

On International Uncertainty Links: BART-Based Empirical Evidence for Canada

December 2015

Abstract

We use the indexes of economic policy uncertainty (EPU) developed by Baker et al. (2015) for Canada and seven (China, France, Germany, Japan, Russia, United Kingdom, and United States) other countries to study international uncertainty links by means of a Bayesian Additive Regression Trees (BART) algorithm. Using partial-dependence plots, relative-importance measures, and permutation tests, we shed light on how EPU in Canada is linked to EPU in seven other countries.

JEL classification: C11; C22; D80

Keywords: Economic policy uncertainty; BART algorithm; Canada

Rangan Gupta, Department of Economics, University of Pretoria, Pretoria, 0002, South Africa.

Christian Pierdzioch, Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany

* Marian Risse, Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany

* Corresponding author. E-mail: marian.risse@hsu-hh.de. Phone: +49-40-6541-3378.

1 Introduction

Against the background of the “Great Recession” and recent financial crises, research on the macroeconomic effects of uncertainty shocks (Bloom 2009; Jones and Olson, 2015), on the implications of uncertainty for asset prices (Antonakakis et al., 2013; Ko and Lee, 2015), and on international uncertainty links (Colombo, 2013; Yin and Han, 2014; Klößner and Sekkel, 2015) have mushroomed. The news-based economic policy uncertainty (EPU) indexes developed by Baker et al. (2015) are widely used in this research to measure uncertainty, and international uncertainty links have been often explored by means of causality tests (Ajmi et al., 2014) and variants of the technique proposed by Diebold and Yilmaz (2009). Our contribution to this rapidly growing strand of research is to show that interesting patterns in international uncertainty links can be recovered by means of a machine-learning algorithm known as Bayesian Additive Regression Trees (BART, Chipman et al., 2010).

We focus on Canada to empirically illustrate how the BART algorithm works. Empirical modeling with the BART algorithm uses Bayesian techniques to compute and additively combine regression trees (on regression trees, see Breiman et al., 1983). A regression tree explains the EPU index in Canada by using recursive hierarchical binary splits to subdivide the range of EPU indexes in other countries into non-overlapping regions. Within every region, the response of the Canadian EPU index is fixed at some region-specific constant. International uncertainty links are modeled by implementing several splits. A regression tree renders it possible to trace out even complex nonlinearities in the data and, thus, is a leading candidate for the study of international uncertainty links. Special techniques like partial dependence plots and measures of variable importance are available that make it straightforward to visualize and quantify international uncertainty links. While regression trees have several interesting advantages (see Hastie et al., 2009, page 351), its hierarchical structure makes individual regression trees high-variance predictors. The BART algorithm overcomes this drawback by combining several hundred regression trees using Bayesian techniques.

Our choice of Canada is motivated by the fact that it is a small open economy, and it is likely to be affected by global uncertainty directly and indirectly. The indirect channel operates through the impact of global uncertainty on domestic uncertainty. A result reported by Yin and Han

(2014) and Klößner and Sekkel (2015) that Canada is a net importer of uncertainty shocks indicates that such an indirect channel is likely to be an important element of international uncertainty links. We attempt to trace out this indirect channel because its presence would imply second round effects of global uncertainty, thereby magnifying the impact of the latter on the Canadian economy. Our research complements recent research using the Diebold and Yilmaz (2009) approach as the BART algorithm accounts in a natural way for potential nonlinearities in the data and sheds light on the quantitative importance of international uncertainty links from a different angle.

In Section 2, we describe the BART algorithm. In Section 3, we describe our data and we summarize our empirical results. In Section 4, we conclude.

2 The BART Algorithm

A regression tree uses binary hierarchical recursive splits to partition the space of the predictors (the EPU indexes for other countries than Canada) into non-overlapping regions. Upon implementing several splits a full-fledged tree emerges that consists of a root, interior nodes, and terminal nodes (the “leaves”). When a tree has b leaves, every leaf is characterized by a leaf parameter, $M = (\mu_1, \dots, \mu_b)$. A leaf parameter is the best prediction of the EUP index for Canada if the predictors assume a value in one of the non-overlapping regions resulting from the binary splits. Finding the non-overlapping regions requires the definition of a splitting rule for each interior node. A splitting rule implies, for example, that leaf parameter μ_1 predicts the EPU index for Canada when $x_i < c$ and leaf parameter μ_2 predicts the EPU index for Canada when $x_i \geq c$, where x_i is one of n predictors and c is a split point.

In order to overcome the data-sensitivity of individual trees that results from the recursive hierarchical ordering of splits, the BART algorithm is constructed in a way such that the EPU index for Canada is modeled using the sum of individual tree predictions. The basic idea underlying the BART algorithm, thus, is to model the EPU index for Canada, y , using an ensemble of m regression trees of the following format:

$$y_t = \sum_{j=1}^m f(x|T_j, M_j) + \epsilon_t, \quad \epsilon \sim N(0, \sigma^2), \quad (1)$$

where $\epsilon =$ a residual term, $T_j =$ a single regression tree, and $f =$ a function that links the leaf parameters of the j -th regression tree to the predictors, x .

The BART algorithm requires specifying prior information on the tree structure (T_j, M_j) and the residual variance (σ^2) to achieve a sufficient amount of regularization. Chipman et al. (1998) set the prior for T_j , which controls for the location of the interior nodes in a tree, equal to $\alpha(1 + d)^{-\beta}$, where $\alpha = 2$ and $\beta = 0.95$ imply (i) moderately-sized trees with 2 or 3 nodes having the highest likelihood, and, (ii) the leaf parameters exhibiting a modest variation. The prior on each leaf parameter is $\mu \sim N(\mu_\mu/m, \sigma_\mu^2)$, with an expected value, μ_μ , of $(y_{max} + y_{min})/2$. In order to prevent σ^2 from getting too small and to reduce the probability of overfitting, we use an Inverse Gamma (IG) distribution, $\sigma^2 \sim IG(\nu/2, \nu\lambda/2)$. The parameter λ is estimated such that there is a priori a q -percentage chance to improve the RMSE of a least squares regression. Finally, the parameters, μ_μ and σ_μ^2 , are specified such that $y_{min} = m\mu_\mu - k\sqrt{m}\sigma_\mu$ and $y_{max} = m\mu_\mu + k\sqrt{m}\sigma_\mu$, where k controls for the tightness of the prior (i.e., increasing k leads to a stronger regularization of μ_μ). In order to prevent model overfitting, we choose a conservative setup advocated by Chipman et al. (2010): $m = 200$, $k = 5$, $q = 0.75$, and $\nu = 10$. Varying the calibration of the parameters as in Chipman et al. (2010, Table 2) leaves our results qualitatively unchanged.

With the priors specified, the BART algorithm proceeds using an iterative Bayesian MCMC backfitting algorithm to sample from the posterior distribution of the sum-of-trees function given in Eq. (1). An in-depth description of the technical details of the Bayesian MCMC backfitting algorithm is given by Chipman et al. (2010). Kapelner and Bleich (2014) discuss its implementation in the add-on package “bartMachine” for the programming environment R (R Core Team 2015), which we use in this research. The algorithm works as follows:

1. Fit tree, T_j , given the other $m - 1$ trees, to the partial residuals that result when T_j is excluded. Use a Metropolis-Hastings algorithm to propose a new tree structure.
2. Given the new structure, sample from the posterior of the leaf parameters, completing the new estimate of T_j .
3. Repeat Steps 1 and 2 iteratively for all regression trees using the corresponding updates of the partial residuals.

4. Once the m -th tree has been updated, use the sum-of-trees function to compute the residuals. Use the residuals to sample from the posterior of the variance parameter, σ^2 .

Simulating Steps 1 to 4 several times gives a large number of draws from the posterior for the sum-of-trees function. We use 3,000 simulation runs of the Bayesian MCMC backfitting algorithm and discard the first 2,000 burn-in runs. Averaging over the draws gives an estimate of the sum-of-trees function (i.e., the posterior mean estimate of the EPU index for Canada), where the dispersion of the draws is used for inferential purposes.

3 Empirical Analysis

We study the EPU indexes developed by Baker et al., (2015) for the following countries: Canada (the predictand) and China, France, Germany, Japan, Russia, United Kingdom, and United States (the predictors).¹ The sample period runs from February 1997 to August 2015.² The EPU indexes are available at a monthly frequency. Like Ajmi et al., (2014), we study the natural log of the EPU indexes (which are found to be stationary), and we use demeaned data. In order to create the uncertainty index, Baker et al. (2015) search for newspaper articles containing the terms “uncertainty” or “uncertain”, the terms “economic” or “economy”, and any (monetary or fiscal) policy-related terms. For inclusion in the index, the newspaper article must include terms in all three categories pertaining to uncertainty, the economy, and policy.

The partial dependence plots depicted in Figure 1 summarize the marginal effect of a predictor on the EPU index for Canada, holding the other predictors fixed. An increase in the EPU index for the United States leads to an almost linear but less than proportional increase in the EPU index for Canada. Similarly, the international uncertainty link with the United Kingdom is linear. For the other countries in our sample, we observe asymmetric international uncertainty links. For example, the Canadian EPU does not change significantly for a negative log EPU

¹Data source: <http://www.policyuncertainty.org>. For an analysis of alternative measures of uncertainty, see Strobel (2015).

²The sample period starts with the earliest available index value for China and ends with the latest available index value for Japan.

index for China, but the Canadian uncertainty index significantly increases as the log of the Chinese index becomes positive. A similar asymmetric link can be observed for France, while the asymmetry is somewhat less pronounced for Germany, Japan, and Russia.

– Please include Figures 1 and 2 about here. –

Figure 2 summarizes the relative importance (in percent) of the predictors. Relative importance measures the average use of a predictor as a splitting variable and is computed as the mean of the average use computed across all posterior samples (Chipman 2010, Bleich et al. 2014). A key finding is that the EPU indexes for the United Kingdom and the United States are the top splitting variables. Their relative importance is estimated around 21% and 19%. This result finding corroborates results documented by Klößner and Sekkel (2015) insofar as they find, using a variant of the Diebold and Yilmaz (2009) approach, that the United Kingdom and the United States are important exporters of international uncertainty spillover effects (see also Yin and Han 2014). The relative importance of the EPU index for France assumes a value of approximately 15%. The relative importance of the EPU indexes for China, Germany, Japan, and Russia is estimated to be about 10% to 13%.

Table 1 summarizes the significance of the predictors and the overall fit of the model. The significance of a predictor is assessed using a permutation test. The null hypothesis of the permutation test is that a predictor does not have explanatory power for the predictand. The permutation tests yield significant results for the United States and the United Kingdom (<1%) and France and China (<5%). The permutation tests lead to marginally significant results for Japan and Russia (<10%) and an insignificant result for Germany (~16%). A joint permutation test for all predictors yields a highly significant result. Results of a Shapiro-Wilk test imply that normality of the residuals cannot be rejected. The BART algorithm produces a model with a reasonable model fit as demonstrated by the Pseudo- R^2 , which assumes the value 0.81.

– Please include Table 1 about here. –

For the sake of comparison, we also report results for a standard OLS model. The p -values of standard t -tests indicate that the OLS model has no problem to detect the linear effects in

the case of the United Kingdom and the United States. In contrast, the OLS model does not recover the nonlinear links in the cases of, for example, France and China. The overall fit of the OLS model falls short of the fit of the BART model.

– Please include Figure 3 about here. –

Figure 3 shows results for a recursive and a rolling-estimation window (of length 5 years). Uncertainty spillover links are estimated to be relatively stable over time. Unlike the Diebold and Yilmaz (2009) approach used in earlier research, the BART algorithm detects a moderate temporary increase in the relative importance of the United States following the 2008 financial market turmoil rather than a sharp temporary spike in international uncertainty links.

4 Concluding Remarks

We have argued that the BART algorithm is a useful technique to study international uncertainty links. Partial dependence plots show how uncertainty in one country is linked in a potentially nonlinear way to a variation in uncertainty in another country. In addition, the BART algorithm can be used to quantify the relative importance of uncertainty links, and significance testing is straightforward. While we have studied the international links of the EPU index for Canada to illustrate how the BART algorithm works, our analysis can be extended in several directions. For example, it is interesting to extend the approach to study alternative uncertainty measures and to incorporate lags of the predictors to shed light on the dynamic interplay between uncertainty measures.

References

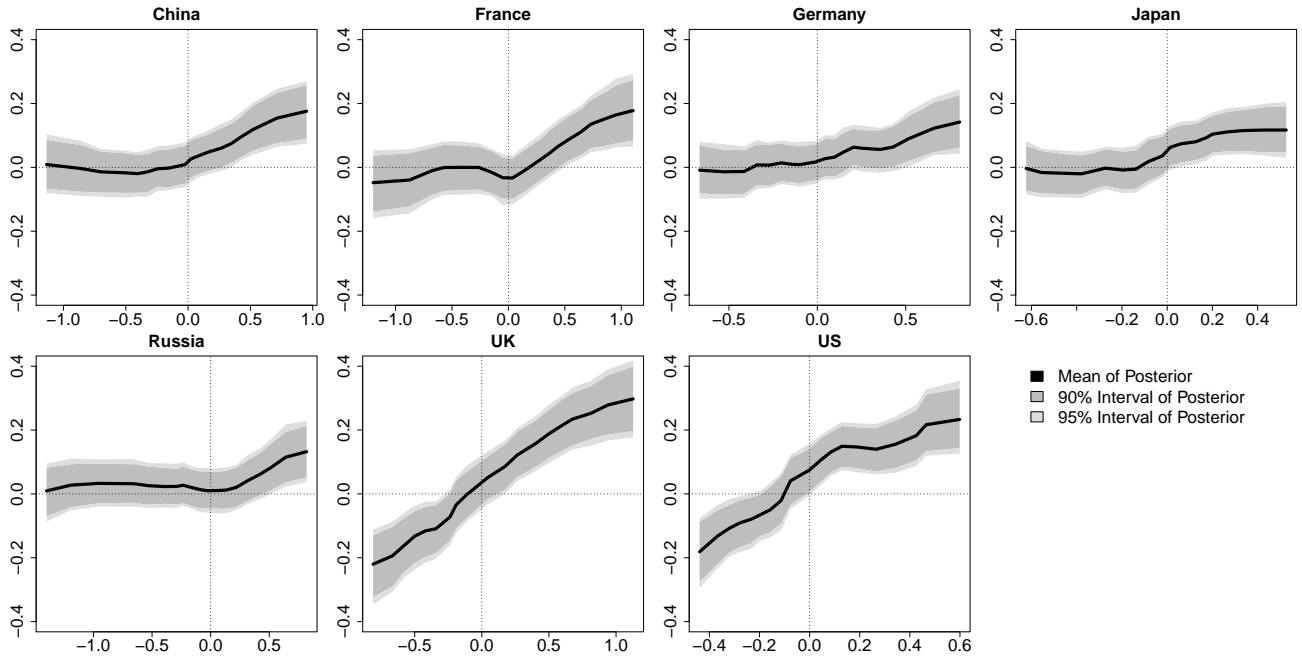
- Ajmi, A. N., Gupta, R. and P.T. Kanda (2014). Causality between economic policy uncertainty across countries: Evidence from linear and nonlinear tests. *Frontiers in Finance and Economics* 11: 73–102.
- Antonakakis, N., Chatziantoniou, I. and G. Filis (2013). Dynamic co-movements of stock market returns, implied volatility and policy uncertainty. *Economic Letters* 120: 87–92.

- Baker, S., Bloom, N. and S. Davis (2015). Measuring economic policy uncertainty. NBER Working Paper No. w21633.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica* 77, 623–685.
- Breiman, L., Friedman, J.H., Oshen, R., and C. Stone (1983). *Classification and Regression Trees*. New York: Chapman & Hall.
- Chipmann, H.A., George, E. I. and R.E. McCulloch (1998). Bayesian CART model search. *Journal of the American Statistical Association* 93: 935–960.
- Chipmann, H.A., George, E. I. and R.E. McCulloch (2010). BART: Bayesian Additive Regression Trees. *Annals of Applied Statistics* 4: 266–298.
- Colombo, V. (2013). Economic policy uncertainty in the US: Does it matter for the Euro area? *Economics Letters* 121: 39–42.
- Diebold, F.X. and K. Yilmaz (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal* 119: 158–171.
- Hastie, T., Tibshirani, R. and J. Friedman (2009). *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. Springer Science and Business Media, LLC, New York.
- Jones, P.M. and E. Olson (2013). The time-varying correlation between uncertainty, output, and inflation: Evidence from a DCC-GARCH model. *Economics Letters* 118: 33–37.
- Kapelner, A. and J. Bleich (2014). *bartMachine: Machine Learning with Bayesian Additive Regression Trees*. ArXiv e-prints.
- Klößner, S. and R. Sekkel (2014). International spillovers of policy uncertainty. *Economics Letters* 124: 508–512.
- Ko, J.-H. and C.-M. Lee (2015) International economic policy uncertainty and stock prices: Wavelet approach. *Economics Letters* 134: 118–122.
- Yin, L. and L. Han (2014). Spillovers of macroeconomic uncertainty among major economies. *Applied Economics Letters* 21: 938–944.

R Core Team (2015). R: A language and Environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>.

Strobel, J. (2015). On the different approaches of measuring uncertainty shocks. *Economics Letters* 134: 69–72.

Figure 1: Partial Dependence Plots



Notes: Intervals are based on 1,000 draws from the posterior. The calibrated parameters are $m = 200$, $k = 5$, $q = 0.75$, and $\nu = 10$.

Figure 2: Relative Importance (in Percent)





Figure 3: Recursive and Rolling Window Estimates



Notes: The relative importance over time assuming a recursive (left) and rolling (right) estimation window. The (initial) window size is 5 years. The calibrated parameters are $m = 200$, $k = 5$, $q = 0.75$, and $\nu = 10$.

Table 1: Significance of Predictors and Model Fit

Country	<i>p</i> -value BART	<i>p</i> -value OLS
China	0.0348	0.2660
France	0.0498	0.2684
Germany	0.1592	0.3409
Japan	0.0597	0.0821
Russia	0.0896	0.3242
United Kingdom	0.0000	0.0000
United States	0.0000	0.0000
Overall significance	0.0000	0.0000
Shapiro-Wilk (p-value)	0.8347	0.0652
(Pseudo-/Adjusted-)R ²	0.8130	0.7494

Notes: For the BART algorithm, p-values are computed by permuting the country-wise data vector over the full data matrix (200 times). Overall significance summarizes the result of a joint permutation test for the BART results. The null hypothesis of the Shapiro-Wilk test is that the residuals have a normal distribution. OLS-standard errors of t-tests are robust to heteroscedasticity and autocorrelation. Overall significance for the OLS model summarizes the results of an F-test.