

Mendes, Rhys; Murchison, Stephen; Wilkins, Carolyn A.

Working Paper

Monetary policy under uncertainty: Practice versus theory

Bank of Canada Staff Discussion Paper, No. 2017-13

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Mendes, Rhys; Murchison, Stephen; Wilkins, Carolyn A. (2017) : Monetary policy under uncertainty: Practice versus theory, Bank of Canada Staff Discussion Paper, No. 2017-13, Bank of Canada, Ottawa,
<https://doi.org/10.34989/sdp-2017-13>

This Version is available at:

<https://hdl.handle.net/10419/200449>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:


Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Staff Discussion Paper/Document d'analyse du personnel 2017-13

Monetary Policy Under Uncertainty: Practice Versus Theory



by Rhys Mendes, Stephen Murchison and Carolyn A. Wilkins

Bank of Canada Staff Discussion Paper 2017-13

November 2017

Monetary Policy Under Uncertainty: Practice Versus Theory

by

Rhys Mendes, Stephen Murchison and Carolyn A. Wilkins

Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
rmendes@bankofcanada.ca
smurchison@bankofcanada.ca
cwilkins@bankofcanada.ca

Acknowledgements

Many colleagues provided helpful comments on drafts of this paper. We are particularly grateful to Bob Amano, Gino Cateau, Jose Dorich, Oleksiy Kryvtsov, Tim Lane, Sylvain Leduc and Césaire Meh for their detailed comments. Laura Murphy provided excellent research assistance. All errors are our own.

Abstract

For central banks, conducting policy in an environment of uncertainty is a daily fact of life. This uncertainty can take many forms, ranging from incomplete knowledge of the correct economic model and data to future economic and geopolitical events whose precise magnitudes and effects cannot be known with certainty. The objective of this paper is to summarize and compare the main results that have emerged in the literature on optimal monetary policy under uncertainty with actual central bank behaviour. To this end, three examples are studied in which uncertainty played a significant role in the Bank of Canada's policy decision, to see how closely they align with the predictions from the literature. Three principles emerge from this analysis. First, some circumstances—such as when the policy rate is at risk of being constrained by the effective lower bound—should lead the central bank to be more pre-emptive in moving interest rates, whereas others can rationalize more of a wait-and-see approach. In the latter case, the key challenge is finding the right balance between waiting for additional information and not falling behind the curve. Second, the starting-point level of inflation can matter for how accommodative or restrictive policy is relative to the same situation without uncertainty, if there are thresholds in the central bank's preferences associated with specific ranges for the target variable, such as the risk of inflation falling outside of the inflation control range. Third, policy decisions should be disciplined, where possible, by formal modelling and simulation exercises in order to support robustness and consistency in decision making over time. The paper concludes with a set of suggested areas for future research.

Bank topics: Monetary policy; Uncertainty and monetary policy

JEL codes: E52, E58, E61, E65

Résumé

Pour les banques centrales, la conduite de la politique monétaire dans un contexte d'incertitude est une réalité quotidienne. Cette incertitude peut prendre de nombreuses formes, qui vont d'une connaissance incomplète du modèle et des données économiques appropriés aux phénomènes économiques et géopolitiques futurs dont il est impossible de cerner l'ampleur et les effets avec précision. Ce document vise à résumer les grandes conclusions que l'on trouve dans la littérature sur ce qui constitue la politique monétaire optimale en contexte d'incertitude, et à les comparer au comportement des banques centrales dans les faits. À cette fin, nous étudions trois exemples de situations où l'incertitude a joué un rôle important dans la décision de la Banque du Canada en matière

de politique monétaire, afin de voir dans quelle mesure la décision correspond aux prédictions de la littérature. Trois principes émergent de cette analyse. Premièrement, dans des circonstances particulières – par exemple, quand le taux directeur risque d’être soumis à la contrainte de la valeur plancher –, certaines sources considèrent que la banque centrale devrait adopter une attitude préventive dans l’établissement des taux d’intérêt, alors que d’autres justifient une approche attentiste. Si l’on choisit cette dernière approche, le principal défi consiste à trouver le juste équilibre, c’est-à-dire attendre des renseignements supplémentaires tout en évitant de réagir en décalage. Deuxièmement, le niveau de l’inflation au point de départ peut influencer sur la décision d’appliquer une politique monétaire plus ou moins expansionniste ou restrictive dans une même situation, en fonction de la présence ou de l’absence d’incertitude, si les préférences de la banque centrale comportent des seuils pour ce qui est des fourchettes particulières établies pour la variable cible (comme le risque que l’inflation sorte de la fourchette visée). Troisièmement, les décisions de politique monétaire devraient être encadrées, autant que possible, par des modélisations et des simulations en bonne et due forme, afin d’assurer la rigueur et la cohérence de la prise de décisions au fil du temps. Nous proposons, enfin, des pistes de sujets qui pourraient faire l’objet de recherches dans l’avenir.

Sujets : Politique monétaire ; Incertitude et politique monétaire

Codes JEL : E52, E58, E61, E65

"The so-called 'certainty equivalence' principle that uncertainty should be ignored and only the mean of each variable or parameter be used is essentially wrong."

Martin Feldstein, Jackson Hole 2003

1 Introduction

One of the pillars of modern thinking is that, if the right method is followed and the facts are observed, how the world works will be uncovered. This thinking led to many achievements: Newton's discoveries in physics, Hume's application of the scientific method to human behaviour, and the birth of modern economics with Adam Smith's *Wealth of Nations*. Yet, even as physical and social scientists hold a lingering desire for certainty, the spirit of Heisenberg's seminal uncertainty principle can be applied well past quantum physics.¹

Indeed, central banks have always had to address uncertainty in the conduct of monetary policy. For this reason, the Bank of Canada has continuously sought to reduce uncertainty by: (i) increasing the clarity around the objectives of monetary policy and the framework to achieve the objectives (i.e., the inflation targeting regime); (ii) leveraging increased computing power to invest heavily in more realistic economic models used for forecasting and policy analysis; and (iii) expanding sources of data and information from businesses, households and market participants.²

These efforts have been critical to supporting the ability of the Bank of Canada to achieve its inflation objectives and maintain its policy credibility. Yet no amount of research effort can completely eliminate uncertainty because of the sheer complexity of the real world. And while some of this uncertainty can be reasonably measured (e.g., errors around a near-term inflation forecast), some cannot be (e.g., risk of a geopolitical event). This latter case is known as "Knightian uncertainty." It is for this reason that central banks, including the Bank of Canada, apply a combination of evidence from data and formal models and informed judgment when making policy decisions. This

¹German physicist Werner Heisenberg articulated a principle in 1927 that the position and the velocity of an object could not both be simultaneously known.

²See Kozicki and Vardy (2017) for a fulsome review of approaches employed by the Bank of Canada to address uncertainty, including those related to communications and transparency.

approach strikes a balance between those who believe that central banks should base their decisions mainly on a simple policy rule that abstracts from uncertainty, and those who believe that economic models are of little use and central bankers must be dependent on judgment.

Rather than making "the best the enemy of the moderately useful" (Blinder 1998), this paper studies how the most influential parts of the literature on monetary policy under uncertainty can inform policy decisions in practice. To do this, we first review briefly the sources of uncertainty that a central bank faces (section 2). We then outline the standard approach in the literature in which uncertainty can be ignored by policy-makers (known as "certainty equivalence"), as well as the conditions under which this approach is not appropriate (section 3). We also study three relevant examples from the past decade and a half, in which uncertainty played a significant role in the Bank of Canada's policy decision, to see how closely the logic of the Bank's policy approach in these examples aligned with what this literature would advise (section 4). The final section draws conclusions and explores areas for further research.

Our analysis confirms that uncertainty can have profound effects on monetary policy, leading policy-makers to deviate substantially from what a typical monetary policy rule used by central banks would suggest. Some sources of uncertainty suggest that monetary policy should be more aggressive, while others can justify a wait-and-see approach. The art is in correctly judging the situation. The specific examples we study lead to the following three guiding principles for decision making under uncertainty in practice:

1. Context matters for whether uncertainty should lead monetary policy to be more preemptive or more gradual. On the one hand, the risk of being constrained by the effective lower bound (ELB) in the future, all else being equal, should lead policy-makers to lower interest rates more aggressively prior to encountering the ELB. Uncertainty regarding the effectiveness of unconventional monetary policy (UMP) contributes to explain this increased aggressiveness. On the other hand, heightened uncertainty about the outlook, especially the type that is more difficult to estimate and control for, can lead to a wait-and-see approach. One reason would be there are real-economy costs to policy reversals that outweigh the costs of delaying policy action. However, a wait-and-see approach can be taken too far. The crucial policy challenge is to find the right balance between waiting for additional information and not falling behind the curve. For large shocks or an accumulation of smaller ones,

the cost of waiting may simply be too high and decisive action will be appropriate, whereas for smaller ones a desire to avoid a policy reversal can justify patience.

2. The starting point of inflation can create a non-linearity in monetary policy strategy. For example, a central bank that is disproportionately more concerned about inflation misses that fall outside of a certain range (such as the target range of 1 to 3 per cent) and that has some desire to smooth interest rates will choose a more accommodative (restrictive) level for the policy interest rate when the starting point for inflation is well below (above) the target, relative to the case where the future path of inflation is certain.
3. While necessary, judgment applied to the policy decision should be disciplined, where possible, by formal modelling and simulation exercises to support robustness and consistency in decision making over time. Policy models, while imperfect, could be used more intensively. For instance, policy rules could be developed that are robust to parameter, model and data measurement uncertainty. More attention could be paid to potential non-linearities in the key structural relationships in the economy, as well as the potential for non-quadratic preferences among policy-makers.

This analysis highlights the importance of several areas for future research for the Bank of Canada in order to add more rigour to the framework and better inform the trade-offs associated with a risk management approach to monetary policy. These include efforts to:

- Improve policy models – this includes work to better articulate the nexus between the real and financial sectors in the economy, better account for the effects of heterogeneity among households and businesses, as well as the interactions between monetary, fiscal and macroprudential policies. A longer-term objective relates to incorporating the effects of uncertainty that households and businesses face in their decision making.
- Employ better and more timely data – this includes efforts to help us understand and monitor important structural changes in the economy such as those related to the digital economy. The Bank is also exploiting new sources of data, including so-called "big data," that could help to improve nowcasting and other analysis.

- Reduce the uncertainty related to the ELB – this work involves refining the design of unconventional monetary policy tools and better understanding their transmission.
- Develop robust policy rules – this includes exercises to assess the robustness of policy rules across different parameterizations of a single model as well as efforts to develop rules that are robust across models.
- Incorporate non-quadratic central bank preferences – this line of work involves more fully exploring the implications of departures from conventional specifications of central bank preferences.

The remainder of the paper is organized as follows. Section 2 reviews the types and sources of uncertainty that central banks face. Section 3 provides a brief and selective review of the literature on monetary policy under uncertainty. In section 4, we study three examples of the Bank of Canada’s behaviour under uncertainty and compare it to the prescriptions of the relevant literature. Section 5 offers some concluding remarks.

2 Types and Sources of Uncertainty

Even though uncertainty has long been a fact of life for many, not just central bankers, there is a general perception that it has increased in recent years. The upward trend in references to uncertainty in mainstream media, particularly in Canada, is one quantitative indication of this (**Chart 1**). While indicative, these types of metrics alone cannot distinguish between an increased fascination with risk, heightened risk aversion or a fundamental increase in uncertainty. In fact, it is very difficult to find objective measures of uncertainty that can.

In this environment, the Bank of Canada has not made larger forecast errors for gross domestic product (GDP) growth or inflation than in the past (**Charts 2 and 3** from Binette and Tchegotarev (forthcoming)).³ Indeed, Binette and Tchegotarev (forthcoming) show that in most *Monetary Policy Report* forecast horizons (FH2-4), the

³GDP growth forecasts for a given year are recorded for four monetary policy reports, and represent the four forecast horizons FH1-4. For example, the first real GDP growth forecast for 2016 is published in the April 2015 *Monetary Policy Report* (FH1). The second is reported in October 2015 (FH2), and the last two annual forecasts for 2016 growth are made in April and October of the same year, 2016 (FH3 and FH4, respectively).

root-mean-squared prediction errors (RMSPes) for GDP forecasts have been lower in the most recent decade (2006–17) compared with the previous one (1997–2006). This is important context because central banks typically take a risk management approach to monetary policy, which relies heavily on economic projections, in conjunction with an assessment of the related uncertainties, to choose the appropriate path for the policy.

Much of the macroeconomic literature treats uncertainty as quantifiable risk. That is, the policy-maker is assumed to be able to ascertain the range of possible outcomes and associated probabilities.⁴ For example, by assuming that the future will closely resemble the past, we can use models to quantify the risk around an inflation forecast from past forecast errors. Uncertainty is not always quantifiable, however. Knightian uncertainty refers to situations where the policy-maker cannot estimate the probabilities or envisage the range of possible outcomes.


There are numerous sources of uncertainty that central banks have developed techniques to deal with. These are reviewed carefully in Kozicki and Vardy (2017) from the point of view of the Bank of Canada and are summarized in this paper in **Table 1**. First, uncertainty is inherent in measurement of the economic data (e.g., business investment in a digital world) and in unobserved metrics that are used in monitoring and policy models (e.g., the neutral rate of interest) that central banks use to inform monetary policy. Second, there is uncertainty related to the policy models that are used to inform monetary policy decisions, either because the structure of the model does not adequately capture the essential elements of the real world or because the parameters cannot be directly observed and may change over time. Finally, uncertainty arises because of the potential for unforeseen developments or “shocks.” For obvious reasons, these are largely out of the central bank’s control, especially when geopolitical events are involved.

Central banks use a number of techniques that aim to reduce uncertainty from sources that they can influence. For instance, the Bank of Canada invests heavily in the development of more accurate projection models,⁵ as well as novel data sources that allow for a more complete assessment of the current state of the economy. In this regard, big data, advanced optimization algorithms and artificial intelligence represent promising areas for future research to reduce uncertainty associated with now- and near-term forecasting. But even as forecasting techniques continue to improve, ongoing structural change in

⁴This approach is consistent with Savage (1954).

⁵For instance, see the Appendix in the *Monetary Policy Report* (October 2017) for details on recent enhancements to one of the Bank of Canada’s main projection models, ToTEM.

Chart 1: News-based economic policy uncertainty indices, Canada and United States, monthly
Jan 1985–Jun 2017




Source: Baker, Bloom and Davis (2017)

Last Observation: June 2017

economies means that researchers are essentially aiming at a moving target. Thus, there is a certain level of uncertainty that, at given point in time, is considered irreducible. This undoubtedly reflects measurement issues and an incomplete understanding of how economies function.


Chart 2: Monetary Policy Report real GDP growth prediction errors, annual 1997–2016



Source: Binette and Tchegotarev (forthcoming)

Last observation: 2016

Chart 3: Monetary Policy Report inflation prediction errors, annual 2003–2016



Source: Binette and Tchegotarev (forthcoming)

Last observation: 2016

Table 1: Summary of sources of monetary policy uncertainty	
Description and example	
1. Uncertainty related to measurement	
Data	Data are subject to measurement error, particularly when there is a structural change. For instance, investment related to the digital economy may be underestimated in the national accounts because business spending in areas like cloud computing are counted as consumption.
Unobserved metrics	Several important metrics used as inputs to economic models (including monetary policy rules), such as the output gap and the neutral rate of interest, cannot be directly observed and must be estimated.
2. Uncertainty related to forecasting models	
Model specification	No model can capture the complexity of the real world, especially one that is subject to structural change. For example, there may be uncertainty about the correct way to model the underlying source of the pressure on prices of goods and services (e.g., focus on the output market or on the labour market).
Model parameters	Even if all parameters are identified, parameter values cannot be directly observed and may change over time. They are estimated using statistical techniques that are subject to error in finite samples even with a correctly specified model. For instance, there is uncertainty about the strength of the transmissions mechanism of “unconventional” monetary policy tools, which have relatively short track records in the jurisdictions where they have been deployed.
3. Uncertainty related to unforeseen developments	
Economic shocks	Unforeseen developments can have implications for inflation, like a sharp rise in the value of the Canadian dollar or a steep decline in oil prices. Sometimes these are Knightian in nature, although they can also reflect deficiencies in monitoring and forecasting tools that could be addressed through research.
Geopolitical shocks and natural disasters	Government decisions, whether they are related to peace, war or public policies such as trade agreements, can have profound implications for the economic outlook, as can natural disasters. They are nonetheless often the classic example of Knightian uncertainty.

3 A Brief Review of the Literature

There is a rich economic literature on how to deal with irreducible uncertainty in the conduct of monetary policy. Many of the approaches mirror those in the engineering literature, including optimal estimation and control using the Kalman filter, which is used extensively to deal with sources of uncertainty such as air and water currents, as well as errors associated with measurement, all of which are inherent to designing and deploying navigation and guidance systems. The attractiveness of these approaches is that they provide a rigorous framework to identify the optimal path for the monetary policy rate that explicitly factors in uncertainty about the current state of the economy. It also supports consistency in decision making over time.

The good news is that the introduction of different types of uncertainty into the framework requires only modifications to the existing optimization problem in the absence of uncertainty, rather than a fundamentally different approach. We start with the simplest result in the literature in which uncertainty does not matter for the optimal path of monetary policy (known as "certainty equivalence") because it provides a useful benchmark to more complex but realistic situations, and it mirrors the approach typically adopted in central bank projections.⁶ We then turn to more complicated situations. The policy prescription can change in important ways if any of the three key assumptions that underpin the certainty-equivalence result are relaxed: the model that governs the dynamics of the economy, the type of uncertainty faced by the central bank, and the preferences of the central bank.

3.1 Optimal policy design and certainty equivalence

A simple example clearly illustrates the logic behind the certainty-equivalence result. Consider first a linear model that posits a relationship between the target variable, π_t , and the policy instrument, R_t :

$$\pi_t = -\varphi R_t + \varepsilon_t, \tag{1}$$

⁶In most cases, central bank models do not use globally optimal rules for projection purposes, so the resulting policy paths implied by these models are not certainty equivalent, in the strict sense. Here, we mean that the parameters and functional forms of the policy rule, however chosen, do not change as uncertainty is introduced.

where ε_t is a random, additive disturbance with $E_t(\varepsilon_t) = 0$ and variance, σ_ε^2 , that is realized after the setting of the policy instrument, and the central bank loss function is quadratic in the target variable:

$$\mathcal{L} = E_t(\pi_t^2). \tag{2}$$

In this case, the optimal response for the central bank is $\varphi^{-1}E_t(\varepsilon_t)$. The optimal response coefficient φ^{-1} is certainty equivalent since it is independent of the variance of the shock. In other words, the policy-maker can make decisions based solely on the mean and safely ignore the variance of shocks. This result can be generalized (see Appendix B for details) provided that: (i) the underlying economic model is linear, including additive shocks (meaning their impact is independent on the state of the economy); (ii) there is no other source of uncertainty other than additive model shocks;⁷ (iii) the central bank's loss function is quadratic. This type of loss function implies that the central bank cares disproportionately more about large deviations of inflation from target than small deviations, but loss is symmetrical around the target.⁸

3.2 Three departures from certainty equivalence

As mentioned earlier, this result is highly relevant to how most central banks construct projections. For instance, models are typically linear by construction or linearized, the only source of uncertainty typically considered is additive, and the reaction functions are either simple rules whose parameters are estimated on historical data or chosen optimally in the presence of additive uncertainty only. These models are then used to construct a base-case forecast, as well as several risk scenarios that assume a different profile for one or more of the model's shock terms. A balanced or mean scenario can be obtained by taking a weighted average of these various scenarios, which is certainty equivalent in the sense that no other aspect of the distributions other than the mean affect the policy path. While the inclusion of risk scenarios acknowledges that the economy may be more likely to evolve in one direction (relative to the base case) than another, it does so in a way that preserves certainty equivalence. This is because the policy-rate path

⁷A similar but somewhat weaker result is obtained when policy is constrained to respond optimally to a subset of model variables, as would be the case with an optimized Taylor rule. Specifically, relative shock variances will affect the design of policy since they affect the relative importance of the target variables in the central bank's loss function, and therefore the trade-off between stabilization objectives.

⁸These conditions are not necessary in all cases. For instance, Chadha and Schellekens (1999) show that the certainty-equivalence result is still obtained with a more general class of convex loss functions.

is being influenced to the extent that the mean of the target variables like inflation is being shifted, but the policy rule itself is unchanged.

There are three cases studied extensively in the literature that show where optimal policy can no longer conveniently abstract from uncertainty by focusing solely on the mean. Specifically, optimal policy in these cases can respond more or less aggressively to economic developments than under certainty equivalence. Alternatively, it can affect the optimal degree of interest-rate smoothing, the timing of interest rate changes or even the average level of interest rates.

3.2.1 Non-linear model dynamics

Equation (1) provides an example of a linear relationship in which the impact on inflation of a demand shock is proportional to the size of the shock. Because of this linear relationship between demand shocks and inflation, the optimal policy response is also linear and moves in direct proportion with the expected size of the demand shock it faces. But suppose instead the relationship between demand and inflation was non-linear such that excess demand creates a larger inflation movement than excess supply. In such a circumstance, policy would have an incentive to respond more aggressively to positive demand shocks than negative ones. But since policy has to be set prior to observing the shock, it must be based on an expectation. In the case of the linear model, the rule was to simply set the policy rate equal to zero each period, $\varphi^{-1}E_t(\varepsilon_t) = 0$, since the expected value of the shock is zero. But with the non-linear Phillips curve, optimal policy will respond to more than just $E_t(\varepsilon_t)$. Why? Under a policy of $R_t = 0$, the average rate of inflation in the economy will be positive even if the shocks are zero mean and symmetrically distributed, because inflation is more sensitive to excess demand. In this case, the optimal reaction function will include a positive constant or wedge that offsets the bias to inflation, and this wedge will depend on the variance of the demand shocks. As a result, certainty equivalence is broken.

A relevant example from the literature is the "zone-linear" Phillips curve of Orphanides and Wieland (2000b). They assume that the slope of the Phillips curve is zero for some symmetric zone of output gaps. Outside of this zone, the Phillips curve is linear with a positive slope. In other words, small deviations in the output gap around zero do not affect inflation, but larger ones do linearly. Two interesting results emerge from this set-up. First, optimal policy without uncertainty implies a zone of inaction for

price shocks, that is for shocks that create a trade-off between inflation and output gap stabilization. This is due to the flat Phillips curve assumption, which implies that the central bank must generate a certain minimum size output gap to offset any of the price shock. This can be thought of as a fixed price of policy action. For small-enough price shocks, it is less costly (for a central bank that cares about both output gap and inflation stabilization) to not respond. The second interesting result is that allowing for uncertainty generates a larger zone of inaction that depends positively on the variance of the price shocks. This option value of waiting reflects the possibility that a current-period deviation of inflation from target will be partially eliminated in the future by favourable (offsetting) price shocks. Finally, given that this zone of inaction varies with the degree of uncertainty, certainty equivalence is broken.

Other examples related to wage and price dynamics include convex Phillips curves and piece-wise linear relationships that are kinked around an output gap of zero. There are, of course, many other important sources of non-linearity in the real world that would nullify the certainty-equivalence result. In general, optimizing agent models used by central banks give rise to non-linearities in the behaviour of all endogenous variables, but the models are often linearized for the purpose of running projections. Historically this has been because of the dearth of robust algorithms to solve non-linear models featuring rational expectations and uncertainty. Also, economic and geopolitical shocks may enter the model non-linearly, as opposed to being additive as is typically assumed. So for instance, the impact of a shock to financial conditions on household and business behaviour may depend on other aspects of the state of the economy, such as the level of indebtedness. Finally, the interaction between extreme shocks and non-linear models has been recently studied by Kim and Ruge-Murcia (2016). In particular, using a small non-linear New Keynesian model with asymmetric shocks, they find that under both the Taylor and the Ramsey policy, the central bank responds non-linearly and asymmetrically to shocks. Moreover, they conclude that strict price stability is preferred to targeting a positive inflation rate in their model.

3.2.2 Different sources of uncertainty

The literature has studied the implications of multiple sources of uncertainty on monetary policy, including uncertainty related to parameters within a given model, uncertainty across models, and uncertainty about the data, including unobserved metrics, used in

models. The results show that optimal policy under these forms of uncertainty may be either more or less aggressive than under certainty equivalence.

Parameter uncertainty Even in the best case where the central bank has the correctly specified model and all parameters are identified, the underlying structural parameter values will generally have to be inferred based on limited historical data. As a result, parameter estimates will be subject to sampling error. Parameter uncertainty is non-linear in the sense that it interacts with policy actions, and this interaction means that the choice of policy rule can mitigate or exacerbate the effects of parameter uncertainty on the economy. To illustrate a particular case of parameter uncertainty, in which policy should be less aggressive, we consider the famous Brainard example (Brainard 1967). In this set-up, the parameter relating the instrument to the target is not known with certainty (unlike equation (1)), so the central bank's model is characterized by

$$\pi_t = -(\varphi - v)R_t + \varepsilon_t = -\varphi R_t + \varepsilon_t + vR_t. \quad (3)$$

In this case there are two shocks in the model from the vantage point of the central banker: the demand shock (ε_t) that is observed by the central bank prior to setting policy, and v , which is a random variable that reflects the central bank's uncertainty about φ . The multiplier on v interacts with the nominal interest rate (vR_t), reflecting the idea that any error in the estimate of φ translates into a policy error. If the central bank implements the same optimal policy derived under certainty equivalence, the variance of inflation will be $\frac{\sigma_\varepsilon^2 \sigma_v^2}{\varphi^2} > 0$. In the absence of uncertainty, it would be zero since the demand shock is assumed to be observed prior to setting policy each period. However, policy could do better, in terms of smaller deviations of inflation from target, by responding less aggressively to inflation. The optimal response parameter, accounting for uncertainty, in this model is $\frac{\varphi}{\varphi^2 + \sigma_v^2} < \varphi^{-1} \quad \forall \sigma_v^2 > 0$.

Intuitively this means that as parameter uncertainty increases, policy should become less aggressive to avoid causing excessive volatility to inflation. In the simple model used by Brainard, the coefficient relating the policy instrument to the target variable is a composite of the coefficient linking the policy instrument to real activity (such as the output gap) and that linking real activity to inflation, meaning the conservatism result should hold for either, at least for this particular model.

A timely example of parameter uncertainty relates to UMP, such as negative interest

rates, extraordinary forward guidance and quantitative easing, the use of which by central banks has been limited relative to conventional policy. As a result, there is much greater uncertainty about the magnitude of the real economy’s response to these tools compared with conventional policy, as well as any unintended consequences. Applying the Brainard conservatism principle would suggest an UMP response that is less aggressive than with conventional policy in the same circumstances. It may also rationalize a higher bar to employ UMP measures.

The Brainard result has received a lot of attention from central banks (see Blinder 1998) since it seems to provide a rationale for why optimal policy rules (or optimized simple rules) tend to be much more aggressive than estimated rules. However, the Brainard result turns out to be a rather special case. Much of the subsequent literature finds the opposite result, in that central banks should respond more aggressively when they are uncertain about their model than when certainty equivalence applies.⁹ This result has been obtained through two main modifications to Brainard’s simple set-up: first, researchers have included uncertainty about multiple parameters of their model instead of just one; and second, more realistic model features such as endogenous expectations have been added (Craine 1979; Söderström 2002; Kimura and Kurozumi 2007). One specific example of a more realistic model feature that can help convey the intuition of this finding is the presence of lagged inflation in a modified New Keynesian Phillips curve:

$$\pi_t = \alpha\pi_{t-1} + (1 - \alpha)\pi_{t+1} + \gamma y_t, \tag{4}$$

where π_t is the quarterly rate of inflation minus the inflation target, and y_t is the output gap. Suppose we don’t know the extent to which lagged inflation matters for current inflation, governed by α , but we do know that the parameter is normally distributed with a mean of $\bar{\alpha}$ and a variance of σ_α^2 . It turns out that even if the distribution of α is normal, the distribution of losses to the central bank stemming from this is heavily skewed to the right. As a result, the loss associated with underpredicting α is greater than that associated with overpredicting α . If the true value of α is greater than the estimated value $\bar{\alpha}$, the central bank’s policy rule should be more aggressive in order to control inflation. Conversely, if the true value is smaller, then the rule should be less aggressive. It is better to behave as if α is high and be wrong than vice versa. So in

⁹For example, see Cateau and Murchison (2010).

contrast to the Brainard result, parameter uncertainty in this instance leads policy to respond more aggressively.

Model uncertainty While central banks often use a reference model for forecasting, there is considerable uncertainty about the correct specification of its key behavioural relationships. In dynamic stochastic general equilibrium (DSGE) models, for instance, there is often disagreement about the correct way to specify many features: household preferences, the production technology, the information sets of private agents, and the degree of heterogeneity in the economy, all of which give rise to model uncertainty. Svensson and Williams (2005) demonstrate that model uncertainty in general breaks certainty equivalence and therefore policy must focus on the whole distribution of possible outcomes, rather than simply the mean. There are two main approaches to doing this in the literature, and they differ only in terms of how each possible model specification is weighted. The first is a Bayesian approach already discussed that weights each possibility by a measure of its likelihood of being the true model. The second is a "minimax" approach, which assigns a weight only to the model that would produce the highest loss. An important example of this type of approach is the robust control methods introduced into economics by Hansen and Sargent (see, for example, Hansen and Sargent 2007). The resulting policy rule in the first case will tend to do well on average, whereas in the second case, policy is geared more towards robustness against the worst-case scenario. The minimax approach is particularly relevant for situations characterized by Knightian uncertainty where the decision maker may be unable or unwilling to assign weights to the different models he/she considers plausible (e.g., after the financial crisis it was known that the standard New Keynesian Phillips curve (NKPC) models had limitations, a number of real-financial-linkages models were being proposed but that past data would not be informative to validate new candidate models). In such situations, a minimax criterion would make sense. These approaches can also be combined. For example, using a framework that nests both the Bayesian and the "minimax" approach, Cateau (2007) finds that a small degree of aversion to the risk associated with facing multiple models and measurement error configurations can generate an optimal rule that matches the empirical Taylor rule.

Measurement uncertainty A third important source of uncertainty stems from measurement issues associated with data used to infer the state of the economy. The common

example cited in the literature is the output gap, which is unobservable and therefore must be estimated based on a number of observable variables. Estimates are often subject to large revisions through time because of the use of two-sided filtering algorithms that rely on future data that are not available in real time. Orphanides (2001) considers the case of optimal policy in a setting when policy-makers are faced with a signal extraction problem. Specifically, if the time (t) observation of the output gap is equal to the unobserved true output gap plus a noise component as in

$$y_{t|t} = y_t + \eta_t, \tag{5}$$

then the magnitude of the optimal response to $y_{t|t}$ will be declining in the variance of the noise component, σ_η^2 , and certainty equivalence will no longer hold. In effect, the optimal response coefficient accounts for the fact that the time- t observed output gap is more volatile than the true output gap and will therefore attenuate the response. Equivalently, if we separate out the estimation and control problems, the former would allow for an optimal estimate of the true output gap, $E_t y_t = \frac{\sigma_y^2}{\sigma_y^2 + \sigma_\eta^2} y_{t|t}$ (using the Kalman filter). If policy were instead allowed to respond to $E_t y_t$ instead of $y_{t|t}$, certainty equivalence would be retained (see Swanson 2004 and Svensson and Woodford 2003). In other words, by separating out the estimation and control stages of the problem, certainty equivalence can be preserved when the only source of uncertainty is related to data measurement. Finally, if the problem is specified as one in which the true output gap contains an orthogonal component omitted from the optimal estimate

$$y_t = y_{t|t} + \eta_t, \tag{6}$$

then η_t won't enter the reaction function and certainty equivalence holds.

To summarize, the link between certainty equivalence and measurement uncertainty comes down to whether the optimal estimate of the unobserved variable enters the policy rule. If yes, policy will remain certainty equivalent in that the response coefficient associated with the optimal estimate will be invariant to the level of uncertainty, σ_η^2 . Otherwise, there will generally be an attenuated response to the noisy indicator that is a function of the relative size of the noise component.

3.2.3 Non-standard central bank preferences

Central bank loss is non-convex Not surprisingly, how a central bank takes uncertainty into account in the conduct of monetary policy depends critically on what it cares about. The popularity of quadratic preferences—which are a cornerstone of the certainty-equivalence result—is likely because they are relatively tractable, and because they can be rationalized as a second-order approximation to the underlying (non-linear) household welfare function.¹⁰ Allowing for more general central bank preferences permits us to analyze the impact of additive uncertainty in environments, for instance, where the central bank is disproportionately more concerned with tail events (relative to a quadratic specification):

$$\mathcal{L} = E \left| \pi_t^\beta \right|; \quad \text{for } \beta > 2, \quad (7)$$

where the limiting case $\beta \rightarrow \infty$ corresponds to preferences whereby the central bank concerns itself only with shocks that generate the largest deviations of inflation from target. Another special case corresponds to loss that is linear in the absolute value of inflation deviations ($\beta = 1$). In this case, the measure of central tendency of the shock distribution of ε_t will now be the median, as opposed to the mean when $\beta = 2$. If the distribution of ε_t is skewed, certainty equivalence will no longer hold since the median, not the mean, will be the appropriate measure of central tendency. Similarly, with $\beta = 0$, the relevant measure will be the mode and any skew (that is, the balance of risks around inflation) would be ignored altogether when policy is set.

The preceding examples demonstrate that in the case of an asymmetric shock distribution, the central bank's β will determine the relevant measure of central tendency when setting policy, and only when $\beta = 2$ will certainty equivalence hold. Orphanides and Wieland (2000b) consider an alternative in which central bank preferences have thresholds, such that loss is very small or zero for a certain range of outcomes, but quadratic when outside of that range. This would be consistent with a central bank that targets a range, rather than a specific level, of inflation. In such a case, the central bank may be indifferent to movements of inflation within the target range, but care increasingly about inflation when it falls outside of the range. The authors show that, without uncertainty, this type of loss function generates a zone of inaction, meaning policy would not respond to shocks that do not push inflation outside the target zone. However, once additive

¹⁰Chadha and Schellekens (1999) extend the certainty-equivalence result to a more general class of convex preferences.

uncertainty is introduced, policy would respond to all shocks, regardless of size, since any deviation of inflation from the midpoint of the target range increases the likelihood that subsequent shocks will cause a breach of the range. Certainty equivalence is nevertheless broken since policy would respond by less to small disturbances relative to the case where preferences are quadratic for all deviations, but becomes disproportionately more aggressive as the deviation from the midpoint of the target zone increases.

Svensson (2003) considers what is essentially the opposite set-up: loss is quadratic for small and moderate deviations of the target variables, and constant for extreme events. Using such a set-up, he shows that optimal monetary policy completely ignores extreme events. While such a result is intuitively appealing—central banks may find it easier to communicate policy that places no weight on low-probability, extreme events—the idea that loss becomes independent of the severity of the event beyond a certain point is inconsistent with standard assumptions about household utility, which is typically the basis for the central bank’s loss function.

Deviations from expected utility The literature reviewed so far has mainly focused on how various potential economic fluctuations would affect the central bank’s loss function. With the exception of minimax preferences discussed briefly in the model uncertainty section, we’ve assumed that each possible outcome is weighted by its respective probability of occurring in order to compute expected loss over the entire distribution of outcomes (Bayesian averaging). More formally, we have assumed that the central bank calculates the mathematical expectation across the various possible outcomes for the target variable:

$$E(\mathcal{L}(\pi)) = \int_{-\infty}^{\infty} \mathcal{L}(\pi(\varepsilon)) \cdot f(\varepsilon) d\varepsilon, \quad (8)$$

where in this case, ε is the additive shock term, $\mathcal{L}(\pi(\varepsilon))$ is the loss associated with a given realization of the shock, and $f(\varepsilon)$ is the density function of ε .

The final deviation from the standard assumptions that underpin certainty equivalence that we consider is related to this weighting function, since there is no requirement that it correspond to the density function for the underlying variable. Al-Nowaihi and Stracca (2002) study optimal policy assuming the central bank chooses to over- or under-weight certain ranges of events. They show that the combination of such a non-linear weighting and a skewed shock distribution will result in certainty equivalence being broken, despite their assumption that the weighting scheme is symmetric with respect to the

probability distribution of the target variable. For instance, if the central bank places a greater weight on extreme events than suggested by their probabilities, and the risks to the target variable are skewed to the downside, policy would be more accommodative relative to the certainty-equivalence case.

4 Theory versus Practice: Examples of the Role of Uncertainty

It is clear from the brief review in the previous section that the literature is highly informative regarding how best to account for uncertainty under different cases. At the same time, the policy prescriptions from any particular model reviewed may not be robust to all of the sources of uncertainties and related trade-offs that are pertinent in the real world. That may be why, in practice, central banks tend to start with a formal framework that is consistent with certainty equivalence and then apply judgment that attempts to take into account the factors and uncertainties that are viewed as missing from the framework. This is part of a broader risk management approach that attempts to weigh (at least qualitatively) the ability of different policy paths to achieve the inflation target while incorporating risks to the real economy and financial stability.

This suggests that central banks view the certainty-equivalent approach as providing a reasonable approximation to the optimal policy path that would prevail if all potential deviations from certainty equivalence were explicitly considered. There are instances, however, in which deviations from certainty equivalence are factored in to policy deliberations. It is therefore instructive to attempt to map this practice to science, in order to assess whether it conforms to the prescriptions in the literature. To do this, we have selected three examples of the Bank of Canada's behaviour in the face of uncertainty that are particularly relevant in the current context. While these episodes pertain to the Canadian economy, they have a universal quality: other central banks have faced similar situations in the past and will likely face them again in the future.

There are two caveats to the analysis that should be underscored: (i) the examples we study are by no means exhaustive, and in fact are not orthogonal to each other; and (ii) the Bank of Canada makes monetary policy decisions by consensus of a committee of six people (chaired by the Governor), all of whom may have somewhat different loss functions and may put different weights on the risks to the projection. As a result,

any inferences regarding a single, unified intent behind policy actions must be viewed as indicative.

Our analysis of these examples demonstrates that judgment can often be justified (qualitatively) by a formal framework that accounts for uncertainty and deviations from certainty equivalence. At the same time, the logic behind the judgment often also has its limits because there are trade-offs that may not be properly considered. This reinforces the importance of augmenting, as much as possible, the art of policy-making with formal modelling frameworks for thinking about policy-making under uncertainty.

4.1 The effective lower bound

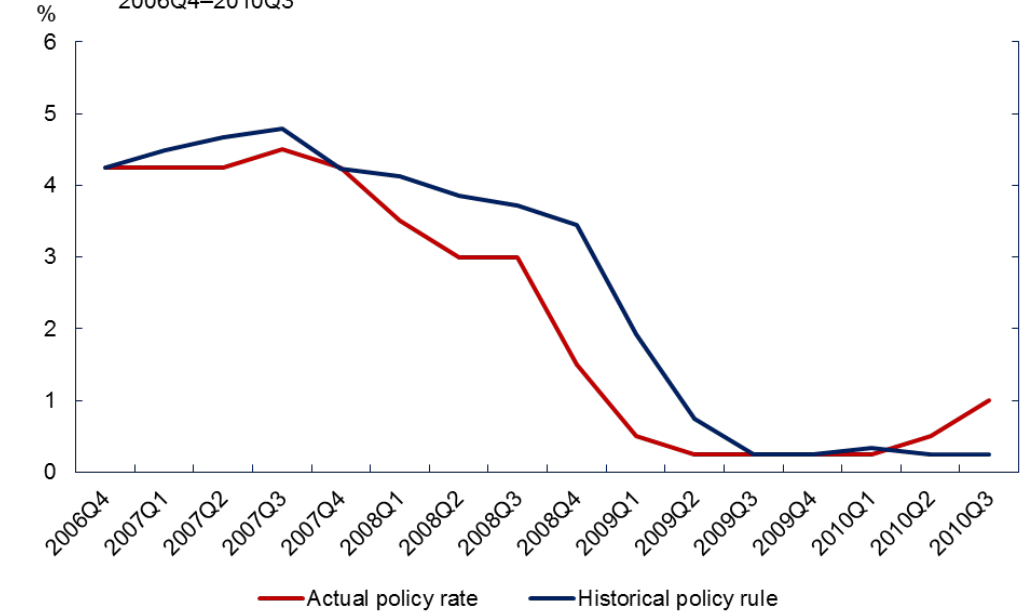
As we saw in section 3, non-linearities in the structure of the economy can invalidate the certainty equivalence result. Over the past decade, the ELB on nominal interest rates has been the most obvious non-linearity. Like many other central banks, the Bank of Canada encountered the ELB during the global financial crisis. The Bank’s policy rate fell to 25 basis points—which was then regarded as the ELB—in April 2009. Moreover, the probability of being constrained by the ELB has increased markedly because of the estimated decline in the neutral rate of interest over the last decade and a half (Dorich et al., forthcoming). Based on our assessment that the ELB in Canada is around -50 basis points, the probability of being at the ELB is around 8 per cent, about five times higher than it was 15 years ago.¹¹

The discussion of non-linearities in the previous section suggests that as the Bank’s policy rate approached the ELB and the risk of being constrained by the ELB rose, the Bank’s behaviour could have been expected to change in order to properly account for the uncertainty about the effectiveness of unconventional monetary policies. Indeed, there is evidence that the Bank’s behaviour did change near the ELB. **Chart 4** shows the path of the policy rate that would have been prescribed by a historically estimated policy rule along with the actual path of the policy rate.¹² Clearly, monetary policy was eased more aggressively during that period than suggested by the historical rule.

¹¹Some estimates for the United States are as high as 40 per cent. Much of this difference with Canada can be explained by differences in methodology rather than fundamentals. For instance, raising the assumption of the ELB to zero and reducing the degree of interest rate smoothing to zero in the policy rule increases the estimate of the probability to around 25 per cent for Canada. For a more detailed explanation of how the probability of being at the ELB is calculated, see Dorich et al. (forthcoming).

¹²The path associated with the historically estimated rule was calculated using real time data, so that the difference between actual and prescribed policy actions reflects only the additional aggressiveness.

Chart 4: Aggressiveness of policy near the effective lower bound, quarterly data
2006Q4–2010Q3



Source: Bank of Canada calculations

Last observation: 2010Q3

This apparent increased aggressiveness is a textbook example of optimal policy in the presence of the specific non-linearity related to the ELB, as discussed in section 3. The intuition behind this is demonstrated the most clearly in Kato and Nishiyama (2005). They show that the optimal policy, expressed as a reaction function, contains both a standard certainty-equivalent component as well as a component that is strictly a function of the possibility that the lower bound may constrain policy in the future as follows:

$$R_t = \varphi_1 \pi_t + \varphi_2 y_t + \varphi_3 \sum_{i=0}^{\infty} \varphi_4^i E_t \Theta_{t+i}; \quad R_t > ELB. \quad (9)$$

The first two terms constitute a Taylor-style rule with optimal parameter settings, whereas the third term is a function of the expected probabilities of being constrained by the ELB, and therefore is indirectly a function of the additive stochastic shocks to the Phillips and investment/savings (IS) curves in the model. The authors go on to demonstrate that optimal policy in this model differs from certainty-equivalent policy in two critical ways. First, policy is more *expansionary*, meaning the stochastic mean of the policy rate is less than or equal to that in the absence of the ELB. So the pres-

ence of the ELB introduces an inflationary bias into policy as a safeguard against being constrained.¹³ Second, policy is more *aggressive*, meaning the response to inflation and output gap movements will be greater than or equal to that which would prevail without the ELB, and the extent of the difference will depend on the expected proximity to the ELB. An implication of this is that policy will respond more aggressively to shocks that drive the economy below its equilibrium level than to those that push it above.

The intuition for this result is straightforward. When the current state of the economy depends partly on previous states (the model contains lagged state variables), optimal policy will be more aggressive as the probability of being constrained by the ELB in future periods increases. Maintaining our assumption of a quadratic loss function, welfare will be higher if the central bank boosts demand and inflation above their target levels prior to the ELB binding, since (through lagged effects) this will imply the same variables don't fall as far below their targeted levels should the constraint bind in the future. In effect, with quadratic preferences, a series of small misses is better than a single large miss.¹⁴

While this may be the optimal strategy given this type of uncertainty, central banks would be able to do better (e.g., lower volatility of inflation and the output gap, less buildup of financial vulnerabilities from credit) if they were confident about the effectiveness of unconventional policy tools. This highlights the importance of further refinement of the design of unconventional policy tools, as well as our understanding of the strength of their transmission to the real economy and potentially undesirable longer-term side effects.


4.2 The option value of waiting

Pauses—relatively long periods without any change in the policy interest rate—are typical of central bank behaviour. This is evident in **Chart 5**, which shows pauses for policy rates for several central banks even outside the more recent period when policy was

¹³This result has been more recently confirmed by Kiley and Roberts (2017) using a small DSGE model and FRB-US.

¹⁴Other things being equal, greater uncertainty about the future state of the economy will increase the probability that the ELB constrains policy, thereby giving the central bank an incentive to preemptively run the economy hotter. Blinder (2000), Goodfriend (2000), Reifschneider and Williams (2000), Orphanides and Wieland (2000a), and Hunt and Laxton (2003) all argue the more general point that policy will be more aggressive and preemptive in the presence of uncertainty and a lower bound on nominal interest rates.

Chart 5: Monetary policy rates, monthly
 % February 1999–October 2017



Source: Haver Analytics

Last observation: October 2017

constrained by the ELB. The numerous pauses in this figure contrast with the optimal policy prescriptions described in section 3, which call for more frequent adjustment of the policy rate.

To understand what might motivate this type of behaviour, it is instructive to consider a specific example. The Bank of Canada left the policy rate unchanged at 1 per cent for four years (September 2010 to January 2015). This was an unusually long period of inaction relative to history. It wasn't until January 2015, after oil prices had plummeted about 60 per cent, that the Bank of Canada cut its policy interest rate by 25 basis points and then another 25 basis points in July of that year. In this context, Poloz (2017) stated that uncertainty can lead the Bank to leave the policy rate unchanged in the face of small shocks while reacting to larger shocks:

All of these sources of uncertainty define the zone in which we can be reasonably assured that policy is on track. Factors that increase uncertainty—such as geopolitical risks—can widen this zone temporarily. Conversely, resolution of uncertainties can narrow it.

...[U]ncertainty does not equal indecision. It is true that the notion of a zone generated by uncertainty can create a degree of tolerance for small shocks. At the same time, a large shock—or, perhaps, an accumulation of smaller shocks—can tilt the balance of risks to projected inflation and prompt policy action.

Other central banks have also cited uncertainty as a rationale for waiting. For example, the European Central Bank (quoted in Al-Nowaihi and Stracca 2002) stated that it "confirms its position of 'wait and see' with regard to its monetary policy stance. In an environment of increased uncertainty over the global economy and its impact on the euro area, the Governing Council is carefully assessing whether and to what extent upward risks to price stability will continue to decline." Similarly, the Sveriges Riksbank (1998) noted that "The element of uncertainty in the inflation assessment can accordingly influence monetary policy's construction. A high degree of uncertainty can be a reason for giving policy a more cautious turn."

While these explanations are intuitively appealing, it is not immediately obvious how to reconcile this type of "wait-and-see" behaviour with the standard optimal policy prescriptions described in section 3. In particular, the standard approaches generally call for continuous adjustment of the policy rate regardless of the level of uncertainty and are therefore not able to explain the pauses that are typical of actual central bank behaviour. One exception to this discussed in section 3.2.1 is under the assumption of a zone linear Phillips curve (Orphanides and Wieland 2000b). However, this result is specific to price shocks. Moreover, a more general case for pauses cannot be justified by introducing an aversion to volatility in interest rates or to volatility in changes in interest rates into a quadratic loss function. Such an aversion would lead to smoothing interest rate adjustments, but this smoothing would be independent of the degree of uncertainty. In general, no change to the central bank's loss function that maintains its quadratic form can motivate the tendency of central banks to pause.

Another way of generating an option value of waiting that gets us closer to a more general explanation but is largely outside of the policy under uncertainty literature is to treat the central bank's policy decision in a manner analogous to a firm's investment decision. The investment literature emphasizes that, when a firm faces uncertainty about the future rate of return on an investment project and a fixed cost of undertaking or reversing the project, it will sometimes choose to exercise its option to wait.

To illustrate this point, consider a firm in the oil and gas extraction sector. Suppose the firm has a lease that gives it the right to extract oil from a plot of land for two years. For simplicity, assume that the firm knows with certainty that it could extract 100,000 barrels of oil each year and that it costs \$50 to extract each barrel. The decision to extract the oil is irreversible in the sense that if the firm starts extraction in year 1, it must continue in year 2. However, it has the option to wait until year 2 to begin extraction.

Suppose the price of oil in year 1 is \$55 and there is a 50 per cent probability that the price of oil will be \$10 higher in year 2 and a 50 per cent probability that it will be \$10 lower. Assuming that the firm has a discount rate of zero, then the expected net present value (NPV) of starting extraction in year 1 is¹⁵

$$(\$55 - \$50) \times 100,000 \times 2 = \$1,000,000.$$

On the other hand, if the firm exercises its option to wait until year 2 to decide whether or not to begin extraction, its expected NPV is

$$0.5 \times (\$65 - \$50) \times 100,000 = \$750,000.$$

The expected NPV of waiting reflects the fact that, conditional on waiting until year 2, the firm will only extract if the price rises and the probability of the price rising is 0.5. Under these assumptions it is optimal for the firm to begin extraction in year 1.

But the option value of waiting hinges on the degree of uncertainty about future oil prices. Suppose that instead of rising or falling by \$10, the price of oil can rise or fall by \$20 in year 2. This increase in price uncertainty does not affect the expected NPV of beginning extraction in year 1. In contrast, the expected NPV of waiting changes to

$$0.5 \times (\$75 - \$50) \times 100,000 = \$1,250,000.$$

This higher expected NPV of waiting results from the fact that, if the firm waits, it extracts the oil only if the price rises to \$75. Waiting allows the firm to avoid producing

¹⁵The expected NPV of starting extraction in year 1 is derived as follows. Profit in year one is $(\$55 - \$50) \times 100,000$. If the price goes up in year 2, then profit in year 2 will be $(\$65 - \$50) \times 100,000$. If the price goes down, then year 2 profit will be $(\$45 - \$50) \times 100,000$. Weighting the two possible year 2 outcomes by their probabilities implies that the expected year 2 profit is $(\$55 - \$50) \times 100,000$. Summing the year 1 profit with the expected year 2 profit yields the result in the main text.

at a loss if the price falls in year 2. Comparing expected NPVs, it is clear that with greater uncertainty it is optimal for the firm to wait until year 2.

This suggests that one could rationalize the tendency of central banks to pause by introducing considerations that lead to an option value of waiting. One possible approach, pursued by Lei and Tseng (forthcoming), is to assume the central bank faces a fixed cost of adjusting its policy rate. They show that, with a fixed adjustment cost, increased uncertainty makes the central bank more reluctant to change its target interest rate and generates long pauses. They argue that this helps to explain recent observed deviations from the Taylor rule. But this approach suffers from two limitations. First, it is difficult to motivate the large adjustment costs that Lei and Tseng (forthcoming) find are necessary to generate pauses. More importantly, while fixed costs of adjustment explain why a central bank would not react to small shocks, they do so at the cost of predicting counterfactually large adjustments when the policy rate is changed. In practice, central banks tend to make many small adjustments in the same direction rather than making one large adjustment.

A more promising alternative, inspired by the notion of irreversibility in the investment literature, is to assume that the central bank faces some fixed cost of *reversing* course. That is, the central bank incurs a fixed cost if it lowers the policy rate when its last change was an increase, or vice versa. This avoids imposing a cost for small adjustments in the same direction, so it can be consistent with the tendency of central banks to adjust policy rates gradually. At the same time, it imposes a cost on changes in direction and can therefore motivate pauses. Such a cost is also easier to motivate as an aversion to reversals.

As Blinder (2006) notes, "Although the basic logic of optimization suggests that such policy reversals should not be uncommon, central bankers seem to avoid them like the plague." This aversion is evident in the data. Defining a policy reversal as a change in direction within three months, Blinder finds that reversals are exceedingly rare for many central banks. In Canada, since the adoption of fixed announcement dates in November 2000, the Bank of Canada has changed its policy interest rate 46 times, but only 2 of these were reversals. Widening the window for reversals to six months only increases the total number of reversals to 4. Blinder points to concerns about central bank credibility and financial market stability as possible explanations of this apparent aversion to reversals.

To demonstrate how reversal aversion may influence policy decisions, we develop a

simple two-period model. We use the model to analyze how an aversion to reversals on the part of the central bank (CB) can lead to optimal waiting in the presence of uncertainty (see Appendix A for the details of our proofs). Let i_t^* be the optimal policy rate setting in the absence of reversal costs. Assume that i_t^* follows a random walk:

$$i_t^* = i_{t-1}^* + \varepsilon_t. \quad (10)$$

The shocks, ε_t , are a reduced-form way of representing all the shocks that could affect the optimal setting of the policy rate. Prior to period $t = 1$ the economy was in steady state (i.e., $\varepsilon_t = 0$ and $i_t^* = i_{ss}$ for all $t < 1$). In periods 1 and 2, the shocks, ε_t , are drawn from a normal distribution with mean 0 and variance σ^2 . Period 2 is the terminal period for this economy. So we have

$$\begin{aligned} i_t^* &= i_{ss} \quad \forall t \leq 0, \\ i_1^* &= i_{ss} + \varepsilon_1, \\ i_2^* &= i_{ss} + \varepsilon_1 + \varepsilon_2. \end{aligned} \quad (11)$$

The actual policy rate, i_t , may or may not be equal to i_t^* in periods 1 and 2. In period 1, the CB chooses whether to take action by setting $i_1 = i_1^*$ or to wait by leaving the policy rate at its steady-state level, $i_1 = i_{ss}$. In period 2, the CB must again choose whether to take action by setting $i_2 = i_2^*$ or to wait by leaving the policy rate at its period 1 level, $i_2 = i_1$. But the CB faces an additional consideration in period 2. If it took an action in period 1, then the CB will incur a cost, k , if it takes an action of the opposite sign in period 2. The cost, k , is measured in terms of the value of the CB's loss function and is intended to capture the idea of reversal aversion in a simple reduced-form manner.

Define the deviation of the actual policy rate from i_t^* as

$$z_t \equiv i_t - i_t^*. \quad (12)$$

The flow cost of deviations is z_t^2 , so the CB's loss is

$$L = z_1^2 + \beta (z_2^2 + I_2 k), \quad (13)$$

where k is a fixed cost of reversing the period 1 action and I_2 is an indicator variable:

$$I_2 \equiv \begin{cases} 1 & \text{if } \text{sign}(i_1 - i_{ss}) \neq \text{sign}(i_2 - i_1) \text{ and } i_1 \neq i_{ss} \\ 0 & \text{if } \text{sign}(i_1 - i_{ss}) = \text{sign}(i_2 - i_1) \text{ or } i_1 = i_{ss} \end{cases}. \quad (14)$$

This indicator variable captures the fact that the CB incurs the cost k only if it took action in period 1 ($i_1 \neq i_{ss}$) and it takes an action of the opposite sign in period 2.

In this environment, we can prove three key results:

1. The fixed cost of reversing implies that the CB will only reverse direction in period 2 for sufficiently large shocks. That is, there is a zone of inaction in period 2 when the period 2 shock is of a different sign from the period 1 shock.
2. The risk of needing to reverse itself and incur the fixed cost causes the central bank to take action in period 1 only for sufficiently large shocks. That is, even though the CB faces no direct costs of action in period 1, there is still a zone of inaction in period 1.
3. The absolute magnitude of the period 1 shock required to induce a policy response from the CB is increasing in the variance of the period 2 shocks. That is, the width of the period 1 inaction zone is increasing in the degree of uncertainty about future shocks.

The first result is a fairly obvious consequence of the fixed cost of reversals. The second and third results are less obvious and more interesting. In particular, our framework can motivate a zone of inaction that depends on the degree of uncertainty while still being consistent with the observed tendency of central banks to make many small changes in policy in the same direction.

One possible criticism of the analysis above is that there is no second period cost of first period inaction. For example, one might think that by waiting in period 1, the central bank could "fall behind the curve" and therefore incur a greater loss in period 2. We can capture this type of idea in a reduced-form way by modifying the loss function to penalize large movements in the policy rate:

$$L = z_1^2 + c(i_1 - i_{ss})^2 + \beta \left(z_2^2 + c(i_2 - i_1)^2 + I_2 k \right), \quad (15)$$

where c is the relative weight on this penalty term. This modification does not materially change our key results.

While the central bank's evident aversion to reversals can rationalize a wait-and-see approach to policy, it is unclear how costly such reversals actually are for the real economy. One possibility is that the perceived costs are self-reinforcing in that the rarity of reversals causes them to be viewed as highly unusual when they do occur. This could, in turn, increase market volatility and be viewed as reflecting a policy error rather than a reaction to new information. Reversal aversion is an area that merits additional research to better understand whether indeed the costs to the real economy are as high as feared. At least one study (Battellino, Broadbent and Lowe 1997) suggests that the costs in terms of increased financial market volatility in Australia are small. This is important because this result raises the possibility that central banks may overly limit their ability to act as a shock absorber.

4.3 The starting point

In July 2014, a link was made in the Bank's policy statement between the starting point for inflation, which was well below target at the time, and the stance of monetary policy in a risk management framework:

At the time, even though we presented a projection in which the forecast risks around inflation were roughly balanced, we gave greater weight to the downside risks than the upside ones. Because inflation had been persistently below target, a negative shock would have mattered more than a positive shock, as inflation would have been driven even farther away from target. In short, we attached greater weight to the downside risks because inflation's starting point was already relatively far from home, and had been for some time.

More recently, Governor Poloz drew a direct link between this asymmetry and the conduct of monetary policy:

Since inflation has been so consistently in the lower half of the target band, our risk-management approach to monetary policy led us to pay greater attention to forces pushing inflation down. This is because when inflation is

already low, a negative shock to the outlook for inflation has more significant policy consequences than a surprise on the upside. Throughout, we wanted to be sure our policy would be sufficiently stimulative to get the economy home. (Poloz 2017)

This latter quote suggests an increased relative concern with downside risks implies that policy will be more accommodative than in the case without uncertainty. This approach is not consistent with the prediction based on the certainty-equivalence framework described in section 3. That is, uncertainty about future shocks should not affect the policy response; policy will be more accommodative if inflation is predicted to be below the target at the relevant policy horizon, but the degree of accommodation should be the same as if the outlook were known with certainty. This point is best illustrated by imagining a situation in which inflation is initially below target because the output gap is negative. For simplicity, we'll suppose that inflation stabilization is the only policy objective, and that there is an n -period lag between policy actions and inflation. In this situation, optimal policy would be to aim to return inflation from its starting point to target in n periods, and in the absence of any future shocks, this is what would occur. However, the central banker knows that in the next period, there will be a shock to the output gap, the sign and magnitude of which are unknown. If the shock is negative inflation will undershoot the target at the n -period horizon, whereas if it's positive but of equal size, the target will be overshoot by the same amount. Since the deviation of inflation from target (in absolute value) is the same in both cases, a symmetric loss function implies that the optimal policy response to this shock should be identical (in absolute value). By this logic, the introduction of symmetric uncertainty about the future evolution of output and inflation should not alter the optimal path of the policy rate. It's true that the realization of the downside risk will cause loss to be higher in the periods preceding period n , relative to an upside risk of equal magnitude, but this is beyond the control of the central bank because of the assumed lag in the transmission mechanism.

In the simple example just presented, the central bank doesn't have to weigh competing objectives, whereas in reality this may be the case. For example, the optimal path for the policy rate may imply that inflation is still below the target at the relevant policy horizon because the central bank has to consider the impact of lower interest rates on financial stability. In this case, the visualization is more difficult because the loss func-

tion is two-dimensional, but the certainty-equivalence result still holds. The realization of positive shock to the output gap in a future period will still elicit the same policy response as a negative shock of equal size. By tightening policy in the face of a positive output gap shock, inflation may still fall short of the target but financial stability risks will be further reduced.

There are deviations from certainty equivalence that could give rise to this kind of asymmetric policy response. For instance, policy-makers may be concerned about the ELB, or that large or persistent deviations of inflation may affect inflation expectations disproportionately. While reasons are consistent with an asymmetric policy response, they also imply a negative skew in the balance of risks to inflation. In contrast, the quote that we refer to above describes the risks as roughly balanced.

A motivation that does not imply a skew in the balance of risks is linked to central bank preferences. This is most clearly evident in Governor Poloz's most recent statement that "...when inflation is already low, a negative shock to the outlook for inflation has more significant policy consequences than a surprise on the upside." (Poloz 2017) In this case, the risk of inflation falling outside of the 1 to 3 per cent control range could be the predominant concern when the starting point for inflation is far from the midpoint of the range. A non-linear loss function, such as the zone-quadratic set-up described in section 3.2.1 (Orphanides and Wieland 2000b), combined with a trade-off faced by the central bank in achieving multiple policy objectives, would be consistent with greater policy stimulus (tightening) when inflation is initially below (above) the target, relative to the case in which there is no uncertainty. Such a policy trade-off is required to rationalize an optimal path for inflation that would deviate from the target over the projection horizon, absent uncertainty. For instance, a trade-off could arise from a desire on the part of the central bank to smooth interest rates.^{16,17}

The 1 to 3 per cent target range, as opposed to just a point target of 2 per cent, primarily reflects the degree of imprecision associated with inflation targeting. For instance, even if the policy path is chosen perfectly to bring inflation to target over the medium

¹⁶A central bank could, in theory, adjust interest rates in every period by enough to perfectly stabilize inflation. Even if this were possible in practice it would likely induce an unacceptably high volatility in interest rates.

¹⁷A lag between policy and inflation is insufficient, because while it would rationalize a period of time during which inflation is expected to deviate from the target over the projection, it would not justify a different path for policy in the presence of uncertainty, since varying the path of the policy rate, by definition, would have no impact on inflation over that interval.

term, unforeseen developments (such as oil price shocks) in the interim will mean inflation never hits the target exactly (see Poloz 2017). However, once such a target range is chosen, there is naturally an expectation that the central bank will achieve it under normal circumstances. This doesn't mean the central bank is indifferent to movements within the range, but rather their loss function becomes more convex, meaning loss increases more disproportionately, once inflation is outside the range. Preferences such as this, combined with a trade-off between policy objectives and uncertainty about the future path of inflation, could explain a greater loosening (tightening) bias associated with inflation that is initially well below (above) the target. From the viewpoint of those who wish to better understand a central bank's motivations for choosing a particular path for policy, information regarding preferences would be useful where possible.

5 Conclusion

Central banks have always needed to chart the right course for monetary policy in the face of uncertainty. The Bank of Canada has devoted considerable effort to reduce uncertainty by being clear about the objectives of monetary policy, investing in state-of-the-art forecasting models and expanding the sources of data and information we can use. Despite these efforts, there will always be an inherent degree of uncertainty that is irreducible. This is especially true for sources of uncertainty that are outside the central bank's control, such as geopolitical risk and natural disasters. Although policy models are imperfect, they are an indispensable starting point for assessing how best to respond to uncertainty.

Our analysis of relevant examples of Bank of Canada policy confirms that uncertainty does have profound effects on monetary policy, leading policy-makers to deviate substantially from what a typical monetary policy rule would suggest. Consistent with the rich literature on conducting monetary policy under uncertainty, some sources of uncertainty mean that monetary policy should be more aggressive, while others can justify a wait-and-see approach. We have articulated some practical principles to provide a guide to how uncertainty can affect monetary policy decisions. We recognize that judgment will always be needed in the decision-making process, because a single rule that is robust to the complexity of real-world situations does not yet exist. At the same time, judgment should be informed by a policy framework that uses formal modelling exercises when possible.

Our analysis highlights the importance of several areas for future research in order to add more rigour to the framework and better inform the trade-offs associated with a risk management approach to monetary policy. These include efforts to broaden our sources of data and information, improve policy models (especially related to real-financial linkages and heterogeneity), incorporate non-quadratic central bank preferences, develop policy rules that are robust across models, and improve our understanding of unconventional monetary policy tools.

A Appendix A: The Simple Analytics of Reversal Costs

We use a simple two-period model to analyze how an aversion to reversals on the part of the central bank (CB) can lead to optimal waiting in the presence of uncertainty. Let i_t^* be the optimal policy rate setting in the absence of reversal costs. Assume that i_t^* follows a random walk:

$$i_t^* = i_{t-1}^* + \varepsilon_t. \quad (16)$$

The shocks, ε_t , are a reduced-form way of representing all the shocks that could affect the optimal setting of the policy rate. Prior to period $t = 1$ the economy was in steady state (i.e., $\varepsilon_t = 0$ and $i_t^* = i_{ss}$ for all $t < 1$). In periods 1 and 2, the shocks, ε_t , are drawn from a normal distribution with mean 0 and variance σ^2 . Period 2 is the terminal period for this economy. So we have

$$\begin{aligned} i_t^* &= i_{ss} \quad \forall t \leq 0 \\ i_1^* &= i_{ss} + \varepsilon_1 \\ i_2^* &= i_{ss} + \varepsilon_1 + \varepsilon_2. \end{aligned} \quad (17)$$

The actual policy rate, i_t , may or may not be equal to i_t^* in periods 1 and 2. In period 1, the CB chooses whether to take action by setting $i_1 = i_1^*$ or to wait by leaving the policy rate at its steady-state level, $i_1 = i_{ss}$. In period 2, the CB must again choose whether to take action by setting $i_2 = i_2^*$ or to wait by leaving the policy rate at its period 1 level, $i_2 = i_1$. But the CB faces an additional consideration in period 2. If it took an action in period 1, then the CB will incur a cost, k , if it takes an action of the opposite sign in period 2. The cost, k , is measured in terms of the value of the CB's loss function. It is intended to capture the idea of reversal aversion in a simple reduced-form

manner.

Define the deviation of the actual policy rate from i_t^* as

$$z_t \equiv i_t - i_t^*. \quad (18)$$

The flow cost of deviations is z_t^2 , so the CB's loss is

$$L = z_1^2 + \beta(z_2^2 + I_2 k), \quad (19)$$

where k is a fixed cost of reversing the period 1 action and I_2 is an indicator variable:

$$I_2 \equiv \begin{cases} 1 & \text{if } \text{sign}(i_1 - i_{ss}) \neq \text{sign}(i_2 - i_1) \text{ and } i_1 \neq i_{ss} \\ 0 & \text{if } \text{sign}(i_1 - i_{ss}) = \text{sign}(i_2 - i_1) \text{ or } i_1 = i_{ss} \end{cases}. \quad (20)$$

This indicator variable captures the fact that the CB incurs the cost k only if it took action in period 1 ($i_1 \neq i_{ss}$) and it takes an action of the opposite sign in period 2.

In this environment, we can prove three key results:

1. The fixed cost of reversing implies that the CB will only reverse direction in period 2 for sufficiently large shocks. That is, there is a zone of inaction in period 2 when the period 2 shock is of a different sign from the period 1 shock.
2. The risk of needing to reverse itself and incur the fixed cost causes the central bank to take action in period 1 only for sufficiently large shocks. That is, even though the CB faces no direct costs of action in period 1, there is still a zone of inaction in period 1.
3. The absolute magnitude of the period 1 shock required to induce a policy response from the CB is increasing in the variance of the period 2 shocks. That is, the width of the period 1 inaction zone is increasing in the degree of uncertainty about future shocks.

The first result is a fairly obvious consequence of the fixed cost of reversals. The second and third results are less obvious and more interesting.

A.1 Result 1: Period 2 zone of inaction

To prove the first result, we begin by analyzing the situations the CB could face in period 2. If $\text{sign}(i_1^* - i_{ss}) = \text{sign}(i_2^* - i_1)$ or no action is taken in period 1 ($i_1 = i_{ss}$), then $I_2 = 0$. In this case, it will be optimal to set $z_2 = 0$. If $\text{sign}(i_1^* - i_{ss}) \neq \text{sign}(i_2^* - i_1)$ and action is taken in period 1 ($i_1 \neq i_{ss}$), then it may or may not be optimal to incur the reversal cost. If the reversal cost is incurred, then the CB will set $z_2 = 0$, so the period 2 loss will be k . If the reversal cost is not incurred, then the period 2 loss will be $z_2^2 = (i_1 - i_{ss} - \varepsilon_1 - \varepsilon_2)^2 = \varepsilon_2^2$. The CB will be indifferent between incurring and not incurring the reversal cost for values of ε_2 that satisfy

$$\varepsilon_2^2 = k. \quad (21)$$

Denote the solution to this quadratic equation by ε_2^* . Then,

$$\varepsilon_2^* = \pm\sqrt{k}. \quad (22)$$

The positive root is the relevant threshold if the period 1 action was a rate cut. The negative root is the relevant threshold if the period 1 action was a rate hike.

For example, if $\varepsilon_1 > 0$ and a tightening action is taken in period 1 ($i_1 > i_{ss}$), then action will be taken in period 2 iff:

$$\varepsilon_2 > 0 \quad \text{or} \quad \varepsilon_2 < -\sqrt{k}. \quad (23)$$

Alternatively, if $\varepsilon_1 < 0$ and an easing action is taken in period 1 ($i_1 < i_{ss}$), then action will be taken in period 2 iff:

$$\varepsilon_2 < 0 \quad \text{or} \quad \varepsilon_2 > \sqrt{k}. \quad (24)$$

These expressions make clear that there will be a zone of inaction in period 2, if the value of i_2^* calls for a reversal. The width of this zone will be \sqrt{k} . Note that since there is no uncertainty beyond period 2, uncertainty does not affect the width of the period 2 zone. But, as we will show below, uncertainty will affect the CB's behaviour in period 1.

A.2 Result 2: Period 1 zone of inaction

To prove the second result, we need to characterize the expected losses faced by the CB when making its period 1 decision. If $\text{sign}(i_1^* - i_{ss}) = \text{sign}(i_2^* - i_1)$ or no action is taken in period 1 ($i_1 = i_{ss}$), then, as noted above, the CB will set $i_2 = i_2^*$. So, in this case, the expected period 2 loss is trivially

$$E_1 [z_2^2 | \text{sign}(i_1^* - i_{ss}) = \text{sign}(i_2^* - i_1) \text{ or } i_1 = i_{ss}] = 0. \quad (25)$$

The case in which $\text{sign}(i_1^* - i_{ss}) \neq \text{sign}(i_2^* - i_1^*)$ and action is taken in period 1 ($i_1 \neq i_{ss}$) is more complicated. For expositional purposes, we will focus on the case of a positive period 1 shock, but the results hold for a negative shock, *mutatis mutandis*. In this case, the expected period 2 loss can be written as

$$\begin{aligned} E_1 [z_2^2 | \text{sign}(i_1^* - i_{ss}) \neq \text{sign}(i_2^* - i_1^*) \text{ and } i_1 \neq i_{ss}] & \quad (26) \\ &= p \cdot 0 + qk + (1 - p - q) E \left[(i_1 - i_{ss} - \varepsilon_1 - \varepsilon_2)^2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] \\ &= qk + (1 - p - q) E \left[\varepsilon_2^2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right], \end{aligned}$$

where

$$\begin{aligned} p &= \Pr(\varepsilon_2 > 0) & (27) \\ q &= \Pr(\varepsilon_2 < -\sqrt{k}). \end{aligned}$$

If we assume that i_t^* follows the random walk process in (16) and ε_t is normally distributed with mean zero and variance σ^2 , then we can write p and q as

$$\begin{aligned} p &= 1 - \Phi(0) = \frac{1}{2} & (28) \\ q &= \Phi\left(-\frac{\sqrt{k}}{\sigma}\right), \end{aligned}$$

where Φ is the standard normal cdf.

Recall that, if the central bank takes action in period 1, then $z_1 = 0$; if the CB does

not take action in period 1, then $z_1 = \varepsilon_1$. So, the expected losses are

$$L^a = \beta \left[qk + (1 - p - q) E \left[\varepsilon_2^2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] \right] \quad (29)$$

$$L^{na}(\varepsilon_1) = \varepsilon_1^2. \quad (30)$$

To derive the expectation of ε_2^2 above, we can use the formulas for the expectation and variance of a doubly truncated normal random variable together with the definition of variance. The relevant formulas yield

$$E \left[\varepsilon_2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] = \sigma \frac{\phi \left(-\frac{\sqrt{k}}{\sigma} \right) - \frac{1}{\sqrt{2\pi}}}{\frac{1}{2} - \Phi \left(-\frac{\sqrt{k}}{\sigma} \right)} \quad (31)$$

$$Var \left[\varepsilon_2 \mid 0 > \varepsilon_2 > -\sqrt{k} \right] = \sigma^2 \left[1 + \frac{-\frac{\sqrt{k}}{\sigma} \phi \left(-\frac{\sqrt{k}}{\sigma} \right)}{\frac{1}{2} - \Phi \left(-\frac{\sqrt{k}}{\sigma} \right)} - \left(\frac{\phi \left(-\frac{\sqrt{k}}{\sigma} \right) - \frac{1}{\sqrt{2\pi}}}{\frac{1}{2} - \Phi \left(-\frac{\sqrt{k}}{\sigma} \right)} \right)^2 \right]. \quad (32)$$

Note that the definition of variance implies

$$Var \left[\varepsilon_2 \mid 0 > \varepsilon_2 > -\sqrt{k} \right] = E \left[\varepsilon_2^2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] - \left(E \left[\varepsilon_2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] \right)^2. \quad (33)$$

Thus, the expectation we wish to compute is

$$E \left[\varepsilon_2^2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] = Var \left[\varepsilon_2 \mid 0 > \varepsilon_2 > -\sqrt{k} \right] + \left(E \left[\varepsilon_2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] \right)^2. \quad (34)$$

Substituting in the formula given above for a doubly truncated normal random variable yields

$$E \left[\varepsilon_2^2 \mid -\sqrt{k} < \varepsilon_2 < 0 \right] = \sigma^2 \left[1 + \frac{-\frac{\sqrt{k}}{\sigma} \phi \left(-\frac{\sqrt{k}}{\sigma} \right)}{\frac{1}{2} - \Phi \left(-\frac{\sqrt{k}}{\sigma} \right)} \right]. \quad (35)$$

So the expected losses are

$$L^a = \beta \left\{ qk + (1 - p - q) \sigma^2 \left[1 + \frac{-\frac{\sqrt{k}}{\sigma} \phi \left(-\frac{\sqrt{k}}{\sigma} \right)}{\frac{1}{2} - \Phi \left(-\frac{\sqrt{k}}{\sigma} \right)} \right] \right\} \quad (36)$$

$$L^{na}(\varepsilon_1) = \varepsilon_1^2. \quad (37)$$

The CB will choose to take action if $L^{na}(\varepsilon_1) > L^a$. Thus, we can define the critical value, ε_1^* , as

$$L^{na}(\varepsilon_1^*) \equiv L^a \quad (38)$$

or,

$$(\varepsilon_1^*)^2 = \beta \left\{ qk + (1-p-q)\sigma^2 \left[1 + \frac{-\frac{\sqrt{k}}{\sigma}\phi\left(-\frac{\sqrt{k}}{\sigma}\right)}{\frac{1}{2} - \Phi\left(-\frac{\sqrt{k}}{\sigma}\right)} \right] \right\}. \quad (39)$$

This implies

$$\varepsilon_1^* = \pm \sqrt{\beta \left\{ \frac{1}{2}\sigma^2 + \Phi\left(-\frac{\sqrt{k}}{\sigma}\right)(k - \sigma^2) - \sigma\sqrt{k}\phi\left(-\frac{\sqrt{k}}{\sigma}\right) \right\}}. \quad (40)$$

Thus, the central bank will choose to take action in period 1 only if the period 1 shock exceeds ε_1^* in absolute value. This proves the second result.

A.3 Result 3: The width of the period 1 zone of inaction depends on the degree of uncertainty

To prove the third result, we can take the derivative of ε_1^* with respect to σ :

$$\frac{\partial \varepsilon_1^*}{\partial \sigma} = \frac{1}{2} (\varepsilon_1^*)^{-1} \left\{ \beta \left[\sigma - 2\sigma\Phi\left(-\frac{\sqrt{k}}{\sigma}\right) - 2\sqrt{k}\phi\left(-\frac{\sqrt{k}}{\sigma}\right) \right] \right\}. \quad (41)$$

For this derivative to be positive we need the term in brackets to be positive. Equivalently, we need

$$\Phi\left(-\frac{\sqrt{k}}{\sigma}\right) + \frac{\sqrt{k}}{\sigma}\phi\left(-\frac{\sqrt{k}}{\sigma}\right) < \frac{1}{2}. \quad (42)$$

Letting $x \equiv \sqrt{k}/\sigma$ we can write

$$f(x) \equiv \Phi(-x) + x\phi(-x) < \frac{1}{2}. \quad (43)$$

We want to prove that this inequality holds for arbitrary positive values of x . Note that

$$\lim_{x \rightarrow 0} f(x) = \frac{1}{2}. \quad (44)$$

So proving that the inequality holds boils down to proving that $f'(x) < 0$. Note that, using the formula for the derivative of the normal pdf, we can write f as

$$f(x) = \Phi(-x) - \phi'(-x). \quad (45)$$

Then:

$$\begin{aligned} f'(x) &= -\Phi'(-x) - \phi''(-x) \\ &= -\phi(-x) - (x^2 - 1)\phi(-x) \\ &= -\phi(-x) - x^2\phi(-x) + \phi(-x) \\ &= -x^2\phi(-x). \end{aligned}$$

So, $f'(x)$ is negative and the inequality holds for any positive value of $x \equiv \sqrt{k}/\sigma$. This proves the third result—the width of the period 1 zone of inaction is an increasing function of the degree of future shock uncertainty, σ .

B Appendix B: Generalizing the Principle of Certainty Equivalence

Dennis (2005) discusses a linear (or linearized), forward-looking system of equations given by

$$\mathbf{A}_0 \mathbf{y}_t = \mathbf{A}_1 \mathbf{y}_{t-1} + \mathbf{A}_2 E_t \mathbf{y}_{t+1} + \mathbf{A}_3 \mathbf{x}_t + \mathbf{A}_4 E_t \mathbf{x}_{t+1} + \mathbf{A}_5 \mathbf{v}_t, \quad (46)$$

in which \mathbf{y}_t is an $(n \times 1)$ vector of endogenous variables, \mathbf{x}_t is a $(p \times 1)$ vector of policy instruments and $\mathbf{v}_t \sim iid(\mathbf{0}, \mathbf{\Omega})$ is an $(s \times 1, s \leq n)$ vector of disturbances. The loss function can also be generalized to multiple target variables as well as transformations of the policy instruments:

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t (\mathbf{y}'_t \mathbf{W} \mathbf{y}_t + \mathbf{x}'_t \mathbf{Q} \mathbf{x}_t), \quad (47)$$

where \mathbf{W} and \mathbf{Q} are symmetric positive semi-definite weighting matrices. Dennis shows that optimal policy under commitment yields a forward-looking system of linear equations in the endogenous variables, instruments and Lagrange multipliers, λ_t . This system can then be solved to eliminate leads using a variety of methods and the resulting

backward-looking system is given as

$$\begin{bmatrix} \lambda_t \\ \mathbf{y}_t \\ \mathbf{x}_t \end{bmatrix} = \begin{bmatrix} \theta_{11} & \theta_{12} & \mathbf{0} \\ \theta_{21} & \theta_{22} & \mathbf{0} \\ \varphi_1 & \varphi_2 & \mathbf{0} \end{bmatrix} \begin{bmatrix} \lambda_{t-1} \\ \mathbf{y}_{t-1} \\ \mathbf{x}_{t-1} \end{bmatrix} + \begin{bmatrix} \theta_{13} \\ \theta_{23} \\ \varphi_3 \end{bmatrix} [\mathbf{v}_t]. \quad (48)$$

Importantly, the matrices φ_1 and φ_2 depend on the structural parameters of the underlying model and the policy preference parameters, but not on the covariance matrix of model shocks, $\mathbf{\Omega}$. Furthermore, this result is independent of the variables that enter the loss function. So, for instance, a central bank that wishes to avoid instrument volatility or policy reversals will have an incentive to smooth interest rates, but the degree of smoothing will not depend on the level of uncertainty as measured by $\mathbf{\Omega}$.

A similar but somewhat weaker result obtains when policy is constrained to respond optimally to a subset of model variables, such as would be case with an optimized Taylor rule. Specifically, *relative* shock variances will affect the design of policy since they affect the relative importance of the target variables in the central bank's loss function, and therefore the trade-off between stabilization objectives. As a very simple example, consider an economy with just two shocks: a demand shock that pushes output and inflation in the same direction, and a supply shock that moves them in opposite directions. Also assume that while the central bank seeks to stabilize output and inflation, the policy interest rate responds only to inflation. In this set-up, the optimal response to a demand shock will be larger than the optimal response to a supply shock, since the policy response to a supply shock pushes output away from potential output. Therefore, the optimal response to inflation in the policy rule will depend on the relative importance of demand versus supply shocks in the economy.

References

- Al-Nowaihi, A. and L. Stracca. 2002. "Non-standard Central Bank Loss Functions, Skewed Risks and Certainty Equivalence." European Central Bank Working Paper No. 129.
- Baker, S., N. Bloom and S.J. Davis. 2017. "Measuring Economic Policy Uncertainty." Accessed July 27, 2017, www.policyuncertainty.com.
- Bank of Canada, 2014. "Release of the *Monetary Policy Report*." Press Release, 16 July.
- Bank of Canada, 2017. *Monetary Policy Report*, October.
- Battellino, R., J. Broadbent and P. Lowe. 1997. "The Implementation of Monetary Policy in Australia." Reserve Bank of Australia Research Discussion Paper No. 9703.
- Bergo, J. 2006. "Projections, Uncertainty and the Choice of Interest Rate Assumptions in Monetary Policy." Economic Bulletin 1/2006 p. 16, Norges Bank.
- Binette, A. and D. Tchegotarev, forthcoming. "Evaluating real GDP growth forecasts in the Bank of Canada Monetary Policy Report," Bank of Canada Staff Analytical Note.
- Blinder, A. 1998. *Central Banking In Theory and Practice*. Cambridge: MIT Press.
- Blinder, A. 2000. "Monetary Policy at the Zero Lower Bound: Balancing the Risks." *Journal of Money, Credit and Banking* 32(4): 1093–1099.
- Blinder, A. 2006. "Monetary Policy Today: Sixteen Questions and about Twelve Answers." CEPS Working Paper No. 29.
- Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review Papers and Proceedings* 57, 411–25.
- Brock, W., S. Durlauf and K. West. 2003. "Policy Evaluation in Uncertain Economic Environments" (with discussion). *Brookings Papers on Economic Activity* 1: 235–322.

- Cateau, G. 2007. "Monetary Policy under Model and Data-Parameter Uncertainty." *Journal of Monetary Economics* 54(7): 2083–2101.
- Cateau, G. and S. Murchison. 2010. "Monetary Policy Rules in an Uncertain Environment" *Bank of Canada Review* (Spring): 27–39.
- Chadha, J. and P. Schellekens. 1999. "Monetary Policy Loss Functions: Two Cheers for the Quadratic," *Bank of England working paper* 101.
- Craine, R. 1979. "Optimal Monetary Policy with Uncertainty." *Journal of Economic Dynamics and Control* 1(1): 59–83.
- Dennis, R. 2005. "Optimal Policy in Rational Expectations Models: New Solution Algorithms." Federal Reserve Bank of San Francisco Working Paper 2001–09.
- Dorich, J., N. Labelle St-Pierre, V. Lepetyuk and R. Mendes, forthcoming. "Could a Higher Inflation Target Enhance Macroeconomic Stability?" Bank of Canada Staff Working Paper.
- Felstein, M. 2003. "Monetary Policy in an Uncertainty Environment". Speech at Jackson Hole, Wyoming, 30 August.
- Goodfriend, M. 2000. "Financial stability, deflation, and monetary policy." *Monetary and Economic Studies* 19 (S-1), 143–167.
- Hansen, L.P. and T.J. Sargent. 2007. *Robustness*. Princeton University Press.
- Hunt, B. and D. Laxton. 2003. "The Zero-Interest Rate Floor and its Implications for Monetary Policy in Japan." In *Japan's Lost Decade: Policies for Economic Revival*, edited by : T. Callen and T. Ostry J.D. International Monetary Funds, Washington, DC: 179–205.
- Kato, R. and S. Nishiyama. 2005. "Optimal Monetary Policy when Interest Rates are Bounded at Zero." *Journal of Economic Dynamics and Control* 29(1-2): 97–133.
- Kiley, M. and J. Roberts. 2017. "Monetary policy in a low interest rate world." *Brookings Papers on Economic Activity- BPEA Conference Drafts*.
- Kim, J. and F. Ruge-Murcia. 2016. "Extreme Events and Optimal Monetary Policy." *CIREQ Cahiers de recherche* 09–2016.

- Kimura, T. and T. Kurozumi. 2007. "Optimal Monetary Policy in a Micro-Founded Model with Parameter Uncertainty." *Journal of Economic Dynamics and Control* 31(2): 399–431.
- Kozicki, S. and J. Vardy. 2017. "Communication, Monetary Policy and Uncertainty." Bank of Canada Staff Discussion Paper.
- Lei, X. and M. Tseng, forthcoming. "‘Wait and See’ Monetary Policy". *Macroeconomic Dynamics*.
- Orphanides, A. 2001. "Monetary Policy Rules Based on Real-Time Data," *American Economic Review* 91(4), 964–985.
- Orphanides A. and V. Wieland, 2000a. "Efficient Monetary Policy Design Near Price Stability." *Journal of the Japanese and International Economies* 14: 327–365.
- Orphanides A. and V. Wieland, 2000b. "Inflation Zone Targeting." *European Economic Review* 44(7): 1351–1387.
- Poloz, S. 2017. "The Meaning of ‘Data Dependence’: An Economic Progress Report." Speech at St. John’s Board of Trade, St. John’s, Newfoundland and Labrador, 27 September.
- Reifschneider, D. and J.C. Williams, 2000. "Three Lessons for Monetary Policy in a Low-Inflation Era." *Journal of Money, Credit and Banking* 32(4): 936–966.
- Savage, L.J. 1954. "The Foundations of Statistics." Wiley, New York. Dover edition 1972.
- Söderström, U, 2002. "Monetary Policy with Uncertain Parameters." *The Scandinavian Journal of Economics* 104(1), 125–145.
- Svensson, L. 2003. "Optimal Policy with Low-Probability Extreme Events." *NBER Working Papers* 10196.
- Svensson, L. and N. Williams. 2005. "Monetary Policy with Model Uncertainty: Distribution Forecast Targeting." NBER Working Papers 11733.
- Svensson, L. and M. Woodford. 2003. "Indicator Variables for Optimal Policy." *Journal of Monetary Economics* 50(3): 691–720.

Sveriges Riksbank, 1998. "Inflation Report". March: 31.

Swanson, E. 2004. "Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rules," *Macroeconomic Dynamics* 8(1): 27–50.