BANK OF CANADA
BANQUE DU CANADA

Discussion Paper/Document d'analyse
2014-6

Integrating Uncertainty and Monetary Policy-Making: A Practitioner's Perspective

by Stephen S. Poloz

www.bank-banque-canada.ca

Bank of Canada Discussion Paper 2014-6

October 2014

Integrating Uncertainty and Monetary Policy-Making: A Practitioner's Perspective

by

Stephen S. Poloz

Governor of the Bank of Canada
Ottawa, Ontario, Canada K1A 0G9

ISSN 1914-0568

Acknowledgements

An earlier working version of this paper was presented to 2014 Moneco-Econtro Summer Outlook Conference, hosted by the Canadian Association for Business Economics (CABE) in August in Kingston, Ontario. This paper benefited from the discussion at that conference and from comments from Bank of Canada colleagues, including Carolyn Wilkins, Timothy Lane, Agathe Côté, Lawrence Schembri, Lynn Patterson, and Stephen Murchison. All remaining errors and omissions are my own responsibility.

Abstract

This paper discusses how central banking is evolving in light of recent experience, with particular emphasis on the incorporation of uncertainty into policy decision-making. The sort of post-crisis uncertainty that central banks are dealing with today is more profound than that which is typically subjected to rigorous analysis and does not lend itself easily to formal modelling. As a practical matter, the policy-maker is dependent on macro models to develop a coherent monetary policy plan, and this burden of coherence means that fundamental uncertainty must be incorporated explicitly into the policy formulation process. As suggested here, doing so transforms policy formulation from an exercise in reverse engineering to one of risk management, one consequence of which is to inject a little more realism about uncertainty into the policy narrative, while trusting markets to wrestle with the data flow and deliver two-way trading. The evolution is likely to be a long one—researchers are encouraged to keep focusing on developing a practical understanding of how the economy works, one that admits that rules around economic behaviour are not cast in stone, but are almost certainly subject to variation through time and events.

JEL classification: C50, E37, E5, E61
Bank classification: Economic models; Financial stability; Monetary policy framework; Uncertainty and monetary policy

Résumé

Le présent article examine l'évolution que connaissent les banques centrales à la lumière de l'expérience récente et met l'accent sur l'intégration de l'incertitude dans le processus d'élaboration de la politique monétaire. L'incertitude à laquelle les banques centrales sont confrontées dans un contexte d'après-crise est plus profonde que celle faisant généralement l'objet d'analyses rigoureuses et ne se prête pas aussi facilement à la modélisation formelle. Dans la pratique, le décideur dépend des modèles macroéconomiques pour la mise au point d'un plan cohérent en matière de politique monétaire, et en raison de cette exigence de cohérence, l'incertitude fondamentale doit être explicitement prise en compte dans l'élaboration de cette politique. Ce faisant, la formulation de la politique ne repose plus sur un modèle purement technique, mais plutôt sur un processus de gestion des risques. C'est donc dire qu'il faut tenir compte de l'incertitude de façon un peu plus réaliste dans le message de la Banque, tout en se fiant aux marchés pour jongler avec les flux de données et s'ajuster dans les deux sens. L'évolution sera sûrement longue. On encourage les chercheurs à poursuivre leurs efforts pour en arriver à une compréhension pratique des rouages de l'économie, et ainsi reconnaître que les règles entourant les comportements économiques ne sont pas immuables, mais varient presque inévitablement au fil du temps ou des événements.

Classification JEL : C50, E37, E5, E61
Classification de la Banque : Modèles économiques; Stabilité financière; Cadre de la politique monétaire; Incertitude et politique monétaire

1 Introduction

Economists can disagree on just about anything, even when they are actually in violent agreement. But one thing on which they do seem to agree is that the world really has changed in the wake of the 2007–08 global financial crisis. This calls into question the reliability of the tools of the policy-maker's trade. Picture the sailors of yore, accustomed to guidance from the stars in the northern hemisphere, enduring a three-day storm that takes them into the southern hemisphere, where the stars are all different.

Some would say that "uncertainty" has always been an integral part of practical economics and policy-making, so these observations are not really new. Economic models, by their very nature, are overly simplified abstractions in theory and rough approximations in practice. But to categorize the sort of experience we have been through in the past five or six years as just the typical realizations of error terms as defined in our models is to trivialize the situation. Certainly, it does not capture the issue as seen through the eyes of this practitioner.

The purpose of this paper is to offer some early thoughts on how central banking is evolving in light of our recent experience, totally through the practitioner's lens. There is an extensive body of literature on dealing with uncertainty in policy-making, and I will not pretend to do it justice here. I will acknowledge that it provides insight into some of the potential pitfalls we are facing, but will suggest that the actual practice of monetary policy is a little more difficult than this literature makes it appear. In effect, I believe that the sort of uncertainty we are dealing with today is more profound than that which is typically subjected to rigorous analysis—that it simply does not lend itself as easily to either mathematical or empirical analysis, or any real sort of formalization. In short, uncertainty today is more Knightian in nature.

2 The Ideal Situation: Policy-Making as Engineering

It is useful to fix ideas by describing the ideal that lies behind much of the thinking around policy. It has been a sacrosanct belief for some 40 years that monetary policy should be constrained by an explicit target or objective, and a policy rule that achieves that target (as opposed to pure policy discretion). The idea is to provide an explicit nominal anchor for inflation expectations, and then to achieve it, thereby improving decision-making throughout the economy and preventing time inconsistency in policy-making. Central bank rules have evolved from monetary targets, and sometimes exchange rate targets, to a preponderance of inflation targets. Fundamentally, the use of a rule is meant to reduce uncertainty for private sector decision-making by stripping out discretion.

Of course, defining monetary policy with a rule or a target does not make policy mechanical, or remove the need for judgment on the part of the central banker. Even the first policy rules that sought to keep monetary growth on target required that the central banker look ahead and ask how the money supply might be expected to evolve in the absence of further policy action. This is because the very notion that a central bank can "choose" the money supply at each point in time is an oversimplification. With inflation targeting, this point is even more obvious since the policy-maker's influence is even more indirect. Central banks can really only influence most measures of the money supply, and especially inflation, with a lag. They can alter interest rates,

thereby affecting private sector decisions with a lag, and those decisions will affect various measures of the money supply and the core behavioural drivers of inflation such as the output gap and wage behaviour and, eventually, prices. This transmission mechanism is widely believed to take several quarters from policy move to inflation impact. In effect, the policy target translates into a reaction function for interest rates, and the actual target outcomes depend on a host of other factors.

Consequently, understanding the economy is fundamental to the policy-maker's job, even if policy is defined as a rule and that understanding is generally approximated by a model. Whether large or small, models decompose economic developments into the part that is understood and the noise. The latter may be taken as a proxy for the policy-maker's "uncertainty," but in the ideal world, this notion of uncertainty might better be described as "benign." This is because, in the ideal situation, the error terms of the policy-maker's model are additive and are believed to be white noise, so the policy-maker can act as if they simply do not exist. If the error terms are white noise, they will cancel out on average and, given that they enter the model additively, so will their impact on the variables the policy-maker cares about.

Accordingly, in this ideal situation, monetary policy-making can be approximated in a fairly formulaic fashion. The monetary authority takes a model of the economy that is considered reliable. A forecast of the economy is developed using all available data and analytical judgment, and the model is solved to produce a path for interest rates that will deliver the policy target, conditional on the remainder of the forecast. One can decompose the process into two steps if preferred: run the forecast forward assuming a no-change scenario for monetary policy, and if that scenario implies that inflation is likely to differ from its target between now and two years from now, then an alternative scenario is developed using the model, one that is engineered with a new path for interest rates that is more likely to deliver the desired policy outcome. Normally, the objective would be to aim for the midpoint of the target range over a reasonable time frame (say between six to eight quarters) so that if another shock occurs along the way there is still an acceptable range of outcomes around that policy goal. When new data are released that show that the central bank has missed its target, the process of updating is akin to a course correction by an airline pilot who must fly around a thunderstorm and calculate a new flight plan—the computer will get them there and tell them their new arrival time.

3 Policy Reality: Less than Ideal

The actual situation differs quite a lot from the ideal one. The macroeconomic model is not just an approximation in theory; it is a living thing with all kinds of error terms in the recent period, many of which have demonstrated persistence. The forecasters who use the model must deal with these error terms in some way right from the start. For example, suppose the model has been predicting more exports than have occurred for several quarters in a row; the forecaster must investigate and create a story around this error term and effectively forecast its evolution into the future in order to run the model forecast.

The natural way to do this may be to assume that the model will reassert itself in due course, so the error term will decline steadily to zero over the first few quarters of the forecast exercise. Alternatively, the persistence of a forecast error could mean that the model is fundamentally

misspecified, in which case, forcing the error term to converge to zero will produce a biased projection. This type of uncertainty is not benign, but fundamental. This is an entirely different matter, for the nature of the structural model is that almost everything in it depends on almost everything else, so how the forecaster deals with the uncertainty has effects elsewhere in the model as well.

Indeed, it is the interdependencies of the variables in the macro economy that make a structural model both interesting and a necessary tool for the policy-maker. Fundamentally, the model imposes desirable constraints on the various components of a forecast. It represents our best understanding of how the variables fit together structurally, at least eventually. Accordingly, one cannot form an independent judgment on the outlook for, say, GDP and another independent judgment on the outlook for inflation, since the two projected paths must be coherent with one another, connected by the underlying economic structure. This coherence applies to a multitude of variables in the economy. This is important because taking this interdependence on board is the only way the policy-maker can develop an interest rate path that will interact through the structure of the model to deliver the inflation path that is desired.

Some forecasters place less value on this internal consistency between macroeconomic variables than others. This may be because they are primarily interested in a short-term forecast. Many work without the benefit of a complete structural macro model, instead working with pure statistical models of individual variables, and some work without formal models entirely. Those forecasters can then generate individual forecasts for key macro variables—output and inflation, for example—that are not constrained or connected in any way by an underlying macroeconomic structure. The policy-maker simply cannot do this, given the requirement that the appropriate policy path consistent with the rest of the projection be engineered from the model over a longer horizon. What this means is that any source of bias in the policy-maker's model—such as a persistent error term and the forecaster's projection of it—will affect the policy-maker's forecasts of such variables as output and inflation through the model's complex interdependencies.

This extra burden of coherence with a macroeconomic structure makes the comparison of forecast performances very difficult, since it places extra value on "theory" as opposed to "statistical fit." More importantly, however, the burden of coherence means that the policy-maker must treat fundamental uncertainty much more seriously than forecasters generally do. What we have called "benign uncertainty" entails expressing a range of possible outcomes around a central forecast. In contrast, the more fundamental sort of uncertainty related to shortcomings with model structure can have explicit policy implications, so must be carefully built into the decision-making process.

4 And the World Has Changed

I asserted earlier that the world has changed in the wake of the 2007–08 financial crisis, and economic models are struggling to cope with those changes. Some may resist this assertion, preferring instead to believe that many of our basic models of economic behaviour are deeply parameterized and not so easily disrupted. But, at a minimum, most would concede that the global economic cycle that the crisis fostered has been so big and so prolonged that we have

moved sufficiently far away from normal, and for a sufficiently long time, that we should be skeptical about existing macro models. A certain amount of remodelling seems inevitable in light of our experience, if only to acknowledge the possibility of some degree of hysteresis in various areas of the economy. Furthermore, even in cases where macro models seem to work as well as before, our post-crisis experience is likely to alter the policy-maker's approach to using models and other information permanently.

By way of illustration, we now believe that the models used for policy-making in the past contained fewer and less detailed real-financial linkages than ideal. Elaborating on those linkages and understanding them better has been a significant research preoccupation since the crisis. Much of that energy has gone into adding richer financial or regulatory structure to existing dynamic stochastic general-equilibrium models, which have become the mainstay of the profession. This work appears promising, and yet the approach clearly has limits—even simple financial frictions produce models of extreme complexity, parameters must often be calibrated as opposed to estimated, and it is very difficult to capture the sort of self-reinforcing behaviour and apparent disequilibria that are generally observed in the lead-up to a crisis. Parallel research, using smaller models to focus on specific financial phenomena observed during crises or capturing the effects of financial sector regulation on behaviour, seems to offer more insight, even if it does not result in the perfect single policy model for the future.

This practitioner's take from this literature has been and remains that two models are better than one, and several models even more so. Our experience since the financial crisis has certainly proved the point again. The Bank has invested in a variety of models to complement its main one, Terms-of-Trade Economic Model, or ToTEM, which is relatively aggregated but extremely theoretically robust, and this should make its parameters less susceptible to shifts in the macro environment. The Large Empirical and Semi-structural (LENS) model, for example, is more disaggregated and imposes less structure on short-term dynamics than ToTEM, providing a tighter fit to the short-term data. It has become very helpful in thinking about alternative policy scenarios. The Bank also has a structure-free data-filtering framework, Canada's Short-Term Indicator (CSI), designed to weight new data in continuously updating our near-term GDP forecast. There are also various models designed to capture more real-financial linkages, including allowing for changes in regulatory architecture. These complementary models help inform the judgments that inevitably get built into any policy decision.

Given the availability of this suite of high-powered tools, and a growing belief that none of them is perfectly suited to the world in which we find ourselves today, it might be fair to say that we are all Bayesians now. And if we were Bayesians already, the weights put on softer information have gone up. The Bank has invested considerably in its gathering of business-related information through its *Business Outlook Survey*, with staff across the country in continuous conversation with companies. These results are cross-checked with a variety of other meaningful business surveys from other organizations. And, Bank management regularly sit down with various industry associations and small groups of company CEOs to get an updated pulse check on the economy, and to test our interpretations of the data against business reality.

These points can be illustrated with two specific issues that arose during 2013–14. The first example concerns the task of forecasting inflation. This is typically based on a variant of the

expectations-augmented Phillips curve, where inflation expectations are built up within the macro model, in effect capturing the notion of rational expectations. However, after some 20 years of successful inflation targeting, it has become increasingly difficult to parameterize the relationship between the output gap and inflation here in Canada; perhaps a case of Goodhart's law or the Lucas critique in action. During the prolonged post-crisis global downturn, inflation first exhibited downside stickiness relative to model forecasts, and then declined even further than predicted later on. This experience was common to several other countries as well.

This experience resulted in an extensive re-examination of alternative inflation specifications, including working with more disaggregated inflation data. The work included comparisons of similar models of other countries, consultations with major Canadian retailers concerning their pricing practices and testing of the importance of a variety of other underlying assumptions. The forecasters then judgmentally developed an error-term profile informed by this research, and ran the new model in the usual way, thereby producing a more appropriate model-based path for monetary policy.

The second example relates to the forecasting of exports. Foreign GDP and exchange rates figure prominently in the explanatory variables in our export model, but when global GDP began to recover, Canada's non-energy exports recovered at first but then began to underperform steadily relative to our model prediction. A more granular analysis of 31 export subsectors, complemented by extensive conversations with Canadian companies in those sectors, helped shed additional light on the puzzle. A subset of export sectors was found to have been underperforming for a much longer period. Indeed, the post-crisis global downturn had caused a large number of Canadian exporting firms in those sectors to close their doors, suggesting that any export recovery from those parts of the economy would be very gradual, and that some loss of market share would almost certainly be permanent. While the forecast puzzle was not fully solved, this investigation allowed the Bank to develop a more realistic forecast for exports that was significantly different from that suggested by its base model. Again, this shift in judgment had a direct impact on the model-consistent suggested path for policy interest rates.

Looking ahead, most of us would acknowledge that the sort of economic and financial cycle we are living in now is almost certainly going to have other legacy effects on economic behaviour. With a shock this persistent, we will spend little time debating the possibility of hysteresis in product or labour markets, instead focusing on how large those effects might be, and whether there are conditions in which they might actually be reversed over time. In particular, a business cycle of six to seven years' duration, perhaps longer, will inevitably leave scars on young people who tried unsuccessfully to enter the workforce during that window. Perhaps they will invest more in education, or perhaps they will collect unpaid work experiences, or perhaps they will live in their parents' basements longer than they otherwise would. Meanwhile, companies that might have survived a typical recession lasting one to two years simply by hunkering down cannot do so for six or seven years. They exit the marketplace, reducing the economy's overall capacity; restoring that capacity once a recovery is in train is a bigger, riskier decision because it requires the creation of a brand-new firm.

How will our models cope with these important economic shifts? Can forecasters be expected to take them into account via judgment alone? It seems to me that these various considerations

cannot be thought of as "benign uncertainty," as defined above. They are more profound than the ones we usually think of as error terms in our models. In effect, as policy-makers, we have little choice but to add them to the elements that we normally park in those confidence intervals or, as I prefer to refer to them, "zones of ignorance."

5 Policy as Risk Management

The above arguments should by now have convinced the reader that policy-making must have moved about as far from a mechanical engineering exercise as one can get. And yet, there is a contradiction here—the complexity of the macro economy, and the multiplicity of channels of interaction that it summarizes, make the use of a model an absolute necessity for policy deliberations. The risks associated with what we used to call "leaning against the wind" or "successive approximation" rule-of-thumb monetary policies are legion. Our models, and the inflation targets that we overlay on them, are an essential accountability framework for the policy-maker. And, the engineering stage—where the policy-maker combines his or her best judgment with a model and a policy target and develops an interest rate path that should deliver the desired target—remains at the heart of the policy conversation, despite its various potential shortcomings.

In recent years, however, increasing amounts of energy have been invested in the analysis of specific risks around that base case. We can begin our deliberations with a base that simply assumes we are in the ideal situation described above. But then we must ask ourselves how the various sources of uncertainty we have identified should influence those deliberations. Of course, we have in mind not the simple "benign uncertainty" that can reasonably be assumed away, except in the sense that policy should remain cautious and gradual. Rather, we have in mind the more fundamental kind of uncertainty in which model shortcomings build biases into the projection. Model interdependencies in turn mean that those biases infect many other key variables, and the coherence imposed on the policy process means that those biases will influence the path for interest rates that has been calculated to deliver the inflation target in undesirable ways, putting the target itself at risk.

Accordingly, we develop key risks that we think should be understood better from a policy viewpoint. We develop ways of acknowledging those risks in our model, and we recalculate the path for interest rates that would be expected to deliver the inflation target. We also draw on other models to develop a richer understanding of those risks. Rather than doing this for a long list of risks, and generating a "cloud" of possible policy paths, we focus on understanding a smaller number of central risks more fully.

The upshot is that for any given base-case economic projection, there will be a wide range of interest rate paths that could ultimately be broadly consistent with the inflation target. The policy-maker weighs the various risks that inform that risk zone, and chooses a path for interest rates that takes a nuanced balance of risks into account—and that interest rate path may depart from the engineered one calculated as the starting point of the discussion.

As complicated as this may seem from a policy-formulation viewpoint, there is more. As noted in the introduction, a new set of risks has moved to the front burner in the wake of the global

financial crisis. These are categorized as "financial stability risks." These run the gamut from excessive leverage in households or financial institutions, to bubbles in equity or bond markets, to real estate overheating, to weaknesses in financial market infrastructure, and so on. In some countries, these risks are actually the responsibility of the central bank, but in Canada, most of these lie outside the central bank's mandate. Indeed, the Bank of Canada perceives monetary policy to be the fourth line of defence against such risks; ahead of monetary policy, we would have (i) the behaviour of individual borrowers and financial intermediaries; (ii) regulatory oversight, which in Canada is the responsibility of the Superintendent of Financial Institutions and the Minister of Finance; and (iii) macroprudential policy, which refers to adjustments to rules around financial system leverage or other parameters to lean against the financial cycle. The Bank of Canada is a key adviser in this space, works with its colleagues at the Office of the Superintendent of Financial Institutions (OSFI), the Financial Consumer Agency of Canada (FCAC), Canada Mortgage and Housing Corporation (CMHC) and the Department of Finance to reach a common understanding of the issues, and monitors the situation carefully, as summarized in its semi-annual *Financial System Review* (FSR). The FSR has become an essential complement to our quarterly *Monetary Policy Report* (MPR).

When it comes to the monetary policy decision, then, financial stability issues amount to a set of risks that are "taken into account" by the Bank. They are not generally seen as a significant constraint on monetary policy actions, but when the flexibility inherent in getting inflation back to target within a reasonable time frame permits the Bank to do so, it can choose its tactics in such a way that it does not unintentionally make financial stability concerns worse through its actions. In extreme conditions, of course, where financial stability risks were actually believed to be constraining monetary policy from achieving its inflation target over a reasonable time frame, there would be grounds for considering some macroprudential adjustments.

Accordingly, the Bank has found its way to deliberating its policy tactics within a risk-management framework, which is defined through its understanding of the various risks associated with economic forecasting and modelling, as well as financial stability risks. Our understanding of these risks, how they might impinge on the core policy objective, and how they might interact with one another is informed through a number of smaller, specialized models, anecdotal and survey evidence, and judgment. In this framework, any number of specific issues might come into play, and the weights placed on them could rise or fall depending on the situation. The essence of this framework is that acting to reduce one type of risk would probably work to increase another, and there is sufficient uncertainty that it is better to do nothing until the picture is clearer.

6 A Case Study: Summer 2013 to Summer 2014

It may be helpful to walk the reader through a familiar example to illustrate how this risk-management framework works in practice. I offer Figure 1 as an informal guide to the discussion. It divides the monetary policy decision-making space into four quadrants—low and high inflation-target risk and below-normal and above-normal financial stability risk—and defines a "neutral zone" where the risks to inflation are tilted to one side or the other. However, financial stability risks are such that the remedy on the inflation front could worsen financial stability risks, and it is better to hold policy unchanged until the outlook is clearer.

Characterizing this balance as a "zone" inserts policy gradualism, which helps avoid the need to reverse policy decisions, which would create even more uncertainty.

To set ideas, let's begin the story in late 2012 and early 2013 (represented by the first dot in Figure 1). The global economic recovery appeared to be on track, Canada's economy was growing but had a fair amount of excess capacity, and inflation was around the low end of the 1 to 3 per cent target range. The Bank's forecast was that exports would pick up through 2013 and significantly in 2014, the output gap would close steadily, and inflation would start to drift up. Consequently, the Bank was indicating that interest rates would begin their inevitable upward adjustment toward more normal levels reasonably soon. Markets acknowledged this forward guidance by, among other things, pricing rate increases into the forward curve. One result was that the Canadian dollar traded higher, around parity with the U.S. dollar. Meanwhile, although the government had undertaken some macroprudential adjustments to take some of the heat out of the housing market, the sector remained hot, and the Bank's expressed tightening bias almost certainly helped to foster the moderations in the housing sector that were to come.

Figure 1: Policy Formulation as Risk Management

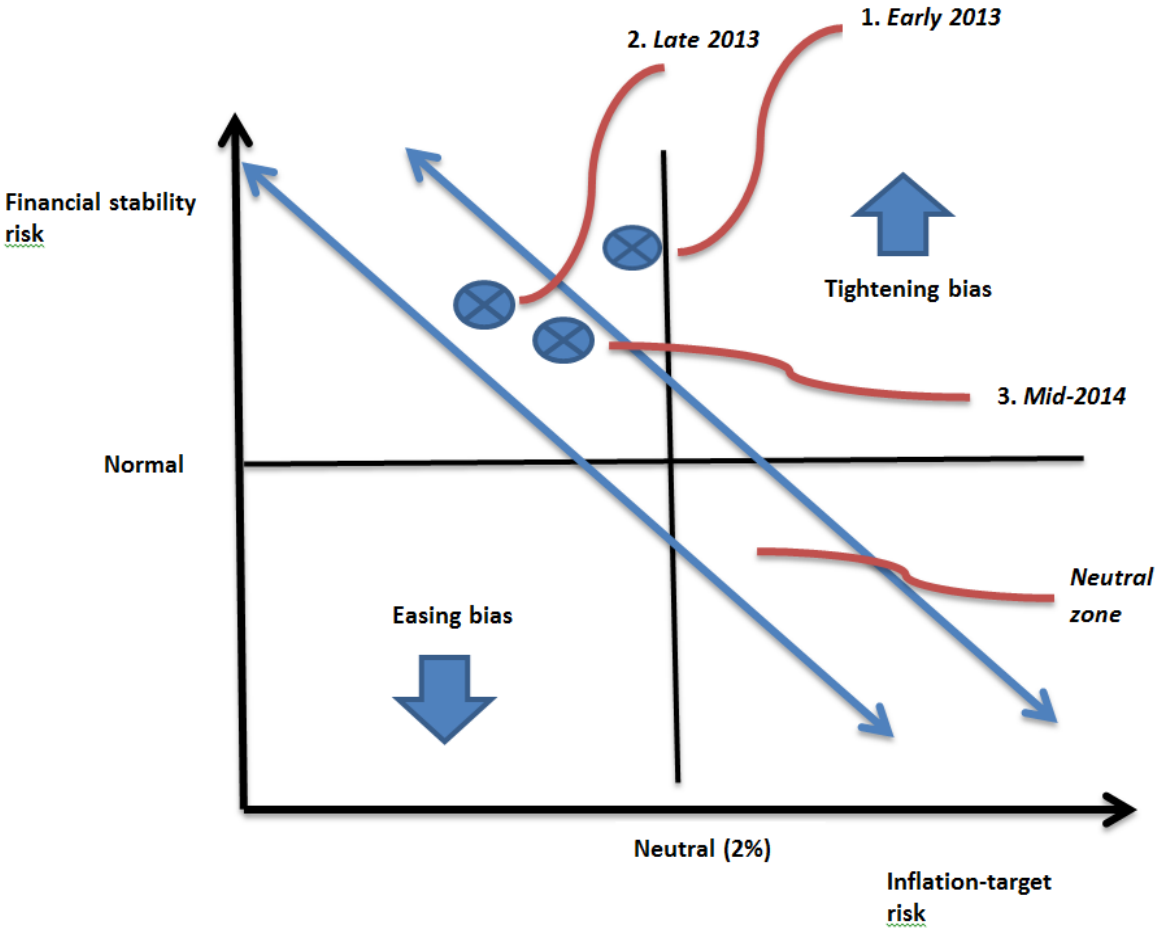

As we moved through the summer of 2013, however, it became increasingly clear that the global expansion would prove more tepid than previously assumed and further, that Canada's exports were underperforming, even relative to that more subdued track for global growth. As our disappointment over the fundamental outlook for Canadian growth—and therefore the projected duration of the output gap and the downward pressure on Canadian inflation—grew, inflation drifted even further below target.

With reference to Figure 1, our situation drifted to the left and downward, into what we describe as the "neutral zone," as indicated by the second dot. Accordingly, in the autumn of 2013, we began to incorporate the notion of risk management explicitly into our policy narrative. In our *Monetary Policy Report*, we set out forecast ranges for potential output and economic growth, to complement our traditional use of confidence intervals around the inflation projection. Our summary judgment was that although we had growing concerns about the downside risks to inflation, we felt that we would still get inflation back to target in about two years (from that point, implying a delay in our return to target compared with earlier projections), and therefore could keep rates stable, thereby avoiding the risk of adding to the heat in the housing market through an actual cut in interest rates. This led markets to downgrade Canada's projected short-term interest rate profile, and, as a consequence, the Canadian dollar moved to a lower range. In effect, the market was reacting to the disappointing data, by simply calculating that the risks to the inflation profile were tilting to the downside.

As this process continued into the autumn of 2013, inflation fell even lower, and we introduced a new concept to our narrative, which we have referred to as "starting-point risk." The idea was simple, if a little tricky to explain: although an updated forecast would, in principle, offer a balanced (symmetric) risk around the inflation profile, if that profile was beginning well below target and would remain below target for some time, the realization of a downside risk to inflation (or any of its fundamental determinants) would matter much more to the policy-maker than an upside risk. Underneath this concern was the risk of expectational drift—that the firmly held expectation that the Bank would achieve its target might erode if inflation remained too long below target.

Therefore, the Bank indicated that it was becoming more preoccupied with the downside risks to inflation, but held interest rates where they were because (i) our analysis suggested that our return to 2 per cent inflation could still be achieved in a reasonable time frame, even if the arrival date was being revised further into the future; and (ii) the possibility of adding to the financial stability risks coming from the household sector was worth taking on board. Expressing this tilt in risks within the risk-management zone led the market to push out the implied path for interest rates further, and the Canadian dollar fell to an even lower level.

As we moved through the first half of 2014, however, measured inflation began to drift back up, and by early summer was quite close to target. Because this occurred in the context of further disappointment with the global and U.S. economies, and no recovery in non-energy exports, our strong belief was that the uptick in inflation was largely due to temporary effects, not domestic fundamentals. These included tight supplies of beef and pork, higher electricity prices, a leap in gasoline prices, and, of course, some measure of exchange rate pass-through. Our fundamental analysis then told us that although the projected date when we would see a closed output gap and

have inflation sustainably at 2 per cent had been pushed out slightly further, it remained within a reasonable time frame. Seeing through the temporary inflation effects, and witnessing a further modest constructive evolution of household imbalances, we judged that we were still in the "neutral zone" (the third dot in Figure 1) and needed to maintain our level of monetary stimulus in order to close the gap. Otherwise, inflation would simply slip back down well below target once the temporary effects had been washed out of the data. Again, the choice was made to hold course, as making a move would serve to worsen some risks and could result in a problem if we were wrong. But we enhanced our communications strategy at this point, offering a sense of our deliberations in our opening remarks at the MPR press conference—in effect, filling a gap between the MPR and its accompanying press release—a practice that will continue.

Some will undoubtedly argue that this potentially makes for slippage around the inflation target. Indeed, there can be little doubt that the projected return of inflation to 2 per cent has been pushed out several times as economic growth has disappointed us. But, provided that inflation expectations continue to be well anchored, and the Bank is aiming to hit the 2 per cent target over a reasonable time frame, it is sensible to remain flexible about the actual speed of reconvergence to target. This flexibility can be of particular value when special circumstances dictate a more gradual approach to policy, such as when financial stability risks are elevated.

The point of this case study is that uncertainty of one form or another is inevitable in policy-making, and the practitioner must take it explicitly on board and be transparent about how risks are impacting their thinking. All of this must be carefully anchored on the mandated inflation target, ensuring that the framework retains its essential, credibility-enhancing logic. To be concrete, managing uncertainty in this way does not constitute a shift in the central bank's certainty-equivalent reaction function. Rather, by being transparent, it shifts some of the policy-maker's uncertainty into the markets for resolution.

7 Forward Guidance: Another Way to Reduce Uncertainty?

As mentioned earlier, monetary policy rules are intended to reduce at least one form of uncertainty faced by market participants; that is, the uncertainty flowing from monetary policy. They also provide a rigorous framework in which markets can evaluate new disturbances—in effect, these rules define a policy reaction to most news, especially data releases. Being transparent about the risks and about any evolution in the Bank's reaction function is, of course, a solid means of reducing uncertainty in the economy.

But during the extreme conditions experienced in the wake of the 2008 financial crisis, central banks have dipped into their toolkits and made use of unconventional monetary policy. Among these tools, one that is relevant to the current topic is forward guidance.

Forward guidance refers to indications given by the central bank to markets about its intentions regarding interest rates, and the factors it will take into account when setting them. Guidance clearly removes a key source of uncertainty from the marketplace. Indeed, in the extreme case, the central banker takes all sources of uncertainty on board, collapses them down to zero, and informs the market of its plan for policy rates. Under certain conditions, this can enhance the power of the single policy instrument to achieve its goal. In particular, when short-term interest

rates are at the zero lower bound, and the yield curve is upward sloping, reflecting a consensus view on the future of monetary policy, the central bank may be able to flatten the yield curve even further by reassuring markets about the horizon over which it expects rates to remain at zero. The idea is that this will enhance the economy's response to low interest rates. The Bank of Canada temporarily exercised this form of guidance in the wake of the financial crisis, and to good effect.

Forward guidance works by taking certain possibilities off the table. For example, the central bank might tell the market that interest rates will not be raised from their current level for at least 12 months. This is tantamount to giving the market a one-way bet. The market responds by positioning itself around that one-way bet, with consequences for other markets that the central bank desires. It is natural that much of this market positioning will be leveraged, the more so the longer the forward guidance is in place. One should not be surprised, therefore, that markets will see some significant volatility when they perceive that the central bank is preparing to change its guidance: the volatility of a return to two-way trading is the future price of successful forward guidance today. That does not invalidate the use of forward guidance, but it underscores that the lunch may not be entirely free.

The other potential downside of forward guidance is that it is inevitably conditional on all the assumptions and forecasts that the central bank must bring to the table. Accordingly, central banks cannot help but insert caveats into their forward guidance. While this is appropriate, what it does is to create a fragile market equilibrium in which every new data point can be interpreted as a potential caveat, and markets may need repeated doses of reassurance. In short, forward guidance can become addictive for markets if it is overly precise or heavily weighted with caveats.

For these reasons, it is my belief that forward guidance should be seen as a useful tool in the central banker's kit, but one that should be reserved primarily for use at the zero lower bound, as a form of additional insurance that the economy will return to equilibrium. Our experience during the second half of 2013, as we moved from the first to the second dot in Figure 1, and, in the process, dropping the previously announced tightening bias, suggests that offering instead full transparency on the risks that the central bank is weighing causes the market to assess new information more or less as the central bank does; and because every data point can give rise to a debate between economists, the market remains two-way and less vulnerable to unusual leveraging and volatile shifts in sentiment. Essentially, the net effect of dropping forward guidance is to shift some of the policy uncertainty from the central bank's plate back onto the market's plate, a more desirable situation in normal times. I would offer the same arguments against those who suggest that the ultimate form of forward guidance, one that is useful in all situations, is for the central bank to publish its interest rate projection. There are both benefits and costs to such an approach, but the main issue is that it shifts too much of the burden of equilibrating *tâtonnement* away from the market and onto the shoulders of the central bank.

8 Conclusion

As economists and policy-makers we know that uncertainty is everywhere and that it has worsened in the aftermath of the global financial crisis. Our tools can only manage some of this.

That is the nature of our craft, yet our constituents usually do not think that way. They demand decimal points in our forecasts, and they demand that we explain why our decimal points differ from someone else's decimal points. And they believe that those decimal points must naturally translate into precision around monetary policy. Accordingly, if a new data point differs from its forecast, they expect that there is a clear, and ultimately quantifiable, implication for monetary policy.

Helping people to appreciate the underlying reality and the limitations of our craft without invalidating our core value proposition is a challenging task. More importantly, the business of central banking is being reinvented in real time in reaction to these realities. At the Bank of Canada, some of the key manifestations of this evolution, as I have tried to motivate above, are:

(i) explicitly building forecast ranges or scenario modelling around key assumption variables, such as potential output, the neutral interest rate and the world price of oil, into our public policy dialogue;
(ii) pointing to key elements of fundamental uncertainty, analyzing the associated policy risks carefully and openly, and laying out complementary research as we learn more about those risks;
(iii) investing more in consultations with Canadian business people and financial market participants, both in the form of surveys and in frank, face-to-face conversations around alternative interpretations of the macroeconomic data;
(iv) bringing a more fulsome narrative to the policy decision-making process, based on a risk-management framework rather than the more conventional policy engineering model; and,
(v) bringing to the table more research on real-financial linkages and financial stability risks to generate a richer set of considerations that influence day-to-day policy thinking.

Some may see this evolution as an erosion in the accountability of central banking. Others may see it as an attack on economic models. Neither of those perspectives is valid. The idea is simply to inject a little more realism about uncertainty into the narrative, while trusting markets to wrestle with the data flow and deliver two-way trading. The idea is also to encourage our researchers to keep focusing on developing a practical understanding of how the economy works—an understanding that admits that rules around economic behaviour are not cast in stone, but are almost certainly subject to variation through time and events.

The bottom line? For the policy practitioner, uncertainty is not abstract, it is a daily preoccupation. Uncertainty, and the policy errors it can foster, must not only be embedded in our decision-making processes *ex ante*, they must be worn like an ill-fitting suit *ex post*—that is, with humility. We know that the sailors of yore who drifted into the southern hemisphere coped. They adapted to the new constellations and found their way—as will we.