

NBER WORKING PAPER SERIES

TFP, NEWS, AND "SENTIMENTS:"
THE INTERNATIONAL TRANSMISSION OF BUSINESS CYCLES

Andrei A. Levchenko
Nitya Pandalai-Nayar

Working Paper 21010
<http://www.nber.org/papers/w21010>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2015

We are grateful to Rudi Bachmann, Christoph Boehm, Jordi Gali, Lutz Kilian, Eric Sims, and Jaume Ventura for helpful suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Andrei A. Levchenko and Nitya Pandalai-Nayar. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

TFP, News, and "Sentiments:" The International Transmission of Business Cycles
Andrei A. Levchenko and Nitya Pandalai-Nayar
NBER Working Paper No. 21010
March 2015
JEL No. E32,F41,F44

ABSTRACT

We propose a novel identification scheme for a non-technology business cycle shock, that we label "sentiment." 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 "sentiment" – from the US to Canada. The US sentiment shock produces a business cycle in the US, with output, hours, and consumption rising following a positive shock, and accounts for the bulk of short-run business cycle fluctuations in the US. The sentiment 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 up to 50% of the forecast error variance of Canadian GDP and about 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.

Andrei A. Levchenko
Department of Economics
University of Michigan
611 Tappan Street
Ann Arbor, MI 48109
and NBER
alev@umich.edu

Nitya Pandalai-Nayar
Department of Economics
University of Michigan
611 Tappan Street
Ann Arbor, MI 48109
npnayar@umich.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. (For 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). 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).¹

This paper investigates empirically the relative importance of the cross-border transmission of both technology and non-technology shocks. It uses US and Canada as a laboratory to study these issues.² These 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.

We begin by identifying three types of US shocks in a structural vector auto-regression (VAR) setting. The first is a shock to contemporaneous TFP. This shock is identified as the reduced-form TFP shock, assuming that the TFP series is ordered first. New to the study of the international business cycle, the TFP series we use is adjusted for unobserved input utilization. Basu, Fernald, and Kimball (2006) show that the utilization adjustment has a large impact on both the properties of the TFP series itself, and on the impulse responses of US macroeconomic aggregates to the TFP shock. The second shock is a news shock about future TFP (Beaudry and Portier, 2006), identified

¹An 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 the US monetary policy explain only very limited share of forecast error variance of other countries’ output, while Ilzetzki and Jin (2013) show that even the sign of the impact is not stable over time. In a similar vein, Eickmeier and Ng (2011), Helbling et al. (2011), and Kollmann (2013) show that the share of variance of other countries’ GDP accounted for by the US credit shocks and bank shocks is small as well.

²This approach has been adopted by Schmitt-Grohé (1998).

following Barsky and Sims (2011) as the shock that has no contemporaneous TFP impact and explains the maximum of the forecast error variance of the utilization-adjusted TFP series.

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 the forecast (or consumer confidence) series. Because the shock is identified explicitly from data on expectations after controlling for shocks to current and future TFP, we label this shock “sentiment” as an homage to Angeletos and La’O (2013). It is important to underscore that this is only a short-hand, as we do not identify the precise mechanism behind the fluctuations posited in Angeletos and La’O (2013). Our shock can be driven by anything that makes agents expect better/worse times, conditional on available information about current and future productivity.

We identify the three shocks in the US data. We then estimate the impact of these shocks on Canadian macro aggregates, included as non-core variables in the baseline VAR. The results can be summarized as follows. The sentiment shock generates a US business cycle and accounts for an important share of forecast error variance in the US macro aggregates. GDP, consumption, and hours (as well as expectations) all increase on impact, peak within a year, and revert back to pre-shock values in the medium run. These dynamics are consistent with the sentiment shock being a transitory “demand” shock. The sentiment shock drives the bulk of short-run fluctuations in the US. It accounts for 70-80% of the forecast error variance in GDP at short frequencies (one year or less). At short frequencies, it also accounts for about one-third of the forecast error variance of consumption and 75-85% of the forecast error variance in total hours. The finding that a non-technology shock is responsible for a large share of short-run fluctuations is of course consistent with results 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. On the rest of the Canadian variables, the three identified shocks have very different impacts.

The common theme is that in the short run, Canadian aggregates react much more strongly to the sentiment shocks than to the surprise and news TFP shocks. Following a sentiment shock, Canadian GDP rises instantaneously and peaks within one year. By contrast, the response of Canadian GDP to the US surprise TFP or news shocks is positive but takes place with a lag of 2-3 quarters. Canadian consumption and hours follow the same pattern. Canadian exports to the US and US exports to Canada both rise instantaneously following a sentiment shock, peak at 1 or 2 quarters, and then fall back to steady state. By contrast, 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. There is suggestive evidence that Canadian imports from and exports to the US actually fall on impact in response to a US news shock.

Among the three identified US shocks, the sentiment shock is by far the most important in accounting for the forecast error variance of the Canadian variables. At short frequencies, it accounts for 20-50% of the forecast error variance of Canadian GDP, 10-15% of Canadian consumption, 20-40% of Canadian hours, and a 25-40% of Canada-US trade flows. By contrast, 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 10 years, and for less than 10% of Canadian consumption. The (lack of) importance of the US news shock is similar at short frequencies, though the news shock does become more important for Canadian output and consumption at frequencies longer than 2 years.

Finally, we examine the role of the three US shocks in business cycle comovement between US and Canada, by means of computing conditional correlations between the variables due to each shock following the approach in Galí (1999). The correlation of the US and Canadian GDP conditional on surprise TFP shocks is 0.47. The surprise TFP shock actually generates a negative correlation in consumption (-0.13) and hours (-0.47) between US and Canada. Conditional GDP correlations due to news (0.99) and sentiment (0.99) shocks are much higher. These shocks generate positive instead of negative correlations of consumption and hours as well. The sentiment shock generates a conditional correlation in consumption of 0.86 and in hours of 0.99, which is substantially higher than that generated by the news shock.

The bottom line is that at short frequencies, the non-technology shocks generate a much stronger cross-border impact of US shocks and account for a higher share of Canadian fluctuations. The sentiment 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 innovation central to most IRBC models is actually a shock that does not generate substantial comovement.

Our analysis is most closely related to 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 sentiment shocks unrelated to TFP.

Our paper draws heavily on the recent closed-economy empirical and theoretical literature on “demand”-driven fluctuations (see, among others, Galí, 1999; Beaudry and Portier, 2006; Lorenzoni, 2009; Barsky and Sims, 2011; Angeletos and La’O, 2013; Blanchard, L’Huillier, and Lorenzoni, 2013). Two recent papers in particular identify shocks that are interpreted as sentiments. Angeletos, Collard, and Dellas (2014) extract a shock that explains the most of the forecast error variance of key macroeconomic aggregates, and show that it has the properties consistent with being a confidence shock. Angeletos, Collard, and Dellas (2014) and Milani (2014) structurally estimate fully-specified DSGE models that incorporate sentiment shocks, and show that the sentiment shocks identified within the structure of those models can explain a large fraction of the US business cycle fluctuations. Our empirical strategy complements both of these approaches. In contrast to both of these alternatives, we are explicit about separating the TFP news shock from a strictly non-technology sentiment shock. Relative to the data-driven exercise in Angeletos, Collard, and Dellas (2014), our identification strategy is based on explaining the variation only in an explicit

expectational variable (e.g., the GDP forecast). Our strategy thus “ties our hands behind our back” to a much greater extent, as we are not looking for a shock that by construction explains the bulk of fluctuations in the key macro aggregates. We complement the fully structural DSGE estimation approach by performing a more data-driven exercise. It is reassuring that our findings regarding the importance of “sentiments” in the US business cycle are consistent with these alternative 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 and estimation methods. Section 3 describes the data. Section 4 presents the main results, while Section 5 discusses interpretation and relates our analysis to the literature. Section 6 concludes.

2 Identification

It is important to assess the international transmission of different identified shocks because different shocks could produce different cross-border spillovers, particularly for small open economies closely linked to the US. A shock that leads to an increase in US consumption would simultaneously increase demand for Canadian exports.³ The general equilibrium effects of such a shock would feed through the Canadian economy. The news shock about a future improvement in US TFP would therefore signal a period of lasting, increased demand for Canadian exports, as well as future deterioration of the US terms of trade. The demand shock properties of the sentiment shock would also lead to increased production in and exports from Canada, though the effect would be expected to be shorter-lived.

The reason we consider the reaction of Canadian aggregates in particular is related to our identification of the news and sentiment shocks. We rely on the shocks being US-specific, i.e. they are not being driven by movements in Canadian TFP. 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. We also check for the possibility of correlated demand shocks, which are not visible in TFP movements, in Section 5. We find no evidence that the responses shown are due to a correlated shock in Canada.

³Relative price changes would mitigate this response, if Canadian products are substitutable for domestic goods in the US consumer’s basket.

2.1 Method

Our identification strategy builds on Uhlig (2003) and Barsky and Sims (2011). We assume 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, we assume that forecasts of GDP are influenced not only by the surprise innovation in TFP and the anticipated future improvement in TFP, but also by “sentiments,” as the forecasters rationally expect the sentiment shock to lead to a temporary boom in the economy and increase GDP. Forecasters also respond to other changes in the economy that could stimulate GDP, but we assume that the bulk of the variation in forecasts is due to these three shocks. A simple process for forecast F_t that satisfies this assumption is:

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

with ϵ^{sent} the sentiment shock.

Clearly it would not be possible to identify the three shocks of interest from movements in TFP and forecasts 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 $\mathbf{u}_t = A\epsilon_t$

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

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 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 $DD' = I$, AD will provide an identification of the structural shocks.

The forecast error h steps ahead is defined as

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

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. Without loss of generality, assume the second structural shock is the news shock, and thus the second column of AD is its impact vector. 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} .⁵ 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} , Barsky and Sims (2011) choose the vector γ^{news} (the second column of D), such that this second

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

shock maximizes the residual forecast error variance of the TFP process over horizon H^{news} .

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

$$\gamma^{news}(1) = 0 \tag{4}$$

$$\gamma^{news} \gamma^{news'} = 1,$$

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

We next proceed to the identification of the sentiment 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. Further, we impose that this shock does not affect true TFP. The procedure outlined above naturally builds in this assumption: by allowing only the first two shocks to affect TFP, we minimize the impact of the remaining $k - 2$ structural shocks, which includes the sentiment shock.

Let the expectational variable F_t be ordered 5th in the VAR, and without loss of generality assume that the sentiment shock is the 3rd shock. Note that by equating the first reduced-form 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 sentiment shock as the linear combination of the remaining $k - 2$ reduced-form innovations that maximizes the forecast error variance of F_t . Because the sentiment shock is short-run, we select it to maximize the forecast error variance for a 2-quarter horizon ($H^{sent} = 2$). Formally:

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

subject to

$$D(1, i) = 0 \quad \forall i \neq 1 \tag{5}$$

$$D(:, 2) = \gamma^{news}. \tag{6}$$

Both stages of this optimization are conditional on an arbitrary orthogonalization, the Choleski

decomposition \tilde{A} . The identification strategy for both shocks is robust to the reordering of the remaining $k - 1$ variables in the VAR other than TFP. The first restriction – (4) and (5) – common to both problems implies that none of the $k - 1$ structural shocks has a time- t impact on TFP. The second restriction (6) ensures that identification of the sentiment shock holds identification of the news shock constant: we expect the surprise innovation in TFP and the news shock, as informative about true fundamentals, to explain the bulk of the movements in the forecast of GDP. The sentiment shock identified in this manner simply captures patterns in the residual variance of the forecast of GDP, once supply-side fundamentals are accounted for.⁶

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. Our identification strategy relies on ‘medium-run’ identification. It might appear that the natural identification of the sentiment 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).

2.2 Identification of International Transmission

We estimate the impact of the 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. This might seem a strong assumption, but we believe it is reasonable given the small size of the Canadian economy relative to the US (Canadian GDP is about one-tenth that of the US).

We estimate the reduced-form VAR with estimated generalized least squares (EGLS) using a

⁶In 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 – over several horizons.

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 restrictions here simply imply that $\forall C_j, C_j(1 : k - 1, k) = 0$. 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} then is:

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

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 with the assumption of Gaussian errors, the estimator for \mathbf{b} is the same as the maximum likelihood estimator. With these estimates of \mathbf{b} , it is then straightforward to compute the impulse response functions of the each Canadian aggregate to the three shocks of interest. Note that the identification of the shocks is unaffected by this procedure. 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).

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 is not available quarterly, so the Fernald (2014) series corrects only for variable capital and labor utilization.

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.

For the forecasts of US GDP, we construct an index from the Survey of Professional Forecasters (SPF) forecasts of the growth rate of GDP one quarter ahead. The SPF is provided by the Federal Reserve Bank of Philadelphia, and is conducted quarterly. The series begins in 1968:Q4. For NIPA variables, it provides quarterly forecasts at several horizons as well as longer-term forecasts. We use the one quarter ahead growth rate forecast. The perturbation to US sentiment 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 $H^{sent} = 4, 6$ over which we expect the sentiment shock to contribute to the forecast error variance of the GDP forecast variable, and find no significant change 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.

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. Details of the algorithm are in the Appendix A. 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.

We present the impulse response functions for both the utilization-adjusted TFP series and the implied capital utilization series.⁷

4 Results

Our baseline specification identifies the news shock at a horizon of ten years, the sentiment shock at a horizon of two quarters, and uses the 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 sentiment shocks on the U.S. economy (Figures 1, 2 and 3), followed by an analysis of the transmission to Canada. Section 5 places the results in the context of standard business cycle models and some variants proposed in the literature.

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.

⁷We 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.

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.⁸ It therefore cannot 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 sentiment 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 3. The business cycle generated by the shock lasts approximately three years. These results are consistent with the empirical factor extracted by Angeletos, Collard, and Dellas (2014). That factor is constructed to explain the bulk of the

⁸This 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 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 (Kurmann and Mertens, 2014).

business cycle variation in hours and investment, and our results provide one direct structural interpretation for it as fluctuations in GDP forecasts orthogonal to current or future productivity. Angeletos, Collard, and Dellas (2014) also point to the potential for this shock to be a driving force of business cycle fluctuations. A substantial empirical literature beginning with Galí (1999) has previously argued that demand shocks are promising for explaining business cycles, but ours is (to our knowledge) the first paper to directly measure these shocks from forecasts 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 of demand shocks in Section 5.

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 sentiment shocks respectively. At short frequencies, the sentiment shock appears most important. It accounts for 70-80% of the variation in GDP, 32-37% in consumption, and 77-86% 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 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 partly identified one such shock.

International Transmission Figure 4 sets the stage for the remainder of the results. It shows the impulse responses of Canadian utilization-adjusted TFP to the three identified US shocks. None of the three identified shocks have a perceptible impact on Canadian technology (note also the different scale of the y-axis compared to the other figures). The news shock actually leads to a barely visible, though persistent and 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 below is not accompanied by a change in Canadian productivity.

Figure 5 shows that 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 sentiment 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 sentiment 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 sentiment shocks could generate spillovers. As Figures 6 and 7 show, Canadian exports to the US and imports from the US show the strongest responses to the sentiment 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 sentiment shock generates increased demand in the US, this response is unsurprising.

The increased exports do not come through lowered Canadian consumption. Rather, as Figures 8 and 9 show, the factors of production are used more intensively following a US sentiment 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 (Figure 10). Demand for imports increases as well as a result of the higher consumption, and US exports to Canada rise.⁹ The empirical evidence clearly suggests that the sentiment 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 U.S. 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 con-

⁹Of 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 U.S. output to Canada. That paper does not distinguish between the types of shocks that affect the U.S., however.

sumption 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 Q3.¹⁰ 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 sentiment shocks in generating business cycle spillovers. Forecast error variance decompositions provide additional support for the importance of sentiment 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 the shares of forecast error variances of the Canadian macro aggregates accounted for by the three identified US shocks. At short frequencies, the sentiment shock is by a large margin the most important of the three. The sentiment shock contributes substantially to the forecast error variance of US-Canada trade, explaining up to 49% of the variance of Canadian exports and 45% of imports at the one year horizon. It also explains a large fraction of the forecast error variance of Canadian output (49% at one year), hours and utilization (about one third), and consumption (10-15%). The impact of this short run US 'demand' shock on a smaller trading partner is persistent, as it still accounts for 31% of the variance at 10 years. The high TFP variation attributed to this shock at 10 years is likely also due the utilization mis-measurement, 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.

¹⁰Other explanations such as habit formation could also play a role here.

As further evidence on the importance of both sentiment 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 $A(L)$ is the matrix of lag order polynomials of the structural moving average representation of the VAR and h is the horizon. In practice, we compute these correlations for a finite but large maximum horizon of 10000 periods.

The results from this exercise are in Table 4. The sentiment 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 observable 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 sentiment shocks deliver strongly positive consumption correlations, the correlation of hours due to the news shock is too low at only 0.46, but due to the sentiment shock it is too high at 0.99. The sentiment shock comes the closest to explaining the unconditional cross-correlations of exports from Canada with US output.¹¹

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. Sentiment shocks are important for transmission at higher frequencies, while news shocks play a stronger role at medium/long frequencies. In future work, we will estimate a DSGE model of international business cycles that incorporates all three of these shocks to quantify their roles in comovement for a larger set of countries.

We conclude this section by discussing the role of all three shocks in recent business cycles.

¹¹Interestingly, 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 sentiment shocks also match the correlation of output and consumption well, but both undershoot the hours and output correlation.

Figures 11– 14 display the historical decompositions of the key US and Canadian macro aggregates into the components due to the three identified US shocks. While the TFP, sentiment, and news components all contributed to the great recession in the US, the fall in output in Canada appears driven entirely by the sentiment shock. Similar patterns are visible for Canadian consumption and hours, as well as exports and imports, indicating the sentiment shock played a key role in the transmission of the recent recession. The sentiment 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.

4.1 Robustness

We check the responses of all variables to variations in the horizons of identification of the sentiment shock, and find no significant qualitative difference for $h = 4$, $h = 8$ or $h = 16$. We also vary the forecast variable used in identification, using forecasts of GDP two quarters ahead and three quarters ahead. To conserve space, the results are not reported here, but the qualitative shape of the dynamic responses remains the same.¹²

Consumer confidence. To check robustness of the results to the choice of expectational variable, we replace the GDP forecasts the VAR with the E12Y variable from the Michigan Survey of Consumers. The results from this exercise are in Figures 15-23. Reassuringly, the patterns described above appear robust to the forward-looking measure used to identify the sentiment shock.

5 Discussion

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 sentiment shock, which is identified as a fully rational change in forecasts or sentiments orthogonal to TFP. That is, the forecaster/consumer cannot believe this shock is news or a noisy signal of future TFP. Further, Barsky and Sims (2012)

¹²Full results are available on request.

assess the importance of noise or 'animal spirits' shocks and find that they do not account for a substantial portion of the relationship between confidence and output. This supports the notion that the responses shown in Figure 2 are informative of true news.

Demand shocks. Our identified shock could be a combination of several shocks traditionally considered demand shocks such as monetary shocks or preference shocks. The strategy simply relies on the forecasters rationally expecting an increase in GDP, which is not due to news or noise about future TFP, and is unrelated to current TFP innovations. To the extent that an increase in demand leads to an increase in output, this would be identified as a sentiment shock in our framework. However, our shock is not consistent with all of the results associated with demand shocks previously identified in the literature. The shock orthogonal to changes in long run labor productivity in Galí (1999) for instance, leads to a temporary increase in labor productivity. Monetary shocks are also commonly proposed as demand shocks. Angeletos, Collard, and Dellas (2014) find that these 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. The response of US consumption to the sentiment 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 have also shown properties that would appear similar to demand shocks, despite being shocks to technology. The dynamics of output and hours 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, and 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 sentiment shock, increasing our confidence that it is not a productivity shock.

Spillovers or correlated shocks. Our sentiment shock is identified using only US data. Therefore, the possibility that the observed spillovers are due to correlated sentiment shocks affecting both countries simultaneously cannot be ruled out. While this is not a problem for the central message of this paper, we test this hypothesis by identifying a Canadian sentiment shock. Using data from the Conference Board of Canada, we identify a surprise TFP innovation, a news

shock, and a sentiment shock in Canada alone. We then check the correlation of the identified shocks with their US counterparts. The Canadian sentiment series 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". As in Barsky and Sims (2012), we construct the 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. The US and Canadian sentiment shocks are actually moderately *negatively* correlated, while the news shocks are moderately positively correlated. The correlation of the surprise TFP innovations is close to zero.¹³ This suggests that the observed comovement in the impulse responses is due to cross-border transmission in the case of the US sentiment shock. The Canadian response to a US news shock might be partially due to correlated news shocks; however, as Figure 4 shows, a US news shock does not generate a quantitatively meaningful impulse response of Canadian TFP at any horizon.

Relationship to standard neoclassical and new Keynesian models. A large body of work on closed-economy business cycles has established substantively that (1) RBC models driven by a technology shock do well at matching the key moments of US data; and (2) estimated technology shocks do 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. 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 in the suddenly more productive country increases. 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 role of the wealth effect, they do not generate the desired impact comovement in the closed economy. Similar intuition holds for an open economy setting:

¹³Results available upon request.

Kosaka (2013) finds that a model with a low elasticity of substitution between domestic and foreign goods, and no wealth effect on labor supply helps generate a news-driven business cycle. 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 sentiment 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 sentiment shocks.

6 Conclusion

We introduce a novel identification scheme to uncover the effects of surprise TFP innovations, news shocks, and "sentiment" 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 sentiment 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 business cycle synchronization. Future work will include estimating a dynamic two-country model to quantify the effects of the different shocks.

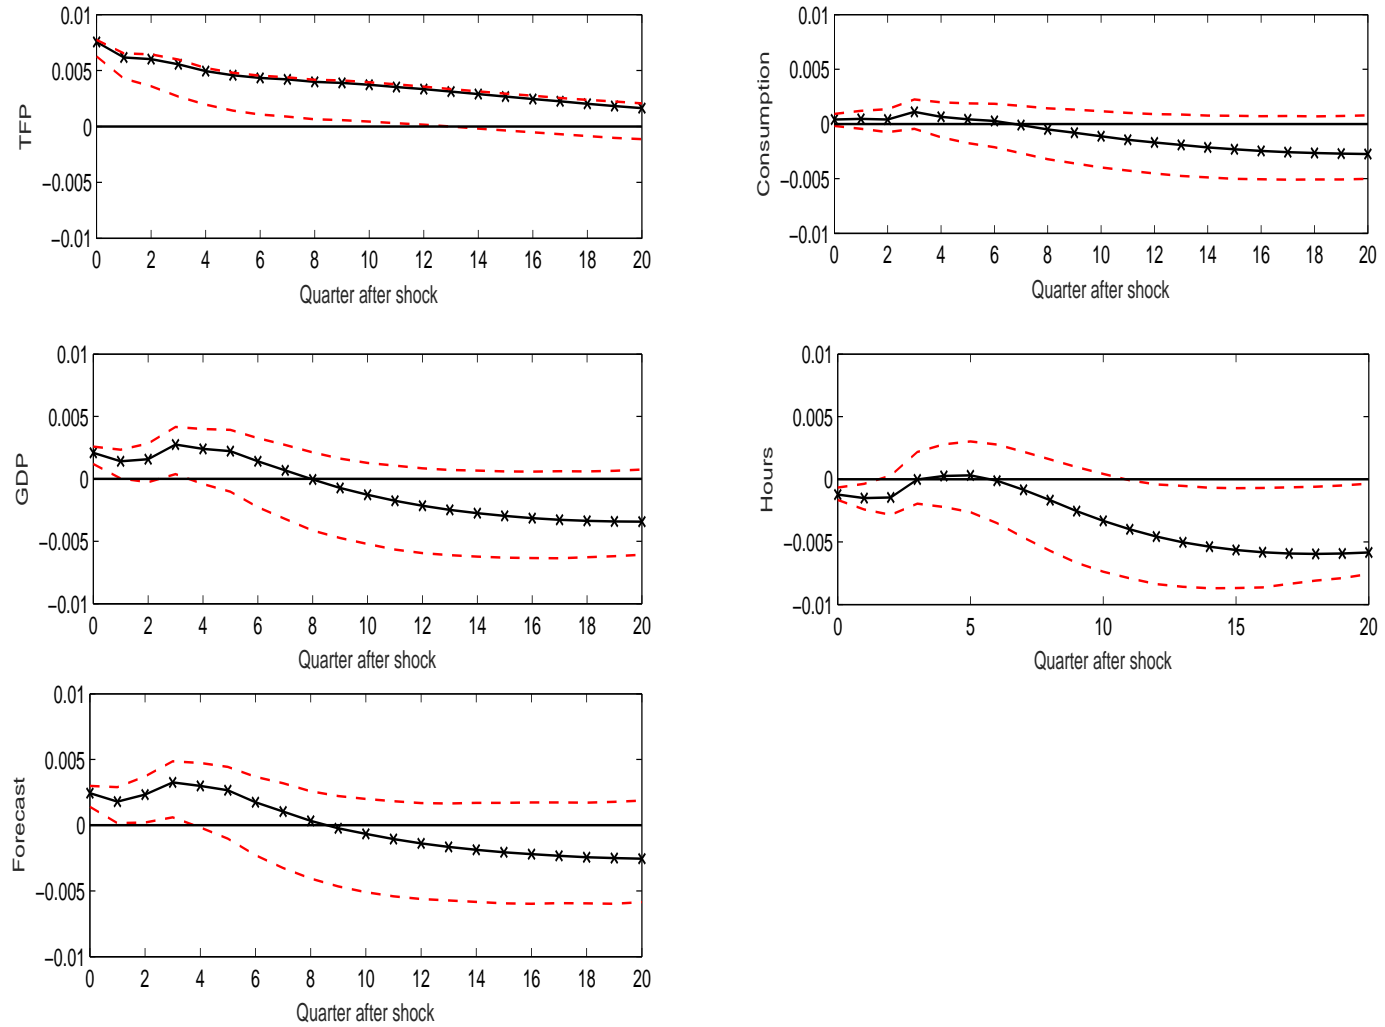
References

- 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.
- Barsky, Robert B., Susanto Basu, and Keyoung Lee. 2014. “Whither News Shocks?” NBER Working Paper 20666.
- 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.
- 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 Danny Quah. 1989. “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review* 79 (4):655–73.
- 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.
- 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.
- Eickmeier, Sandra and Tim Ng. 2011. “How do credit supply shocks propagate internationally? A GVAR approach.” Deutsche Bundesbank, Research Centre in its series Discussion Paper Series 2011,27.
- Engel, Charles and Jian Wang. 2011. “International Trade in Durable Goods: Understanding Volatility, Cyclical, 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.
- Fisher, Jonas D. M. 2006. "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks." *Journal of Political Economy* 114 (3):413–451.
- 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.
- Ilzetzki, 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.
- 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.
- Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *American Economic Review* 99 (5):2050–84.

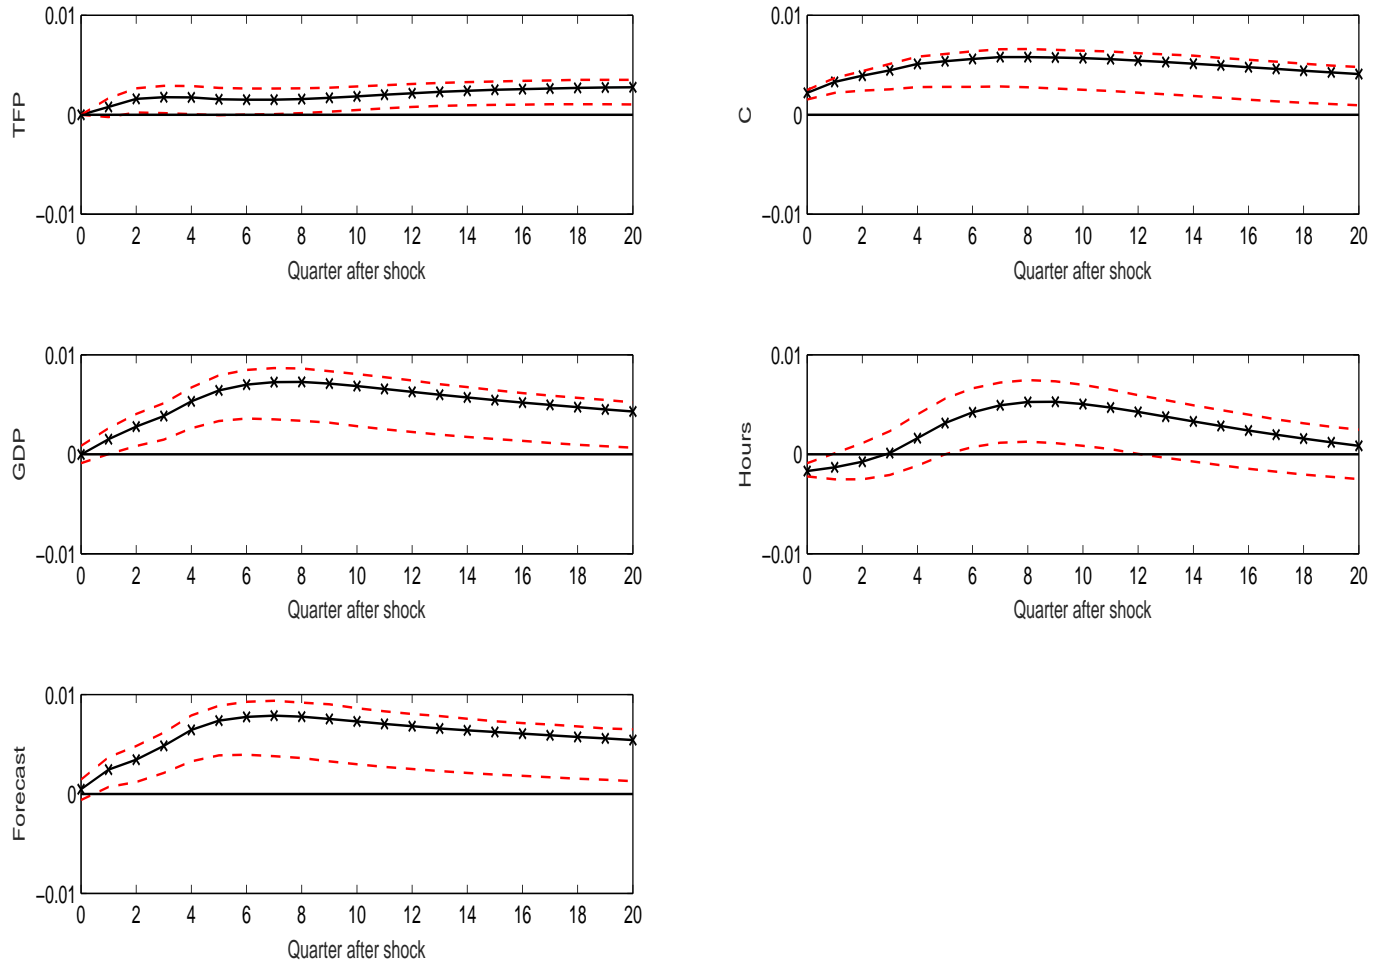
- Lütkepohl, Helmut. 2005. *New Introduction to Multiple Time Series Analysis*. Berlin Heidelberg: Springer-Verlag.
- Maćkowiak, Bartosz. 2007. “External shocks, U.S. monetary policy and macroeconomic fluctuations in emerging markets.” *Journal of Monetary Economics* 54 (8):2512–2520.
- Milani, Fabio. 2014. “Sentiment and the U.S. Business Cycle.” Mimeo, UC Irvine.
- 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.
- 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.
- 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



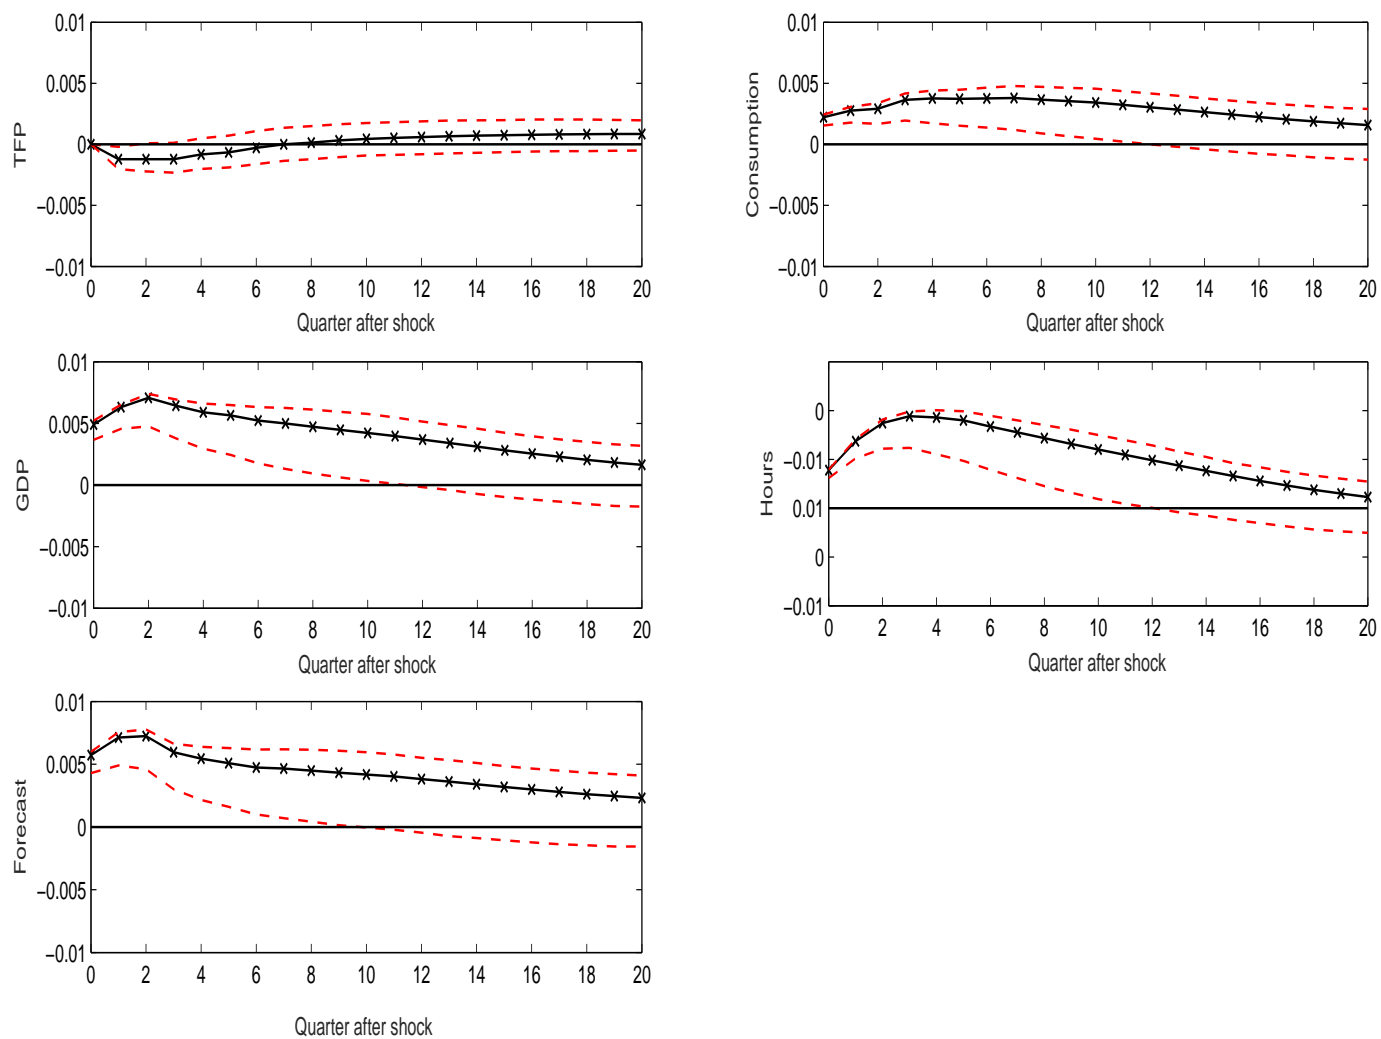
Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the surprise TFP shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

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



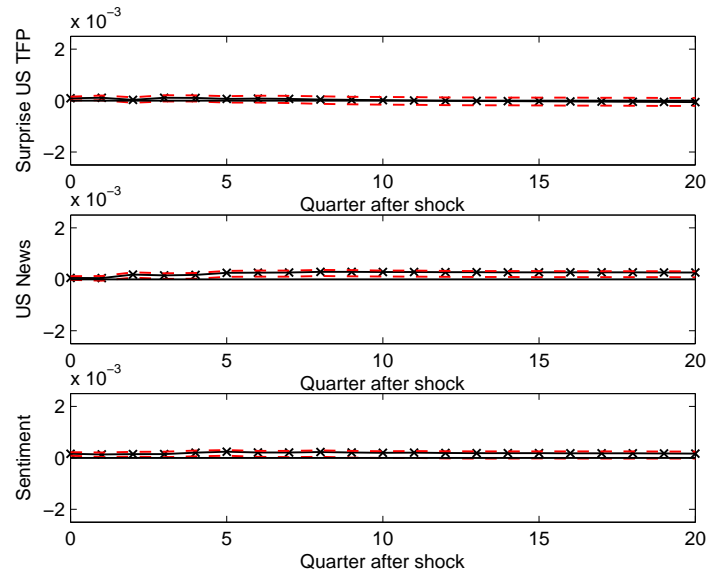
Notes: This figure plots the impulse responses of US TFP, GDP, consumption, hours, and the forecast of US GDP in response to the news shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 3: The Impulse Responses to the US Sentiment Shock




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

Figure 4: The Impulse Responses of Canadian TFP to the Three US Shocks




Notes: This figure plots the impulse responses of Canadian utilization-adjusted TFP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 5: The Impulse Responses of Canadian GDP to the Three US Shocks




Notes: This figure plots the impulse responses of Canadian GDP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 6: The Impulse Responses of Canadian Exports to the US to the Three US Shocks




Notes: This figure plots the impulse responses of Canadian exports to the US to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 7: The Impulse Responses of Canadian Imports from the US to the Three US Shocks




Notes: This figure plots the impulse responses of US exports to Canada to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 8: The Impulse Responses of Canadian Hours to the Three US Shocks




Notes: This figure plots the impulse responses of Canadian total hours to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 9: The Impulse Responses of Canadian Utilization to the Three US Shocks




Notes: This figure plots the impulse responses of Canadian utilization to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 10: The Impulse Responses of Canadian Consumption to the Three US Shocks




Notes: This figure plots the impulse responses of Canadian consumption to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 11: Historical Decompositions, Part 1




(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 12: Historical Decompositions, Part 2




(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 13: Historical Decompositions, Part 3




(a) US Hours




(b) Canadian Hours

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 14: Historical Decompositions, Part 4



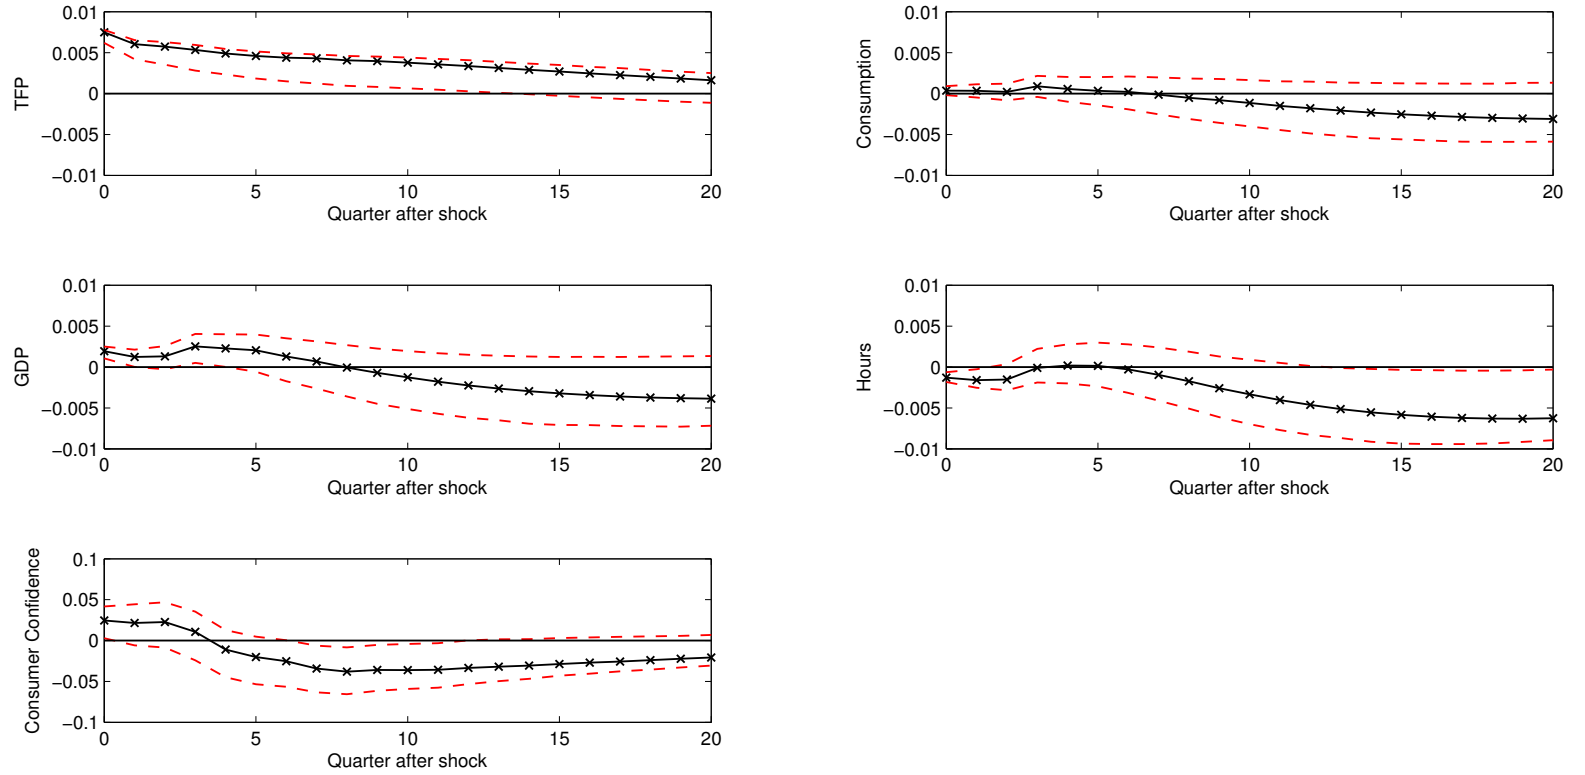
(a) US Consumption



(b) Canadian Consumption

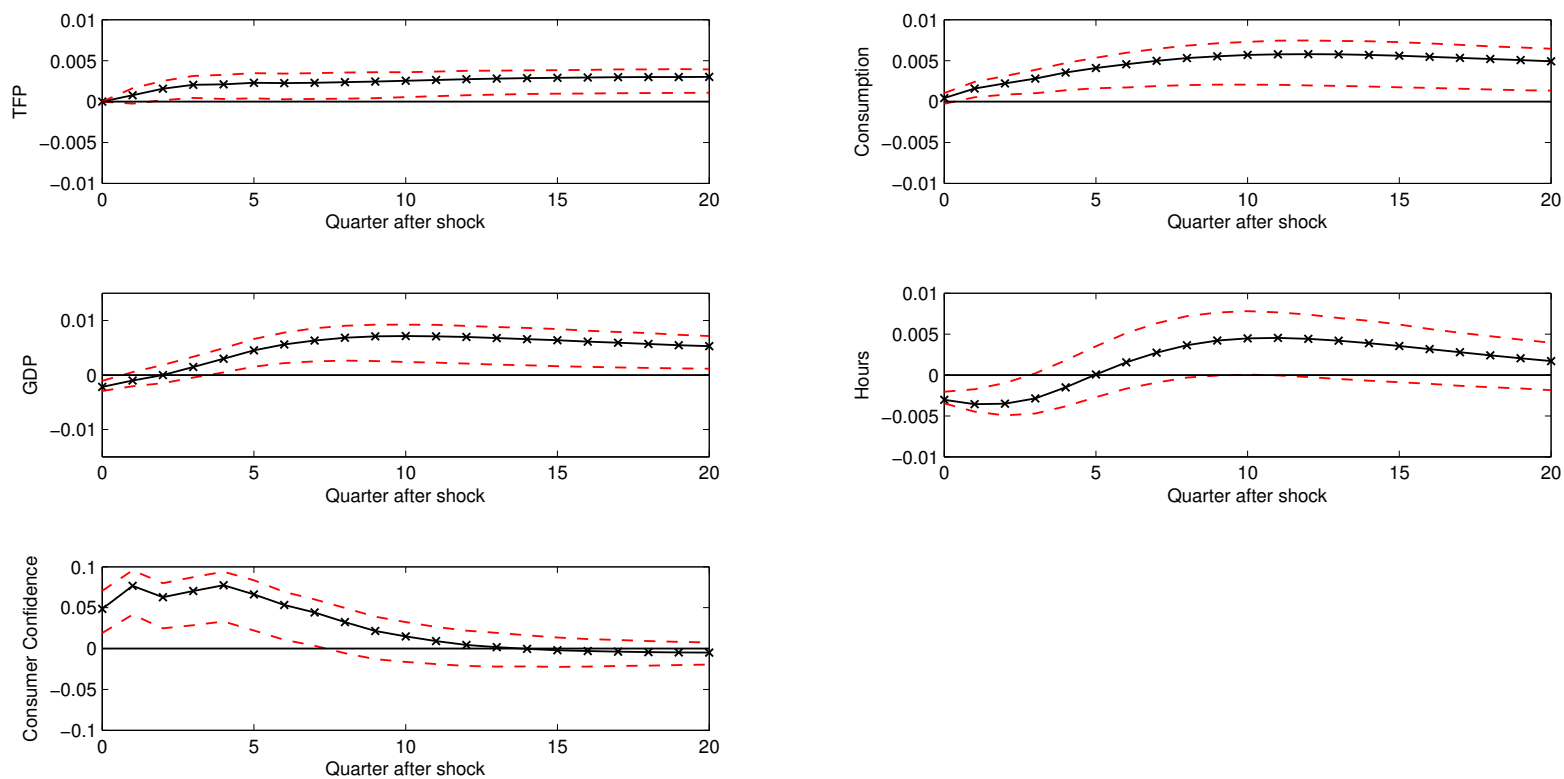
Notes: These figures show the decomposition of historical data into components due to the three identified shocks. The shaded areas are US NBER recession dates.

Figure 15: The Impulse Responses to the US Surprise TFP Shock, Using US Consumer Confidence



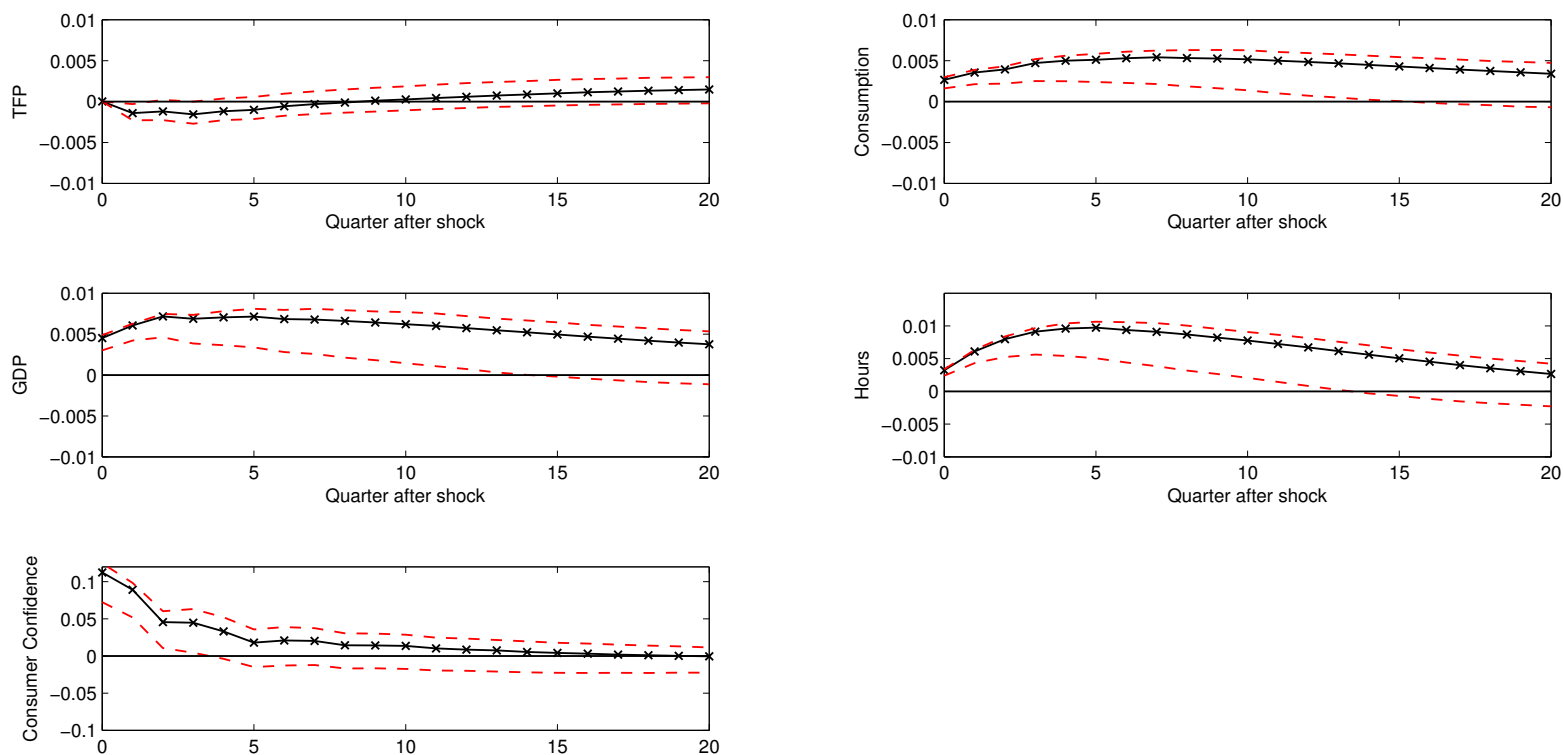
Notes: This figure plots the impulse responses of the core US variables to a surprise TFP innovation, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 16: The Impulse Responses to the US News TFP Shock, Using US Consumer Confidence



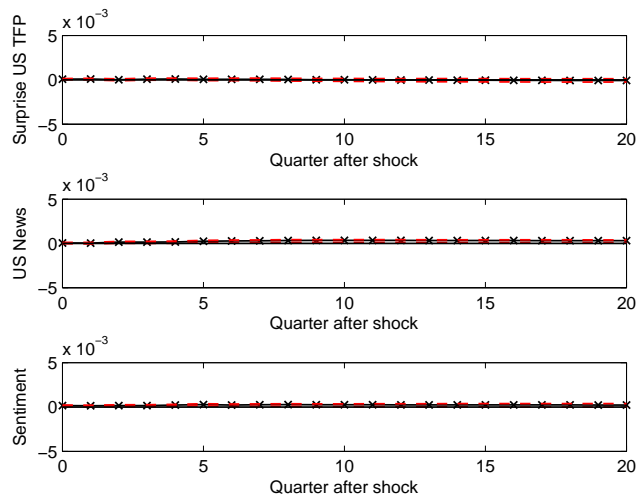
Notes: This figure plots the impulse responses of the core US variables to the news shock, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 17: The Impulse Responses to the US Sentiment Shock, Using US Consumer Confidence



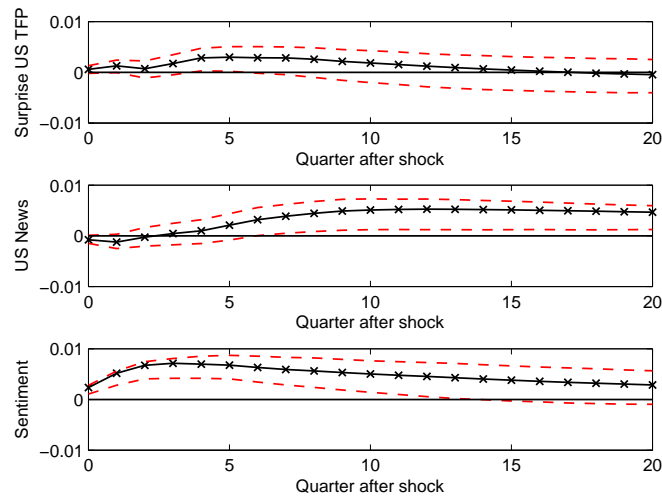
Notes: This figure plots the impulse responses of the core US variables to the sentiment shock, identified in a VAR with the Michigan Consumer Confidence indicator ordered fifth. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 18: The Impulse Responses of Canadian TFP to the Three US Shocks, Using US Consumer Confidence




Notes: This figure plots the impulse responses of Canadian utilization-adjusted TFP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 19: The Impulse Responses of Canadian GDP to the Three US Shocks, Using US Consumer Confidence




Notes: This figure plots the impulse responses of Canadian GDP to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 20: The Impulse Responses of Canadian Exports to the US to the Three US Shocks, Using US Consumer Confidence




Notes: This figure plots the impulse responses of Canadian exports to the US to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 21: The Impulse Responses of Canadian Imports from the US to the Three US Shocks, Using US Consumer Confidence



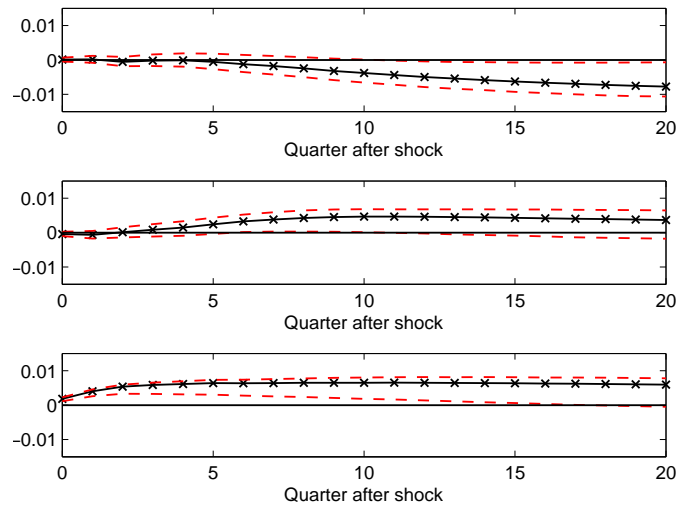
Notes: This figure plots the impulse responses of US exports to Canada to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 22: The Impulse Responses of Canadian Hours to the Three US Shocks, Using US Consumer Confidence




Notes: This figure plots the impulse responses of Canadian hours per worker to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 23: The Impulse Responses of Canadian Utilization to the Three US Shocks, Using US Consumer Confidence



Notes: This figure plots the impulse responses of Canadian capital utilization to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

Figure 24: The Impulse Responses of Canadian Consumption to the Three US Shocks, Using US Consumer Confidence



Notes: This figure plots the impulse responses of Canadian consumption to each of the three shocks in the US: surprise TFP shock, news shock about future US TFP, and US sentiment shock, identified in the VAR with the consumer confidence series. Standard errors are bias-corrected bootstrapped 90 percent confidence intervals.

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: Sentiment Shock: Variance Decomposition

Panel A: US					
Horizon	TFP	Output	Consumption	Hours	Forecast
1Q	0.00 (0.00)	0.77 (0.19)	0.37 (0.13)	0.69 (0.14)	0.71 (0.17)
2Q	0.02 (0.02)	0.85 (0.15)	0.36 (0.11)	0.80 (0.12)	0.73 (0.12)
1Y	0.03 (0.03)	0.86 (0.11)	0.32 (0.11)	0.72 (0.12)	0.62 (0.11)
2Y	0.02 (0.04)	0.67 (0.12)	0.28 (0.13)	0.45 (0.14)	0.37 (0.12)
5Y	0.02 (0.06)	0.44 (0.12)	0.21 (0.12)	0.30 (0.16)	0.25 (0.12)
10Y	0.04 (0.07)	0.37 (0.12)	0.16 (0.12)	0.24 (0.16)	0.19 (0.12)

Panel B: Canada							
Horizon	Output	Consumption	Hours	Exports	Imports	TFP	Utilization
1Q	0.21 (0.06)	0.11 (0.05)	0.23 (0.06)	0.25 (0.07)	0.29 (0.06)	0.08 (0.04)	0.20 (0.06)
2Q	0.37 (0.07)	0.15 (0.06)	0.35 (0.08)	0.40 (0.08)	0.40 (0.07)	0.08 (0.05)	0.35 (0.07)
1Y	0.49 (0.09)	0.12 (0.07)	0.41 (0.09)	0.49 (0.09)	0.45 (0.09)	0.10 (0.06)	0.45 (0.09)
2Y	0.43 (0.10)	0.10 (0.07)	0.38 (0.11)	0.46 (0.10)	0.40 (0.10)	0.17 (0.08)	0.38 (0.11)
5Y	0.35 (0.10)	0.05 (0.06)	0.36 (0.11)	0.39 (0.10)	0.32 (0.10)	0.23 (0.10)	0.27 (0.12)
10Y	0.31 (0.10)	0.04 (0.06)	0.36 (0.11)	0.37 (0.10)	0.28 (0.10)	0.22 (0.10)	0.20 (0.11)

Notes: This table shows the contribution of the sentiment 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	Sentiment
US, Canada Output	0.73	0.47 (0.28)	0.99 (0.07)	0.99 (0.04)
US, Canada Consumption	0.51	-0.13 (0.47)	0.94 (0.12)	0.86 (0.23)
US, Canada Hours	0.68	-0.47 (0.56)	0.46 (0.47)	0.99 (0.31)
Exports from Canada, US Output	0.66	0.95 (0.21)	0.97 (0.17)	0.89 (0.09)
Canadian Imports, US Output	0.58	0.79 (0.32)	0.91 (0.29)	0.93 (0.11)
US Output, US Consumption	0.84	0.97 (0.05)	0.99 (0.03)	0.95 (0.04)
US Output, US Hours	0.87	0.81 (0.39)	0.20 (0.44)	0.83 (0.32)

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.

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 is 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 .¹⁴
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. We acknowledge this procedure is not as state of the art as methods used today to construct utilization-adjusted TFP in the US. Future work will examine possibilities for improving the measurement of utilization-adjusted TFP in other countries.

¹⁴The derivation of this expression uses the household’s optimization problem and can be found in Imbs (1999).