BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel—2025-28

Last updated: October 31, 2025

Risk Scenarios and Macroeconomic Forecasts

Kevin Moran
Université Laval
CIRANO
kevin.moran@ecn.ulaval.ca

Stéphane Surprenant
Canadian Economic Analysis Department
Bank of Canada
SSurprenant@bankofcanada.ca

Dalibor Stevanovic
Université du Québec à Montréal
CIRANO
dstevanovic.econ@gmail.com

DOI: https://doi.org/10.34989/swp-2025-28 | ISSN 1701-9397

Abstract

This paper discusses the usefulness of risk scenarios—forecasts conditional on specific future paths for economic variables and shocks—for monitoring the Canadian economy. To do so, we use a vector autoregressive (VAR) approach to produce macroeconomic forecasts conditional on four risk scenarios: high oil prices, a US recession, a tight labor market, and a restrictive monetary policy. The results show that these scenarios represent significant risk factors for the evolution of the Canadian economy. In particular, the high-oil-price scenario is beneficial for the Canadian economy, while a US recession induces a significant slowdown. The very tight labor market scenario leads to additional price increases relative to an unconditional forecast, and the restrictive monetary policy scenario increases the unemployment rate while lowering the inflation rate slightly.

Topics: Econometric and statistical methods, Business fluctuations and cycles, Monetary policy
JEL codes: E32, F41, F44

Résumé

Dans cette étude, nous examinons l'utilité des scénarios de risque – qui sont des prévisions conditionnelles à certaines hypothèses concernant la trajectoire des variables et des chocs économiques – dans un contexte de surveillance de l'économie canadienne. Pour ce faire, nous utilisons une approche vectorielle autorégressive (VAR) pour produire des prévisions macroéconomiques conditionnelles à quatre scénarios de risque : des prix du pétrole élevés, une récession aux États-Unis, un marché du travail tendu et une politique monétaire restrictive. Les résultats montrent que ces scénarios représentent des facteurs de risque importants pour l'évolution de l'économie canadienne. Plus précisément, le scénario prévoyant des prix du pétrole élevés est favorable à l'économie canadienne, tandis qu'une récession aux États-Unis entraîne un ralentissement important. Comparativement aux prévisions inconditionnelles, le scénario supposant un marché du travail très tendu entraîne des hausses de prix supplémentaires, et le scénario de politique monétaire restrictive fait augmenter le taux de chômage tout en réduisant légèrement le taux d'inflation.

Sujets : Méthodes économétriques et statistiques, Cycles et fluctuations économiques, Politique monétaire
Codes JEL : E32, F41, F44

1 Introduction

Economic policy and decisions made in central banks or government departments often rely on economic forecasts. These forecasts represent the best estimates of the future trajectory of economic variables of interest that analysts can establish using available data. These analyses often focus on a *baseline* (or *reference*) scenario, which generally corresponds to the expected future values of the variables because of the attractive statistical properties of expectations.[1]

However, prudent decision-making suggests that it is important to take into account the risks associated with forecasts, i.e., the possibility that the future evolution of the economy deviates from the central trend identified by the baseline forecast. In this context, central banks and governmental officials may also be interested in forecasts related to different *risk scenarios*, which explore the sensitivity of forecasts to specific uncertainties.

For example, the Bank of Canada recently examined the sensitivity of its forecasts by posing the following question: "Although a pronounced global economic slowdown is not the most probable outcome in the coming months, what would be the consequences of such a slowdown on the Canadian economy and the inflation rate if it were to occur nonetheless?" This scenario is analyzed in the April 2023 publication of the *Monetary Policy Report* and reflects the risk management practices of that institution, engaged at the time in a process of monetary policy tightening.[2]

These examples illustrate the interest of public decision-makers in exploring the sensitivity of forecasts to specific eventualities that, while not the most probable, remain possible. Moreover, a scenario could also focus on the future trajectory of the economy in the event that public policy itself undergoes significant changes in the future, in that case reflecting the desire of public decision-makers to explore the consequences of policy

[1]When the goal is to minimize the mean squared error of forecasts, the expected future values—considering the available information in the present—provide the optimal forecasts.

[2]See `https://www.bankofcanada.ca/publication/mpr/`. Alternatively, recent work by Desjardins Economic Studies calculates how housing prices and affordability in the Greater Toronto Area would react to a severe recession in Ontario, although such an outcome was not a baseline forecast of the institution: `https://www.desjardins.com/content/dam/pdf/en/personal/financial-advice/economic-studies/toronto-housing-september-5-2023.pdf`.

changes. As argued by McCracken and McGillicuddy (2019), the range of possible applications for these conditional forecasts is wide, including academic research, banking sectors' stress-testing or monetary policy management.

The formulation of such conditional forecasts should ideally use a systematic methodological framework that produces realistic forecasts that are also consistent across variables. Since the observed historical evolution of economic variables reflects both the various shocks affecting the economy and the interactions linking the variables together, a scenario regarding the future trajectory of a variable should therefore be constructed and analyzed taking both shocks and interactions into account.

This paper shows how these forecasting scenarios can be implemented in a coherent and systematic way when analyzing the Canadian economy. Given the importance of energy in Canada's exports and our economy's significant openness with the United States, the first two scenarios describe how to modify a baseline forecast by making it conditional first on the possibility that oil prices remain high in the coming months and, second, on the possibility that the U.S. economy enters a recession soon. The third scenario reflects important issues in the labor market and shows how forecasts are modified if the recent resilience of this market — likely caused by significant labor shortages in certain sectors — continues. Finally, the fourth scenario reflects uncertainty about the future conduct of the Bank of Canada and the likely magnitude and persistence of the ongoing monetary tightening. To do this, this last scenario calculates how forecasts are modified in the event that Canadian monetary authorities further tighten their policy over the next few quarters.

The analytical framework we employ is that of vector autoregressive (VAR) models, a statistical tool commonly used to produce forecasts. Two important methodological contributions examine the best way to modify these forecasts to make them compatible with risk scenarios (Waggoner and Zha, 1999; Baumeister and Kilian, 2014), and several empirical contributions demonstrate the interest of these methods by computing forecasts compatible with these scenarios (Jarociński, 2010; Giannone et al., 2014; Bańbura et al., 2015; McCracken and McGillicuddy, 2019). The use of this methodological framework also helps emphasize the close relationship between the scenario and the statistical approach: the scenario, its sources, and its consequences are analyzed through the lens of one unified statistical approach.

The results confirm that the scenarios described above represent significant risk factors for the Canadian economy. The scenario of high oil prices and that of a U.S. recession generate paths reminiscent of those resulting from a sequence of positive demand shocks (for high oil prices) and negative demand shocks (for the U.S. recession). In other words, a scenario in which the price of oil remains high is beneficial for the overall Canadian economy, while a U.S. recession induces a severe slowdown. Persistent low unemployment rates generate paths where the inflation rate and housing prices are higher than in the baseline forecast, prompting a more aggressive monetary policy that raises interest rates further. Finally, the scenario of a more restrictive monetary policy than expected does not substantially alter the forecasts, although it results in a slight increase in the unemployment rate and a slight decrease in the inflation rate, relative to the baseline.

This report is structured as follows. Section 2 presents the methodology and the concepts of unconditional and conditional forecasts. Section 3 discusses our contribution, defining the VAR model and presenting the data used for estimation. Section 4 details the four scenarios analyzed, and Section 5 then presents the conditional forecasts resulting from these scenarios. Section 6 checks the robustness of these results to certain statistical assumptions, and Section 7 concludes by suggesting possible avenues for future research.

2 Methodology

2.1 Vector Autoregressive Models

The analytical framework used in this paper is that of VAR models. VARs are statistical tools commonly employed to synthesize the joint evolution of a set of time series, to identify the factors causing fluctuations in these variables, and to perform economic forecasts.

Generally, a VAR is represented by an equation such as

$$\mathbf{y}_t = \nu + \mathbf{A}_1\mathbf{y}_{t-1} + \mathbf{A}_2\mathbf{y}_{t-2} + \cdots + \mathbf{A}_p\mathbf{y}_{t-p} + \mathbf{u}_t, \tag{1}$$

where $\mathbf{y}_t$ denotes a vector containing the K time series used in the analysis. Equation (1) thus constitutes a multivariate autoregression with K variables and p lags. Additionally, ν is a vector of constants; the matrices $\mathbf{A}_1$, $\mathbf{A}_2$, $\mathbf{A}_p$ contain coefficients expressing how different variables interact; and $\mathbf{u}_t$ is a vector of statistical (white noise) innovations (a vector of shocks). These innovations are uncorrelated over time but likely contemporaneously correlated with each other. The notation (1) represents the reduced form of the VAR, in the sense that the shock vector $\mathbf{u}_t$ represents statistical innovations to the value of the variables in $\mathbf{y}_t$, which are not assigned any particular economic interpretation.

The 'companion representation' allows rewriting (1) in a form that encompasses all p lags in one. This representation is as follows:

$$\begin{bmatrix} \mathbf{y}_t \\ \mathbf{y}_{t-1} \\ \vdots \\ \mathbf{y}_{t-p+1} \end{bmatrix} = \begin{bmatrix} \nu \\ \mathbf{0}_{K\times 1} \\ \vdots \\ \mathbf{0}_{K\times 1} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_1 & \dots & \mathbf{A}_{p-1} & \mathbf{A}_p \\ \mathbf{I}_K & \dots & \mathbf{0}_K & \mathbf{0}_K \\ \vdots & \ddots & \vdots & \vdots \\ \mathbf{0}_K & \dots & \mathbf{I}_K & \mathbf{0}_K \end{bmatrix} \begin{bmatrix} \mathbf{y}_{t-1} \\ \mathbf{y}_{t-2} \\ \vdots \\ \mathbf{y}_{t-p} \end{bmatrix} + \begin{bmatrix} \mathbf{u}_t \\ \mathbf{0}_{K\times 1} \\ \vdots \\ \mathbf{0}_{K\times 1} \end{bmatrix}.$$

Alternatively, in a more concise manner:

$$\mathbf{Y}_t = \mu + \mathbf{A}\mathbf{Y}_{t-1} + \mathbf{U}_t, \tag{2}$$

where the uppercase vector $\mathbf{Y}_t$ now contains all lags of the initial vector $\mathbf{y}_t$ and all parameters of the system are now contained in the vector μ and the matrix $\mathbf{A}$.

A VAR like (1) or (2) aims to synthesize the macroeconomic evolution of the economy under study. Choosing which and how many variables to include in $\mathbf{y}_t$ is an important decision that seeks to balance two opposing effects: including more variables allows for incorporating more information and reducing problems related to omitted variables, but it requires estimating an increasing number of parameters (those of ν and matrices $\mathbf{A}_1$, $\mathbf{A}_2$, ..., $\mathbf{A}_p$). The estimation of (1)-(2) is generally performed using ordinary least squares and therefore does not pose a technical problem.

2.2 Forecasts: Baseline Scenario

Consider that we are at time t, and the VAR has been estimated using data covering the period from 1 to t. Furthermore, suppose that our interest lies in using these estimation results and the available data to forecast the future values of all variables y_{t+h}, with $h = 1, \dots H$, where H represents the maximum horizon of interest.

The first step is to construct the *baseline* forecasts, i.e., the conditional expectation of the variables for the sequence of horizons $h = 1, \dots, H$. To do this, we first use the companion representation (2) recursively to obtain the value of $\mathbf{Y}_{t+h}$ for any value of h:

$$\mathbf{Y}_{t+h} = \left(\sum_{j=0}^{h-1} \mathbf{A}^j \right) \mu + \mathbf{A}^h \mathbf{Y}_t + \sum_{j=0}^{h-1} \mathbf{A}^j \mathbf{U}_{t+h-j}. \tag{3}$$

The baseline forecast is then the expectation of the vector of variables at period $t + h$ given the information available at time t, i.e., $\mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^{t}\right)$. By using (3), we obtain

$$\mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^{t}\right) = \left(\sum_{j=0}^{h-1} \mathbf{J}\mathbf{A}^j \mathbf{J}' \right) \nu + \mathbf{J}\mathbf{A}^h \mathbf{Y}_t, \tag{4}$$

where $\mathbf{J} := \begin{bmatrix} \mathbf{I}_K & \mathbf{0}_K & \dots & \mathbf{0}_K \end{bmatrix}$ is a selection matrix and the calculations use the assumption that the expectation of all future values for the innovations $\mathbf{u}_{t+h}$ is zero $\forall\, h \geq 1$, which cancels the third term on the right-hand side of (3). $\mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^{t}\right)$ represents the baseline scenario and is the optimal forecast, assuming a quadratic loss function, of future values of $\mathbf{y}_{t+h}$ given the data available up to period t and the estimated model.[3]

[3] See Hamilton (1994) for a discussion on the properties of the conditional expectation as a forecast.

2.3 Conditional Forecasts

Although $\mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^{t}\right)$ corresponds to the optimal forecast, public institutions may also be interested in forecasts compatible with risk scenarios. As mentioned above, these scenarios explore the sensitivity of forecasts to specific eventualities that, while not the most probable, remain possible. The methods of Waggoner and Zha (1999) and Baumeister and Kilian (2014), WZ and BK hereafter, can be used to produce these alternative scenarios.

2.3.1 The Waggoner and Zha (1999) Method

This approach first entails defining a sequence of future values for one or more variables within the model and then computing forecasts for the remaining variables of the model that are contingent upon and consistent with these specified trajectories. For instance, such a scenario might seek to determine a forecast for future gross domestic product (GDP) given the condition that the inflation rate remains at its current level.

These conditional forecasts will differ from those of the baseline for the following reason: a scenario imposing a given trajectory for certain variables in the vector $\mathbf{y}_t$ is equivalent to imposing a specific future evolution for the model's shocks. In other words, the alternative trajectory for the variable targeted by the scenario can only arise because shocks have combined to make it possible, and these shocks will influence the forecasts of all other variables in the model.

To show this, we can use equations (3) and (4) to compute the difference between the actual future values of $\mathbf{y}_{t+h}$ and the forecasts derived from the baseline scenario:

$$\mathbf{y}_{t+h} - \mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^{t}\right) = \sum_{j=0}^{h-1} \mathbf{J}\mathbf{A}^j\mathbf{J}'\mathbf{u}_{t+h-j}. \tag{5}$$

This expression shows that imposing that future values of certain variables in the vector $\mathbf{y}_t$ be different from the baseline forecasts (left-hand side of the equation) is equivalent to imposing that a specific evolution of the error terms has occurred (right-hand side).

Blake and Mumtaz (2017) show that this logic results in a restriction of the form $\mathbf{Ru} = \mathbf{r}$. The vector $\mathbf{r}$ is $MH \times 1$ for $0 < M < K$, the number of variables constrained over H periods, and is equivalent to the difference between the constrained values and the baseline forecasts for these M variables: the left-hand side of (5). The matrix $\mathbf{R}$ corresponds to the right-hand side of (5) and contains the coefficients of moving averages, i.e. the elements of $\mathbf{\Phi}_h = \mathbf{J}\mathbf{A}^h\mathbf{J}'$ for horizons $h = 1, \ldots, H$, since these represent the change in values relative to the baseline forecast. Finally, the vector $\mathbf{u}$ contains all future innovations stacked vertically, i.e., $\mathbf{u} := \begin{bmatrix} u_{1,t+1} & \ldots & u_{K,t+1} & \ldots & u_{K,t+H} \end{bmatrix}'$. Note that both $\mathbf{R}$ and $\mathbf{r}$ are known once the VAR has been estimated.

There is no unique solution of shocks that respects the constraint $\mathbf{Ru} = \mathbf{r}$: intuitively, this is because several possible combinations of shocks could have generated a trajectory compatible with the postulated scenario. However, by invoking Doan et al. (1984), we can find the shock path that respects the constraint and minimizes the sum of squares for this equation, so that $\hat{\mathbf{u}} = \mathbf{R}' \left(\mathbf{R}'\mathbf{R}\right)^{-1} \mathbf{r}$. This $\hat{\mathbf{u}}$ represents the most likely combination of shocks that could have generated the trajectory postulated by the scenario. Once identified, this shock path $\hat{\mathbf{u}}$ can then be added to the right-hand side of (5) and the conditional forecasts are constructed by adding this correction to the unconditional ones.

2.3.2 The Baumeister and Killian (2014) Method

This approach differs from that of WZ in two important ways. First, the structural representation of the VAR is used, which requires that a structural or economic interpretation be given to the shocks affecting the VAR. The VAR (1) is rewritten as:

$$\mathbf{y}_t = \nu + \mathbf{A}_1\mathbf{y}_{t-1} + \mathbf{A}_2\mathbf{y}_{t-2} + \cdots + \mathbf{A}_p\mathbf{y}_{t-p} + \mathbf{D}\epsilon_t, \tag{6}$$

where the statistical innovations $\mathbf{u}_t$ have been replaced by the term $\mathbf{D}\epsilon_t$, with ϵ_t the structural shocks, which in addition to being white noise like $\mathbf{u}_t$ are also assumed to be uncorrelated. The matrix $\mathbf{D}$, which needs to be identified, shows how the statistical innovations $\mathbf{u}_t$ in (1) are actually weighted sums of structural shocks ($\mathbf{u}_t = \mathbf{D}\epsilon_t$).

The second difference is that the conditional forecast is obtained by making assumptions about future sequences of the structural shocks ϵ_t, rather than assumptions made about the future values of the variables $\mathbf{y}_t$ themselves, as in WZ.

As before, the method starts with the baseline forecast $\mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^t\right)$. However, it now computes the relevant modification to this forecast when a specific scenario for the future trajectory of the structural shocks ϵ_t is analyzed. We thus have

$$
\begin{aligned}
&\mathbb{E}\left(\mathbf{y}_{t+h} | \left\{\epsilon_{t+h-j} = \epsilon_{t+h-j}^{scenario}\right\}_{j=0}^{h-1}, \{\mathbf{y}_s\}_{s=1}^t\right) = \\
&\mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^t\right) + \left[\mathbb{E}\left(\mathbf{y}_{t+h} | \left\{\epsilon_{t+h-j} = \epsilon_{t+h-j}^{scenario}\right\}_{j=0}^{h-1}, \{\mathbf{y}_s\}_{s=1}^t\right) - \mathbb{E}\left(\mathbf{y}_{t+h} | \{\mathbf{y}_s\}_{s=1}^t\right)\right].
\end{aligned}
$$

A solution for the unconditional forecast (the first term of the expression) has already been derived using (4) above, so an expression for the second term remains to be obtained. By inspecting equations (3) and (4) and using the equivalence $\mathbf{u}_t = \mathbf{D}\epsilon_t$, it can be shown that the bracketed second term, which needs to be subtracted, is equal to

$$
\sum_{j=0}^{h-1} \mathbf{A}^j \mathbf{U}_{t+h-j} = \sum_{j=0}^{h-1} \mathbf{J}\mathbf{A}^j \mathbf{J}' \mathbf{u}_{t+h-j} = \sum_{j=0}^{h-1} \mathbf{J}\mathbf{A}^j \mathbf{J}' \mathbf{D}\epsilon_{t+h-j} := \sum_{j=0}^{h-1} \mathbf{\Theta}_j \epsilon_{t+h-j},
$$

where $\mathbf{\Theta}_j := \mathbf{J}\mathbf{A}^j\mathbf{J}'\mathbf{D}$ corresponds to the matrix of impulse response functions of the model variables to structural shocks, for horizon j. It suffices then to introduce the specific scenario for the shock sequence $\epsilon_{t+h-j} = \epsilon_{t+h-j}^{scenario}$ to obtain the second term.

As mentioned above, this approach requires the identification of the structural impact matrix $\mathbf{D}$, which poses the following identification problem. While an estimate for the covariance matrix of $\mathbf{u}_t$ is obtained during the estimation process of the VAR, as $\hat{\mathbf{\Sigma}}_u := \sum_{t=1}^T \hat{\mathbf{u}}_t \hat{\mathbf{u}}_t' / (T - Kp - 1)$ where $\hat{\mathbf{u}}_t$ represents the ordinary least squares residuals, that matrix is symmetric and thus contains only $K(K+1)/2$ distinct elements. In contrast, the matrix $\mathbf{D}$ contains K^2 distinct elements but is related to $\mathbf{\Sigma}_u$ by $\mathbf{\Sigma}_u := \mathbb{E}\mathbf{u}_t\mathbf{u}_t' = DD'$. Some assumptions must therefore be made to add the missing equations that are necessary to identify the matrix $\mathbf{D}$ from $\hat{\mathbf{\Sigma}}_u$.

This paper employs a short-term recursive approach to complete the identification. This approach assumes that the matrix $\mathbf{D}$ has a lower triangular form, which provides the $K(K-1)/2$ missing restrictions to complete the identification. Intuitively, this imposes that the shock at position k affects variables at positions $1, \ldots, k-1$ in the vector $\mathbf{y}_t$ only with a delay, but variables at positions $k, \ldots, K$ contemporaneously. We estimate $\mathbf{D}$ by the Cholesky decomposition of $\hat{\mathbf{\Sigma}}_u$. Practically, this identification strategy requires careful selection of the order in which variables enter the vector $\mathbf{y}_t$, so that the timing restrictions in the causal order are plausible.

2.4 Variance Decomposition

Once the structural shocks affecting the evolution of a VAR have been identified, it is possible to obtain a measure of their relative importance for each of the variables, by performing a decomposition of the forecast errors' variances. This decomposition measures, for each of the identified shocks, each of the variables, and any forecast horizon, the importance of the shock in the variability of the variable.

This decomposition is performed as follows. Using the notation of Kilian and Lütkepohl (2017), let us note that the forecast error in the vector of variables $\mathbf{y}$ at horizon h is related to the structural shocks by $\mathbf{y}_{t+h} - \mathbb{E}(\mathbf{y}_{t+h}) = \sum_{j=0}^{h-1} \mathbf{\Theta}_j \epsilon_{t+h-j}$. Hence, we have the following expression for the mean squared forecast error:

$$MSPE(h) = \mathbb{E}\left((\mathbf{y}_{t+h} - \mathbb{E}(\mathbf{y}_{t+h}))\,(\mathbf{y}_{t+h} - \mathbb{E}(\mathbf{y}_{t+h}))' \right) = \sum_{j=0}^{h-1} \mathbf{\Theta}_j \mathbf{\Theta}_j'.$$

Since the response of variable k to shock l at horizon h is given by $\theta_{k,j,h}$, the (k,l) element of $\mathbf{\Theta}_h$, the contribution of shock l to the variance of the forecast error of variable k at horizon h is given by

$$MSPE_{k,l}(h) = \theta_{k,l,0}^2 + \cdots + \theta_{k,l,h-1}^2.$$

The total value (corresponding to the aggregated variability for variable k) simply sums over the contribution of all shocks, that is, $MSPE_k(h) = \sum_{l=1}^{K} MSPE_{k,l}(h)$. Finally, the *relative* decomposition sought is obtained by dividing both sides by $MSPE_k(h)$.

3 Model Specification and Data Used

The model employed for the analysis comprises seven variables and operates at a monthly frequency. The autoregressive order of the VAR, denoted as p in Equation (1), was determined to be 3 based on the Akaike criterion. To account for Canada's substantial exposure to both global markets and the United States, the first three variables included in the VAR consist of the world oil price, an index of U.S. economic activity, and the exchange rate between the U.S. dollar and the Canadian dollar. The remaining four variables encapsulate various dimensions of Canadian economic activity: the consumer price index and a real estate price index to capture price dynamics, the unemployment rate to reflect labor market conditions, and finally, a measure of the stance of Canadian monetary policy.

More specifically, the West Texas Intermediate benchmark price for oil is utilized and U.S. economic activity is represented by the industrial production index, both of which are sourced from the FRED-MD database introduced by McCracken and Ng (2016). Conversely, the Canadian variables are drawn from the extensive macroeconomic database developed and updated by Fortin-Gagnon et al. (2022): the consumer price index corresponds to the CPI, the monetary policy stance is gauged by the official discount rate, and the exchange rate denotes the quantity of Canadian dollars per U.S. dollar (wherein a decline in the exchange rate signifies an appreciation of the Canadian currency).

The period analyzed spans from January 1992 to December 2022, with the starting date selected to mitigate the influence of significant shifts in Canadian monetary policy during the early 1990s on our analysis. The variables are transformed to ensure covariance-stationarity; consequently, we utilize the first differences of the logarithm of each variable, except for the unemployment rate and the discount rate, which are included in levels.

The sequencing of variables within the vector $\mathbf{y}_t$ is pivotal for the aforementioned short-term recursive identification. In this regard, our methodology broadly aligns with established practices in several studies utilizing VAR methodology to scrutinize the Canadian economy (Kim and Roubini, 2000; Bhuiyan and Lucas, 2007; Li et al., 2010; Boivin et al., 2010; Moran et al., 2023): global or U.S. variables are prioritized at the forefront of the vector $\mathbf{y}_t$, followed by Canadian-specific variables, with variables deemed particularly responsive to events – such as financial asset prices in Li et al. (2010), or the exchange rate in our study or in Kim and Roubini (2000) – placed last.

The following ordering of variables is therefore used: oil price, U.S. industrial production, the consumer price index, the real estate price index, the unemployment rate, the discount rate, and the exchange rate. As previously mentioned, the positioning of international variables at the beginning of a VAR applied to Canadian data reflects Canada's status as a small open economy, as we don't expect that Canadian-specific economic developments (the Canadian shocks) will have strong impacts on oil prices or U.S. industrial production. Price variables (CPI and real estate price index) precede the Canadian activity variable (unemployment rate), reflecting the premise that these prices exhibit relative rigidity and do not contemporaneously respond to demand shocks, unlike those affecting the unemployment rate and monetary policy. Lastly, the discount rate, measuring the stance of monetary policy, follows economic activity but precedes the exchange rate, reflecting the conjecture that if the Bank of Canada reacts to economic developments within the period, financial markets, represented by the exchange rate, are assumed to exhibit forward-looking behavior and are even more responsive than the discount rate.

4 Scenarios

We consider four scenarios through which the future evolution of the Canadian economy might diverge from its unconditional forecast: (1) a rapid increase in the world oil price followed by a sustained high plateau for several months, (2) an economic slowdown in the United States evolving into a substantial recession, (3) persistently low unemployment rates in Canada, and (4) a further tightening of Canada's monetary policy stance.

The construction of the specific numerical trajectories for these scenarios is guided by the historical evolution of variables, particularly during past economic slowdowns. A graphical summary of the considered scenarios is presented in Figure 1: for each scenario, the observed evolution of the relevant variable up to December 2022 is shown in green, its forecasted evolution under the baseline scenario is in black, and its future evolution under the scenario is represented in yellow. As an example, the top-left panel of the figure presents the expected evolution of the oil price in black and the projected evolution according to the scenario in yellow.

To construct the scenario related to oil prices, we note that the price peaked at around 134 U.S. dollars in June 2008, following several increases spread over a few months. More recently, an increase of similar magnitude drove the price up to 115 U.S. dollars in June 2022. Our scenario is calibrated to be comparable to these events and assumes that the price of oil, starting from an initial level around 80 U.S. dollars, rises to 120 U.S. dollars over a period of 7 months and remains at this plateau for the rest of the forecast period.

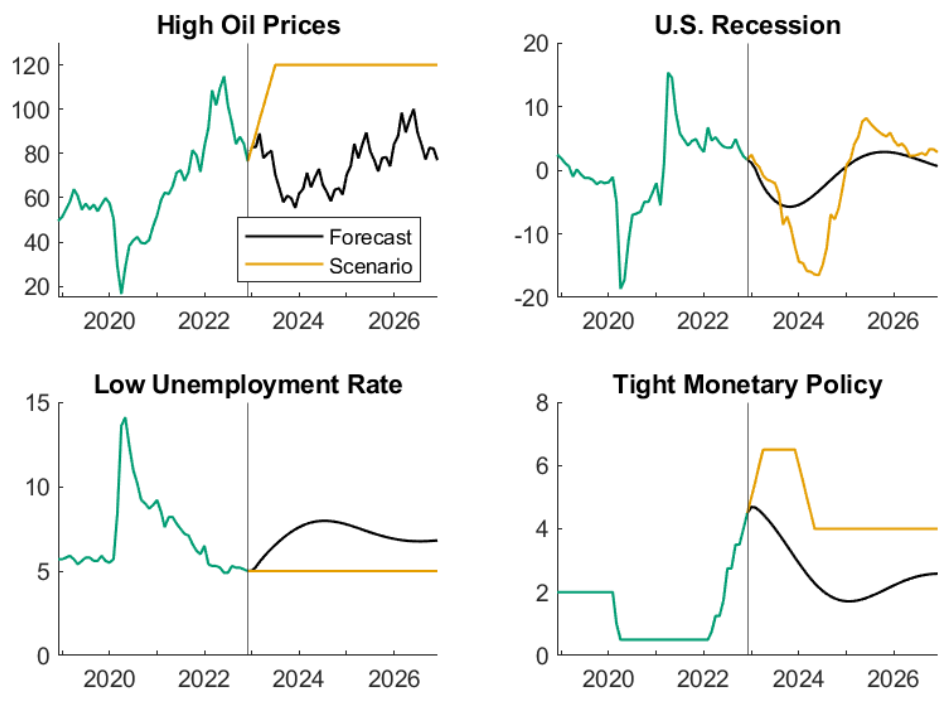

Figure 1: Summary of scenarios

The second scenario is related to U.S. economic activity. To simulate a situation where the United States enters a recession, we suppose that the growth rate of U.S. industrial production follows the same pattern as during the 2008-2009 financial crisis: hence, the scenario duplicates the growth rates displayed by this variable for the 48 months following January 2008.

The third scenario concerns the Canadian unemployment rate: we assume it remains very low, setting it at its value in December 2022 –a very low figure by historical standards– for the entire forecast period. This scenario can be interpreted as reflecting an unexpected and persistent resilience in the labor market, perhaps caused by demographic factors.

Lastly, for monetary policy, we employ a sequence of rate hikes, followed by a pause and a potential decrease in interest rates, similar to those envisioned by some observers predicting that the Bank of Canada's tightening stance will persist for longer than generally expected. In this scenario, the Bank of Canada gradually raises the rate to 6% at a pace of 50 basis points per month, pauses for 8 months, and eventually lowers it back to 4%, where it remains.[4]

We then compute conditional forecasts for all variables under each scenario, using the WZ approach, for the 48-month period starting after the last observation, i.e., from January 2023 to December 2025. Recall that this involves predicting the future evolution of all variables while ensuring that the evolution of the constrained variable follows the path proposed by the scenario, as per equation (5) and shocks $\hat{\mathbf{u}} = \mathbf{R}' \left(\mathbf{R}'\mathbf{R}\right)^{-1} \mathbf{r}$.

5 Results

Figure 2 presents the VAR forecasts for Canadian variables: recall that those are growth rates for prices, growth rates for the real estate price index, the level of the unemployment rate, the level of the Bank of Canada discount rate, and the growth rate for the exchange rate. For each variable, the baseline forecast is shown in black while the forecasts from the

[4]This scenario is valid for the approach of Waggoner and Zha (1999): we describe how to represent a monetary tightening scenario according to the approach of Baumeister and Kilian (2014) below.

scenarios are in purple (oil price scenario), orange (U.S. economic activity scenario), blue (unemployment rate scenario), and yellow (Canadian monetary policy scenario). Additionally, Figure 3 replicates these forecasts but displays them in levels (rather than growth rates) for real estate prices and the exchange rate.[5]

Analyzing the results in Figure 2 is helped by a variance decomposition exercise similar to that analyzed in Section 2.4. This decomposition, reported in Table 1 of Appendix A, highlights the importance of oil price shocks in explaining variability in the inflation rate (first column of the table), the unemployment rate (third column), and the exchange rate (last column). In this context, we expect a conditional forecasting exercise that features oil prices, as in our first scenario, to result in significant deviations between the unconditional forecast and that derived from the scenario. The results below confirm this intuition. In contrast, the second panel of Table 1 suggests that shocks to U.S. economic activity are less important for overall variability in our VAR, implying that deviations between the baseline scenario and a conditional forecast under a U.S. economic slowdown will be more modest.

Scenario 1: Oil Prices Remain High

In this scenario, the world price of oil experiences a sudden increase and remains higher than its level according to the baseline scenario.

The contrast between the economy's evolution under the scenario (purple lines) and the baseline forecast (black lines) shows the responses to be consistent with the idea that, for a net oil exporter like Canada, elevated oil prices are akin to demand shocks. Indeed, the figure shows, on one hand, that the inflation rate and the growth rate in real estate prices decline more slowly than implied by the baseline scenario, while on the other hand, the unemployment rate rises more slowly. Relative to the baseline case, we thus have an expanding economy with prices increasing more rapidly. In response, the Bank of Canada discount rate remains high for a longer period, indicating that the tightening of Canadian

[5] The forecasting ability of the VAR model is examined in Appendix B.

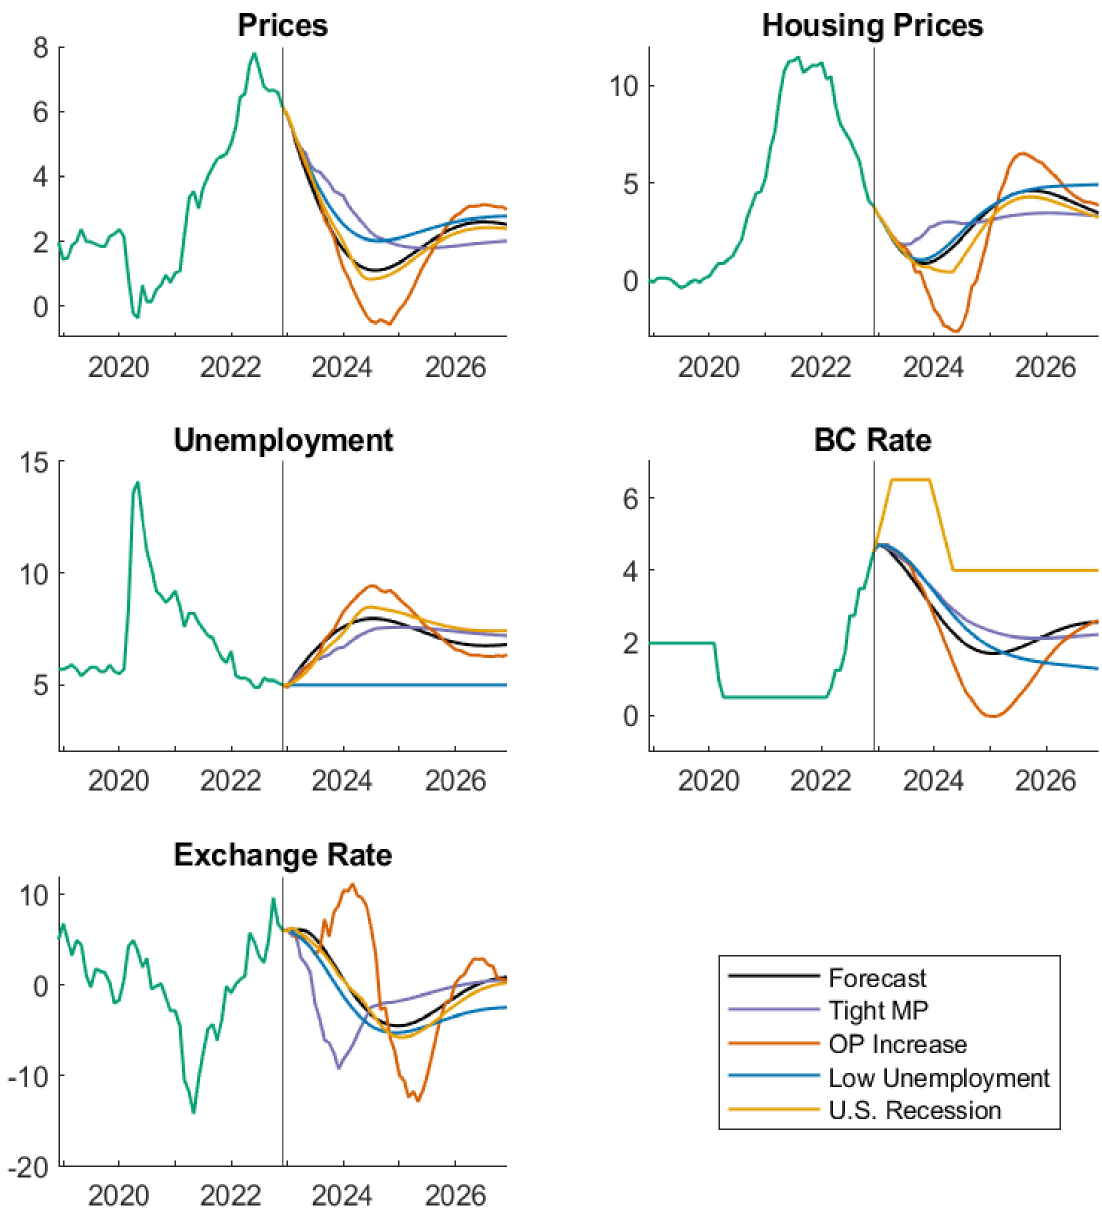

Figure 2: Forecasts for Canadian variables: baseline case and scenarios

Note: Colors indicate different scenarios. The acronyms are Bank of Canada (BC), monetary policy (MP), and oil prices (OP).

monetary policy is prolonged in response to the unexpected rise in oil prices. Finally, the exchange rate follows a substantially lower trajectory than that according to the baseline case, indicating that the Canadian dollar appreciates relative to the U.S. dollar following this shock. Overall, these differences between the trajectories of macroeconomic variables under the scenario and those under the baseline forecast align well with previous studies analyzing the impacts of oil price shocks on the Canadian economy (Hou et al., 2016).

Scenario 2: A U.S. Recession

The second scenario analyzes the consequences that a recession in the United States, of similar magnitude to that of 2008 and occurring in the coming months, would have on the Canadian economy. Figure 2 reports that in this eventuality, a significant contraction would also affect the Canadian economy. The figure indicates that the Canadian unemployment rate would then experience a rapid and substantial increase (dark orange line), relative to the baseline forecast (black line). Relatedly, the growth in prices would be slowed down, with the figure showing decreases in both the inflation rate and the rate of growth in real estate prices. In response to these slowdowns, the stance of monetary policy would transition more rapidly from the current tightening period to a fairly significant easing, highlighted by the substantial and persistent decrease in the discount rate seen in the figure. Finally, the exchange rate would depreciate sharply in the first few months following the shock, before firming up by the 2025-2026 horizon. Overall, these results align well qualitatively with those of other studies analyzing how adverse shocks in the United States also affect the Canadian economy: see Miyamoto and Nguyen (2017) (adverse technological shocks in the United States), Bedock and Stevanovic (2017) (credit shocks), Moran et al. (2022) (uncertainty shocks) or Moran et al. (2023) (shocks to American confidence).

This scenario can be interpreted as a sequence of negative demand shocks over several quarters, before a recovery - positive demand shocks - sets in around 2025. Therefore, a natural similarity with the results associated with the first scenario - oil price increases

- is expected, as they too can be interpreted as demand shocks. However, the forecast deviations between scenario and baseline now appear quantitatively larger, which may seem surprising given the results in Table 1 that suggested shocks from the U.S. economy are less important than oil shocks for Canada (first and second panels of the table). One possible explanation is to recall that we are using a reduced form representation and as such, our U.S. recession scenario may also partially capture impacts from oil shocks: some of what our exercise considers shocks to U.S. economic activity could, in fact, be the response of this economy to an oil price shock. This would be consistent with the interpretation of the oil price scenario and the identification strategy employed in Table 1, which allows U.S. production to immediately respond to oil shocks but imposes that the oil price can only respond to shocks to U.S. production with a delay.

Scenario 3: Unemployment Near Record Lows

The third scenario posits that the very low unemployment rates recently recorded in Canada will persist in the coming months. This could reflect the continuation of the labor shortages that have been reported in several sectors over recent years. This shock, possibly stemming from demographic factors, could be interpreted as negatively affecting labor supply. Standard economic intuition therefore suggests that following this type of shock, real wages would rise, generating upward pressure on prices, but that the evolution of potential GDP would be slowed down.

Figure 2 displays results largely in line with this intuition. The model now anticipates that the inflation rate will be significantly higher during the forecast period (blue line) than its trajectory according to the baseline forecast (black line). This also holds for the growth in real estate prices, although the gap between the two forecasts is initially smaller in this case, before a more significant divergence appears later.

Since the scenario implies that the inflation rate remains higher than in the baseline case, it is expected that the discount rate will also remain high, as the Bank of Canada extends its tightening policy to counteract these additional price increases. Under interest rate parity, it is also expected that this tightening will lead to an appreciation of the Canadian dollar. Overall, the results of Figure 2 are in line with this intuition, particularly at the 2025 horizon at the end of the forecast period. However, the placement of the unemployment rate in the third position in the model's variable vector and the result of the variance decompositions in Table 1 of Appendix A, which report that oil shocks are the predominant source of fluctuations, suggest that the model may assign to oil shocks and U.S. production a portion of the fluctuations originating in the Canadian labor market.

Scenario 4: Tight Monetary Policy

The last scenario considers the central bank aggressively increasing the discount rate before pausing and partially lowering it again. This scenario thus corresponds to an additional tightening in monetary policy, beyond the one already present in the baseline scenario.

Figure 2 shows that the trajectories of the variables following this additional tightening (yellow line), relative to those established by the baseline forecast, are consistent with the extensive empirical literature on shocks to Canadian monetary policy (Kim and Roubini, 2000; Bhuiyan and Lucas, 2007; Li et al., 2010; Champagne and Sekkel, 2018). The unemployment rate increases more than in the baseline case, prices decline further, and the exchange rate experiences a temporary appreciation (possibly followed by a depreciation in more distant periods). However, the deviations between these conditional forecasts and those obtained in the baseline case are small, despite the significant discrepancy in the evolution of the discount rate itself. One possible explanation for this is again found in Table 1 of Appendix A: that table shows (panel D) that monetary policy shocks account for a small fraction of the overall volatility in macroeconomic variables but dominate the evolution of the discount rate itself. Thus, the sequence of shocks that causes the discount rate to deviate from the forecast path in this scenario (the yellow line in the figure) is primarily composed of monetary policy shocks that have only a modest influence on the rest of the variables.

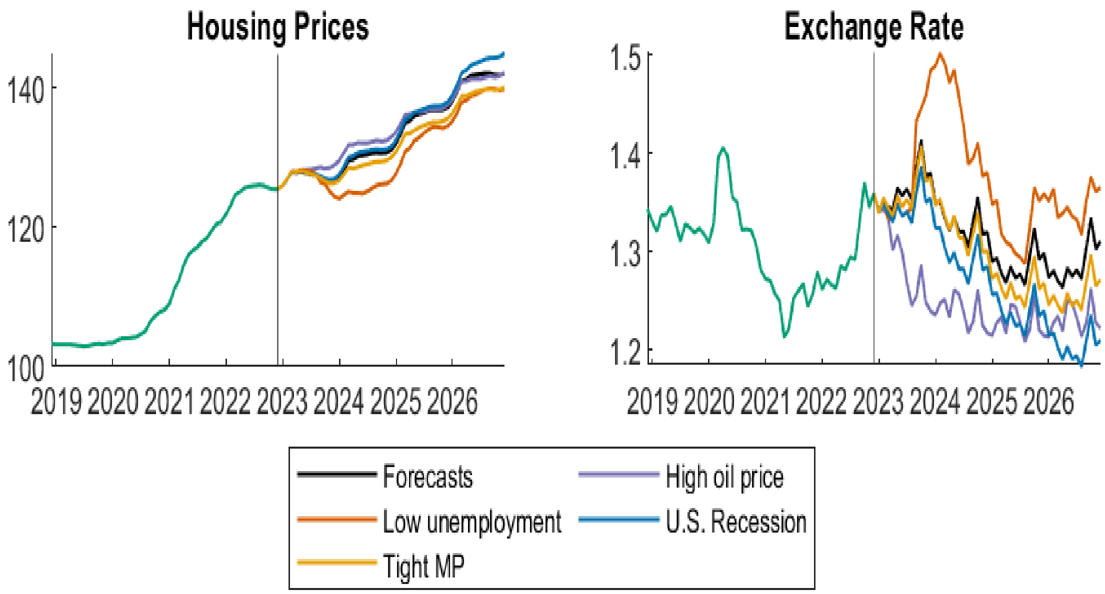

Figure 3: Forecasts for Canadian variables: baseline and scenario levels

Note: Colors indicate different scenarios. MP refers to monetary policy.

Finally, recall that the house price index and the exchange rate are used in growth rates in our VAR specification, while the level of those variables is more amenable to an economic interpretation of the evolution of these variables. To help this interpretation, Figure 3 reports the trajectories of these two variables by accumulating the growth rates from the VAR, which provides a measure in levels of the variables. The results analyzed above are naturally repeated: the evolution of house prices is moderated by the slowdown in U.S. economic activity and the tightening of monetary policy (orange and yellow lines), while the oil price shock leads to a slight increase in this index. In contrast, the oil shock leads to a sharp decline in the exchange rate (an appreciation of the Canadian dollar), while the U.S. slowdown is associated with a significant depreciation.

6 Robustness of Monetary Policy Shocks

Section 2 presented two ways of constructing conditional scenarios: the method of WZ, which starts from a postulated trajectory for a variable and finds shocks that are compatible with this trajectory, and that of BK, which instead assumes to choose a sequence of shocks before considering their implications for the variables. These methodological differences are potentially important for the analysis of a scenario regarding monetary policy, where the evolution of the variable (the discount rate in our case) incorporates both shocks to monetary policy, but also the usual response of the discount rate to shocks affecting other variables (the reaction of monetary authorities to economic developments).

Our results reveal that the estimated monetary policy shocks since the late 1990s are infrequent and of very modest magnitude. To build our sequence of monetary policy shocks and use them in the method of BK, we re-estimate the model on a sample ending in January 1982 to have a time period during which these shocks were more frequent and more significant. Once this is done, we choose the 48-month period displaying the greatest variability in these estimated shocks, and the period between 1988 and 1991 emerges from this analysis. The estimated structural shocks for these 48 months are thus our input for calculating the conditional forecasts. The selected sequence of shocks, as well as the resulting discount rates over these 48 months, are reported in Figure 6 in Appendix C.

The results are presented in Figure 4, which contains the forecasts for Canadian variables under the baseline (black lines), the conditional forecast (monetary policy tightening) using the approach of WZ already studied above (yellow lines), and the tightening forecast according to the approach of BK (purple lines).

The panel containing the evolution of the discount rate ("BC Rate") illustrates the difference between the two methods: the evolution of the rate according to WZ is smooth and precise because it corresponds precisely to the postulated scenario, while the evolution according to BK is more erratic, due both to the estimated shocks and to the fact that the variable continues to react to developments in other variables. Overall, however, we see that the monetary policy tightening is rapid and transitory according to the

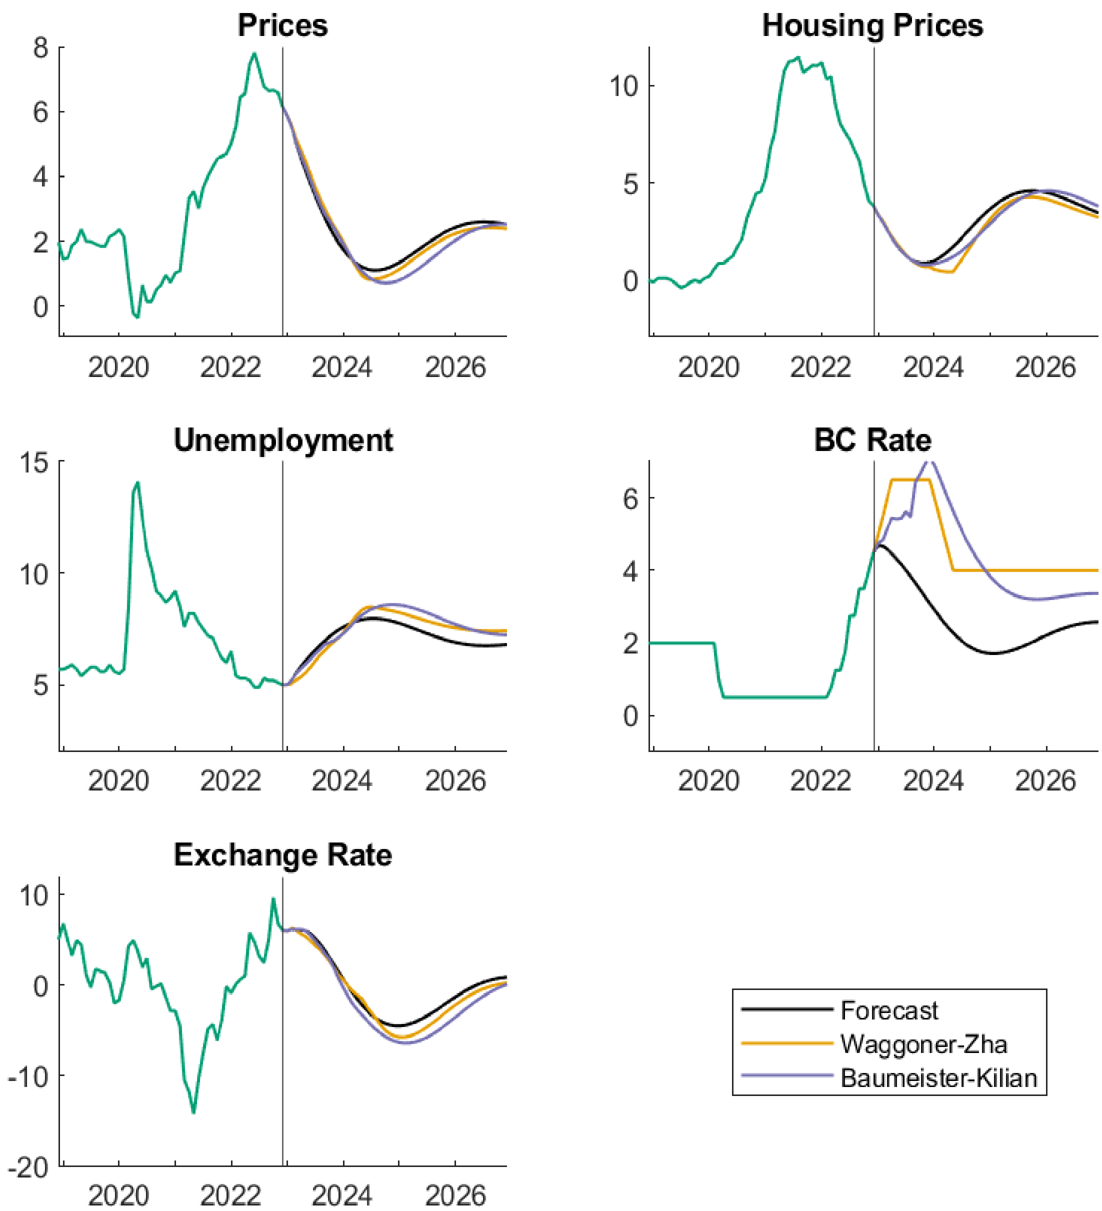

Figure 4: Predictions for Canadian variables: robustness regarding the identification of monetary policy shocks (WZ and BK)

WZ method, while according to BK it occurs more slowly. However, Figure 4 indicates that the differences between the trajectories obtained using these two methods are very modest. Considering these modest quantitative differences, it seems preferable to use the WZ method, since the scenarios developed using this method are expressed in terms of observable variable values rather than random shocks subject to identification restrictions.

7 Conclusion

This paper argues that conditional forecasts used for monitoring the Canadian economy should employ a consistent and systematic methodological framework that generates coherent forecasts across variables. To illustrate our argument, we formulate and estimate a VAR using long time series of macroeconomic variables, so that the risk scenarios analyzed reflect both the methodological framework employed and the historical data utilized. We then use our VAR to investigate the sensitivity of Canadian economic forecasts to four different risk scenarios: (1) high oil prices, (2) a U.S. recession, (3) a tight labor market, and (4) restrictive monetary policy.

The findings reveal that these scenarios represent significant risk factors for the Canadian economy. Specifically, the scenario of high oil prices resembles a sequence of positive demand shocks for a net oil-exporting country, resulting in increased real activity and higher prices. Conversely, a U.S. recession yields opposite forecasts. A tight labor market manifests in higher prices, while the scenario of highly restrictive monetary policy is accompanied by a slight increase in the unemployment rate and a decrease in the inflation rate.

Note that the empirical methodology employed in this study is not the sole option for constructing conditional forecasts. Several alternatives could be considered. For instance, Foroni et al. (2022) adjusted forecasts during the COVID-19 recession based on forecast errors observed during the 2007-09 recession, providing a non-parametric way of incorporating a historical scenario. Another approach is to directly target the tails of the distribution of the variables of interest. For example, Adrian et al. (2019) use a quartile predictive regression and show that a deterioration in financial conditions predicts lower percentiles of real GDP growth rate in the United States. Finally, a dynamic general equilibrium model would offer a theoretical alternative for conducting counterfactual analyses.

A Variance Decomposition for the Estimated VAR

Table 1: Variance decomposition analysis

	Percentage of variability caused by each shock to each Canadian variable				
Horizon	Inflation	Housing Prices	Unemployment Rate	Discount Rate	Exchange Rate
Panel A: Shock to oil prices					
3 months	44	5	19	2	20
	[33 55]	[1 12]	[5 34]	[0 8]	[10 31]
12 months	60	3	22	3	14
	[38 70]	[1 14]	[5 41]	[0 17]	[7 27]
24 months	49	5	17	3	14
	[27 62]	[2 18]	[5 36]	[0 20]	[7 28]
48 months	45	5	16	3	14
	[24 59]	[2 20]	[5 34]	[1 18]	[7 27]
Panel B: Shock to U.S. economic activity					
3 months	1	0	6	1	1
	[0 5]	[0 4]	[0 3]	[0 6]	[0 5]
12 months	1	2	7	7	2
	[0 7]	[0 12]	[0 25]	[0 21]	[0 12]
24 months	1	8	6	10	3
	[0 10]	[0 24]	[1 22]	[1 29]	[0 16]
48 months	4	11	9	9	3
	[1 14]	[1 27]	[2 23]	[1 30]	[1 15]
Panel C: Shock to Canadian unemployment rate					
3 months	0	0	74	0	1
	[0 1]	[0 2]	[46 90]	[0 4]	[0 4]
12 months	1	2	66	1	1
	[0 8]	[0 10]	[34 80]	[0 10]	[0 8]
24 months	4	2	52	8	1
	[0 16]	[0 15]	[27 65]	[0 25]	[0 10]
48 months	4	3	46	16	2
	[1 17]	[1 18]	[23 59]	[1 35]	[1 12]
Panel D: Shock to Canadian monetary policy					
3 months	1	0	1	96	0
	[0 3]	[0 1]	[0 5]	[85 98]	[0 2]
12 months	1	0	1	87	0
	[0 6]	[0 6]	[0 8]	[63 91]	[0 6]
24 months	2	1	5	71	1
	[0 13]	[0 15]	[1 17]	[41 82]	[0 10]
48 months	3	1	7	62	1
	[1 15]	[0 17]	[1 22]	[31 77]	[0 12]

Note: The table reports the percentage of variability of each of the five Canadian variables explained by each of the four shocks. The decomposition is performed for transformed variables and 95% confidence intervals (obtained with 2000 bootstrap replications) are in square brackets.

B Forecasting Ability of the VAR

In this section, we assess the forecasting ability of the VAR model presented in equation (1) in an out-of-sample forecasting exercise. We perform a comparison with an autoregressive model (AR) for each variable:

$$y_{k,t} = \nu_k + a_1 y_{k,t-1} + ... + a_p y_{k,t-p} + u_{k,t}. \tag{7}$$

Like the VAR model, this model can be written in a companion form of equation (2), and average forecasts can be formed in the same way using equation (4). Note that an AR model for each equation corresponds to imposing diagonal coefficient matrices ($\mathbf{A}_1, ..., \mathbf{A}_p$) in the VAR model. The lag length is chosen by the Akaike information criterion.

The out-of-sample period runs from January 2005 to December 2022. The choice of model order and estimation is done recursively using an expanding window approach. For instance, we use data from January 1992 to January 2003 to choose the lag length of the VAR and each of the AR models. Then, we estimate the selected specifications on this sample and use the estimates to form forecasts between February 2003 and January 2005. We add observed values for February 2003 and restart the exercise to obtain forecasts between March 2003 and February 2005. The target variables are the same as in the rest of the report and have been transformed in the same way. We consider forecasts from 1 to 24 months ahead.

We use three performance metrics to evaluate the models. First, denote by $y_{k,t}^{(h)} := \frac{1}{h}\sum_{j=1}^{h} y_{k,t-j+1}$ the average observed value of a variable over the most recent h periods. We define the mean forecast error as:

$$e_{k,t}^{(h,m)} = \frac{1}{h}\sum_{j=1}^{h} y_{k,t-j+1}^{(h)} - \frac{1}{h}\sum_{j=1}^{h} \hat{y}_{k,t}^{(h,m)}, \tag{8}$$

where $\hat{y}_{k,t}^{(h,m)}$ is the predicted value h periods ahead for variable k at period t for model

m. The root mean squared error (RMSE) and mean absolute error (MAE) are given by

$$RMSE(k,h,m) = \frac{1}{\bar{t}-\underline{t}} \sum_{t=\underline{t}}^{\bar{t}} \left(e_{k,t}^{(h,m)}\right)^2 \tag{9}$$

$$MAE(k,h,m) = MSE(k,h,m) = \frac{1}{\bar{t}-\underline{t}} \sum_{t=\underline{t}}^{\bar{t}} |e_{k,t}^{(h,m)}|. \tag{10}$$

We report the ratios of these quantities in Figure 5. As the AR model serves as the reference, values below unity in the figure indicate when the VAR model is preferred to the AR model.

We also conduct a density forecasting exercise. Let's revisit equation (1) and assume $\mathbf{u}_t \sim N(\mathbf{0}, \mathbf{\Sigma}_u)$. The covariance matrix of the error terms $\mathbf{\Sigma}_u$ is also recursively estimated from the residuals $\hat{\mathbf{u}}_t$ by $\hat{\mathbf{\Sigma}}_u = \sum_{t=1}^{T} \hat{\mathbf{u}}_t \hat{\mathbf{u}}_t' / (T - Kp - 1)$. To obtain density forecasts, we simply proceed by simulation. Specifically, using the estimated values for the VAR model parameters, the latest observations in the sample used to estimate the model, and draws of normal random variables $\mathbf{u}_{t+h} \sim N(\mathbf{0}, \hat{\mathbf{\Sigma}}_u)$ for $h = 1, ..., 24$ months, we can recursively apply equation (1) to obtain a simulated path for each variable over 24 months. We repeat this exercise 5000 times to obtain a distribution at each of the 24 horizons considered, namely $\left\{\mathbf{y}_{t+1}^{(s)}, ..., \mathbf{y}_{t+24}^{(s)}\right\}_{s=1}^{5000}$. A similar exercise is done with each of the AR models to obtain path distributions for each variable. This is done recursively using the same expanding window.

To evaluate density forecasts, we use the continuous ranked probability score (CRPS). One advantage is that it suffices to retain the quartiles of the predicted densities to compute it.[6] Let $\hat{q}_{\tau,k,t}^{(h,m)}$ be the $\tau \in [0,1]$ quantile of variable k at time t and forecast horizon h by model m, and let $\rho_\tau(u) := u\,(\tau - \mathbb{I}(u < \tau))$ be the quantile scoring function. For a quartile

[6] This is one of the advantages mentioned by Gneiting and Raftery (2007). They also show that this evaluation criterion has a minimal expected value when using the true model.

grid $0 \leq \tau_1 < ... < \tau_N \leq 1$, the CRPS is given by:

$$CRPS(k,h,m) = \frac{2}{N-1} \sum_{j=1}^{N} \left[\frac{1}{\bar{t} - \underline{t}} \sum_{t=\underline{t}}^{\bar{t}} \rho_\tau \left(y_{k,t} - \hat{q}_{\tau,k,t}^{(h,m)} \right) \right]. \tag{11}$$

Essentially, it is an average of mean quartile losses. We use an equidistant grid of 19 points between 0.05 and 0.95. Carriero et al. (2022) use the same grid to evaluate a weighted version of this criterion, while Carriero et al. (2020) use a grid of 9 equidistant points between 0.1 and 0.9. We also report the ratios for CRPS in Figure 5. Once again, values below unity indicate that the VAR is preferred over the AR model.

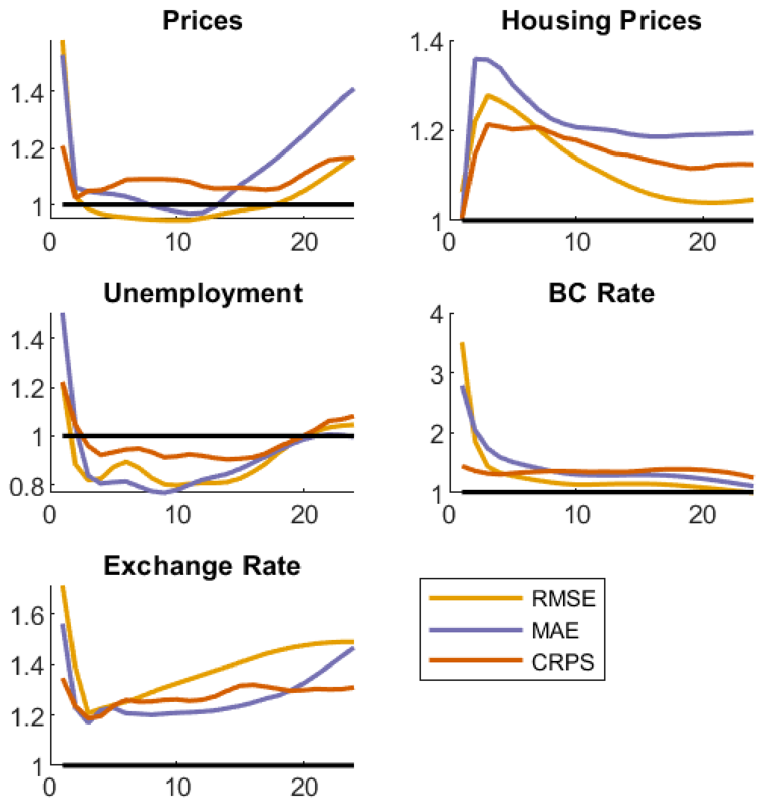

Figure 5: Performance of the VAR model (relative to that of the AR model)

Note: The acronyms are Bank of Canada (BC), root mean squared errors (RMSE), mean absolute errors (MAE) and continuous rank probability score (CRPS).

The VAR model generally performs better for inflation and unemployment, while the AR model produces more accurate forecasts for the other variables. The advantages are generally not substantial. It is worth noting that the objective here is not to construct the most adequate forecasting model but to make conditional forecasts. In other words, we are not interested in $\mathrm{E}(Y_{t+h}|Y_t, Y_{t-1}, ...)$, but rather in $\mathrm{E}(Y_{t+h}|Y_{k,t+h}, ..., Y_{k,t+1}, Y_t, Y_{t-1}, ...)$ for the scenario variables $Y(k,t)$. Furthermore, although the AR model can be written as a restricted VAR model, this set of constraints makes no economic sense because all co-movements would be excluded from the system, and thus no transmission would enable the construction of forecasts under a given scenario.

C Estimated Monetary Policy Shocks (BK Method)

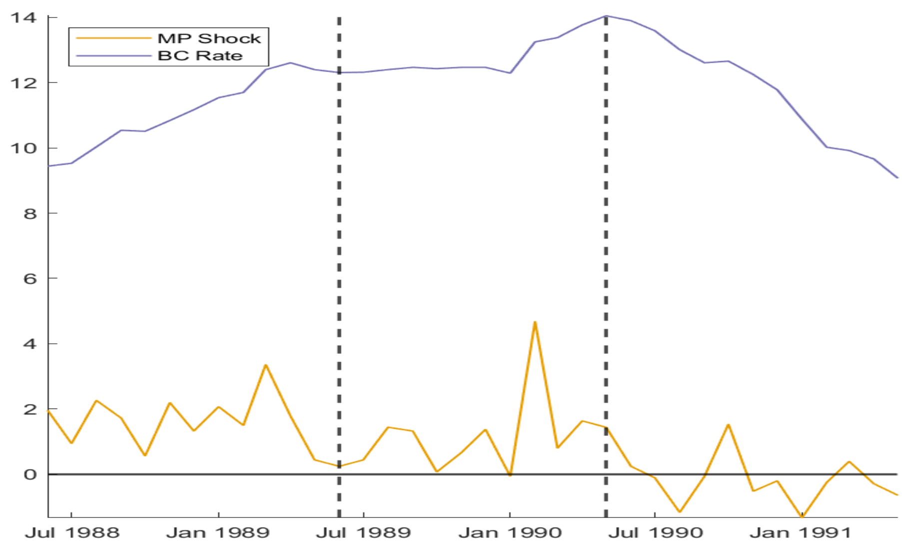

Figure 6: Monetary policy shocks estimated: 1988-1991

Note: The acronyms are Bank of Canada (BC) and monetary policy (MP).

References

T. Adrian, N. Boyarchenko, and D. Giannone. Vulnerable growth. *The American Economic Review*, 109(4):1263–1289, 2019.

C. Baumeister and L. Kilian. Real-time analysis of oil price risks using forecast scenarios. *IMF Economic Review*, 62(1):119–145, 2014.

M. Bańbura, D. Giannone, and M. Lenza. Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections. *International Journal of Forecasting*, 31(3):739–756, 2015.

N. Bedock and D. Stevanovic. An empirical study of credit shock transmission in a small open economy. *Canadian Journal of Economics*, 50(3):541–570, 2017.

R. Bhuiyan and R. F. Lucas. Real and nominal effects of monetary policy shocks. *Canadian Journal of Economics*, 40(2):679–702, 2007.

A. P. Blake and H. Mumtaz. *Applied Bayesian econometrics for central bankers*. Centre for Central Banking Studies, Bank of England, 2017.

J. Boivin, M. P. Giannoni, and D. Stevanovic. *Monetary transmission in a small open economy: More data, fewer puzzles*. Mars, 2010.

A. Carriero, T. E Clark, and M. Marcellino. *Capturing macroeconomic tail risks with Bayesian vector autoregressions*, 2020. ISSN 2573-7953. https://doi.org/10.26509/frbc-wp-202002r.

A, Carriero, T. E. Clark, and M. Marcellino. *Specification choices in quantile regression for empirical macroeconomics*, 2022. ISSN 2573-7953. https://doi.org/10.26509/frbc-wp-202225.

J. Champagne and R. Sekkel. Changes in monetary regimes and the identification of monetary policy shocks: Narrative evidence from Canada. *Journal of Monetary Economics*, 99:72–87, 2018.

T. Doan, R. Litterman, and C. Sims. Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3(1):1–100, 1984.

C. Foroni, M. Marcellino, and D. Stevanovic. Forecasting the Covid-19 recession and recovery: Lessons from the financial crisis. *International Journal of Forecasting*, 38(2): 596–612, 2022.

O. Fortin-Gagnon, M. Leroux, D. Stevanovic, and S. Surprenant. A large Canadian database for macroeconomic analysis. *Canadian Journal of Economics/Revue canadienne d' ̌economique*, 55(4):1799–1833, 2022.

D. Giannone, M. Lenza, D. Momferatou, and L. Onorante. Short-term inflation projections: a Bayesian vector autoregressive approach. *International Journal of Forecasting*, 30:635–644, 2014.

T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102:359–378, 3 2007. ISSN 01621459. doi: 10.1198/016214506000001437.

J. D. Hamilton. *Structural vector autoregressive analysis*. Princeton University Press, 1994.

K. Hou, D. C. Moutain, and T. Wu. Oil price shocks and their transmission mechanism in an oil-exporting economy: A VAR analysis informed by a DSGE model. *Journal of International Money and Finance*, 68:21–49, 2016.

M. Jarociński. Conditional forecasts and uncertainty about forecast revisions in vector autoregressions. *Economics Letters*, 108:257–259, 2010.

L. Kilian and H. Lütkepohl. *Structural vector autoregressive analysis*. Cambridge University Press, 2017.

S. Kim and N. Roubini. Exchange rate anomalies in the industrial countries: A solution with a structural VAR approach. *Journal of Monetary Economics*, 45:561–586, 2000.

Y. D. Li, T. B. Iscan, and K. Xu. The impact of monetary policy shocks on stock prices: Evidence from Canada and the United States. *Journal of International Money and Finance*, 29:876–896, 2010.

M. W. McCracken and J. T. McGillicuddy. An empirical investigation of direct and iterated multistep conditional forecasts. *Journal of Applied Econometrics*, 34:181–204, 2019.

M. W. McCracken and S. Ng. FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, 34(4):574–589, 2016.

W. Miyamoto and T. L. Nguyen. Understanding the cross-country effects of U.S. technology shocks. *Journal of International Economics*, 106:143–1644, 2017.

K. Moran, D. Stevanovic, and A. K. Touré. Macroeconomic uncertainty and the COVID-19 pandemic: Measure and impacts on the Canadian economy. *Canadian Journal of Economics*, 55:379–405, 2022.

K. Moran, D. Stevanovic, and A. K. Touré. Confiance et activité économique : analyse d'impact sur l'économie canadienne, 2023. Rapport du CIRANO.

D. F. Waggoner and T. Zha. Conditional forecasts in dynamic multivariate models. *The Review of Economics and Statistics*, 81(4):639–651, 1999.