. 2018. "Innovation, Productivity, and Monetary Policy." *Journal of Monetary Economics* 93:24–41.

Nakamura, E., and J. Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–1464.

Özden, T. Forthcoming. "Heterogeneous Expectations and the Business Cycle at the Effective Lower Bound." *International Journal of Central Banking.*

Ramey, V. A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics,* edited by J. B. Taylor and H. Uhlig, 2:71–162. Amsterdam: Elsevier.

Ramey, V. A., and S. Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy* 126 (2): 850–901.

Romer, C. D., and D. H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–1084.

Rubbo, E. 2020. "Networks, Phillips Curves and Monetary Policy." Available at https://scholar.harvard.edu/sites/scholar.harvard.edu/files/elisarubbo/files/rubbo_jmp.pdf..

Sims, C. A. 2010. "Rational Inattention and Monetary Economics." In *Handbook of Monetary Economics,* edited by B. M. Friedman and M. Woodford, 3:155–181. San Diego, CA: Elsevier.

Tetlow, R. J. 2022. "How Large is the Output Cost of Disinflation?"

Wei, S.-J., and Y. Xie. 2020. "Monetary Policy in an Era of Global Supply Chains." *Journal of International Economics* 124:103299.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ: Princeton University Press.

———. 2019. "Monetary Policy Analysis when Planning Horizons Are Finite." *NBER Macroeconomics Annual* 33 (1): 1–50.

Woodford, M., and Y. Xie. 2022. "Fiscal and monetary stabilization policy at the zero lower bound: Consequences of limited foresight." *Journal of Monetary Economics* 125:18–35.