

TFP, News, and Expectations: The International Transmission of Business Cycles*

Andrei A. Levchenko
University of Michigan
NBER and CEPR

Nitya Pandalai-Nayar
University of Texas at Austin

October 27, 2017

Abstract

We propose a novel identification scheme for a non-technology business cycle shock, that we label “non-technology expectations” (NTE). This is a shock orthogonal to identified surprise and news TFP shocks that maximizes the short-run forecast error variance of an expectational variable, alternatively a GDP forecast or a consumer confidence index. We then estimate the international transmission of three identified shocks – surprise TFP, news of future TFP, and NTE – from the US to Canada. The US non-technology shock produces a business cycle in the US, with output, hours, and consumption rising following a positive shock, and accounts for the bulk of US short-run business cycle fluctuations. The non-technology shock also has a significant impact on Canadian macro aggregates. In the short run, it is more important than either the surprise or the news TFP shocks in generating business cycle comovement between the US and Canada, accounting for over 40% of the forecast error variance of Canadian GDP and over one-third of Canadian hours, imports, and exports. The news shock is responsible for some comovement at 5-10 years, and surprise TFP innovations do not generate synchronization. We provide a simple theoretical framework to illustrate how US non-technology expectations shocks can transmit to Canada.

Keywords: Non-Technology Shocks, Demand Shocks, News Shocks, International Business Cycles

JEL Codes: E32, F41, F44

*We would like to thank the editor (Dirk Krueger), three anonymous referees, Rudi Bachmann, Regis Barnichon, Kenza Benhima, Christoph Boehm, Fabio Canova, Olivier Coibion, Jordi Galí, Kyle Jurado, Alejandro Justiniano, Lutz Kilian, Jennifer La’O, Franck Portier, Eric Sims, Fabrizio Venditti, Jaume Ventura, and workshop participants at several institutions for helpful suggestions. We are especially grateful to Zhen Huo for his input on the model and for the code that implements the Huo-Takayama model. Email: alev@umich.edu and npnayar@utexas.edu.

1 Introduction

Business cycles in advanced economies exhibit strong positive comovement. A complete empirical and theoretical account of positive cross-border comovement remains elusive. The International Real Business Cycle (IRBC) literature going back to Backus, Kehoe, and Kydland (1992) develops quantitative models in which fluctuations are driven by surprise TFP shocks, and assesses their performance in generating comovement. However, a series of empirical contributions in the closed-economy literature have argued that the bulk of (short-run) business cycle fluctuations is actually accounted for by non-technology shocks, customarily referred to as “demand” shocks.¹ It is thus a natural conjecture that international business cycle comovement can be driven by transmission of non-technology as well as technology shocks across borders. Indeed, international business cycle models are more successful at matching basic moments in the data when augmented with demand shocks (Stockman and Tesar, 1995; Wen, 2007).²

This paper investigates empirically the relative importance of the cross-border transmission of technology and non-technology shocks. It uses US and Canada as a laboratory to study these issues.³ We begin by identifying three types of US shocks in a structural vector auto-regression (SVAR) setting. The first is a shock to contemporaneous TFP, identified as the reduced-form shock assuming that the TFP series is ordered first. The second is a news shock about future TFP (Beaudry and Portier, 2006), identified following Barsky and Sims (2011). Most importantly, we propose a new identification strategy for a non-technology business cycle shock. The VAR includes an expectation variable, alternatively a GDP forecast from the Philadelphia Fed’s Survey of Professional Forecasters or the Michigan/Reuters Consumer Confidence variable. The non-technology shock is identified as the shock orthogonal to both the surprise-TFP and the news-TFP shocks that explains the maximum of the residual forecast error variance of this expectational variable at a short

¹For a number of different approaches that reach this conclusion, see Blanchard and Quah (1989); Galí (1999); Canova and de Nicoló (2003); Basu, Fernald, and Kimball (2006).

²An obvious alternative is that international comovement is generated by transmission of policy or credit shocks. Available evidence suggests that the importance of these shocks in fluctuations is limited. Kim (2001) and Maćkowiak (2007) show that shocks to US monetary policy explain only a very small share of forecast error variance of other countries’ output, while Ilzetzi and Jin (2013) show that even the sign of the impact is not stable over time. In a similar vein, Helbling et al. (2011), Kollmann (2013), and Eickmeier and Ng (2015) show that the share of variance of other countries’ GDP accounted for by US credit shocks and bank shocks is small as well.

³These two economies are closely integrated, and very asymmetric in size. The latter feature implies that identified US shocks are unlikely to be “contaminated” by endogenous US responses to Canadian shocks. This approach has been adopted by Cushman and Zha (1997), Schmitt-Grohé (1998), Justiniano and Preston (2010), and Miyamoto and Nguyen (2014), among others.

horizon. Because the shock is identified explicitly from data on expectations after controlling for shocks to current and future TFP, we label this shock “non-technology expectations,” or NTE for short. The NTE shock accounts for the bulk of US fluctuations at business cycle frequencies (65-75% of the forecast error variance in GDP and hours) and generates positive comovement of GDP, consumption, and hours. These properties are consistent with the NTE shock being a transitory “demand” shock, and are similar to the findings from other ways of identifying “demand” shocks (see, e.g., Galí, 1999; Canova and de Nicoló, 2003, among others).

Our main results concern the cross-border transmission of shocks to Canada. The first important finding that sets the stage for the rest of the results is that Canadian utilization-adjusted TFP does not react to any of the three identified US shocks. This makes us confident that the business cycle impact of US shocks on Canada is not contaminated by an underlying correlation between US shocks and Canadian TFP.

Canadian aggregates react much more strongly to the non-technology shocks than to the surprise and news TFP shocks, and in the short run the NTE shock is by far the most important of the identified US shocks in accounting for fluctuations in Canadian variables. Following an NTE shock, Canadian GDP, hours, and consumption rise instantaneously and peak within one year. The strongest response is of Canadian exports to the US and US exports to Canada, which both rise instantaneously, peak at 1 or 2 quarters, and then fall back to steady state. At short frequencies, the NTE shock accounts for 20-40% of the forecast error variance of Canadian GDP, 8-12% of Canadian consumption, 20-35% of Canadian hours, and 25-44% of Canada-US trade flows. We also compute conditional correlations between the variables due to each shock following the approach in Galí (1999). The NTE shock generates very high conditional correlations in GDP, hours, and consumption between US and Canada.

The responses of Canadian GDP, hours, and consumption to the US surprise TFP or news shocks are positive but take place with a lag of 2-3 quarters. There is not much of a trade response to surprise TFP shocks. US news do not generate a positive trade response for over 1 year following the shock, in fact there is suggestive evidence that Canadian imports from and exports to the US actually fall on impact following a US news shock. The surprise TFP shock accounts for less than 6% of the forecast error variance of Canadian GDP and hours across all frequencies between 1 quarter and 5 years, and for less than 10% of Canadian consumption. The US news shock is

similarly unimportant at short frequencies, though it does become more important for Canadian output and consumption at frequencies longer than 2 years. The surprise TFP shock generates a positive conditional correlation in GDP between US and Canada, but it actually produces negative US-Canada correlations in consumption and hours, whereas in the data those are positive. The conditional US-Canada correlations in GDP, hours, and consumption due to news shocks are similar to those due to NTE shocks.

The bottom line is that at short frequencies, the US non-technology shocks generate a much stronger cross-border impact, and account for a higher share of Canadian fluctuations than technology shocks. The NTE shocks also generate much higher conditional correlations between US and Canadian aggregates than surprise TFP shocks. At the same time, news shocks are also important for international comovement at medium frequencies. An empirical account of observed international comovement therefore requires knowledge of the impact of both types of shocks, coupled with the understanding that the surprise TFP shock central to most IRBC models actually does not generate substantial comovement.

Our identification strategy is inspired by the recent theories of aggregate fluctuations arising from expectational shocks (e.g., Angeletos and La'O, 2013; Benhabib, Wang, and Wen, 2015; Huo and Takayama, 2015). It is important to underscore that our empirical exercise does not necessarily identify the precise mechanisms that produce fluctuations in these theoretical contributions. In principle, our shock can be driven by anything that makes agents expect better/worse times, conditional on available information about current and future productivity. However, we provide two types of evidence that an expectational shock interpretation may be warranted.

First, in a series of checks we show that our NTE shock is not a monetary policy, fiscal policy, oil price, or uncertainty shock. We also add to the VAR a number of variables to increase the information set used to identify shocks: stock prices, consumer prices, the real exchange rate, as well as an estimated factor variable. Adding these variables does not alter the features of the NTE shock or diminish noticeably its importance. Hence, it is difficult to account for our non-technology shock with other standard business cycle shocks, leaving an expectations shock interpretation as one of the few remaining potential explanations. Second, we establish internal validity of our identification strategy. We simulate data from the Huo and Takayama (2015) DSGE model that features persistent TFP and expectational shocks. Our identification procedure applied to model-

simulated data correctly extracts the non-technology expectational shock. In addition, the impulse responses to the expectational shock in model-simulated data are quite similar to the impulse responses to the NTE shock in actual US data. Jointly, these two exercises are suggestive that we may be uncovering an intrinsic shock to expectations, rather than simply picking up a combination of traditional “demand” shocks.

To help understand our empirical findings on international transmission, we set up a simple model of expectation-driven fluctuations following Angeletos and La’O (2013) and Huo and Takayama (2015), and extend it to include a small open economy representing Canada. In this framework, US fluctuations arise from the combination of shocks to agents’ expectations and imperfect information. Canada trades with the US, but itself does not experience the expectational shock. That is, unlike US agents, the Canadian agent observes perfectly both the fundamentals of the economy and the magnitude of the US expectational shock. The key result for our purposes is that Canadian output increases in response to the US non-technology expectations shock, generating output comovement between the US and Canada. This is because Canada knows that its US trading partner will have high output following a positive US expectational shock, and therefore high demand for Canadian output.

Our paper draws on the recent closed-economy literature on “demand”-driven fluctuations (see, among others, Galí, 1999; Beaudry and Portier, 2006; Lorenzoni, 2009; Barsky and Sims, 2011; Blanchard, L’Huillier, and Lorenzoni, 2013). Most closely related are empirical assessments of cross-border transmission of shocks, in particular non-technology shocks. Canova (2005) examines the impact of US supply and demand shocks on Latin America, while Corsetti, Dedola, and Leduc (2014) assess the reaction of externally-oriented variables – such as real exchange rates and foreign assets – to US supply and demand shocks. Both of these papers identify supply and demand shocks using sign restrictions. Our paper contributes a novel identification strategy for supply and demand shocks, based on expectational variables (for demand) and utilization-adjusted TFP (for supply). Importantly, we separate news about future TFP – which can look like a demand shock in the short run – from expectational shocks unrelated to TFP.

Several recent papers identify shocks that are interpreted as sentiments, in both VAR settings and fully specified DSGE models (Angeletos, Collard, and Dellas, 2014; Milani, 2014; Nam and Wang, 2016). We complement these contributions in two main respects. First, we explicitly separate

a strictly non-technology expectations shock from the TFP news shock. Second, our approach is based on explaining the variation only in an expectational variable. Our strategy thus “ties our hands behind our back” to a much greater extent, as we are not extracting a shock that by construction has a particular impact on the key macro aggregates. It is reassuring that the findings of our data-driven exercise regarding the importance of expectational shocks in the US business cycle are consistent with the fully structural DSGE estimation approaches. Substantively, of course, our focus is on the international dimension of shock transmission.

The rest of the paper is organized as follows. Section 2 discusses the empirical strategy, estimation methods, and data. Section 3 presents the main estimation results. Section 4 lays out the theoretical framework, reports the results of an internal validation exercise, and discusses additional interpretation issues. Section 5 concludes. The Appendix collects additional details on data, robustness, and theory.

2 Empirical Strategy

2.1 Identification of Shocks

Our identification strategy builds on Uhlig (2003, 2004) and Barsky and Sims (2011). As an illustration of why it is important to separate non-technology shocks from news TFP shocks, suppose that the TFP process in the US is affected by only two innovations: an unanticipated ‘surprise’ TFP shock and a ‘news’ shock. An example of a process that would satisfy these conditions is:

$$TFP_t = \lambda_1 \epsilon_t^{sur} + \lambda_2 \epsilon_{t-s}^{news}, \quad (1)$$

where ϵ^{sur} and ϵ^{news} are the surprise and anticipated innovations in TFP and the agents learn about the news shock $s > 0$ periods in advance.⁴

Further, assume that expectations of future economic activity are influenced not only by the surprise innovation in TFP and the anticipated future improvement in TFP, but also by ‘confidence,’ as the agents rationally expect a positive shock to expectations to lead to a temporary boom in the economy and increase output. Forward-looking agents also respond to other changes in the economy that could stimulate GDP, but we assume that the bulk of the variation in expectations

⁴This TFP process can clearly be modified to include a persistent component.

of future activity is due to these three shocks. A simple process for expectations F_t that satisfies this assumption is:

$$F_t = \lambda_1^F \epsilon_t^{sur} + \lambda_2^F \epsilon_{t-s}^{news} + \lambda_3^F \epsilon_t^{nte} + \zeta_t, \quad (2)$$

with ϵ^{nte} the non-technology expectations shock.

Expectations of better future economic conditions, controlling for current fundamentals, can be due to either news of high future TFP, or to positive ‘confidence.’ Clearly, in order to extract a non-technology shock from data on expectations, we must control for news of future productivity. It would not be possible to identify the three shocks of interest from movements in TFP and expectations alone. We therefore consider the processes for these variables together with other forward-looking macroeconomic aggregates in a VAR. Let Y_t denote the $k \times 1$ vector of observables in levels. For much of our analysis, this will be US TFP, real GDP, consumption, hours, and forecasts of GDP. The moving average representation of this k -variable VAR is:

$$Y_t = B(L) \mathbf{u}_t,$$

where \mathbf{u}_t is the vector of reduced-form disturbances, L denotes the lag operator and $B(L)$ is the matrix of lag order polynomials.

To identify the structural shocks, we assume that there exists a linear relationship $\mathbf{u}_t = A\epsilon_t$ where ϵ_t is the vector of structural shocks and A is the impact matrix. This implies that the structural representation of the VAR is

$$Y_t = A(L) \epsilon_t,$$

where $A(L) = B(L)A$. Clearly, assuming that the structural shocks each have unit variance, $AA' = \Sigma$, where Σ is the covariance matrix of \mathbf{u} . It is well known that the Choleski decomposition \tilde{A} of Σ provides one candidate for A , but this is just one among many. For any orthonormal $k \times k$ matrix D such that $DD' = I$, $\tilde{A}D$ will provide an identification of the structural shocks.

The forecast error h steps ahead is defined as

$$Y_{t+h} - E_{t-1}Y_{t+h} = \sum_{\tau=0}^h B_\tau \tilde{A}D \epsilon_{t+h-\tau},$$

where B_τ is the reduced-form matrix of lag- τ moving average coefficients. Since the elements of ϵ_t are independent, this equation illustrates that the forecast error variance of a particular variable i at horizon h is the sum of the contributions of the k structural shocks. Let $\Omega_{i,j}(h)$ denote the contribution of shock j to the forecast error variance of variable i at horizon h . The assumption that only two shocks (surprise and news) affect true TFP then implies:

$$\Omega_{1,sur}(h) + \Omega_{1,news}(h) = 1 \quad \forall h. \quad (3)$$

The unexpected TFP innovation ϵ_t^{sur} in (1) is identified as the reduced-form innovation in a VAR with TFP ordered first. By identifying the reduced-form innovation in TFP as the first structural shock, we effectively fix $\Omega_{1,1}(h)$ at all horizons. The news shock ϵ_{t-s}^{news} is true news about future changes in TFP s periods ahead. Of course, in practice (3) is unlikely to hold as an identity for all $h \leq H^{news}$. Thus, given the Choleski decomposition \tilde{A} , the news shock is identified as the linear combination of the remaining VAR innovations that maximizes the residual forecast error variance of TFP, $1 - \Omega_{1,1}(h)$, over a finite horizon H^{news} (Barsky and Sims, 2011).⁵

Without loss of generality, assume the second structural shock is the news shock, and thus the second column of $\tilde{A}D$ is its impact vector. Formally, we select γ^{news} as the solution to the problem:

$$\gamma^{news} = \operatorname{argmax} \sum_{h=0}^{H^{news}} \Omega_{1,2}(h) = \operatorname{argmax} \sum_{h=0}^{H^{news}} \left(\frac{\sum_{\tau=0}^h B_{1,\tau} \tilde{A} \gamma^{news} \gamma^{news'} \tilde{A}' B'_{1,\tau}}{\sum_{\tau=0}^h B_{1,\tau} \Sigma B'_{1,\tau}} \right)$$

subject to

$$D(1, i) = 0 \quad \forall i \neq 1 \quad (4)$$

$$D \text{ is orthonormal,} \quad (5)$$

where the lower-triangular matrix \tilde{A} is the Choleski decomposition (so $\tilde{A}(1, m) = 0 \quad \forall m > 1$).

We next proceed to the identification of the non-technology expectations shock. As this shock cannot be inferred from movements to TFP, our identification will rely on its impact on expectational variables. These will be alternately forecasts of GDP by professional forecasters or consumer confidence. Let the expectational variable F_t be ordered 5th in the VAR, and without loss of generality assume that the NTE shock is the 3rd shock. Note that by equating the first reduced-form

⁵In the empirical implementation we select $H^{news} = 40$, or a ten-year horizon.

shock to the surprise innovation to TFP and then identifying the news shock as in Barsky and Sims (2011), we have in effect fixed $\Omega_{5,1}(h)$ and $\Omega_{5,2}(h)$ at all horizons. We therefore select the NTE shock as the linear combination of the remaining $k - 2$ reduced-form innovations that maximizes the forecast error variance of F_t , where k is the total number of core and non-core variables in the VAR.⁶ Because the NTE shock is short-run, we select it to maximize the forecast error variance for a 2-quarter horizon ($H^{nte} = 2$). Formally:

$$\gamma^{nte} = \operatorname{argmax} \sum_{h=0}^{H^{nte}} \Omega_{5,3}(h) = \operatorname{argmax} \sum_{h=0}^{H^{nte}} \left(\frac{\sum_{\tau=0}^h B_{5,\tau} \tilde{A} \gamma^{nte} \gamma^{sent'} \tilde{A}' B'_{5,\tau}}{\sum_{\tau=0}^h B_{5,\tau} \Sigma B'_{5,\tau}} \right)$$

subject to

$$D(1, i) = 0 \quad \forall i \neq 1 \tag{6}$$

$$D \text{ is orthonormal} \tag{7}$$

$$D(:, 2) = \gamma^{news}. \tag{8}$$

Both the news and NTE identification steps are conditional on an arbitrary orthogonalization, the Choleski decomposition \tilde{A} . The first restriction – (4) and (6) – common to both problems specifies that none of the $k - 1$ structural shocks has a contemporaneous impact on TFP. The second restriction, (5) and (7), states that the matrix D remains orthonormal throughout, and thus the identified shocks are orthogonal to each other. Restriction (8) ensures that identification of the NTE shock holds identification of the news shock constant. We expect the surprise TFP and the news shocks, as informative about true fundamentals, to explain the movements in the forecast of GDP. The NTE shock identified in this manner simply captures patterns in the residual variance of the forecast of GDP, once supply-side determinants are accounted for. The identification strategy for both shocks is robust to the reordering of the remaining $k - 1$ variables in the VAR other than TFP.⁷

⁶Note, we do not allow the reduced form shock to the Canadian variable, ordered k , to affect the identification of the NTE shock. Hence, the NTE shock is identified from $k - 2$ reduced form shocks (as the surprise TFP innovation also does not affect it).

⁷In a recent paper, Angeletos, Collard, and Dellas (2014) adopt a closely related identification strategy to extract a factor that explains most of the business cycle variation in hours and investment at frequencies of 6-32 quarters. In contrast to our approach, that paper obtains an expression for the share of the variance of a variable due to a shock at this frequency through a spectral decomposition, and then chooses a linear combination of shocks that maximizes the variance of the selected variables. TFP is not included in their VAR. In short, they sum across variables, while we maximize the residual forecast error variance of a single, expectational, variable – either GDP forecast or confidence

Note that we do not impose that the NTE shock has no effect on true TFP, except on impact. The procedure outlined above naturally minimizes the TFP impact of the NTE shock (as well as of the other remaining structural shocks), by selecting the news TFP shock with the maximum explanatory power on the TFP series. Still, if the surprise and news TFP shocks do not account sufficiently well for the forecast error variance of the TFP series, there is potentially room for the other shocks to drive TFP. Having identified the NTE shock, we check whether it has a noticeable impact on the TFP series, and find that it does not.

Our strategy relies on ‘medium-run’ identification. It might appear that the natural identification of the non-technology expectations shock would make use of a ‘long-run’ restriction, namely that it has no long-run impact on output or forecasts. We prefer the method here as several papers have emphasized that long-run restrictions are problematic in VARs of finite order, where the coefficient estimates are biased (Faust and Leeper, 1997). Medium-run identification has shown better behavior in finite samples (Francis et al., 2014).⁸

2.2 Estimating International Transmission

We estimate the impact of the US shocks on various Canadian aggregates in turn, treating them as ‘non-core’ variables in the VAR. The Canadian variables are included one at a time and are ordered last in a six-variable VAR with 5 US series. The matrices of coefficients are restricted to allow no current or lagged impact of the Canadian variable on the five US variables. We believe this assumption is reasonable given the small size of the Canadian economy relative to the US (Canadian GDP is about one-tenth that of the US). Section 3.1 shows that the results are robust to allowing lagged Canadian variables to affect US variables.

The impulse responses of Canadian variables to the identified US shocks are interpreted as evidence of cross-border transmission of those shocks to Canada, rather than a correlation of underlying Canadian shocks with the US shocks. A useful check presented below is to construct the impulse responses of Canadian TFP to these identified shocks, and ascertain that Canadian TFP does not comove with the identified US shocks. Section 4.3 also checks for the possibility of cor-

– over several horizons.

⁸An alternative approach to long-run identification in VARs uses the spectral factorization of the variance matrix at frequency zero. This does not circumvent the issues related to long-run restrictions in general, however. We do not pursue the spectral approach in this paper as we are not aware of methods by which we would be able to identify the three shocks, while maintaining the medium-run identification structure.

related non-technology expectations shocks, which would not be visible in TFP movements, and finds little evidence that the impulse responses of Canadian aggregates to US shocks are due to a correlated Canadian shock.

We estimate the reduced-form VAR with estimated generalized least squares (EGLS) using a method adapted from Lütkepohl (2005). The VAR in p lags is:

$$Y_t = C_0 + C_1LY_t + \dots + C_pL^pY_t + u_t$$

where C_j are $k \times k$. If the Canadian variable is ordered last, the restriction that Canadian variables have no impact on US variables amounts to $C_j(1 : k - 1, k) = 0 \forall C_j$. Rewrite the VAR in compact form as $Y = CZ + U$, where $Y = [Y_1, \dots, Y_T]$, $Z_t = [1, Y_t, \dots, Y_{t-p+1}]$, $Z = [Z_0, \dots, Z_{T-1}]$, $C = [C_0, \dots, C_p]$, and $U = [u_1, \dots, u_T]$.

Let the constraints on the coefficients of the six-variable VAR be written as $\beta = \text{vec}(C) = R\mathbf{b} + \mathbf{r}$, where R is a known matrix of rank M , \mathbf{r} is a vector of constants, and \mathbf{b} is the $(M \times 1)$ vector of unknown parameters to be estimated. Appropriately pick R (size $k(kp + 1) \times M$) and \mathbf{r} such that the desired constraints on C_j hold. Clearly, linear restrictions of the type we are interested in can easily be expressed in this form.

The EGLS estimate of \mathbf{b} is then:

$$\mathbf{b} = [R'(ZZ' \otimes \Sigma_u^{-1})R]^{-1}R(Z \otimes \Sigma_u^{-1})z \quad (9)$$

where

$$z = \text{vec}(Y) - (Z \otimes I_K)\mathbf{r}$$

and Σ_u is any consistent estimator of the unknown covariance matrix of $\text{vec}(U)$. We initialize Σ_u as

$$\hat{\Sigma}_u = \frac{1}{T - kp - 1} \hat{U}_{ols} \hat{U}'_{ols}$$

where \hat{U}_{ols} are the residuals from an unconstrained ordinary least squares estimation of the six-variable VAR(p). We use an iterative procedure, in which we compute a new covariance matrix from the first stage EGLS residuals to replace Σ_u in the computation of the next value of \mathbf{b} and iterate to convergence. This procedure is asymptotically more efficient than standard multivariate

least squares, and under the assumption of Gaussian errors the estimator for \mathbf{b} in (9) is the same as the maximum likelihood estimator. Using estimates of \mathbf{b} it is then straightforward to compute the impulse response functions of each Canadian macro aggregate to the three shocks of interest. Note that the identification of the shocks is unaffected by this procedure.

Following the recommendation of Hamilton (1994), the model is specified in levels, since parameter estimates in levels are still consistent even in the presence of cointegration, while the vector error correction model might be misspecified when the cointegration is of unknown form. The baseline implementation uses $p = 4$ lags, the optimal lag length according to the Akaike Information Criterion. All standard errors are constructed from 2000 bias-corrected bootstraps as in Kilian (1998).

2.3 Data

The time period covered by our data is 1968:Q4 to 2010:Q3. All variables are logged. For a measure of US productivity, we use the quarterly, utilization-adjusted TFP series from Fernald (2014). The series is the quarterly version of the annual series developed by Basu, Fernald, and Kimball (2006). That paper constructs a modified Solow residual from industry-level data, allowing for both non-constant returns to scale and varying unobserved capital and labor utilization. The identification of the three structural shocks in our VAR relies on an accurate measure of US technology. Clearly, accounting for measurement issues arising from changes in utilization is crucial. Basu, Fernald, and Kimball (2006) find that the detrended utilization-adjusted TFP is both less correlated with output, and less volatile than the standard Solow residual. Unfortunately the industry-level data required for controlling for non-constant returns to scale are not available quarterly, so the Fernald (2014) series corrects only for variable capital and labor utilization.⁹

⁹The entire TFP series are updated every quarter, and therefore several vintages of the series exist. Kurmann and Sims (2017) document differences in the business cycle properties of the different vintages of the Fernald quarterly series. They also propose a novel way to identify the news shock, which extracts a shock that explains the maximum of the forecast error variance of the TFP series at a fixed, long horizon (80 quarters), rather than cumulatively over all horizons between 1 and 40 quarters as in the original Barsky and Sims (2011) paper. The idea behind taking a fixed and long horizon is that measurement error in the TFP series is unlikely to plague the long-run evolution of TFP, and thus this identification strategy suffers less from measurement error. The Kurmann and Sims (2017) paper has two reassuring findings for our purposes. First, the properties of the Kurmann-Sims shock are not sensitive to which vintage of the Fernald series is used. Second, the Kurmann-Sims shock in fact has very similar behavior to the original Barsky-Sims shock. We implemented the Kurmann-Sims identification strategy instead of Barsky-Sims to extract the news shock, and then extracted a non-technology expectations shock using our approach. The NTE shock obtained conditional on the Kurmann-Sims shock is exceedingly similar to the NTE shock in our baseline analysis (results available upon request).

US population and hours data are from the BLS. For population, we use the civilian non-institutionalized population age 16 and over. Aggregate hours are the total hours of wage and salary workers on non-farm payrolls. For consumption and output, we use the National Income and Product Accounts (NIPA) tables from the BEA. Output is measured as quarterly real GDP, chain-weighted, from NIPA table 1.1.6. As a chain-weighted series for non-durables and services consumption is not available, we construct a series using the Tornqvist approximation (see Whelan, 2000, for details on chain-weighting in the BEA data). For this procedure, we use the nominal shares of spending on non-durables and services from NIPA table 1.1.5. Chain-weighting reduces the dependence of a series on the choice of base year, and is the current standard for macroeconomic series constructed by all major statistical agencies. All variables are converted into per capita terms.

The data on the forecasts of US GDP come from the Survey of Professional Forecasters (SPF), provided by the Federal Reserve Bank of Philadelphia. For NIPA variables, the survey contains quarterly forecasts at several horizons as well as longer-term forecasts. We use the one quarter ahead growth rate forecast. The perturbation to US expectations that we are interested in identifying is not related to true technological progress, and we would expect the effects of this shock to be very short-lived. The survey provides mean and median levels forecasts as well as growth rates. The base year for the levels forecasts changes periodically throughout the survey. To avoid issues related to rebasing the forecasts ex-post, we construct an index of implied GDP levels forecasts from the mean forecast of the one quarter ahead growth rate. We check the sensitivity of our results to using a two- or three-quarter ahead growth rate forecast, as well as different horizons $H^{nte} = 4, 8, 16$ over which we expect the expectational shock to contribute to the forecast error variance of the GDP forecast variable, and find no significant differences in the shape of the responses.

In addition, we re-do the analysis using an index of consumer confidence from the University of Michigan Survey of Consumers instead of the SPF GDP forecast. We use the consumer confidence series E12Y, constructed from the responses to the question *‘And how about a year from now, do you expect that in the country as a whole, business conditions will be better, or worse than they are at present, or just about the same?’*

A consistent measure of quarterly hours for the length of our sample is not easily available for most countries. For Canada, we use a new dataset assembled by Ohanian and Raffo (2012), constructed from the OECD’s Main Economic Indicators database and other sources. Our Canadian

hours measure is the total hours worked in Canada divided by the Canadian population. The population data are taken from CANSIM (the Statistics Canada database), and is the quarterly estimate of total population in all provinces and territories of Canada. Canadian real GDP and consumption are taken from the OECD Economic Outlook and are also converted into per capita terms. For the bilateral exports and imports series, we use data from the IMF's Direction of Trade Statistics (DOTS) database. The series are deflated with a US GDP deflator and deseasonalized using the X-12 ARIMA program developed by the US Census Bureau.

The sample period is constrained on the two ends by different data series availability. In particular, the SPF only starts in 1968. On the other end, we are limited by the availability of Canadian hours in the Ohanian and Raffo (2012) data.

2.3.1 Utilization-Adjusted TFP for Canada

The last critical variable for the analysis is a measure of Canadian TFP. Ideally, we would use a utilization-adjusted series with further adjustments for non-constant returns to scale, similar to the Basu, Fernald, and Kimball (2006) series for the US. Unfortunately, such a series to our knowledge is not available for any other country. The data required to construct such a series are also not available at the quarterly frequency for Canada. Therefore we build our own utilization-adjusted TFP series for Canada, following the approach in Imbs (1999). This method uses a similar insight, namely that with a constant returns to scale production function the first-order conditions for capital and labor are informative about the choices of capital utilization and the workweek of labor. As data on the capital stock are also not available at the quarterly frequency, we use the perpetual inventory method to construct an initial capital stock series, given data on investment from the OECD and a constant depreciation rate. This produces a starting utilization series. We then use an iterative procedure to construct a time-varying depreciation rate, capital stock, and implied utilization series consistent with the observed investment in the data. We construct labor utilization from information on hours worked, wages, and consumption in Canada. The wage data is from the OECD Main Economic Indicators (MEI). The utilization-adjusted Solow residual is then $\log TFP = \log Y_t^{Can} - (1 - \alpha)(\log K_t + \log u_t) - \alpha(\log N_t + \log e_t)$, where e_t is labor utilization, u_t is capital utilization, Y_t^{Can} is output, K_t is capital and N_t is hours worked. Details of the procedure are in Appendix A.

We present the impulse response functions for both the utilization-adjusted TFP series and the implied capital utilization series.¹⁰

3 Results

Our baseline specification identifies the news shock at a horizon of ten years, the NTE shock at a horizon of two quarters, and uses the SPF forecast of GDP one quarter ahead as the fifth variable in the VAR. We begin by discussing the responses to the surprise TFP, news, and NTE shocks on the US economy (Figures 1, 2, and 3), followed by the analysis of the transmission to Canada.¹¹

The surprise TFP innovation signals a deviation in TFP from trend of about 0.8%. The effects of the shock die out slowly, with TFP decreasing but staying significantly above trend for 12 quarters. The responses of other domestic variables to this shock are consistent with other empirical investigations (Basu, Fernald, and Kimball, 2006; Barsky and Sims, 2011). Output increases temporarily before falling below trend after two years. Consumption stays constant on impact, and declines with output.

Our identified news shock signals a slowly building increase in utilization-adjusted TFP, beginning in quarter 2. Consumption increases slightly on impact and continues for two years, after which it exhibits a very slight decline. There is an impact decrease in hours, qualitatively consistent with the results in Barsky and Sims (2011). The response of hours turns positive one year after the shock, peaking at about Q9. There is no significant impact effect on output. Rather, the response of output builds slowly, similar to technology (but stronger). The peak increase is later than for surprise TFP, two years after the shock. Reassuringly, the forecasts of GDP track the responses of actual GDP quite well, with the response of the forecast variable peaking about one quarter before GDP.

Overall, these responses are in line with Barsky and Sims (2011). As in that paper, the impact decrease in hours is consistent with a strong wealth effect, and indicates that the news shock does not solve the impact comovement problem of hours, consumption, and output.¹² It therefore cannot

¹⁰We check the responses of the unmodified Solow residual as well, and find it does not move in response to the shocks. However we think it is still important to correct for utilization, as it is a channel through which the Canadian economy could respond.

¹¹The paper reports traditional VAR-based impulse responses throughout. Impulse responses computed using local projections (Jordà, 2005) are quite similar and available upon request.

¹²This problem has been commonly observed in response to estimated TFP shocks (Galí, 1999), and news shocks were originally discussed as a possible solution. For instance Beaudry and Portier (2006) identify news shocks as

explain the unconditional positive comovement of these variables in the data. As Barsky and Sims (2011) point out, however, the responses to the news shock shown here are consistent with the predictions of a simple neoclassical growth model augmented with news shocks. As the response of hours is eventually positive, our news shock does generate comovement a few periods after impact, indicating that it is an important component of business cycle fluctuations in the medium term. On the other hand, Barsky, Basu, and Lee (2014) argue that it is unclear whether the comovement in the dynamic paths of all variables is due to the news shock itself or the realized productivity growth.

The impulse responses to the NTE shock look noticeably different. There is an impact increase in output, consumption, hours, and the forecast variable. There is a very small and insignificant decrease in measured TFP, which might be due to the quarterly series not perfectly correcting for utilization as discussed in Section 2.3. The business cycle generated by the shock lasts approximately three years. A substantial empirical literature beginning with Galí (1999) has previously argued that demand shocks are promising for explaining business cycles. Ours is (to our knowledge) the first paper to directly measure these shocks based on forecast or confidence data while ensuring they are uncorrelated with both current and future technological change. We discuss the relationship of our identified shock to the literature on demand shocks in Section 4.3.

The top panels of Tables 1-3 report the share of the forecast error variances of the US macro aggregates accounted for by the TFP, news, and NTE shocks respectively. At short frequencies, the NTE shock appears most important. It accounts for 65-75% of the variation in GDP, 18-22% in consumption, and 62-71% in hours at horizons 1 year or less. By contrast, at these frequencies surprise TFP shocks explain less than 8-12% of the variation in GDP, 2% in consumption, and 2-8% in hours. The news shock does a little bit better for consumption (36-48%), but is about equally unimportant for GDP and hours. Not surprisingly, at longer frequencies the news shock increases in importance. Barsky and Sims (2011) reach a qualitatively similar conclusion about the news and surprise TFP innovations, and point out that unexplained shocks were responsible for most of the variation at business cycle frequencies in domestic aggregates. Our analysis has now

the innovation in stock prices orthogonal to current TFP and find that the identified shock does generate positive comovement on impact. The news shocks identified in that paper capture a much longer-term improvement in technology, and therefore dissimilar to those in Barsky and Sims (2011) and our paper. Furthermore, the Beaudry and Portier (2006) identification scheme has been shown to deliver non-unique dynamic paths when extended to several variables (Kurmman and Mertens, 2014).

identified one such shock.

International Transmission. Figures 4-6 report the impulse responses of the Canadian variables to the three identified US shocks. The first panel in each figure sets the stage for the remainder of the results. It shows the impulse responses of Canadian utilization-adjusted TFP. None of the three identified shocks have a perceptible impact on Canadian technology. The news shock actually leads to a barely visible, though persistent and statistically significant increase in Canadian TFP beginning about five quarters ahead. This might indicate the presence of technology spillovers, but the magnitude is quantitatively tiny. Thus, whatever impact of US shocks on Canada that we find is not accompanied by a change in Canadian productivity.¹³

The three shocks lead to very different reactions of Canadian GDP. Neither shock to true TFP leads to an impact increase in GDP. The surprise TFP innovation in the US generates the smallest visible spillovers, with a slight increase in output three quarters after impact. The increase is short-lived, peaking at four quarters, after which Canadian output quickly returns to trend. In contrast, the news shock leads to more persistent Canadian output growth. GDP starts to increase two quarters after impact, lagging one quarter behind its US counterpart. The effects of the shock are more long-lived, with GDP peaking a little over two years after impact. At five years, output is still significantly above steady state.

The most striking is the response to the NTE shock. Canadian GDP jumps on impact, in sync with US output. It increases further for two quarters, before gradually returning to steady state. The effects of the shock are significant for two and a half years, demonstrating that the NTE shock has the potential to generate output comovement at high frequencies.

As it is clear that Canadian TFP is not affected, we propose one channel, consistent with our results, through which US expectational shocks could generate spillovers. As Figure 6 shows, Canadian exports to the US and imports from the US show the strongest responses to the NTE shock. Both series jump on impact, a two percent deviation from trend. They demonstrate a strong hump-shaped pattern: the increase in Canadian exports peaks at one quarter. However they stay significantly above trend for two years. Since the US is Canada's largest trade partner and the

¹³It may seem surprising that Canadian TFP does not react much to US TFP shocks. Whether or not there are noticeable cross-border technology spillovers is an open question, that to our knowledge has not been addressed comprehensively. In ongoing work (Levchenko and Pandalai-Nayar, 2017), we implement the full Basu, Fernald, and Kimball (2006) procedure on 30 countries over a period of 30 years, and also find that utilization-adjusted TFP growth is on average uncorrelated among G7 countries.

NTE shock generates increased demand in the US, this response is unsurprising.

The increased exports do not come through lowered Canadian consumption. Rather, the factors of production are used more intensively following a US NTE shock: Canadian hours increase, as does capital utilization. The increased production for export increases GDP, and generates an income effect which leads to higher consumption on impact. Demand for imports increases as well as a result of the higher consumption, and US exports to Canada rise.¹⁴ The empirical evidence clearly suggests that the NTE shock has the potential to not just generate a domestic business cycle, but explain both international synchronization as well as the positive correlation between exports and imports (Engel and Wang, 2011).

The news shock also generates comovement between Canadian exports and imports, but the impact effect is actually negative. The impact of higher future demand in the US contains both a substitution effect and an income effect. Holding TFP and production constant, the news shock would increase the price of future Canadian output and lead to a substitution effect towards consumption today. That said, cheaper future imports lower the price of future output and induce a negative substitution effect. However, the income effect from the future prolonged period of high export demand should unambiguously increase consumption and decrease hours. Each of these effects cannot be isolated in our framework, but the net effect is a slight decrease in Canadian hours after a US news shock, an insignificant decrease in GDP and a decrease in exports on impact. Consumption does not jump, so the wealth effect is not dominant, but it also begins to increase at about Q3.¹⁵ After one year, there is positive comovement among the key US and Canadian aggregates following a news shock. This implies that news shocks could also be an important component of comovement at medium- to long-term frequencies. It is unclear why US exports fall on impact. One possible explanation is weak demand in Canada coupled with the decreased production in the US.

The discussion above points to the different and complementary roles of the news and NTE shocks in generating business cycle spillovers. Forecast error variance decompositions provide additional support for the importance of NTE shocks at shorter frequencies, and of the news shocks at longer frequencies, internationally as well as domestically. The bottom panels of Tables 1-3 report

¹⁴Of course this is only one plausible channel. Schmitt-Grohé (1998) finds that exports are not a strong enough channel for the transmission of a generic shock to US output to Canada. That paper does not distinguish between the types of shocks that affect the US, however.

¹⁵Other explanations such as habit formation could also play a role here.

the shares of forecast error variances of the Canadian macro aggregates accounted for by the three identified US shocks. At short frequencies, the NTE shock is by a large margin the most important of the three. The NTE shock contributes substantially to the forecast error variance of US-Canada trade, explaining up to 44% of the variance of Canadian exports and 41% of imports at the one year horizon. It also explains a large fraction of the forecast error variance of Canadian output (41% at one year), hours and utilization (over one third), and consumption (8-12%). The impact of this short run US ‘demand’ shock on a smaller trading partner is persistent, as it still accounts for 36% of the variance of output at 10 years. The small share of Canadian TFP variation attributed to the NTE shock at 10 years is likely due “leakage” in the utilization adjustment, as our procedure for Canada is even coarser than the Fernald (2014) method on the US data.

In contrast the news shock is only responsible for very long run variation in TFP, output, and consumption, and does not contribute much to explaining the forecast error variance of other Canadian variables. The surprise TFP shock contributes very little to the forecast error variance of the Canadian aggregates at any frequency.

As further evidence on the importance of both non-technology expectations and news shocks for international comovement, we construct correlations of key variables conditional on only one type of shock. As in Galí (1999), these correlations can be inferred directly from the structural impulse response coefficients. Formally, the correlation of variables j and k conditional on shock i , ρ_{jk}^i , is

$$\rho_{jk}^i = \frac{\sum_{h=0}^{\infty} A_{ji}^h A_{ki}^h}{\sqrt{\sum_{h=0}^{\infty} (A_{ji}^h)^2} \sqrt{\sum_{h=0}^{\infty} (A_{ki}^h)^2}},$$

where A_{ji}^h is the lag- h , (j, i) -th element of the matrix $A(L)$ of lag order polynomials of the structural moving average representation of the VAR, that captures the impulse response of variable j to shock i at lag h . In practice, we compute these correlations for a finite but large maximum horizon of 10000 periods.

The results of this exercise are in Table 4. The NTE and news shocks both generate high correlations (both 0.99) of US and Canadian output, while the surprise TFP innovation delivers a much lower correlation than observed in the data. The surprise TFP shock actually generates a slightly negative US-Canada correlation of consumption (-0.13) and a strongly negative (-0.47) US-Canada correlation of hours. While both news and NTE shocks deliver strongly positive con-

sumption correlations, the correlation of hours due to the news shock is too low at only 0.46, but due to the NTE shock it is too high at 0.98. The NTE shock comes the closest to explaining the unconditional cross-correlations of exports from Canada with US output.¹⁶

In summary, the impulse response functions, variance decompositions, and conditional correlations show that surprise TFP innovations, which are usually assumed to be the key driver of IRBC models, play a negligible role in the international transmission of shocks. NTE shocks are important for transmission at higher frequencies, while news shocks play a stronger role at medium/long frequencies.

We conclude this section by discussing the role of all three shocks in recent business cycles. Figures 7–8 display the historical decompositions of the US and Canadian GDP and trade flows into the components due to the non-technology expectations and the combined technology (surprise plus news) US shocks. While both technology and non-technology shocks contributed to the great recession in the US, the fall in output in Canada appears driven predominantly by the NTE shock. Similar patterns are visible for Canadian exports and imports, indicating the NTE shock played a key role in the transmission of the recent recession. The NTE shock does not appear to contribute equally to all recessions however, with the dips in output and consumption in the 1981–82 recession driven primarily by news.

3.1 Robustness

Consumer confidence. To check robustness of the results to the choice of expectational variable, we replace the SPF GDP forecasts in the VAR with the E12Y variable from the Michigan Survey of Consumers, constructed from the responses to the question *‘And how about a year from now, do you expect that in the country as a whole, business conditions will be better, or worse than they are at present, or just about the same?’* The results from this exercise are in Appendix Figures A1–A6. Reassuringly, the patterns described above are robust to the expectational variable used to identify the NTE shock.¹⁷

¹⁶Interestingly, the surprise TFP innovation does a reasonable job of reproducing the cross-correlations of US output, consumption, and hours, despite the impact impulse responses being inconsistent with closed economy RBC models. As King and Rebelo (1999) point out, data generated by feeding utilization-adjusted TFP into a model with sufficient internal propagation mechanisms does a reasonable job of matching historical US time series. The news and NTE shocks also match the correlation of output and consumption well, but both undershoot the hours and output correlation.

¹⁷One notable difference is that while the SPF GDP forecast does not jump in response to the TFP news shock, consumer confidence does. This could be due to the horizon difference in the two variables: the GDP forecast is

We also examine the properties of the forecast variables used in estimation. If the forecast is a very accurate predictor of future GDP, then the NTE shock by construction would maximize a substantial fraction of the forecast error variance of GDP as well. Table 5 presents the correlations of the forecast and consumer confidence variables, in growth rates, with GDP growth. The forecast of GDP growth is highly correlated with contemporaneous GDP growth (correlation of 0.93), but substantially less correlated with realized future GDP growth one quarter ahead (correlation of only 0.38). The growth of consumer confidence does display the highest correlation with one quarter ahead GDP growth, but this correlation is still very low at 0.25. This is evidence that selecting the NTE shock as the shock that maximizes its contribution to the forecast error variance of the forecast/consumer confidence variables is not equivalent to simply selecting a shock that by construction explains the bulk of the variation in GDP at business cycle frequencies.

Monetary policy shocks. Our empirical strategy permits the separate identification of a monetary policy shock. We extend the core VAR to include the federal funds rate, ordered fifth. The identification of the monetary policy shock is then standard (Christiano, Eichenbaum, and Evans, 1999): we assume that the shock has no contemporaneous impact on the variables ordered above the federal funds rate (TFP, Consumption, GDP, and Hours), but does have an impact on the variables below (the expectational variable). This simply requires certain zero restrictions on the impact matrix $\tilde{A}D$.¹⁸ The identification of the news and NTE shocks then follows as in the baseline with the added monetary policy shock restriction.

Figure 9 plots the impulse responses of the main US macro aggregates and Canadian GDP to the monetary policy shock and the NTE shock side by side. Controlling explicitly for monetary policy shocks does not change the properties of the NTE shock. As in the baseline, consumption, output, and hours all increase on impact in response to the NTE shock. Interestingly, the federal funds rate increases following an NTE shock, pointing to policy tightening in response to increased demand generated by positive expectations. The federal funds rate rises slightly on impact, and then increases sharply further in the second quarter. It then stays flat for about two years, and the

of the next quarter, whereas the confidence variable asks about economic conditions a year from now (the shortest horizon available in confidence data). It could be that the news shock reveals productivity improvements more than one quarter ahead, explaining no response of the quarter-ahead forecast. But the news shock may reveal productivity improvements a year ahead, which will appear in a year-ahead confidence.

¹⁸If the monetary policy variable is ordered in position j , the restrictions simply imply that the j th column of $\tilde{A}D$ must have zeros in rows 1 through $j - 1$.

subsequent decline is slower than following a monetary policy shock.

Figure 9 also makes clear that the identified NTE shock is not a traditional monetary policy shock. In the figure, the monetary policy shock is an unexpected decrease in the federal funds rate, and is thus expansionary. The macroeconomic aggregates respond broadly as expected to a monetary policy loosening (Christiano, Eichenbaum, and Evans, 1999). Output and consumption are flat initially and then rise. Hours decrease for about 5 quarters, and then increase. Turning to comovement, the path of Canadian GDP tracks US GDP in both cases, so both shocks generate spillovers. However, the dynamic responses are different: a negative shock to the Federal Funds rate generates a gradual expansion, while the NTE shock generates a short-lived expansion on impact.

The first column of Appendix Table A1 reports the forecast error variance decomposition for the NTE shock in the specification that includes the monetary policy shock. Augmenting the analysis with the monetary policy shocks does not affect the share of the forecast error variance attributed to the non-technology expectations. The NTE shock is still the dominant shock for the variability of output and hours at time horizons of less than two years.¹⁹

Fiscal policy shocks. We augment the VAR with a fiscal policy variable, namely government spending.²⁰ Identifying fiscal policy shocks in our VAR is more challenging than monetary policy. The best-known identification scheme, due to Blanchard and Perotti (2002), identifies fiscal policy shocks by ordering fiscal policy variables first in the VAR. The assumption is that shocks to output are unlikely to affect government spending within a quarter. Unlike Blanchard and Perotti (2002)'s, our VAR includes TFP, and there is no conventional way of ordering government spending and TFP. Ordering government spending first and TFP second yields an undesirable property that government spending shocks affect TFP. Ordering TFP first and government spending second would imply that shocks to TFP affect government spending within the quarter, but shocks to output do not. We therefore simply augment our core VAR with government spending, and proceed to identify the NTE shock as before. This is also not ideal, as it implies that TFP shocks can have contemporaneous effects on government spending. Nonetheless, under this approach reduced-form innovations in government expenditure can affect the identification of the NTE shock. If

¹⁹The monetary policy shock itself explains most of the variance of the federal funds rate (88% at one quarter) but does not contribute much to variance of the other variables: it explains 5% of the variance of US output and 9% of the variance of Canadian output at a horizon of five years (not reported).

²⁰To be exact, we use the Real Government Consumption Expenditures and Gross Investment (seasonally adjusted, federal, state, and local), NIPA Table 3.9.1.

the identified impulse responses to the NTE shock were similar to the conventional government spending shock, we could plausibly interpret the NTE shock as a fiscal policy shock.

To contrast the impulse responses of the core variables to the NTE shock identified in this manner to a standard government spending shock, we implement the canonical Blanchard and Perotti (2002) VAR and extract the fiscal policy shocks in that conventional setting. We obtain the responses of all core variables to this identified government spending shock.

The impulse responses to the NTE and fiscal policy shocks implied by this procedure are reported in Figure 10. The conclusions that emerge from this exercise are similar to the monetary policy exercise. First, adding a fiscal policy variable does not change substantively the behavior of our NTE shock. And second, macro variables respond to the NTE shock and the fiscal policy shock very differently. In addition, the second column of Appendix Table A1 reports the forecast error variance decomposition for the NTE shock in the specification with government spending. The share of the forecast error variance of the macro variables accounted for by the NTE shock does not change appreciably when we add fiscal policy. Thus, we can be confident that the NTE shock explored in this paper is not a fiscal policy shock.

Finally, we correlate the series of fiscal policy shocks obtained from implementing the Blanchard and Perotti (2002) VAR with our baseline NTE shocks. The correlation between the two series is quite low at 0.128, giving further credence to the notion that our NTE shock is not due to fiscal policy.

Additional Exercises. Appendix B discusses a large battery of additional robustness checks on the main results. In all cases, the main conclusions about the shape and importance of the non-technology expectations shock and its transmission to Canada are robust. The exercises include: (i) adding more variables, such as stock prices, financial conditions, oil prices, or estimating a FAVAR; (ii) examining responses of prices and wages; (iii) checking whether the results are driven by transmission or correlated shocks; (iv) comparing the NTE shocks to uncertainty shocks; and (v) a number of additional robustness checks, such as including Canadian variables as core variables in the VAR, and including multiple Canadian variables in the VAR.

4 Theoretical Framework and Discussion

4.1 A Model of International Transmission of Expectational Shocks

This section presents a formal illustration of how US expectational shocks transmit to Canada. We use a simplified version of the confidence shock model of Huo and Takayama (2015, henceforth HT). For the US economy, we maintain much of the framework in HT. A unit measure of islands in the US meet bilaterally and trade with each other in each period. Each island does not observe its partner's productivity perfectly, and its signal about the trading partner's productivity is affected by an aggregate confidence shock. The US islands do not perfectly observe the confidence shock, and thus face the standard forecasting problem. Aggregate fluctuations in the US are generated by the combination of the expectational shock and imperfect information. HT show that a positive confidence shock generates a US output expansion in this setting.

To this model we add Canada as an island of measure zero.²¹ Canada trades with a randomly selected US island in each period. Canada does not experience the expectational shock, in the sense that it observes perfectly both its partner's productivity and the magnitude of the US confidence shock. The key result is that, nonetheless, Canadian output increases in response to the US expectational shock, generating output comovement between the US and Canada. This is because Canada knows that following a positive US expectational shock, its US trading partner will overestimate Canadian productivity and therefore have high output, increasing its demand for Canadian output.

Preliminaries The model economy comprises of a unit measure of islands corresponding to the US indexed by $i \in [0, 1]$ as in Angeletos and La'O (2013) and HT, plus a single island of measure zero corresponding to Canada and denoted by c . Labor is supplied elastically by households and is the only input in production. Island i has time-invariant TFP of a_i . The cross-sectional productivity distribution is known to all agents, and given by $a_i \sim N(0, \tau_a^{-1})$. A time period contains two stages. In the first stage, islands are randomly bilaterally matched, receive (potentially noisy) signals about their partner's productivity, and decide how much to produce. In the second stage, bilateral trade and consumption take place.

Let island i match with island j . Angeletos and La'O (2013) and HT show that island i 's policy

²¹This assumption is designed to reflect the notion that Canada is a small open economy with respect to the US.

rule for output, expressed as a log deviation from steady state is

$$y_{it} = \alpha_0 a_i + \alpha_1 E_{it} [y_{jt}], \quad i \in [0, 1], i = c \quad (10)$$

where $E_{it}[\cdot] \equiv E[\cdot | \Phi_{it}]$ is island i 's expectation operator, conditional on its time t information set Φ_{it} . We maintain the following assumptions on the constants α_0 and α_1 : (i) $\alpha_0 > 0$ – output increases in own productivity; and (ii) $0 < \alpha_1 < 1$ to ensure i 's positive output response to higher trading partner's expected output (strategic complementarity) and a stable solution.²²

Information structure A US island i receives two signals:

$$x_{it}^1 = a_{jt} + \xi_t, \quad \xi_t \sim N(0, \tau_\xi^{-1}),$$

and

$$x_{it}^2 = \xi_t + u_{it}, \quad u_{it} \sim N(0, \tau_u^{-1}).$$

The first signal x_{it}^1 is a noisy signal of the time- t partner island j 's productivity a_j . The signal contains noise due to the economy-wide confidence shock ξ_t . Island i also receives a noisy private signal x_{it}^2 about the confidence shock ξ_t . Finally, we assume that when a US island gets matched with the Canadian island, it knows that it is meeting the Canadian. This is relevant because higher-order beliefs matter in this environment.

The Canadian island c is not directly affected by the confidence shock. It observes both its current trading partner's productivity and ξ_t . Formally, when Canada matches with a US island j , the signals it receives are²³

$$x_{ct}^1 = a_{jt} + \xi_t,$$

and

$$x_{ct}^2 = \xi_t.$$

²²There are a variety of ways to arrive at (10), and the constants α_0 and α_1 are functions of the parameters governing the physical environment of production and trade. For example, in HT's framework, $\alpha_0 = \frac{1}{1 - \frac{\theta}{(\gamma+1)\omega}}$ and $\alpha_1 = \frac{(1-\omega)}{\frac{(\gamma+1)}{\theta} - \omega}$, where ω governs home bias (i 's preference for its own variety), γ is the inverse Frisch labor supply elasticity, and θ is the labor share.

²³This amounts to assuming that $\tau_u^{-1} = 0$ for Canada.

Note that in this framework, Canada knows that its US match faces the expectational shock, and cannot perfectly observe a_c . This implies that higher-order beliefs will still play a role in c 's policy rule for output y_{ct} . Following HT, we assume that second stage trade and consumption are carried out by shoppers, who die before they can reveal the value of ξ_t to the island's producers. This implies that c cannot inform j of the state of the economy when trading occurs, and the US economy does not learn the Canadian island's true a_c over time.²⁴

Equilibrium Output As the Canadian island and its time- t US trading partner are of measure zero, they do not affect US outcomes. Thus, the behavior of an individual US island matched with another US island, and aggregate US output correspond to the closed economy model analyzed by HT under i.i.d. shocks. It can be shown that the equilibrium policy rule for any US island i that trades with another US island $k \neq c$ is:

$$y_{it} = h_{us}^a a_i + h_{us}^1 x_{it}^1 + h_{us}^2 x_{it}^2. \quad (11)$$

Aggregate US output is therefore:

$$y_t = \int y_{it} = \int h_{us}^a a_i + \int h_{us}^1 x_{it}^1 + \int h_{us}^2 x_{it}^2 = (h_{us}^1 + h_{us}^2) \xi_t. \quad (12)$$

As in a number of papers, including Angeletos and La'O (2013), Benhabib, Wang, and Wen (2015) and HT, the confidence shock generates aggregate fluctuations. It can be shown that a positive expectational shock ξ_t increases aggregate output in the US ($h_{us}^1 + h_{us}^2 > 0$).²⁵

Canadian output and output of the US island matched with Canada must be analyzed separately because of Canada's different information set, and because the US island knows when it is matched to Canada. The coefficients on the policy rules for these agents will in general not coincide with those of the US islands matched to other US islands in equation (11).

²⁴An alternative assumption would be that a_c is i.i.d. over time, with the distribution given by $a_c \sim N(0, \tau_a^{-1})$. In this case, the US island's inference problem is unchanged. Canada will experience productivity-driven output fluctuations, but its response to ξ_t conditional on its own productivity will be the same as that in the Proposition 1 below.

²⁵The detailed derivations are available on request. Note that while the entire history of signals for all matches in $t-1, t-2, \dots$ belong to island i 's information set, all the shocks are i.i.d. Therefore, having traded with c in the past does not affect i 's decisions in t . See HT for the complete treatment of the closed-economy case in which the shock processes are persistent.

Proposition 1 (Canadian Output and its Response to US Expectational Shocks). *The policy rule for Canada is*

$$y_{ct} = h_c^a a_c + h_c^1 x_{ct}^1 + h_c^2 x_{ct}^2, \quad (13)$$

with $h_c^1 + h_c^2 > 0$.

Canadian output increases following the US expectational shock:

$$\frac{dy_{ct}}{d\xi_t} > 0. \quad (14)$$

Proof. See Appendix C. □

Proposition 1 along with (12) provide a model of positive output comovement between US and Canada generated by US expectational shocks. Even though Canada is perfectly informed about its trading partner’s productivity and other fundamentals of its own and the US economies, it increases output following a positive US shock because it knows that the US island will produce a lot in this state of the world and demand for Canadian products will be high.

4.2 Internal Validation

This section describes the results of an internal validation exercise. The question is whether the identification scheme proposed in this paper can correctly extract the expectational shock when applied to model-generated data. To this end, we use the full-fledged DSGE version of the confidence shock model in Section 4.1. This model is developed, solved, and estimated by HT, which can be consulted for the full details. Relative to the simple model above, the full-fledged model is dynamic and enriched with the following features: (i) a persistent aggregate TFP shock; (ii) a persistent expectational shock, with ξ_t following an AR(1) process; (iii) persistent match productivity, with a_{jt} being AR(1) for each i ; (iv) capital accumulation subject to adjustment costs; (v) search frictions in the goods markets following Bai, Rios-Rull, and Storesletten (2017).

We simulate model-generated data for all the variables in our VAR: TFP, Consumption, Hours, GDP, and the GDP forecast for 169 quarters, the same length of time as our actual data.²⁶ We then implement our identification strategy on model-generated data. As the model has only 2 shocks (TFP and ξ_t), while the VAR has 5 variables, we add a modest amount of measurement

²⁶We simulate the DSGE model for 1600 periods, and use the last 169 periods as our sample.

error to each equation in the VAR in order to ensure that the system of regression equations is not overdetermined.²⁷

Appendix Figure A12 reports the impulse responses of the VAR variables to an expectational shock extracted from model-generated data. The impulse responses are virtually identical to the analytical impulse responses in HT (their Figure 7), and quite similar to the empirical impulse responses to the NTE shock reported above. Importantly, our identification strategy correctly separates the TFP from the expectational shock. Appendix Table A3 reports the correlations between the true TFP and expectational shocks with which the model was simulated, and the TFP and NTE shocks that were identified by our SVAR. The correlations are quite high, with the true and VAR-based expectational shock having a 0.934 correlation with each other.

We conclude that (i) our identification procedure correctly recovers the true expectational shock when applied to model-simulated data driven by TFP and expectational shocks; and (ii) the impulse responses to the expectational shock of the variables in the VAR applied to model-simulated data are quite similar to the impulse responses generated by the NTE shock in actual US data.

4.3 Discussion

We now discuss additional conceptual and estimation issues falling outside the scope of the formal model above.

News and noise shocks. Signals about future TFP are likely to be riddled with noise. Blanchard, L’Huillier, and Lorenzoni (2013) point out that news shocks cannot be separated from noise shocks (unfounded signals about future TFP) in a structural VAR setting, since if the econometrician can extract different paths of variables in response to a noise shock, so can the consumer. It is clear that these noise shocks are not related to our non-technology expectations shock, which is identified as a fully rational change in forecasts or confidence orthogonal to surprise and news TFP. That is, in our economy neither the forecaster/consumer nor the econometrician will believe the NTE shock to be either news or a noisy signal of future TFP. Further, Barsky and Sims (2012) assess the importance of noise or ‘animal spirits’ shocks and find that they do not account for a

²⁷We add Normal $(0, \sigma)$ measurement error, where σ is 1/20th of the standard deviation of model-generated GDP growth. Shrinking this measurement error further improves the performance of our identification scheme in extracting the correct expectational shock.

substantial portion of the relationship between confidence and output. This supports the notion that the responses to the news TFP shocks in Figure 2 are informative of the impact of true news.

Demand shocks. Our identified shock could be a combination of several shocks traditionally considered ‘demand’ shocks. The strategy simply relies on the forecasters rationally expecting an increase in GDP that is not due to current TFP innovations or news/noise about future TFP. To the extent that an increase in ‘demand’ leads agents to forecast an increase in economic activity, it would be identified as an NTE shock in our framework. However, our shock is not consistent with all of the results associated with demand shocks previously identified in the literature. For instance, the demand shock identified by Galí (1999) as the shock orthogonal to changes in long run labor productivity leads to a temporary increase in labor productivity. Monetary shocks are also commonly proposed as demand shocks. As we demonstrated in Figure 9, controlling for monetary policy shocks does not significantly change the shape of the impulse responses to the NTE shock. Further, monetary policy and NTE shocks generate very different impulse responses both in the US and Canada.²⁸ Similarly, our NTE shock does not appear to be a fiscal policy shock. The response of US consumption to the NTE shock would be consistent with models with taste shocks. For instance, Stockman and Tesar (1995) show that the addition of a preference shock increases the volatility of consumption.

Investment-specific technology shocks also exhibit properties that would appear similar to demand shocks, despite being shocks to technology. The dynamics of output and hours following our NTE shock do bear a resemblance to those in response to the investment-specific technology shocks as identified in Fisher (2006) post 1982:Q3. The magnitudes are very different however. More importantly, the investment-specific shock generates increases in labor productivity. From the responses of output and hours in Figure 3, this does not appear to be true for our NTE shock, increasing our confidence that it is not the investment-specific technology shock.

Relationship to standard neoclassical and New Keynesian models. A large body of work on closed-economy business cycles has established that (i) RBC models driven by a technology shock do well at matching the key moments of US data; and (ii) estimated technology shocks do

²⁸Angeletos, Collard, and Dellas (2014) find that monetary policy shocks deliver closed-economy business cycle moments very similar to their sentiment shock, but require the assumption that they affect the economy in an implausibly large way.

not deliver impulse responses similar to those in RBC models, calling into question the mechanisms driving the model's success (see also Galí and Rabanal, 2005). While surprise TFP shocks have proven of questionable value in explaining US business cycles, they have been even worse at explaining international transmission. The seminal work of Backus, Kehoe, and Kydland (1992) showed that even with correlated shocks, the gap between theory and data was large (see also Justiniano and Preston, 2010, for a recent statement of this result). Many variants of the original model have been proposed to improve on these results, with limited success. Part of the reason for this failure is that with uncorrelated technology shocks investment increases in the country receiving a positive TFP surprise. As a result, consumption is more highly correlated across countries than output, contrary to the data.

News shocks identified from structural VARs have proven a better fit to the predictions of the neoclassical model. However, due to the wealth effect on labor supply they do not generate the desired impact comovement in the closed economy. A similar outcome appears in an open economy setting: Kosaka (2013) finds that news shocks do help generate an international business cycle, but only when the model is parameterized with a low elasticity of substitution between domestic and foreign goods and no wealth effect on labor supply. Beaudry, Dupaigne, and Portier (2011) also rely on a low elasticity between the goods produced by different countries in order to generate positive comovement in response to a news shock.

Nominal rigidities have also been proposed as an explanation for the drop in hours in response to a technology shock. The intuition is simple: with sticky prices, only a fraction of firms adjust their prices downward in response to a productivity shock. Therefore, aggregate demand (price level) will rise (fall) less than proportionately to the shock, and hours will fall as a consequence (Galí, 1999; Galí and Rabanal, 2005). Galí and Rabanal (2005) also find in their estimated New Keynesian model that a pure preference shock accounts for the bulk of the variation in the output growth, hours, and the nominal interest rate. However, Angeletos, Collard, and Dellas (2014) estimate a medium-scale DSGE model allowing for nominal rigidities and find a shock that generates impulse responses for the US economy alone similar to our NTE shock. Allowing for nominal rigidities helps, but does not seem crucial in their results. At this stage, it is unclear whether sticky prices will be indispensable in a model of the international transmission driven by news and NTE shocks.

VAR invertibility. For a structural VAR to uncover the “true” shocks in an economic model, Fernández-Villaverde et al. (2007) demonstrate that the matrices associated with the state space representation of the model must satisfy a certain invertibility condition. In particular, this is necessary for the VAR to have a moving average representation. This condition can be violated if the structural shocks are not fundamental, i.e. they cannot be recovered from the current and lagged values of the variables. This is a particular concern for news shocks, which are shocks that contain information about future TFP innovations.²⁹ This problem is mitigated if the VAR has sufficient information (Forni and Gambetti, 2014). Beaudry and Portier (2014) discuss this issue in detail, and argue that while VARs with news shocks can be non-fundamental, they don’t have to be. Furthermore, Sims (2012) shows that even if the VAR is non-fundamental, the resulting impulse response functions can be good approximations of the true impulse responses. In our context, the baseline VAR includes several forward-looking variables, and the results are robust to including additional variables such as stock prices or a factor, further mitigating this concern.

5 Conclusion

We introduce a novel identification scheme to uncover the effects of surprise TFP innovations, news shocks, and non-technology expectations shocks. These shocks have very different implications for international comovement in US and Canadian data. The bulk of high-frequency business cycle comovement can be attributed to the non-technology shocks, while the news shocks are important for medium- to long-term synchronization. Surprise TFP innovations, which are the most common driver of IRBC models, are found to be nearly irrelevant for international business cycle synchronization. We provide a theoretical illustration of how expectational shocks in one country can increase output in its trading partner even if it does not experience the confidence shock itself.

²⁹This issue would be of less concern if the true VAR only contained surprise TFP innovations and non-technology expectations shocks, as neither are signals of future shocks. However, it would not be possible to identify NTE shocks without first controlling for news shocks.

Appendix A Data Appendix

We describe the algorithm used to construct a utilization-adjusted TFP series for Canada. Our procedure is adapted from Imbs (1999), as the quarterly data necessary to construct a series with the Fernald (2014) methodology are not currently available for Canada for the requisite time period. The method in Imbs (1999) is in the spirit of Basu, Fernald, and Kimball (2006), in that it also relies on identifying movements in unobserved (aggregate) utilization from observed changes in inputs and output. Unlike Basu, Fernald, and Kimball (2006), this method does not control for sectoral differences or non-constant returns to scale. We briefly describe the steps of the algorithm here, using commonly seen relationships from a firm's profit maximization problem. For a detailed derivation of the equations that follow see Imbs (1999).

1. Construct a starting capital stock series using the perpetual inventory method from official investment series I_t and a quarterly depreciation rate of 0.025. For the initial value of the capital stock we chose $K_0 = \frac{I_1}{r+g_I}$, where g_I is the growth rate of investment in Canada. We tested our results with other choices for the initial capital stock and found no substantive difference.
2. Construct an initial series for utilization u_t using the capital stock series K_t , output Y_t , and values for average depreciation $\bar{\delta}$ and the interest rate r from the equilibrium relationship $u_t = \left(\frac{Y_t/K_t}{Y/K}\right)^{\frac{\bar{\delta}}{\bar{\delta}+r}}$, where Y/K is the average period value.
3. Use the initial utilization series and the assumed relationship between depreciation and utilization $\delta_t = \bar{\delta} u_t^{1+r/\bar{\delta}}$ to construct a time-varying series for δ_t .
4. Together with the official series for investment and the time-varying δ_t , construct a new capital stock using the standard capital accumulation equation.
5. Using the new $\bar{\delta}$ and capital stock, return to step (1) and construct a new utilization series.
6. Iterate until the capital stock and $\bar{\delta}$ converge. Then construct the final implied u_t .
7. Construct a series for the household's labor effort e_t from $e_t = \left((1-\alpha)\frac{Y_t}{C_t}\right)^{f(w_t, N_t, Y_t)}$ using data on consumption C_t , wages w_t , and labor input N_t .³⁰

³⁰The derivation of this expression uses the household's optimization problem and can be found in Imbs (1999).

8. Construct the utilization-adjusted TFP series from the production function $Y_t = X_t (u_t K_t)^\alpha (e_t N_t)^{(1-\alpha)}$.

The only additional data series required for this procedure are data on investment and wages. For consistency with the rest of our data, both series were taken from the Ohanian and Raffo (2012) dataset, which in turn uses data from the OECD Main Economic Indicators along with national databases.

Appendix B Additional Robustness

Additional controls: Stock Prices, Financial Conditions, Oil Prices, FAVAR. Beaudry and Portier (2006, 2014) identify news shocks with a long-run restriction in a VAR with TFP and an index of stock prices. Our identification of news shocks is medium-run and based on the information content in forward-looking real variables. Stock prices are also forward-looking, so we test the robustness of our identified shock to adding stock prices as an additional control. We use the index of stock prices from Beaudry and Portier (2014), ordered second, but we maintain the medium-run identification strategy.³¹ Figures A7 and A8 display the impulse responses of US and Canadian GDP and hours, respectively, to an NTE shock while augmenting the VAR with the stock price variable. The results are very similar to the baseline specification. In particular, the impact effects are almost identical. The addition of stock prices as a control leads a slightly different dynamic path, but the difference is minor.

To control for the role of financial conditions, we augment the VAR with an interest rate spread variable (the Baa-Aaa spread), ordering it fifth after TFP, Consumption, GDP, and Hours. We identify a credit spread shock using a block identification approach, where we assume the slow-moving block (variables ordered above the credit spread in the VAR) respond to the financial shock with a lag. Since the expectational variable is ordered below the spread variable, it can react on impact to the financial shock. The NTE shock is then identified orthogonal to the credit spread shock. The results are presented in Figures A7 and A8. Once again, the addition of this variable does not change the essential properties of the sentiment shock or its transmission to Canada.

Another potential concern is whether our NTE shock might be picking up oil-price shocks. That is, perhaps not including a measure of oil prices would lead to omitted variable bias in our specification. We test this by augmenting the core VAR with the oil price index available from FRED.³² The IRFs of US and Canadian GDP and hours are reported in Figures A7 and A8. The responses of the core variables to the NTE shock are almost unchanged relative to the baseline. While we do not identify the oil price shock formally, below we report the impulse response of the

³¹We can increase H^{news} to an arbitrarily large number to approximate the long-run restriction in Beaudry and Portier (2006), and we do find that the responses of key variables to the news shock approach their findings (results available on request). However, as long-run restrictions can be problematic (Faust and Leeper, 1997), we favor our medium-run approach.

³²We use a seasonally adjusted consumer price index for all urban consumers, fuel oil and other fuels, series ID CUSR0000SEHE.

oil price to the NTE shock. The oil price index shows no impact response to the NTE shock, ruling out the possibility that the NTE shock is an oil price shock.

We also test the responses to adding the US-Canada real exchange rate or US CPI as additional controls. We construct the bilateral real exchange variable using the nominal Canadian-US dollar exchange rate and the US and Canadian consumer price indices from the International Financial Statistics. The units are Canadian basket/US basket, so an increase in the variable is a US appreciation. Figures A7 and A8 report the responses of US and Canadian GDP and hours when including the real exchange rate, and show that the main results are unaffected.

We augment the core VAR with the first factor identified in Forni, Gambetti, and Sala (2014) to increase the information available about the macroeconomy in identifying news and NTE shocks. Including this factor further mitigates the possible omitted variables issues in the VAR.³³ Figures A7 and A8 present the impulse responses of GDP and hours to the NTE shock in the FAVAR. Reassuringly, we find very similar responses of all core variables with the FAVAR, though the point estimates of the dynamic responses are smaller for longer frequencies.

Table A1 reports the share of the forecast error variance of selected core variables attributed to the NTE shock, while including each of the additional controls. The importance of the NTE shock for accounting for the forecast error variance of US GDP and hours and Canadian GDP does not differ appreciably from the baseline in each case.

Response of prices and wages. Figure A9 displays the impulse responses to the three identified shocks of the key price series: US CPI, US stock prices, oil prices, the US-Canada real exchange rate, and US wages.³⁴ The response of the price variables to the three shocks is consistent with theories of news and demand shocks (particularly, demand shocks not driven by shocks to the price variable itself). The US consumer price index (expressed in log levels) increases slightly following the NTE shock, and the increase is persistent. This response supports the notion that the demand shock embodied in the identified NTE shock is inflationary. By contrast, there is no response of US CPI to the surprise TFP shock, and prices fall following a news TFP shock. This difference is further illustration that the NTE shock affects the economy differently from

³³Forni, Gambetti, and Sala (2014) use 107 macroeconomic series to extract factors, which are used to augment a number of VARs to assess the non-fundamentalness problem in identifying news shocks. Details of the data series used are available in the appendix to their paper.

³⁴The response of US and Canadian GDP and hours in the VAR including US CPI are not substantially different from the other robustness exercises reported in Figures A7 and A8 so they are omitted to conserve space.

disturbances to technology.

There is no impact response of oil prices to the NTE shock, ruling out the possibility that the NTE shock is an oil price shock. Two quarters following the NTE shock, oil prices if anything rise modestly, indicating that times of positive NTE do not systematically coincide with low oil prices. In response to the news shock, oil prices fall and stay low, consistent with the decline in inflation documented in Barsky and Sims (2012).

The information content of stock prices is evident in the response to the news shock. On impact, there is a large jump in the stock price index, with a further increase for about five quarters followed by slow reversion. However, at the maximum horizon plotted (20 quarters) the index is still substantially above trend. Stock prices also display an impact increase in response to the NTE shock, but the increase is more muted. This suggests that the NTE shock is indeed a shock to higher-order beliefs about the economy, which are rational though not based on expected changes to TFP.

The bilateral real exchange rate displays an impact increase only in response to the news shock (this is followed by a gradual decline that approximately coincides with the actual increase in TFP). With the NTE shock, there is no response for two quarters, and then a slight but persistent US appreciation. The surprise TFP shock leads to a gradual depreciation of the real exchange rate. This is similar to the results in Nam and Wang (2015), who estimate the response of real exchange rates to news and surprise TFP shocks.

Finally, the last panel of Figure A9 displays the impulse responses of the US real wage to the three shocks. The real wage is constructed by deflating the BLS hourly nonfarm business sector compensation (series PRS85006103) by US CPI. In response to both surprise and news TFP shocks, the real wage increases on impact and stays permanently higher. In stark contrast, there is no impact of the NTE shock on the real wage. This result further supports the notion that the NTE shock is a “demand” shock. In addition, it helps rule out the possibility that our NTE shock is a labor supply shock (Shapiro and Watson, 1988). A labor supply shock should plausibly decrease wages, whereas our NTE shock has no wage impact.

Spillovers or correlated shocks. Our three US shocks are identified using only US data. Therefore, there is the possibility that the observed impulse responses of Canadian variables to US shocks are due to exogenously correlated shocks affecting both countries simultaneously. For

the two technology shocks, this is unlikely to be a problem: Figures 4-6 show that Canadian TFP does not respond to any of the three identified US shocks. However, it may still be that there are exogenous common shocks to US and Canadian non-technology expectations. We evaluate this hypothesis by identifying a surprise TFP innovation, a news shock, and an NTE shock in Canadian data alone. We then check the correlation of these identified shocks with their US counterparts. Note that if there are indeed spillover effects from US to Canada, this is not a clean exercise: Canadian expectations of future Canadian economic activity will rise following a US NTE shock, not because optimism increased exogenously in Canada, but because Canadian agents know that a positive NTE shock in the US will increase Canadian GDP via cross-border transmission. In this sense, identifying a Canadian NTE shock as if it were a closed economy stacks the deck against us, as those shocks might embody Canadian agents' endogenous revisions of expectations following an increase in US confidence.

The Canadian expectational variable is an index constructed from the responses to the question "Do you expect overall economic conditions in Canada six months from now to be: Better/Same/Worse", and comes from the Conference Board of Canada. As in Barsky and Sims (2012), we construct the composite index by subtracting the percentage of responses answering 'worse' from those answering 'better' and adding 100. This series corresponds most closely to the US confidence series from the Michigan Survey of Consumers. Therefore we compare the Canadian shocks identified with these data with those identified from the five variable core US VAR using the consumer confidence series. Table A2 presents the correlations between the US and Canadian shocks identified this way. The correlation of the surprise TFP innovations is 0.16. The US and Canadian NTE shocks are actually slightly positively correlated, while the news shocks are negatively correlated. However, these correlations are low (0.18 for the NTE shock and -0.17 for the news shock), indicating that the spillovers observed in the estimated impulse responses are unlikely to be driven primarily by exogenous common shocks.

As an alternative way to check whether correlated Canadian confidence shocks could drive our impulse responses, we control directly for the Canadian confidence series when computing impulse responses of the Canadian macro aggregates to the US NTE shock. The results (not reported to conserve space) show that controlling for Canadian confidence leads to impulse responses to US NTE shocks that are virtually indistinguishable from the baseline.

Our baseline identification strategy does not allow feedback effects from the Canadian variables to the US variables in the VAR coefficients. These restrictions are testable. For each Canadian variable, we perform likelihood ratio tests comparing restricted (baseline) and unrestricted VARs. When the Canadian variables are TFP, output, hours, exports, or imports, we fail to reject the null that the restricted VAR is the true model. For Canadian consumption, however, the null is rejected. Therefore Figure A10 presents the responses to the identified shocks when Canadian consumption is the sixth variable and there are no restrictions on the lagged coefficients. Substantively, this does not change the baseline results for any shock, indicating the addition of Canadian consumption as a core variable is not extremely informative for the news or NTE shocks.

We also attempted to test an alternative model where the Canadian variables were the core entries in the VAR and the US variables were treated as non-core. However, this model does not converge for any US variable. As we cannot estimate the restricted version of this model, we could not evaluate this alternative setup. This is supportive of our assumption that while shocks to the US matter for Canada, the converse is not true.

Uncertainty. The robustness checks above show that the NTE shock is not a monetary policy shock, fiscal policy shock, or an oil price shock. The NTE shock has characteristics suggesting it is similar to a rational “optimism” shock, where the optimism is not related to a change in productivity. As there may be some relationship between confidence and uncertainty (see, e.g. Ilut and Saijo, 2016), we also examine whether this shock is related to uncertainty (second moment) shocks. We obtain uncertainty shocks by implementing the Bloom (2009) VAR. The correlations of our non-technology expectations shock with the average Bloom (2009) quarterly uncertainty shock and the max uncertainty shock in the previous quarter are -0.0046 and -0.0144 respectively. Thus, our shocks bear little resemblance to the uncertainty shocks as estimated in the literature.

Additional exercises. The baseline analysis adds Canadian variables to the VAR one by one. To assess the robustness of the results to this approach and permit the Canadian variables to interact among themselves, we include the three main Canadian indicators – GDP, consumption, and hours – together as non-core variables in the VAR. As before, we do not allow these Canadian series to affect the US series, but the Canadian variables interact with each other. The impulse responses of the three main Canadian variables to the three US shocks are reported in Figure A11.

It is clear that including several Canadian variables together and letting them interact does not change the basic conclusions.

Finally, we check the responses of all variables to variations in the horizons of identification of the NTE shock, and find no significant qualitative difference for $H^{nte} = 4, 8, \text{ or } 16$. We also vary the forecast variable used in identification, using forecasts of GDP two quarters ahead and three quarters ahead. The qualitative shape of the dynamic responses remains the same. To conserve space the results are not reported here, but are available on request.

Appendix C Proof

Proof of Proposition 1. Denote by j the island that matched to Canada at time t . Guess that the equilibrium policy rules for output for both c and j are linear in their signals:

$$y_{ct} = h_c^a a_c + h_c^1 x_{ct}^1 + h_c^2 x_{ct}^2 \quad (\text{C.1})$$

$$y_{jt} = h^a a_j + h^1 x_{jt}^1 + h^2 x_{jt}^2 \quad (\text{C.2})$$

The US island's expectation of Canadian output is:

$$\begin{aligned} E_{jt}[y_{ct}] &= E_{jt}[h_c^a a_c + h_c^1 x_{ct}^1 + h_c^2 x_{ct}^2] \\ &= h_c^a E_{jt}[a_c] + h_c^1 E_{jt}[a_j + \xi_t] + h_c^2 E_{jt}[\xi_t] \\ &= h_c^a x_{jt}^1 + h_c^1 a_j + (h_c^1 + h_c^2 - h_c^a) \left(\frac{\tau_a}{\tau_\xi + \tau_a + \tau_u} x_{jt}^1 + \frac{\tau_u}{\frac{1}{\sigma_\xi^2} + \tau_a + \tau_u} x_{jt}^2 \right), \end{aligned}$$

where the last line comes from applying Bayes' rule. Therefore,

$$\begin{aligned} y_{jt} &= \alpha_0 a_j + \alpha_1 \left[h_c^a x_{jt}^1 + h_c^1 a_j + (h_c^1 + h_c^2 - h_c^a) \left(\frac{\tau_a}{\tau_\xi + \tau_a + \tau_u} x_{jt}^1 + \frac{\tau_u}{\tau_\xi + \tau_a + \tau_u} x_{jt}^2 \right) \right] \\ &= h^a a_j + h^1 x_{jt}^1 + h^2 x_{jt}^2 \end{aligned}$$

For this equality to hold for any shocks and signals, it must be that

$$h^a = \alpha_0 + \alpha_1 h_c^1 \quad (\text{C.3})$$

$$h^1 = \alpha_1 h_c^a + \alpha_1 (h_c^1 + h_c^2 - h_c^a) \frac{\tau_a}{(\tau_\xi + \tau_a + \tau_u)} \quad (\text{C.4})$$

$$h^2 = \alpha_1 (h_c^1 + h_c^2 - h_c^a) \frac{\tau_u}{(\tau_\xi + \tau_a + \tau_u)}. \quad (\text{C.5})$$

Similarly, c 's expectation of j 's output can be expressed as:

$$\begin{aligned}
E_{ct}[y_{jt}] &= h^a E_{ct}[a_j] + h^1 E_{ct}[x_{jt}^1] + h^2 E_{ct}[x_{jt}^2] \\
&= h^a E_{ct}[a_j] + h^1 E_{ct}[a_{ct} + \xi_t] + h^2 E_{ct}[\xi_t + u_{jt}] \\
&= h^a E_{ct}[a_j] + h^1 E_{ct}[a_{ct}] + h^1 E_{ct}[\xi_t] + h^2 E_{ct}[\xi_t] + h^2 E_{ct}[u_{jt}] \\
&= h^a(x_{ct}^1 - x_{ct}^2) + h^1 a_{ct} + (h^1 + h^2)x_{ct}^2,
\end{aligned}$$

where the last step follows because $E_{ct}[u_{jt}] = 0$. Combining with (10), the Canadian output becomes:

$$y_{ct} = \alpha_0 a_c + \alpha_1 [h^a(x_{ct}^1 - x_{ct}^2) + h^1 a_{ct} + (h^1 + h^2)x_{ct}^2].$$

Therefore, for the policy rule (C.1) to hold, it must be that

$$h_c^a = \alpha_0 + \alpha_1 h^1 \tag{C.6}$$

$$h_c^1 = \alpha_1 h^a \tag{C.7}$$

$$h_c^2 = \alpha_1 (h^1 + h^2 - h^a). \tag{C.8}$$

The equations (C.4)-(C.8) are 6 linearly independent equations in 6 unknowns. Thus, they can be solved for unique $h_c^a, h_c^1, h_c^2, h^a, h^1$ and h^2 . In particular, under the assumption that $0 < \alpha_1 < 1$, the following expressions characterize h^1 and h^2 :

$$h^2 = -\frac{\alpha_0 \alpha_1 \tau_u}{\tau_\xi + \tau_a + (1 - \alpha_1^2) \tau_u} < 0 \tag{C.9}$$

$$h^1 = \frac{\alpha_0 \alpha_1 \tau_\xi - [\tau_\xi + (1 - \alpha_1^2) (\tau_a + \tau_u)] h^2}{(1 - \alpha_1^2) (\tau_\xi + \tau_a + \tau_u)} > 0. \tag{C.10}$$

It is straightforward to add (C.9) and (C.10) to establish that $h^1 + h^2 > 0$. In addition (C.7) and (C.8) imply that:

$$h_c^1 + h_c^2 = \alpha_1 (h^1 + h^2) > 0.$$

This establishes the first claim. Plugging x_{ct}^1 and x_{ct}^2 into the Canadian policy rule (C.1) yields:

$$\begin{aligned}y_{ct} &= h_c^a a_c + h_c^1(a_{jt} + \xi_t) + h_c^2 \xi_t \\ &= h_c^a a_c + h_c^1 a_{jt} + (h_c^1 + h_c^2) \xi_t,\end{aligned}$$

and therefore

$$\frac{dy_{ct}}{d\xi_t} = h_c^1 + h_c^2 > 0.$$

□

References

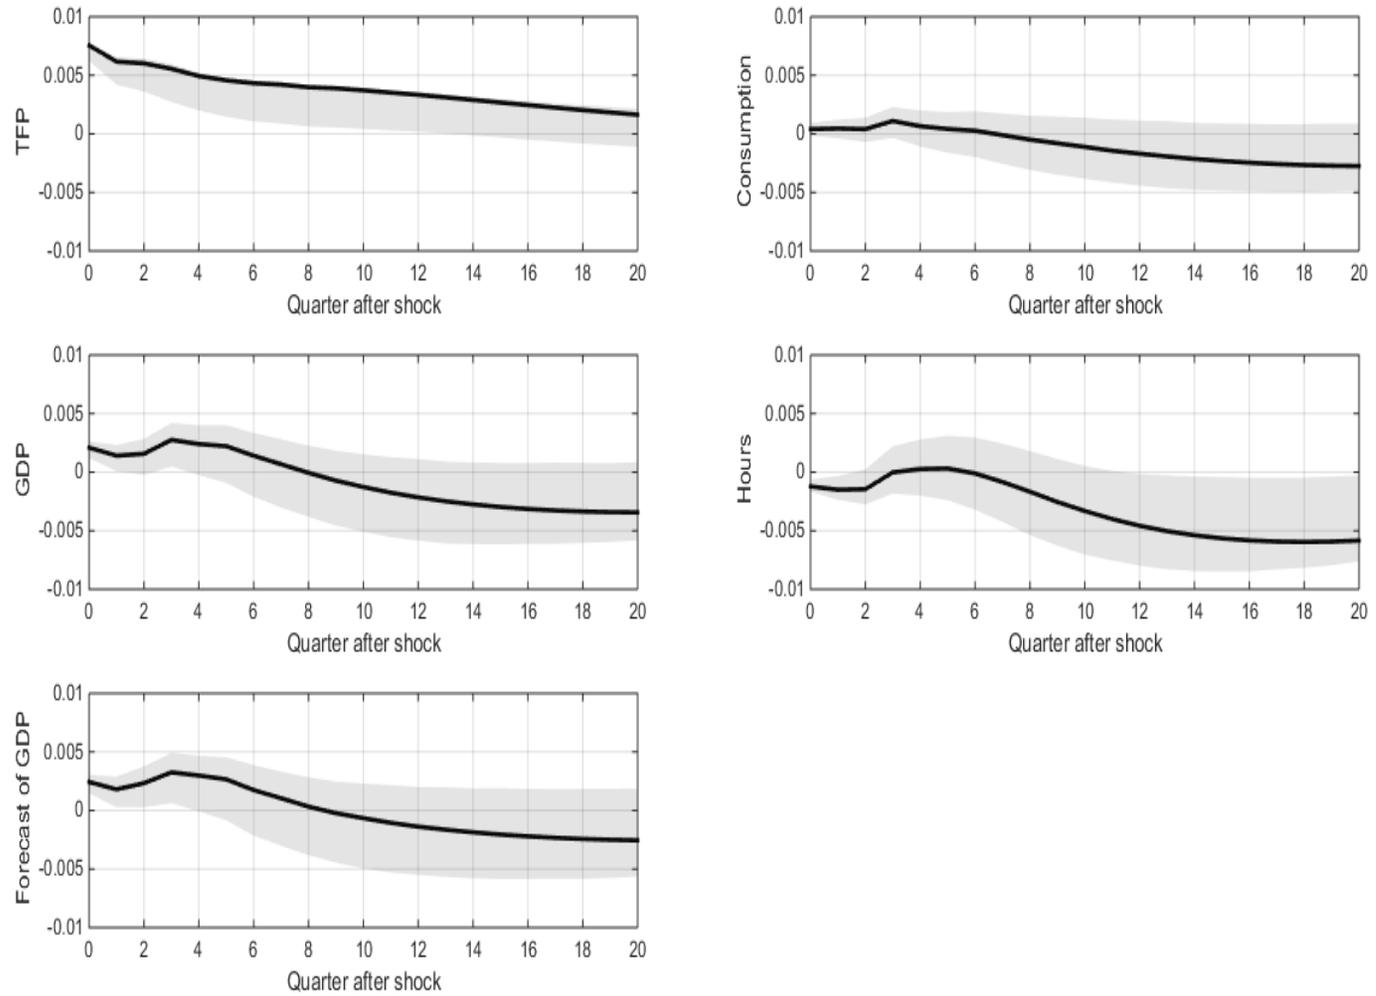
- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2014. “Quantifying Confidence.” NBER Working Paper 20807.
- Angeletos, George-Marios and Jennifer La’O. 2013. “Sentiments.” *Econometrica* 81 (2):739–779.
- Backus, David K, Patrick J Kehoe, and Finn E Kydland. 1992. “International Real Business Cycles.” *Journal of Political Economy* 100 (4):745–75.
- Bai, Yan, Jose-Victor Rios-Rull, and Kjetil Storesletten. 2017. “Demand Shocks as Technology Shocks.” Mimeo, University of Rochester, University of Pennsylvania, and University of Oslo.
- Barsky, Robert B., Susanto Basu, and Keyoung Lee. 2014. “Whither News Shocks?” *NBER Macroeconomics Annual* 29 (1):225–264.
- Barsky, Robert B. and Eric R. Sims. 2011. “News shocks and business cycles.” *Journal of Monetary Economics* 58 (3):273–289.
- . 2012. “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence.” *American Economic Review* 102 (4):1343–77.
- Basu, Susanto, John G. Fernald, and Miles S. Kimball. 2006. “Are Technology Improvements Contractionary?” *American Economic Review* 96 (5):1418–1448.
- Beaudry, Paul, Martial Dupaigne, and Franck Portier. 2011. “Modeling news-driven international business cycles.” *Review of Economic Dynamics* 14 (1):72–91.
- Beaudry, Paul and Franck Portier. 2006. “Stock Prices, News, and Economic Fluctuations.” *American Economic Review* 96 (4):1293–1307.
- . 2014. “News-Driven Business Cycles: Insights and Challenges.” *Journal of Economic Literature* 52 (4):993–1074.
- Benhabib, Jess, Pengfei Wang, and Yi Wen. 2015. “Sentiments and Aggregate Demand Fluctuations.” *Econometrica* 83:549–585.
- Blanchard, Olivier Jean, Jean-Paul L’Huillier, and Guido Lorenzoni. 2013. “News, Noise, and Fluctuations: An Empirical Exploration.” *American Economic Review* 103 (7):3045–70.
- Blanchard, Olivier Jean and Roberto Perotti. 2002. “An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output.” *Quarterly Journal of Economics* 117 (4):1329–1368.
- Blanchard, Olivier Jean and Danny Quah. 1989. “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review* 79 (4):655–73.
- Bloom, Nicholas. 2009. “The Impact of Uncertainty Shocks.” *Econometrica* 77 (3):623–685.
- Canova, Fabio. 2005. “The transmission of US shocks to Latin America.” *Journal of Applied Econometrics* 20 (2):229–251.

- Canova, Fabio and Gianni de Nicoló. 2003. “On the sources of business cycles in the G-7.” *Journal of International Economics* 59 (1):77–100.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 1999. “Monetary policy shocks: What have we learned and to what end?” In *Handbook of Macroeconomics*, vol. 1, edited by John Taylor and Michael Woodford, chap. 2. Elsevier, 65–148.
- Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc. 2014. “The International Dimension of Productivity and Demand Shocks in the US Economy.” *Journal of the European Economic Association* 12 (1):153–176.
- Cushman, David O. and Tao Zha. 1997. “Identifying monetary policy in a small open economy under flexible exchange rates.” *Journal of Monetary Economics* 39 (3):433–448.
- Eickmeier, Sandra and Tim Ng. 2015. “How do credit supply shocks propagate internationally? A GVAR approach.” *European Economic Review* 74:128–145.
- Engel, Charles and Jian Wang. 2011. “International Trade in Durable Goods: Understanding Volatility, Cyclicalities, and Elasticities.” *Journal of International Economics* 83 (1):37–52.
- Faust, Jon and Eric M. Leeper. 1997. “When Do Long-Run Identifying Restrictions Give Reliable Results?” *Journal of Business and Economic Statistics* 15 (3):345–53.
- Fernald, John. 2014. “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity.” Federal Reserve Bank of San Francisco Working Paper 2012-19.
- Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson. 2007. “ABCs (and Ds) of Understanding VARs.” *American Economic Review* 97 (3):1021–1026.
- Fisher, Jonas D. M. 2006. “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks.” *Journal of Political Economy* 114 (3):413–451.
- Forni, Mario and Luca Gambetti. 2014. “Sufficient information in structural VARs.” *Journal of Monetary Economics* 66:124–136.
- Forni, Mario, Luca Gambetti, and Luca Sala. 2014. “No News in Business Cycles.” *The Economic Journal* 124 (581):1168–1191.
- Francis, Neville, Michael T. Owyang, Jennifer E. Roush, and Riccardo DiCecio. 2014. “A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks.” *The Review of Economics and Statistics* 96 (3):638–647.
- Galí, Jordi. 1999. “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review* 89 (1):249–271.
- Galí, Jordi and Pau Rabanal. 2005. “Technology Shocks and Aggregate Fluctuations: How Well Does the RBS Model Fit Postwar U.S. Data?” *NBER Macroeconomics Annual* 19:225–288.
- Hamilton, James D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- Helbling, Thomas, Raju Huidrom, M. Ayhan Kose, and Christopher Otrok. 2011. “Do credit shocks matter? A global perspective.” *European Economic Review* 55 (3):340–353.

- Huo, Zhen and Naoki Takayama. 2015. “Higher Order Beliefs, Confidence, and Business Cycles.” Mimeo, University of Minnesota.
- Ilut, Cosmin and Hikaru Saijo. 2016. “Learning, Confidence, and Business Cycles.” Mimeo, Duke University and UC Santa Cruz.
- Ilizetzi, Ethan and Keyu Jin. 2013. “The Puzzling Change in the International Transmission of U.S. Macroeconomic Policy Shocks.” Mimeo, LSE.
- Imbs, Jean M. 1999. “Technology, growth and the business cycle.” *Journal of Monetary Economics* 44 (1):65 – 80.
- Jordà, Òscar. 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review* 95 (1):161–182.
- Justiniano, Alejandro and Bruce Preston. 2010. “Can structural small open-economy models account for the influence of foreign disturbances?” *Journal of International Economics* 81 (1):61–74.
- Kilian, Lutz. 1998. “Small-Sample Confidence Intervals For Impulse Response Functions.” *The Review of Economics and Statistics* 80 (2):218–230.
- Kim, Soyoung. 2001. “International transmission of U.S. monetary policy shocks: Evidence from VAR’s.” *Journal of Monetary Economics* 48 (2):339–372.
- King, Robert G. and Sergio T. Rebelo. 1999. “Resuscitating real business cycles.” In *Handbook of Macroeconomics*, vol. 1, edited by J. B. Taylor and M. Woodford, chap. 14. Elsevier, 927–1007.
- Kollmann, Robert. 2013. “Global Banks, Financial Shocks, and International Business Cycles: Evidence from an Estimated Model.” *Journal of Money, Credit and Banking* 45 (s2):159–195.
- Kosaka, Michiru Sakane. 2013. “News-driven international business cycles.” *The BE Journal of Macroeconomics* 13 (1):751–793.
- Kurmann, André and Elmar Mertens. 2014. “Stock Prices, News, and Economic Fluctuations: Comment.” *American Economic Review* 104 (4):1439–1445.
- Kurmann, André and Eric Sims. 2017. “Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks.” Mimeo, Drexel University and University of Notre Dame.
- Levchenko, Andrei A. and Nitya Pandalai-Nayar. 2017. “Technology and Non-Technology Shocks: Measurement and Implications for International Comovement.” Mimeo, University of Michigan and University of Texas.
- Lorenzoni, Guido. 2009. “A Theory of Demand Shocks.” *American Economic Review* 99 (5):2050–84.
- Lütkepohl, Helmut. 2005. *New Introduction to Multiple Time Series Analysis*. Berlin Heidelberg: Springer-Verlag.
- Maćkowiak, Bartosz. 2007. “External shocks, U.S. monetary policy and macroeconomic fluctuations in emerging markets.” *Journal of Monetary Economics* 54 (8):2512–2520.

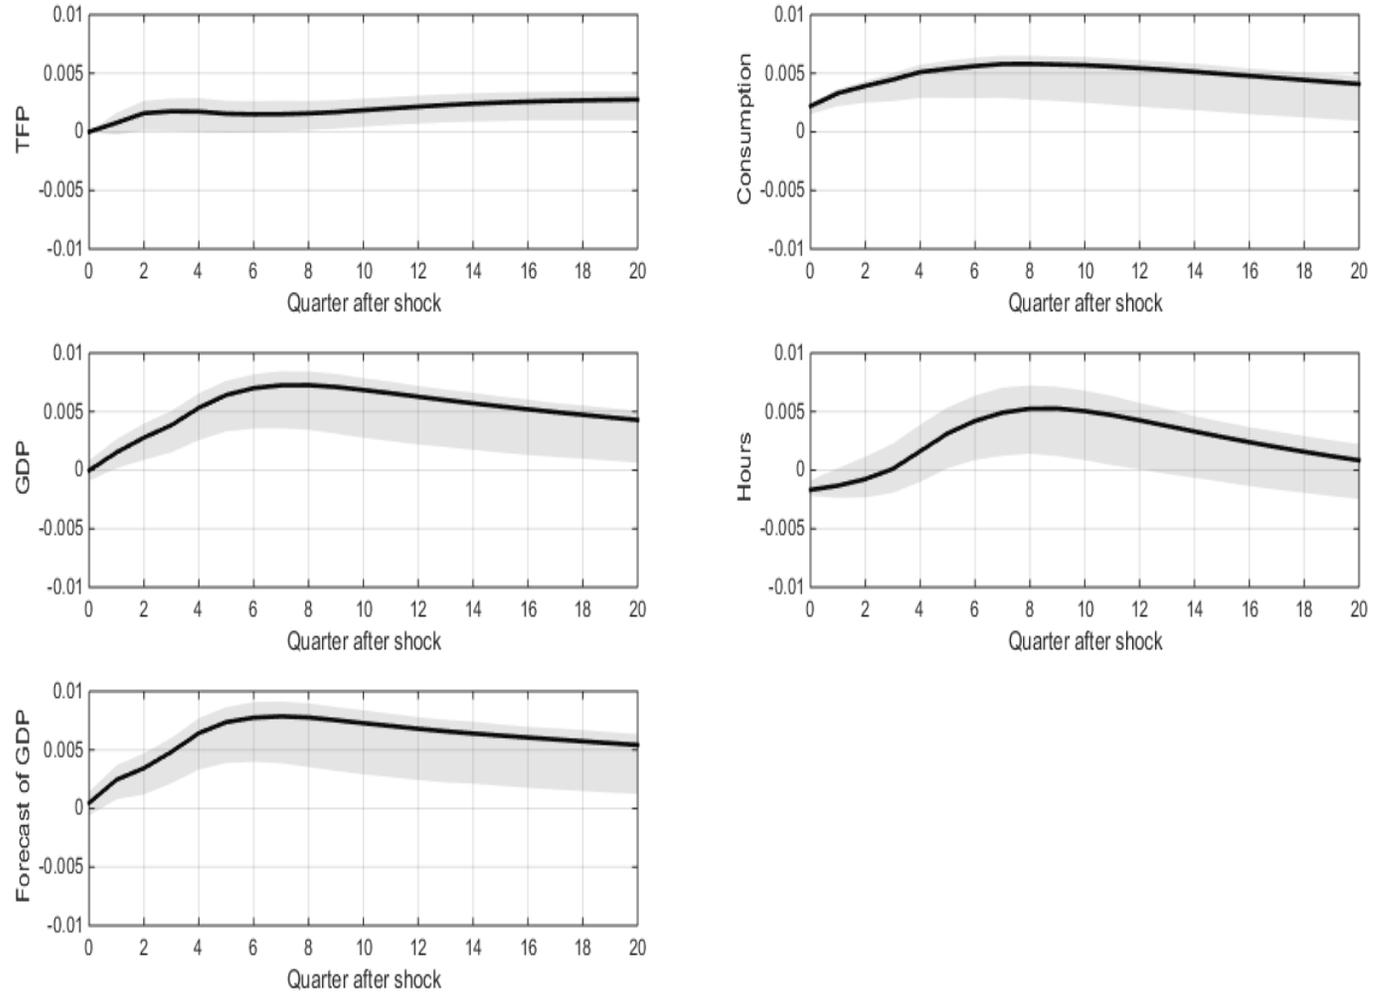
- Milani, Fabio. 2014. "Sentiment and the U.S. Business Cycle." Mimeo, UC Irvine.
- Miyamoto, Wataru and Thuy Lan Nguyen. 2014. "Understanding the Cross Country Effects of U.S. Technology Shocks." Mimeo, Bank of Canada and Santa Clara University.
- Nam, Deokwoo and Jian Wang. 2015. "The Effects of Surprise and Anticipated Technology Changes on International Relative Prices and Trade." *Journal of International Economics* 97 (1):162–177.
- . 2016. "Mood Swings and Business Cycles: Evidence from Sign Restrictions." Mimeo, Hanyang University and CUHK-Shenzhen.
- Ohanian, Lee E. and Andrea Raffo. 2012. "Aggregate hours worked in OECD countries: New measurement and implications for business cycles." *Journal of Monetary Economics* 59 (1):40–56.
- Schmitt-Grohé, Stephanie. 1998. "The international transmission of economic fluctuations: Effects of U.S. business cycles on the Canadian economy." *Journal of International Economics* 44 (2):257–287.
- Shapiro, Matthew and Mark Watson. 1988. "Sources of Business Cycles Fluctuations." *NBER Macroeconomics Annual* 3:111–156.
- Sims, Eric R. 2012. "News, Non-Invertibility, and Structural VARs." *Advances in Econometrics* 28:81–136.
- Stockman, Alan C and Linda L Tesar. 1995. "Tastes and Technology in a Two-Country Model of the Business Cycle: Explaining International Comovements." *American Economic Review* 85 (1):168–85.
- Uhlig, Harald. 2003. "What Drives GNP?" Mimeo, Euro Area Business Cycle Network.
- . 2004. "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association* 2 (2-3):361–371.
- Wen, Yi. 2007. "By force of demand: Explaining international comovements." *Journal of Economic Dynamics and Control* 31 (1):1–23.
- Whelan, Karl. 2000. "A Guide to the Use of Chain Aggregated NIPA Data." Mimeo, Federal Reserve Board of Governors.

Figure 1: The Impulse Responses to the US Surprise TFP Shock



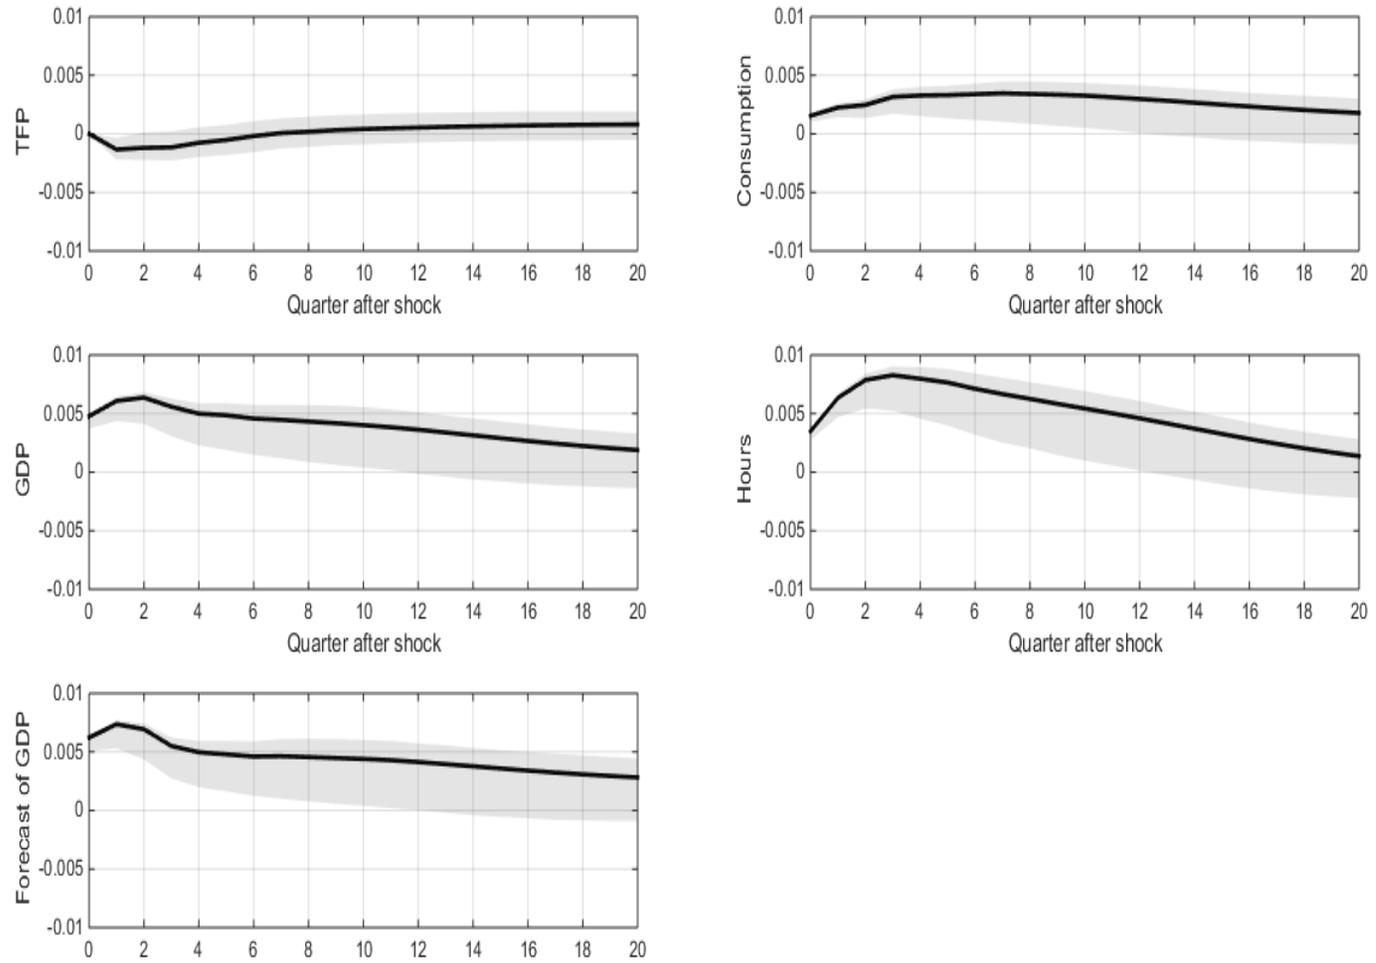
Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the surprise TFP shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 2: The Impulse Responses to the US News TFP Shock



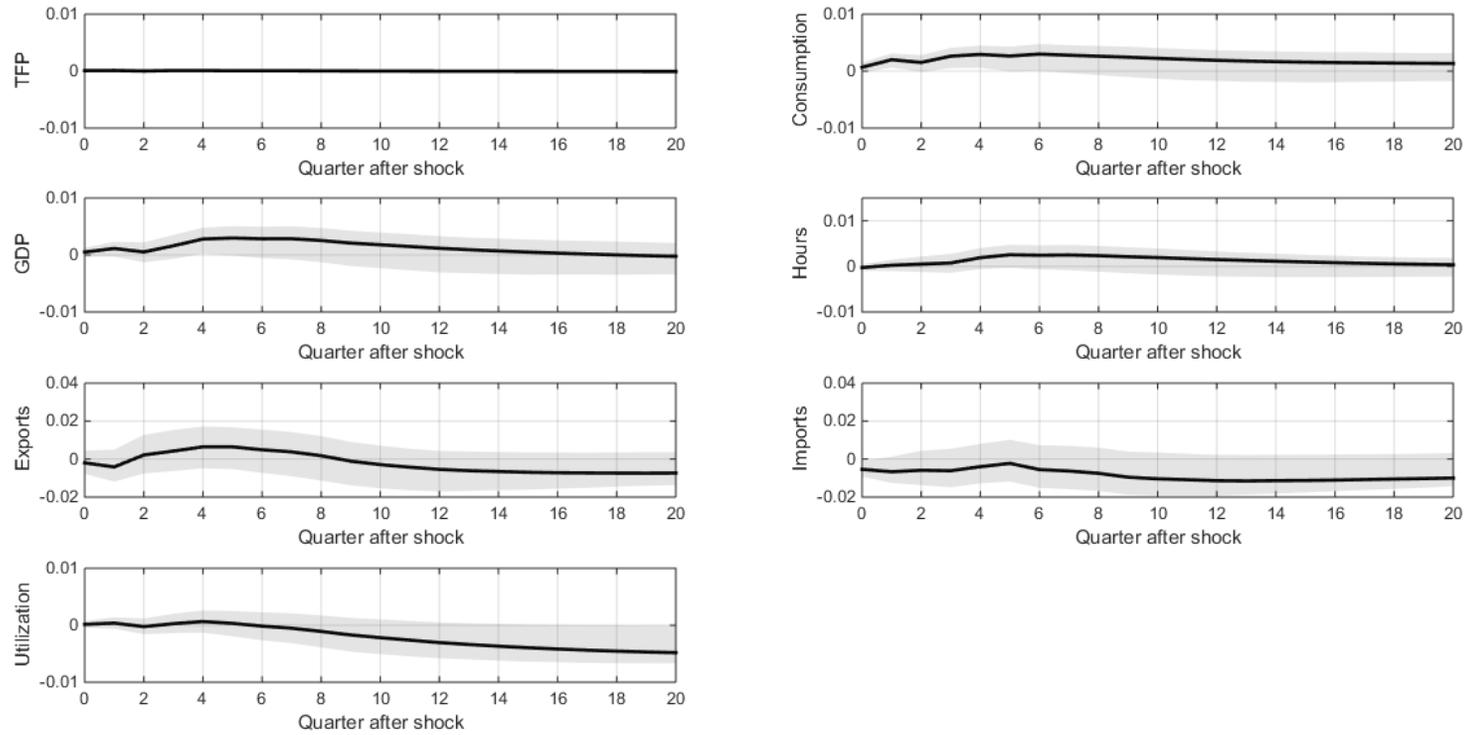
Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the news shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3: The Impulse Responses to the US Non-Technology Expectations Shock



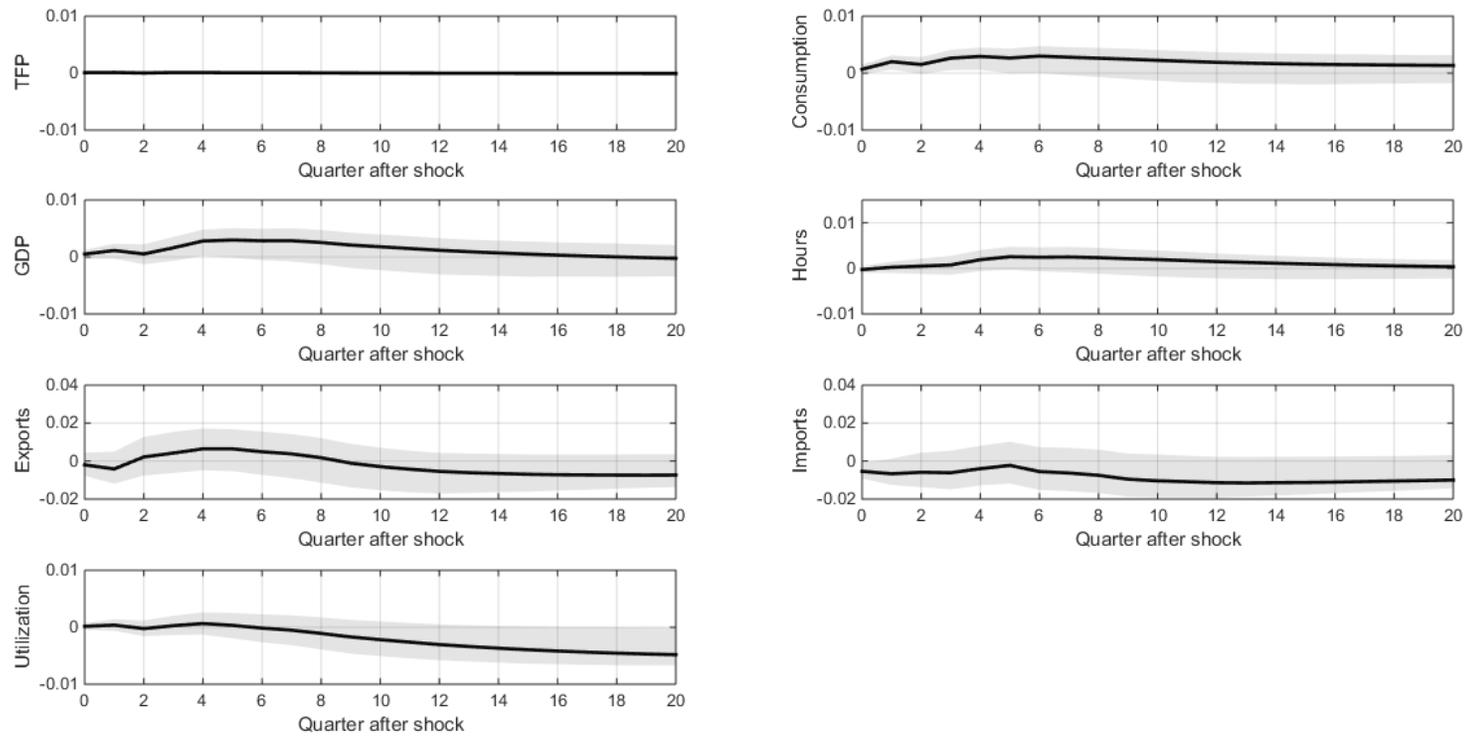
Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the non-technology expectations shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 4: The Impulse Responses of Canadian Variables to the US TFP Shock



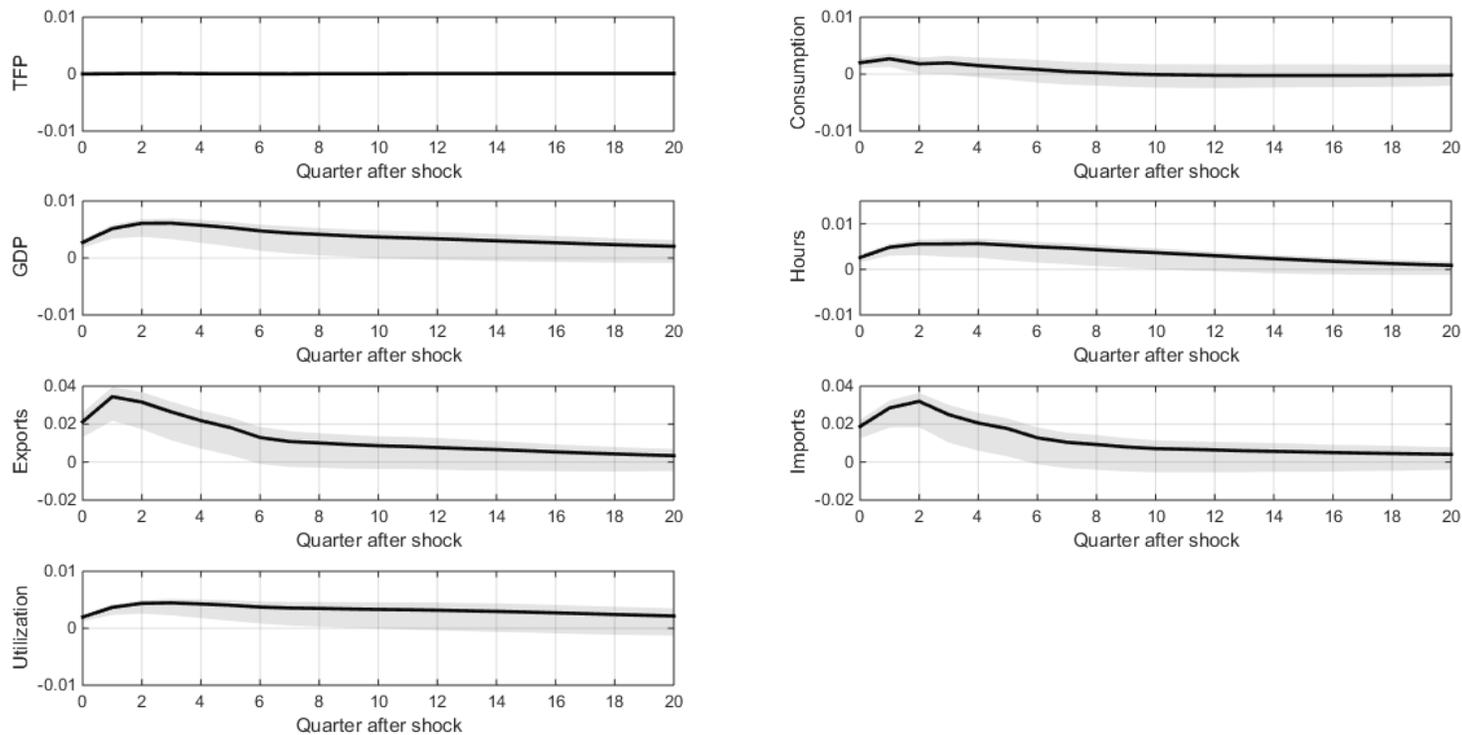
Notes: This figure plots the impulse responses of Canadian macro variables to the US surprise TFP shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 5: The Impulse Responses of Canadian Variables to the US News Shock



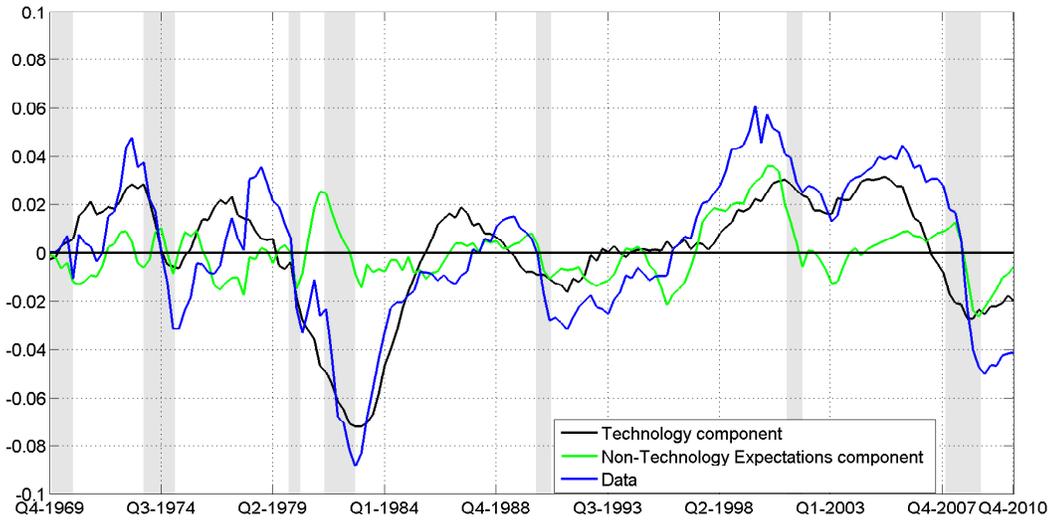
Notes: This figure plots the impulse responses of Canadian macro variables to the US news TFP shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 6: The Impulse Responses of Canadian Variables to the US Non-Technology Expectations Shock

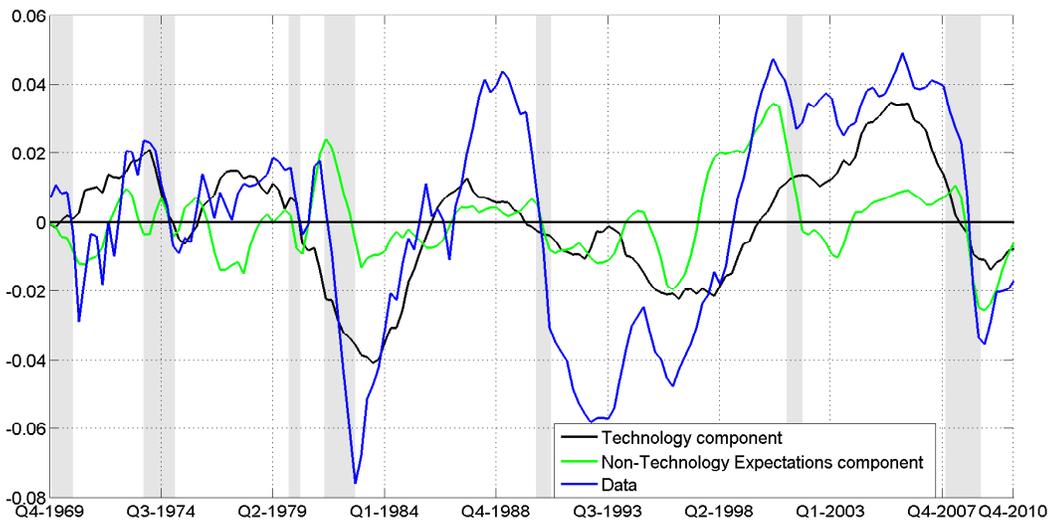


Notes: This figure plots the impulse responses of Canadian macro variables to the US NTE shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 7: Historical Decompositions: GDP



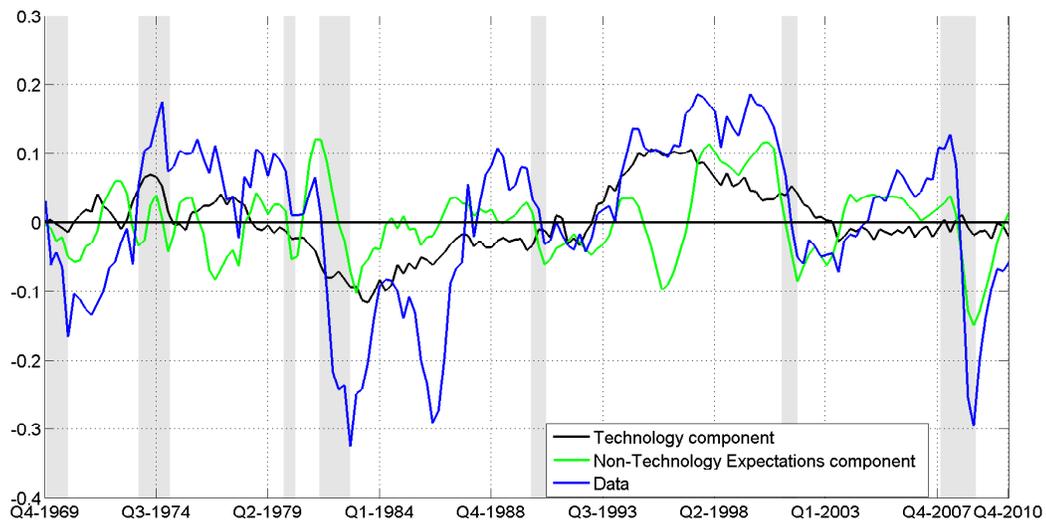
(a) US Output



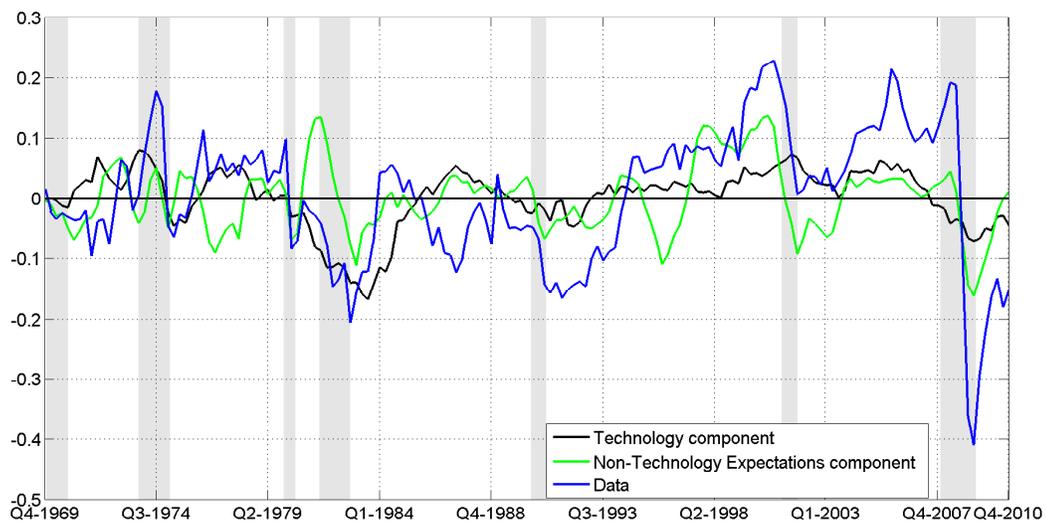
(b) Canadian Output

Notes: These figures show the decomposition of historical data into components due to the three identified shocks. The shaded areas are US NBER recession dates.

Figure 8: Historical Decompositions: Trade



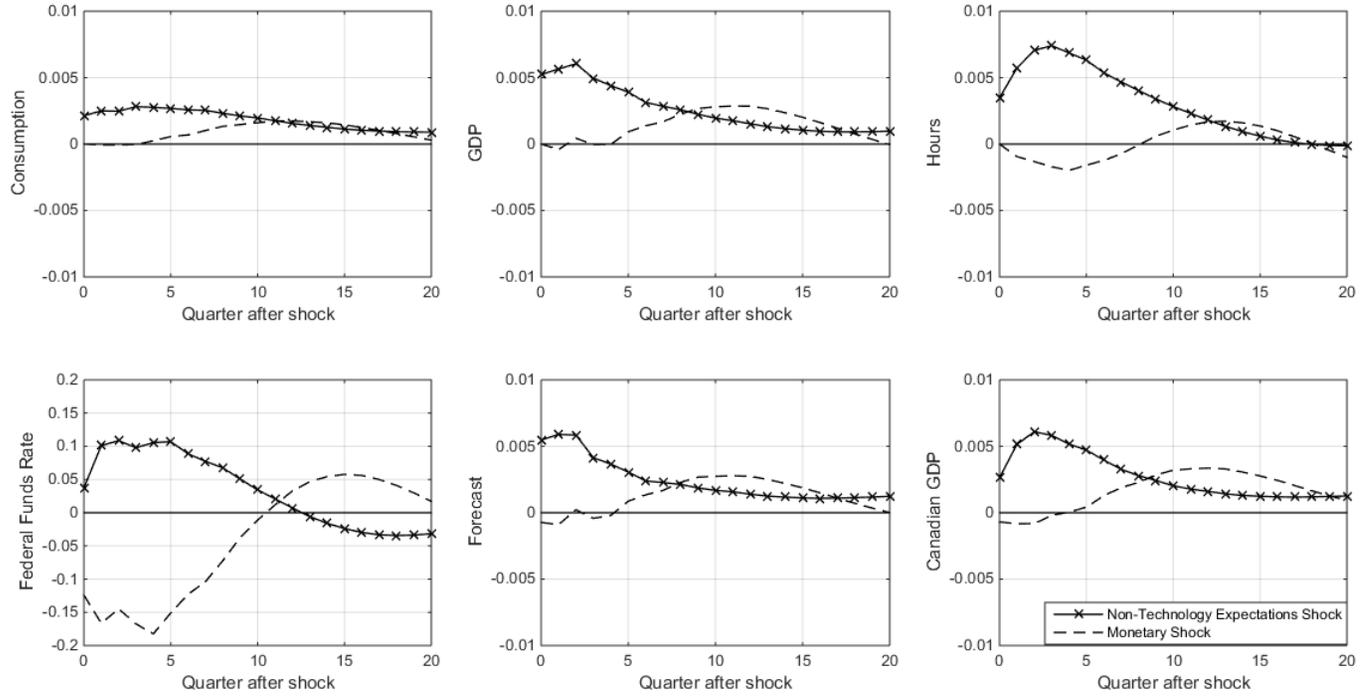
(a) Canadian Imports from the US



(b) Canadian Exports to the US

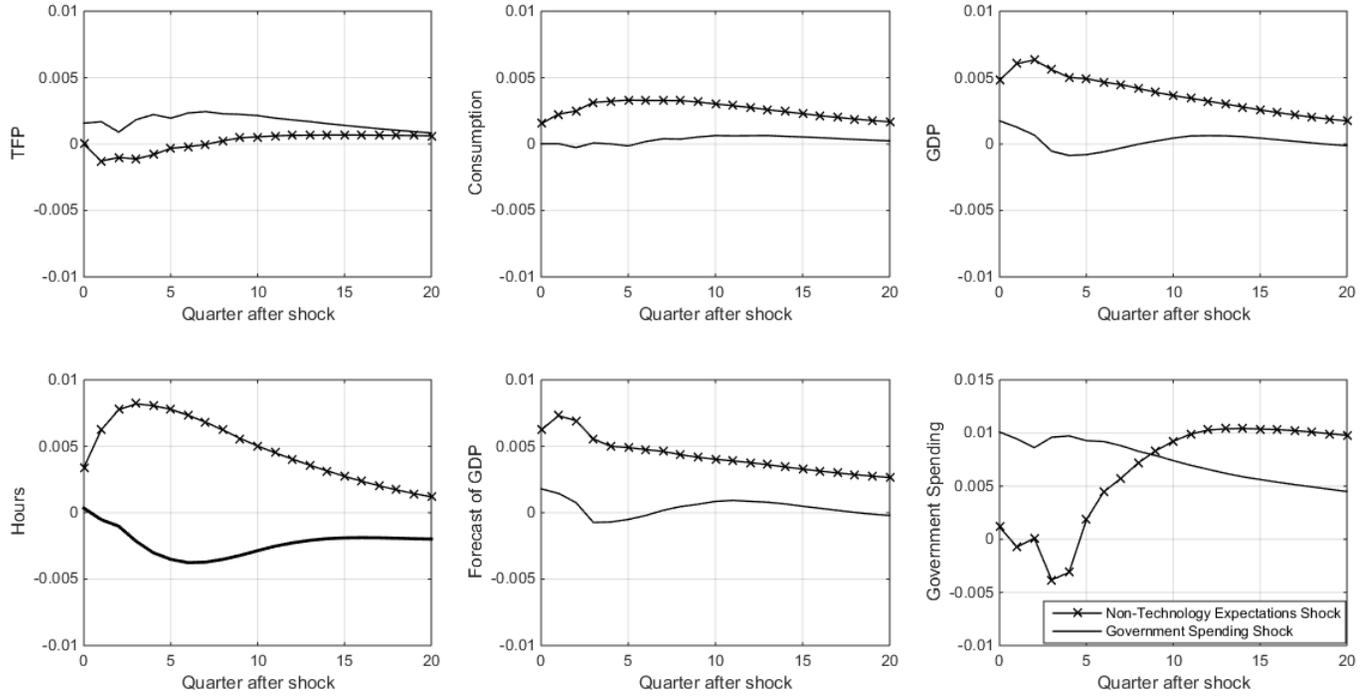
Notes: These figures show the decomposition of historical data into components due to the three identified shocks. The shaded areas are US NBER recession dates

Figure 9: The Impulse Responses of the Core US Variables and Canadian GDP to a Non-Technology Expectations and Monetary Policy Shocks



Notes: This figure plots the impulse responses of US Consumption, GDP, Hours, the Federal Funds rate, the Forecast variable and Canadian GDP to an NTE shock and a monetary policy shock, identified as discussed in Section 3.1.

Figure 10: The Impulse Responses of the Core US Variables and Canadian GDP to a Non-Technology Expectations and Fiscal Policy Shocks



Notes: This figure plots the impulse responses of US Consumption, GDP, Hours, Government Expenditure, the Forecast variable and Canadian GDP to an NTE shock and a fiscal policy shock, identified as discussed in Section 3.1.

Table 1: Surprise TFP Shock: Variance Decomposition

Panel A: US						
Horizon	TFP	GDP	Consumption	Hours	Forecast	
1Q	1.00 (0.00)	0.12 (0.05)	0.01 (0.03)	0.08 (0.04)	0.13 (0.05)	
2Q	0.98 (0.02)	0.08 (0.04)	0.01 (0.03)	0.05 (0.04)	0.08 (0.04)	
1Y	0.93 (0.05)	0.08 (0.05)	0.02 (0.03)	0.02 (0.03)	0.09 (0.06)	
2Y	0.91 (0.07)	0.05 (0.05)	0.01 (0.04)	0.01 (0.03)	0.06 (0.05)	
5Y	0.79 (0.12)	0.08 (0.10)	0.06 (0.11)	0.17 (0.13)	0.05 (0.08)	
10Y	0.59 (0.13)	0.13 (0.12)	0.12 (0.13)	0.28 (0.14)	0.07 (0.10)	

Panel B: Canada							
Horizon	Output	Consumption	Hours	Exports	Imports	TFP	Utilization
1Q	0.01 (0.02)	0.01 (0.02)	0.00 (0.01)	0.00 (0.01)	0.02 (0.02)	0.02 (0.03)	0.00 (0.01)
2Q	0.01 (0.03)	0.05 (0.05)	0.00 (0.01)	0.00 (0.02)	0.02 (0.03)	0.04 (0.04)	0.00 (0.02)
1Y	0.02 (0.03)	0.07 (0.07)	0.00 (0.02)	0.01 (0.02)	0.02 (0.03)	0.04 (0.04)	0.00 (0.02)
2Y	0.06 (0.07)	0.10 (0.10)	0.04 (0.06)	0.01 (0.04)	0.02 (0.05)	0.04 (0.06)	0.00 (0.04)
5Y	0.05 (0.08)	0.10 (0.11)	0.06 (0.09)	0.04 (0.07)	0.11 (0.10)	0.02 (0.07)	0.11 (0.12)
10Y	0.05 (0.09)	0.08 (0.11)	0.06 (0.09)	0.07 (0.08)	0.16 (0.11)	0.05 (0.09)	0.29 (0.16)

Notes: This table shows the contribution of the surprise TFP innovation to the forecast error variance of all variables at different horizons. Standard errors are from 2000 bootstrap repetitions.

Table 2: News Shock: Variance Decomposition

Panel A: US						
Horizon	TFP	GDP	Consumption	Hours	Forecast	
1Q	0.00 (0.00)	0.00 (0.01)	0.36 (0.07)	0.14 (0.06)	0.00 (0.02)	
2Q	0.01 (0.01)	0.03 (0.03)	0.45 (0.08)	0.06 (0.04)	0.05 (0.04)	
1Y	0.04 (0.04)	0.11 (0.07)	0.48 (0.09)	0.02 (0.03)	0.15 (0.07)	
2Y	0.06 (0.06)	0.32 (0.11)	0.52 (0.11)	0.08 (0.06)	0.35 (0.11)	
5Y	0.18 (0.11)	0.45 (0.12)	0.56 (0.13)	0.13 (0.08)	0.46 (0.12)	
10Y	0.36 (0.13)	0.45 (0.13)	0.54 (0.15)	0.12 (0.08)	0.47 (0.15)	

Panel B: Canada							
Horizon	Output	Consumption	Hours	Exports	Imports	TFP	Utilization
1Q	0.00 (0.01)	0.02 (0.03)	0.00 (0.01)	0.02 (0.03)	0.03 (0.03)	0.01 (0.02)	0.01 (0.02)
2Q	0.00 (0.02)	0.01 (0.02)	0.00 (0.02)	0.01 (0.02)	0.02 (0.02)	0.01 (0.02)	0.01 (0.02)
1Y	0.02 (0.04)	0.03 (0.04)	0.00 (0.02)	0.01 (0.02)	0.01 (0.03)	0.07 (0.05)	0.05 (0.05)
2Y	0.11 (0.08)	0.12 (0.08)	0.02 (0.05)	0.05 (0.04)	0.02 (0.03)	0.18 (0.09)	0.16 (0.09)
5Y	0.26 (0.11)	0.32 (0.13)	0.06 (0.07)	0.15 (0.07)	0.06 (0.06)	0.39 (0.12)	0.24 (0.11)
10Y	0.33 (0.13)	0.51 (0.14)	0.06 (0.07)	0.17 (0.08)	0.10 (0.08)	0.47 (0.14)	0.20 (0.11)

Notes: This table shows the contribution of the news shock to the forecast error variance of all variables at different horizons. Standard errors are from 2000 bootstrap repetitions.

Table 3: Non-Technology Expectations Shock: Variance Decomposition

Panel A: US							
Horizon	TFP	Output	Consumption	Hours	Forecast		
1Q	0.00 (0.00)	0.65 (0.06)	0.18 (0.05)	0.62 (0.07)	0.85 (0.05)		
2Q	0.02 (0.02)	0.75 (0.06)	0.21 (0.06)	0.71 (0.07)	0.81 (0.06)		
1Y	0.03 (0.03)	0.61 (0.09)	0.22 (0.07)	0.69 (0.08)	0.62 (0.09)		
2Y	0.02 (0.04)	0.36 (0.10)	0.21 (0.09)	0.52 (0.12)	0.35 (0.10)		
5Y	0.02 (0.05)	0.25 (0.10)	0.18 (0.10)	0.35 (0.13)	0.26 (0.10)		
10Y	0.03 (0.05)	0.21 (0.10)	0.15 (0.10)	0.29 (0.12)	0.22 (0.10)		

Panel B: Canada							
Horizon	Output	Consumption	Hours	Exports	Imports	TFP	Utilization
1Q	0.19 (0.06)	0.08 (0.05)	0.19 (0.06)	0.25 (0.07)	0.25 (0.07)	0.01 (0.01)	0.18 (0.06)
2Q	0.32 (0.08)	0.12 (0.06)	0.29 (0.08)	0.38 (0.08)	0.36 (0.08)	0.01 (0.02)	0.28 (0.07)
1Y	0.41 (0.09)	0.09 (0.06)	0.35 (0.10)	0.44 (0.09)	0.41 (0.09)	0.04 (0.04)	0.33 (0.09)
2Y	0.34 (0.10)	0.05 (0.05)	0.36 (0.11)	0.39 (0.10)	0.36 (0.10)	0.03 (0.04)	0.25 (0.10)
5Y	0.29 (0.10)	0.03 (0.05)	0.37 (0.12)	0.32 (0.10)	0.29 (0.10)	0.06 (0.06)	0.17 (0.10)
10Y	0.26 (0.10)	0.02 (0.05)	0.37 (0.11)	0.30 (0.10)	0.26 (0.10)	0.08 (0.07)	0.12 (0.09)

Notes: This table shows the contribution of the NTE shock to the forecast error variance of all variables at different horizons. Standard errors are from 2000 bootstrap repetitions.

Table 4: Conditional Correlations

	Data	<i>Correlation Conditional on:</i>		
		TFP	News	NTE
US, Canada Output	0.73	0.47 (0.28)	0.99 (0.07)	0.99 (0.05)
US, Canada Consumption	0.51	-0.13 (0.47)	0.94 (0.13)	0.80 (0.23)
US, Canada Hours	0.68	-0.47 (0.56)	0.46 (0.45)	0.98 (0.30)
Exports from Canada, US Output	0.66	0.95 (0.19)	0.97 (0.16)	0.89 (0.09)
Canadian Imports, US Output	0.58	0.79 (0.31)	0.91 (0.28)	0.91 (0.12)
US Output, US Consumption	0.84	0.97 (0.05)	0.99 (0.03)	0.96 (0.04)
US Output, US Hours	0.87	0.81 (0.39)	0.20 (0.44)	0.73 (0.31)

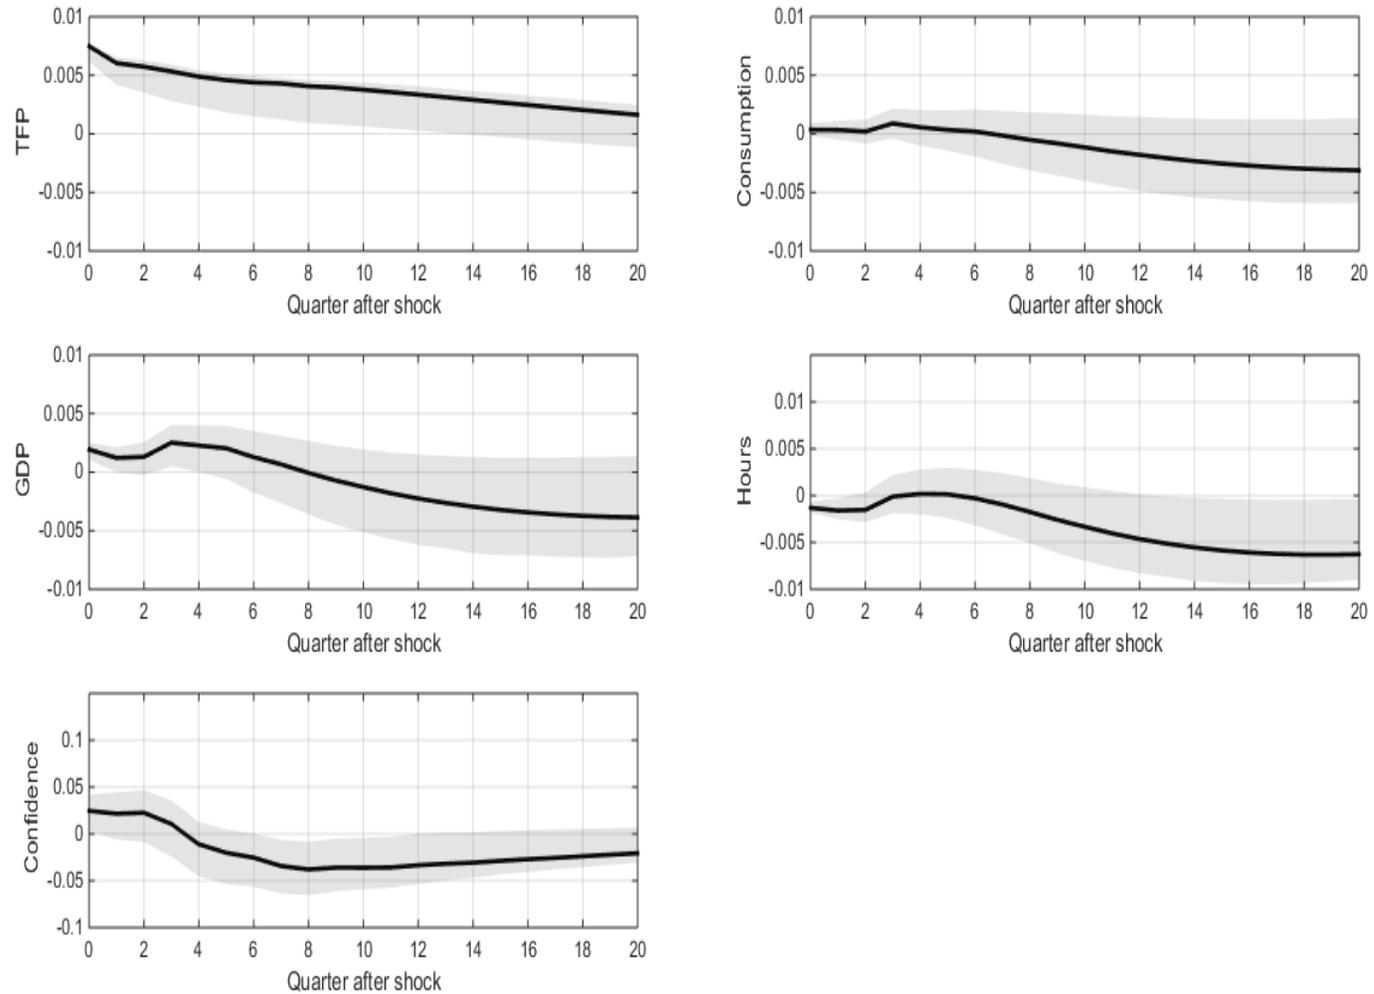
Notes: This table shows conditional correlations of various macroeconomic aggregates in response to the three shocks. Standard errors are from 2000 bootstrap replications. The data column refers to the unconditional correlations from HP-filtered data with a smoothing parameter of 1600.

Table 5: Cross-Correlations of Forecast and Confidence with GDP

Variable	Lags										
	-5	-4	-3	-2	-1	0	1	2	3	4	5
GDP Forecast	-0.07	0.04	0.01	0.18	0.20	0.93	0.38	0.24	0.14	0.11	-0.03
Consumer Confidence	-0.17	-0.12	-0.21	-0.13	-0.15	0.21	0.25	0.18	0.12	0.25	0.03

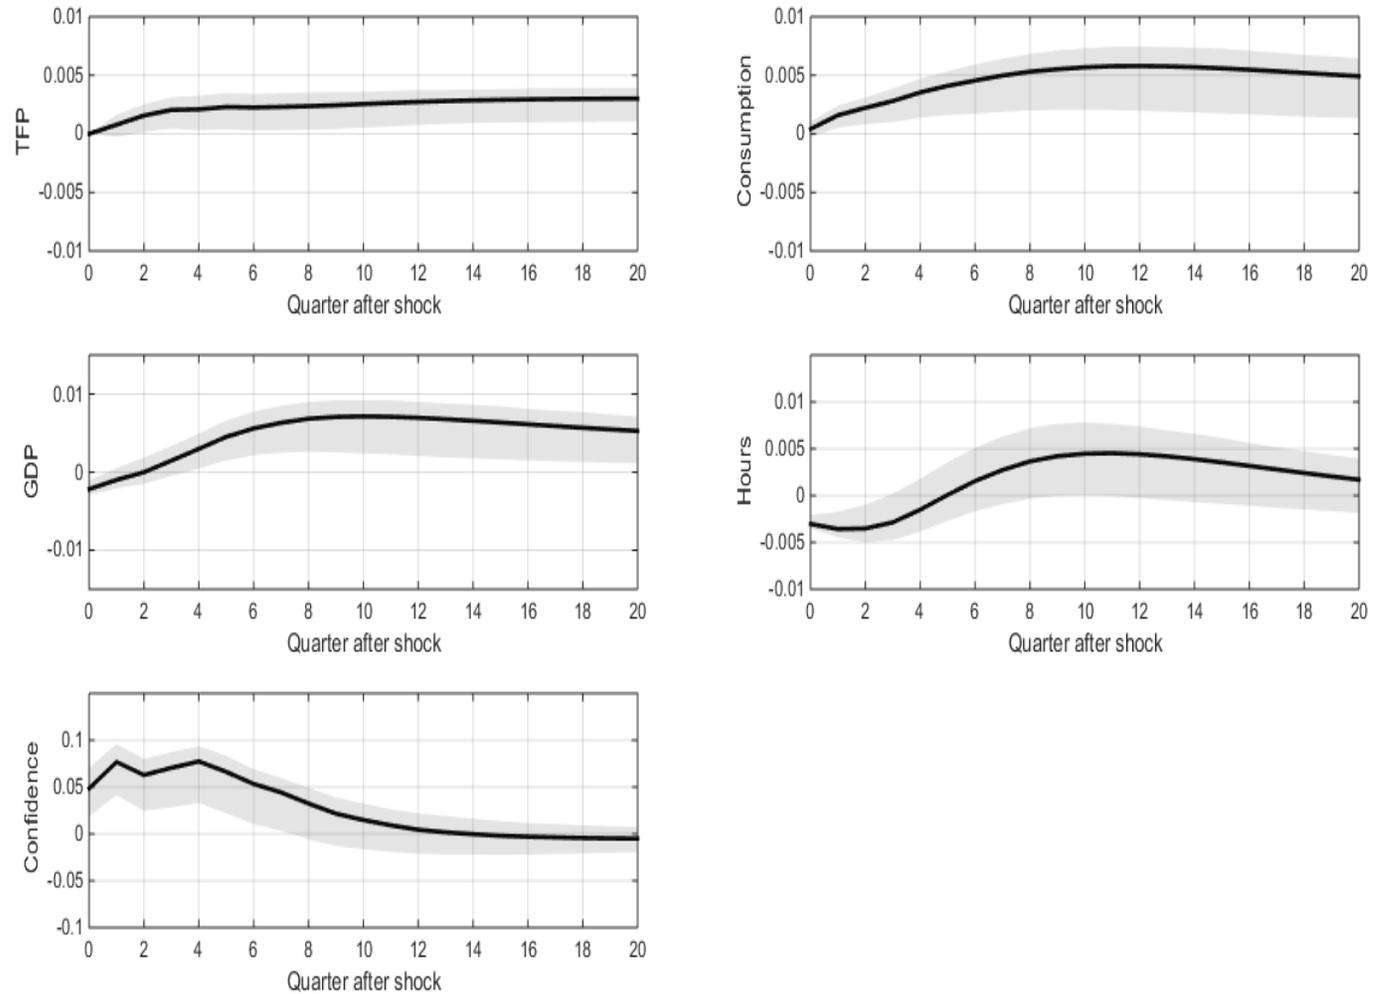
Notes: This table shows the cross-correlation of the GDP forecast and the Consumer Confidence variable with GDP at leads and lags. All variables are in growth rates.

Figure A1: The Impulse Responses to the US Surprise TFP Shock, Using US Consumer Confidence



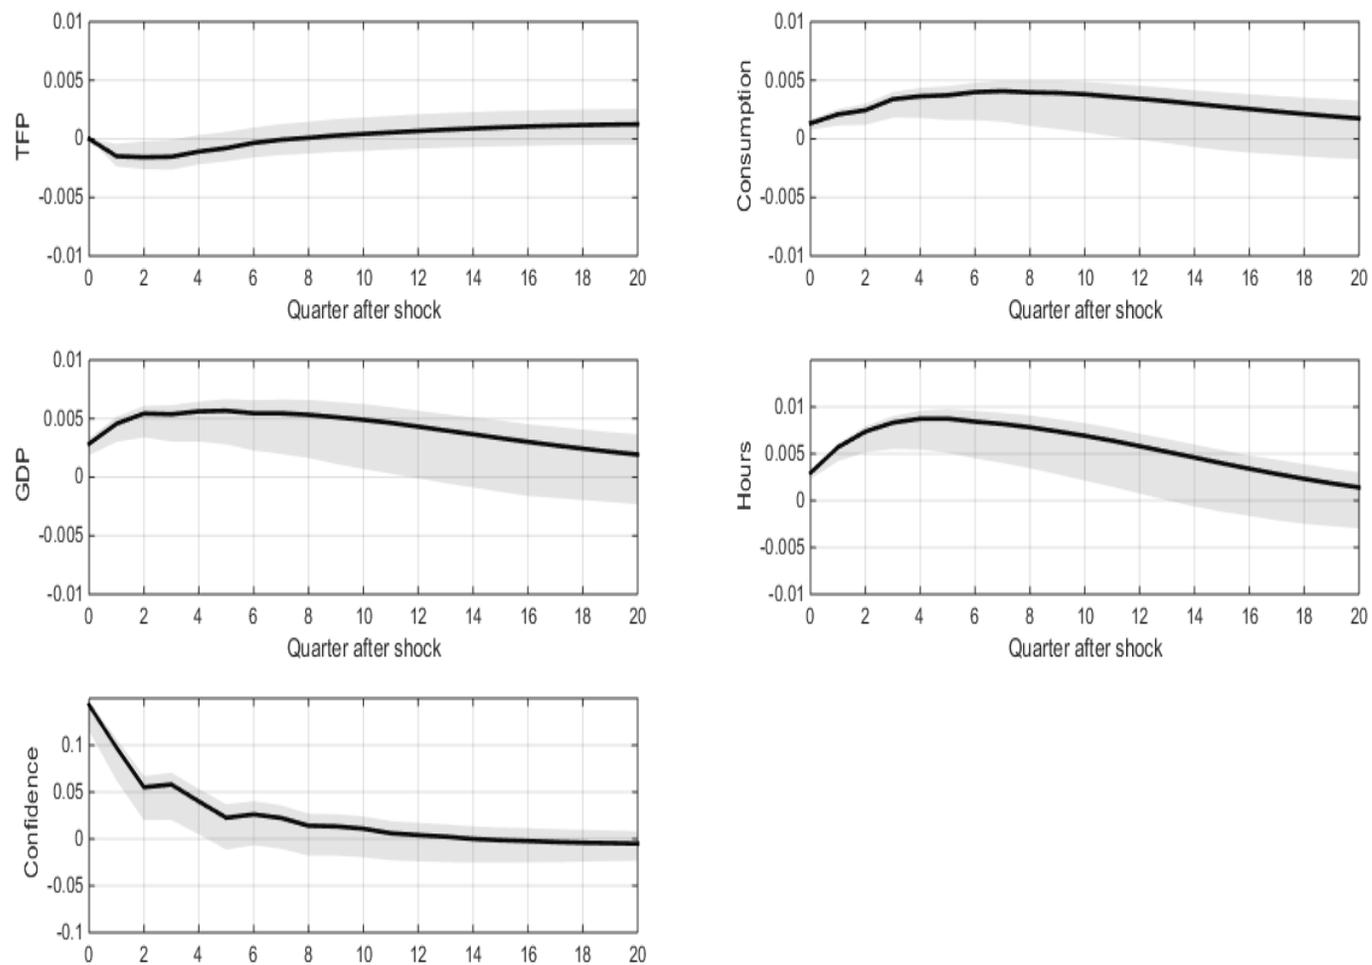
Notes: This figure plots the impulse responses of the core US variables to a surprise TFP innovation, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure A2: The Impulse Responses to the US News TFP Shock, Using US Consumer Confidence



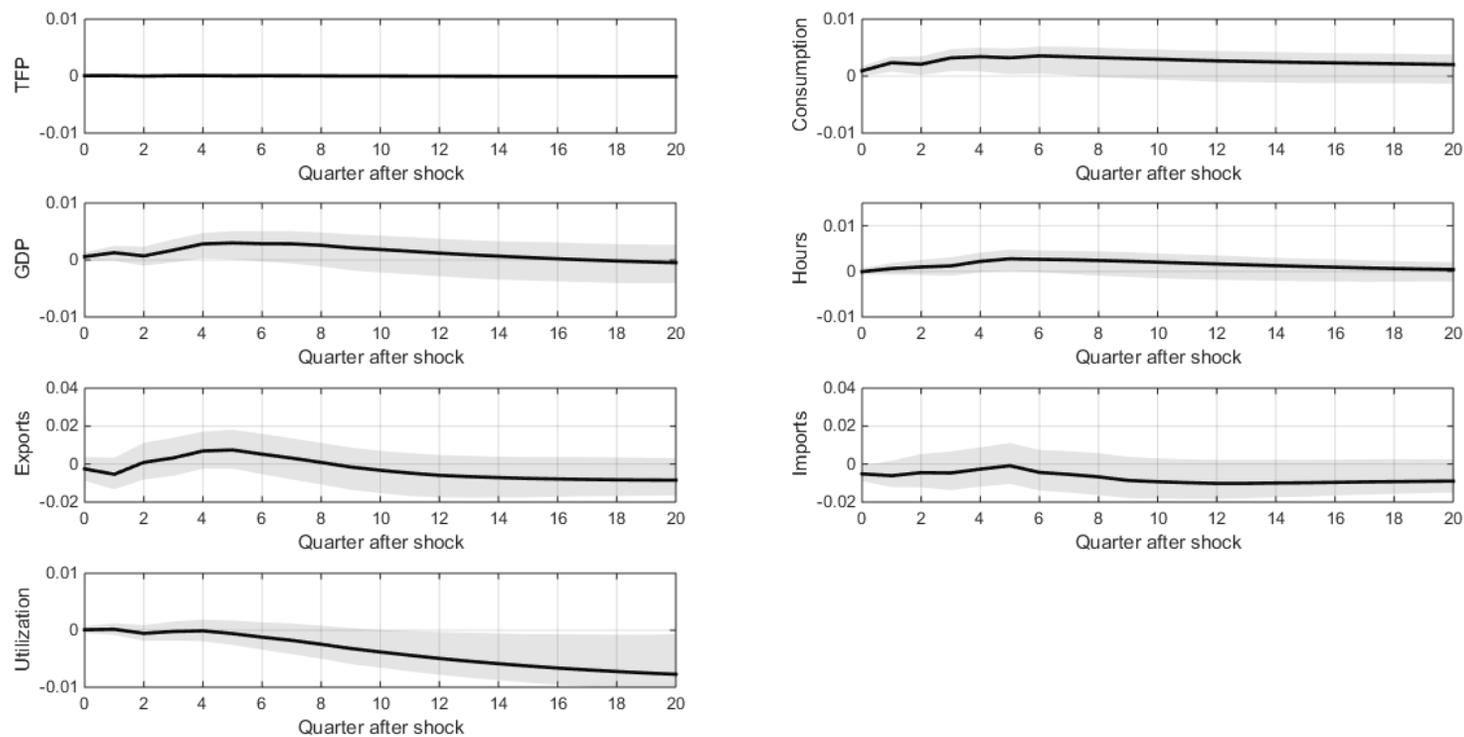
Notes: This figure plots the impulse responses of the core US variables to the news shock, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure A3: The Impulse Responses to the US Non-Technology Expectations Shock, Using US Consumer Confidence



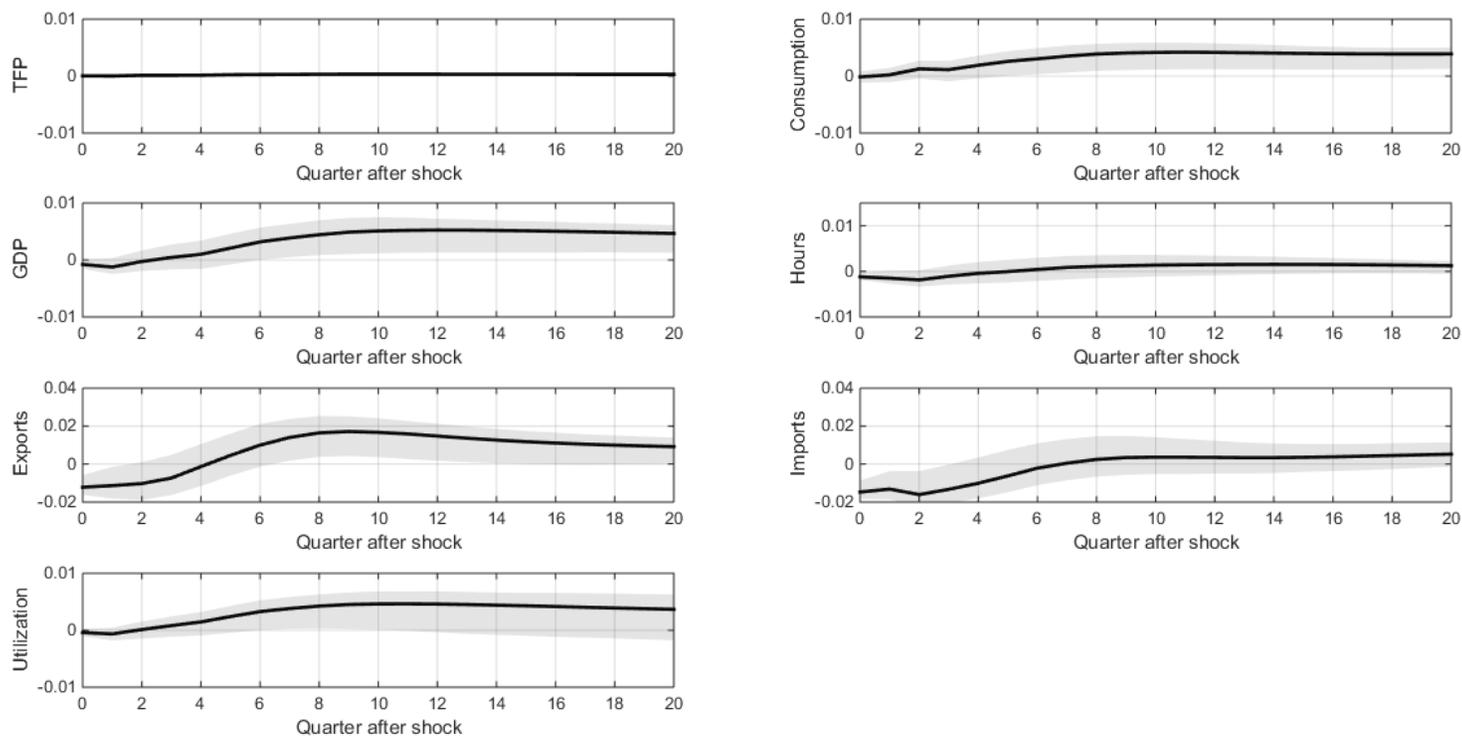
Notes: This figure plots the impulse responses of the core US variables to the NTE shock, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure A4: The Impulse Responses of Canadian Variables to a US TFP Shock, Using US Consumer Confidence



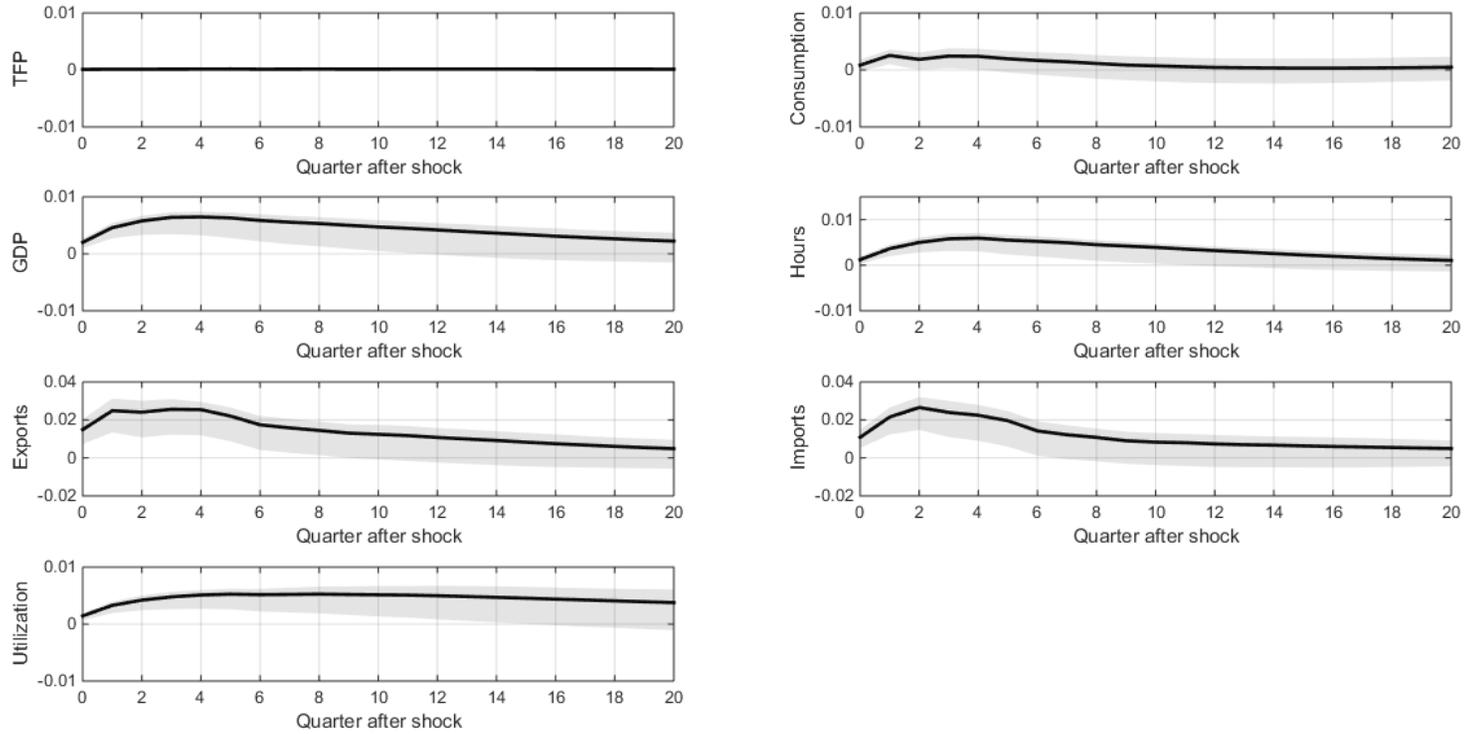
Notes: This figure plots the impulse responses of Canadian macro variables to surprise TFP shock in the US, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure A5: The Impulse Responses of Canadian Variables to a US News Shock, Using US Consumer Confidence



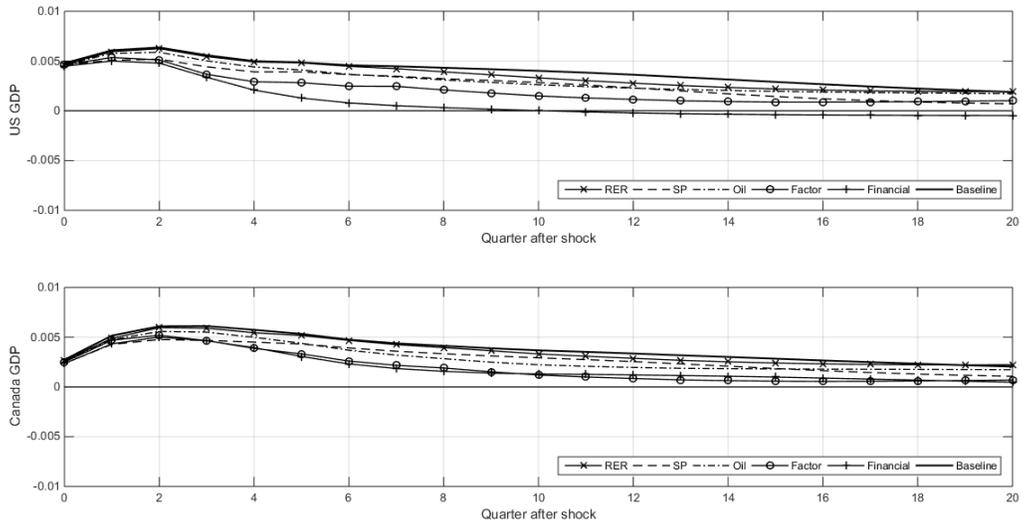
Notes: This figure plots the impulse responses of Canadian macro variables to news of TFP shock in the US, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure A6: The Impulse Responses of Canadian Variables to a US Non-Technology Expectations Shock, Using US Consumer Confidence



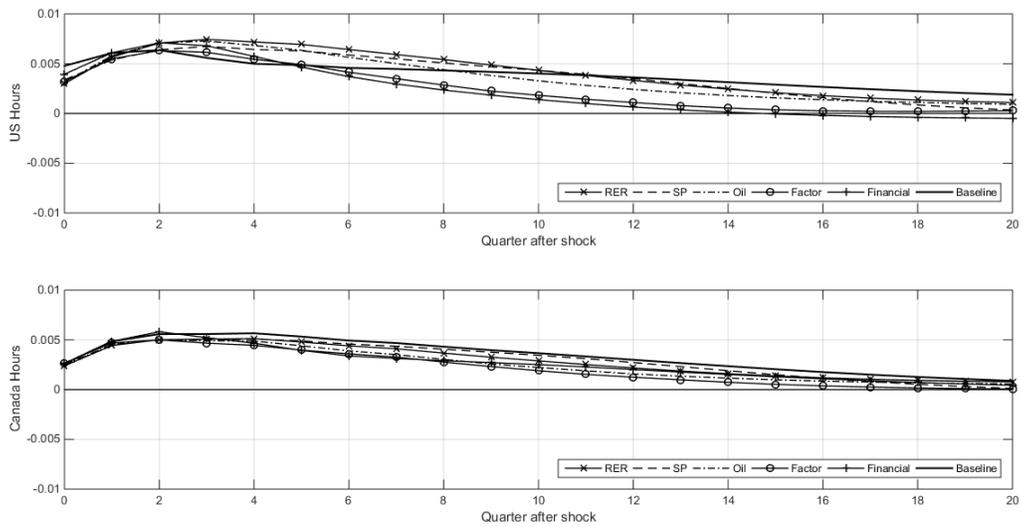
Notes: This figure plots the impulse responses of Canadian macro variables to NTE shock in the US, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure A7: The Impulse Responses US and Canadian GDP to the Non-Technology Expectations Shock in a VAR with Additional Controls



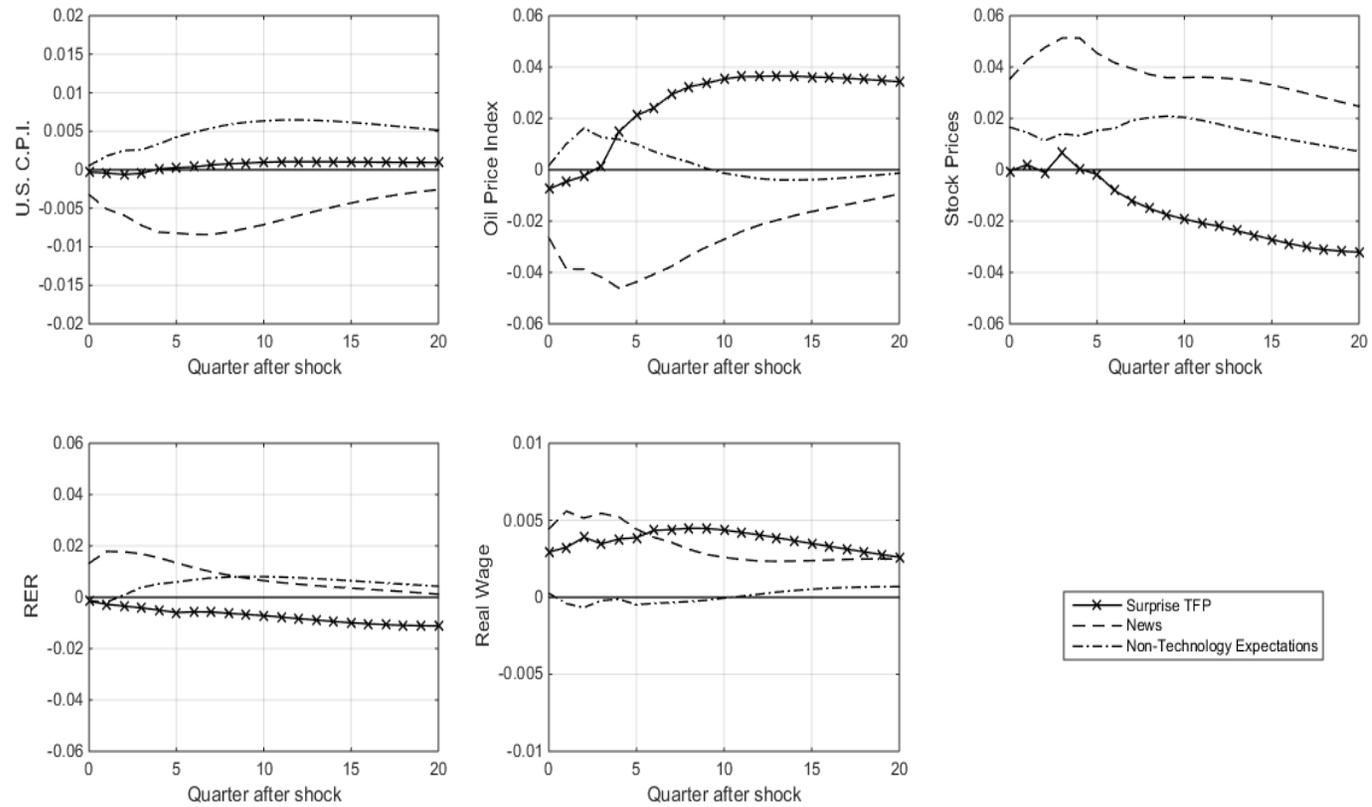
Notes: This figure plots the impulse responses of US and Canadian GDP to the NTE shock, in a VAR with additional controls identified as discussed in 3.1. The additional controls are a measure of stock prices (labeled SP), an oil price index (Oil), the real exchange rate (RER), a US factor (Factor), and the Baa-Aaa credit spread (Financial).

Figure A8: The Impulse Responses of US and Canadian Hours to the Non-Technology Expectations Shock in a VAR with Additional Controls



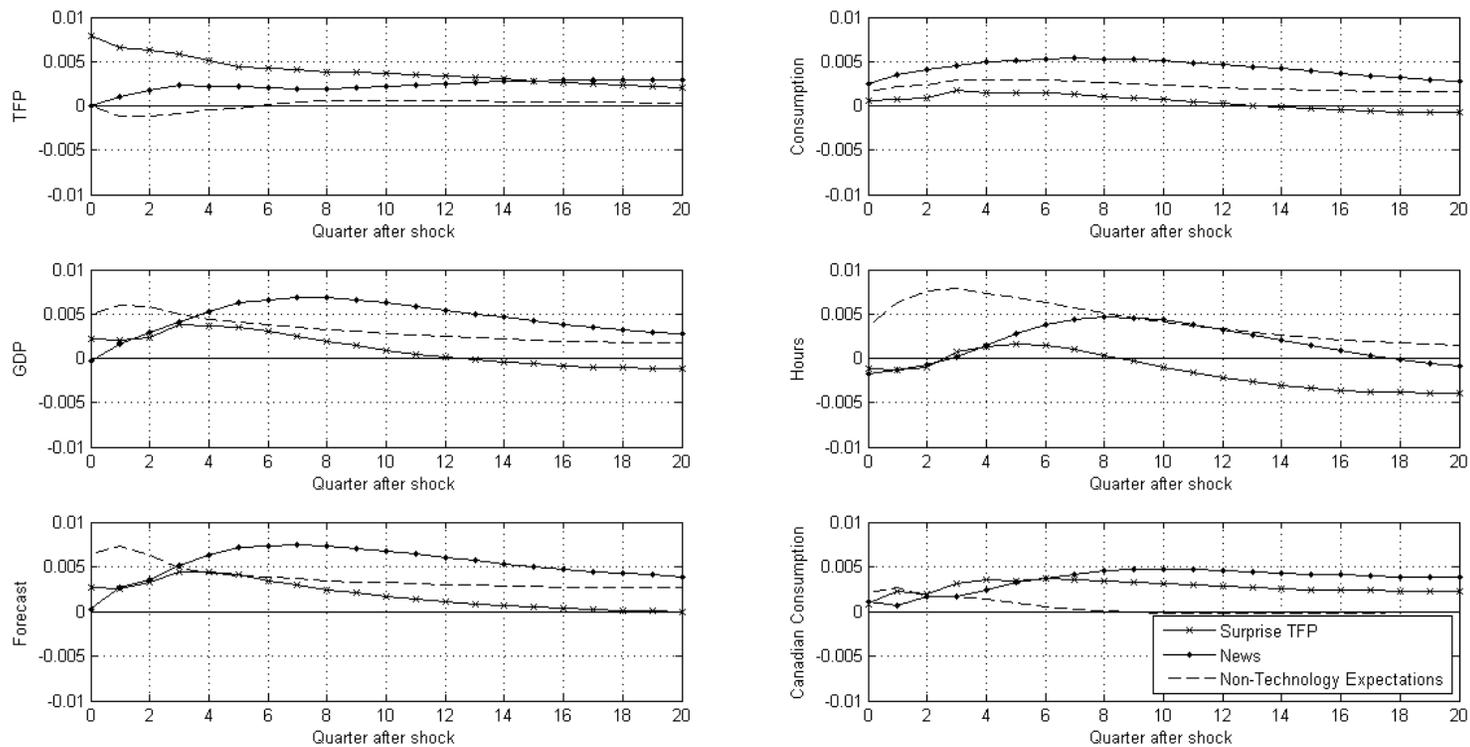
Notes: This figure plots the impulse responses of US and Canadian GDP to the NTE shock, in a VAR with additional controls identified as discussed in 3.1. The additional controls are a measure of stock prices (labeled SP), an oil price index (Oil), the real exchange rate (RER), a US factor (Factor), and the Baa-Aaa credit spread (Financial).

Figure A9: The Responses of Price Variables to the Three Shocks



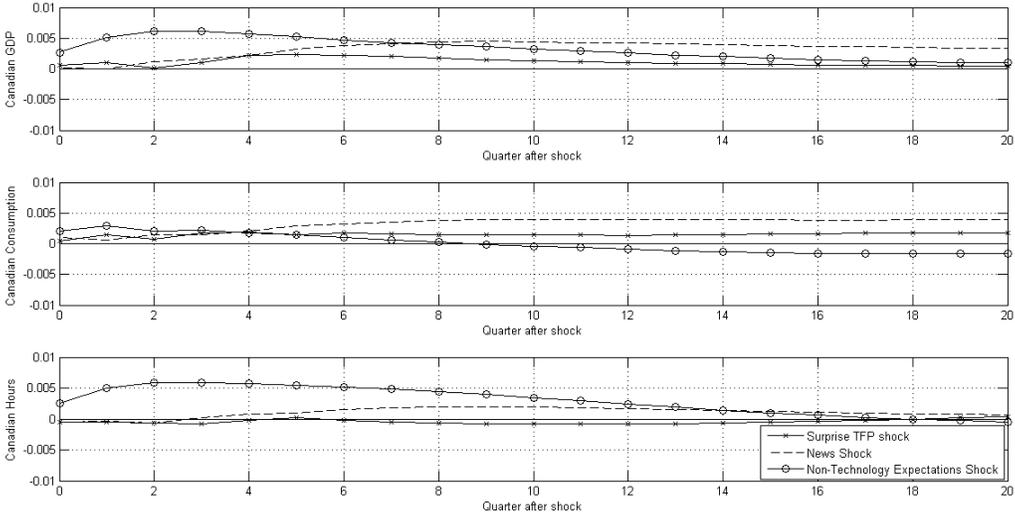
Notes: This figure plots the responses of the US CPI, the oil price index, the stock price index, the US-Canada real exchange rate, and US real wage to the three shocks. Note the CPI variable is in log-levels.

Figure A10: Responses with Canadian Consumption as a Core Variable



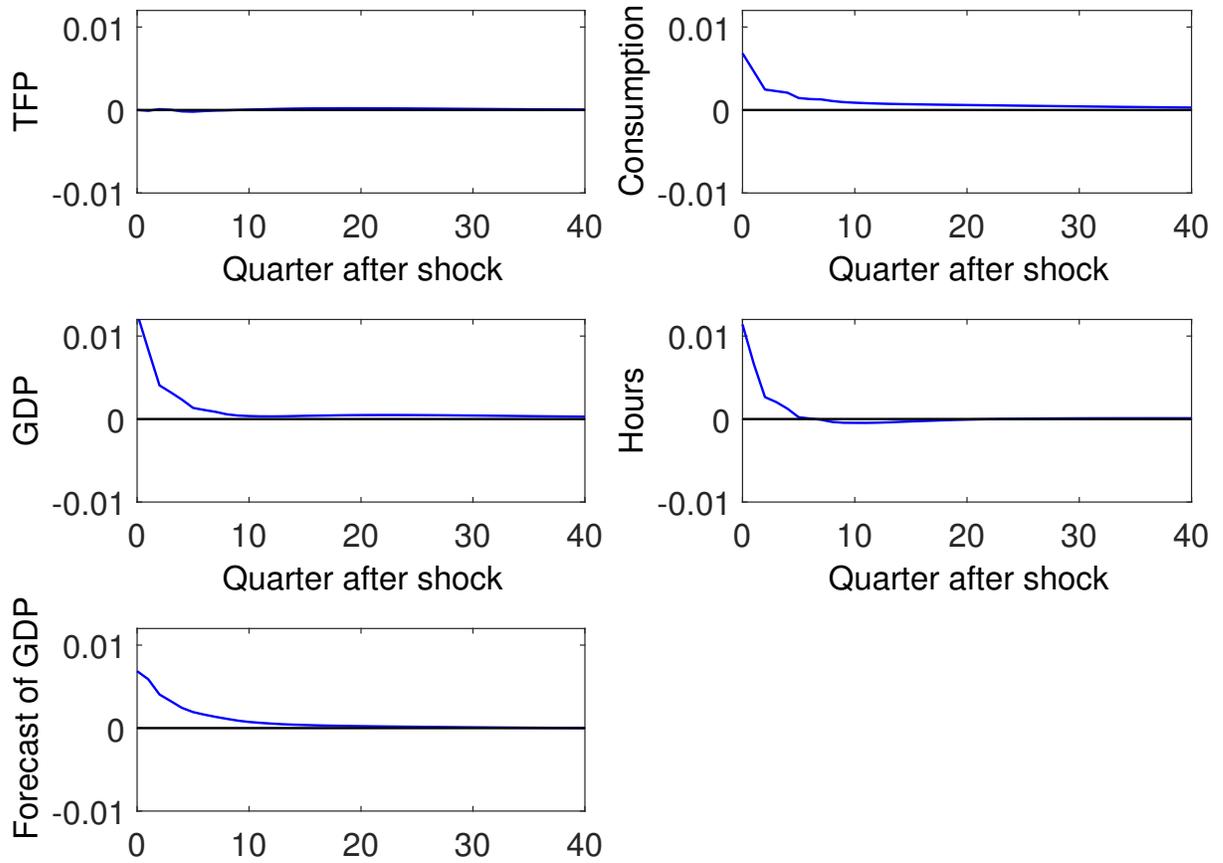
Notes: This figure plots the impulse responses of all key variables to the three shocks in a six-variable VAR where Canadian consumption is treated as a core variable (see section 3.1 for details).

Figure A11: Responses with Canadian GDP, Consumption, and Hours Included in the Same VAR



Notes: This figure plots the impulse responses of the three Canadian variables to the three shocks in a eight-variable VAR in which Canadian GDP, consumption, and hours are included together.

Figure A12: The Impulse Responses to the Non-Technology Expectations Shock: Model-Simulated Data



Notes: This figure plots the impulse responses to the non-technology expectations shock of true TFP, consumption, output, labor and the forecast of output in simulated data from the model, as described in Section 4.2.

Table A1: Other Controls: Non-Technology Expectations Shock Variance Decomposition

Panel A: US GDP							
Horizon	Monetary	Fiscal	Factor	RER	Stock Prices	Oil	Financial
1Q	0.83	0.66	0.65	0.63	0.66	0.64	0.62
2Q	0.83	0.76	0.72	0.76	0.72	0.76	0.57
1Y	0.64	0.62	0.54	0.66	0.49	0.63	0.38
2Y	0.34	0.38	0.23	0.45	0.26	0.34	0.38
5Y	0.17	0.28	0.11	0.27	0.15	0.18	0.07
10Y	0.12	0.27	0.09	0.22	0.12	0.15	0.07

Panel B: US Hours							
Horizon	Monetary	Fiscal	Factor	RER	Stock Prices	Oil	Financial
1Q	0.64	0.62	0.59	0.47	0.66	0.53	0.88
2Q	0.64	0.71	0.70	0.58	0.70	0.65	0.76
1Y	0.61	0.70	0.66	0.58	0.57	0.66	0.60
2Y	0.42	0.53	0.37	0.48	0.36	0.47	0.31
5Y	0.23	0.37	0.14	0.30	0.23	0.24	0.14
10Y	0.17	0.32	0.11	0.24	0.19	0.18	0.11

Panel C: Canadian GDP							
Horizon	Monetary	Fiscal	Factor	RER	Stock Prices	Oil	Financial
1Q	0.18	0.22	0.17	0.19	0.15	0.18	0.20
2Q	0.33	0.36	0.28	0.31	0.25	0.29	0.29
1Y	0.42	0.45	0.33	0.43	0.28	0.39	0.33
2Y	0.33	0.37	0.21	0.39	0.21	0.30	0.19
5Y	0.21	0.30	0.12	0.23	0.16	0.20	0.11
10Y	0.17	0.28	0.10	0.20	0.14	0.18	0.08

Notes: This table shows the contribution of the NTE shock to the forecast error variance of the key variables when the core VAR is extended with various controls, discussed in Section 3.1. Monetary refers to the VAR with an identified monetary policy shock, Fiscal refers to the VAR augmented with government spending, Factor refers to the FAVAR, RER refers to the real exchange rate, and Stock Prices and Oil refer to VARs that include an index of stock prices and an oil price index respectively

Table A2: Correlations of Shocks: US and Canada

Contemporaneous Correlations	US TFP	US News	US NTE
Canada TFP	0.16	0.01	0.17
Canada News	-0.04	-0.02	-0.17
Canada NTE	-0.02	0.11	0.18

Notes: This table shows the contemporaneous correlation of the three identified shocks between the US and Canada. They are identified in separate VARs with only the core variables corresponding to each country.

Table A3: Model Shocks vs. Identified Shocks: Correlations

Shock	TFP	NTE
Correlation	0.792	0.934

Notes: This table reports the correlations of the true surprise TFP and non-technology expectations shocks in the model with the shocks identified using the method described in Section 2. The shocks are identified using a panel with the same dimensions as available in the data. Measurement error with a standard deviation of 1/20th of output is added to each variable prior to estimating the VAR.