

CAMA

Centre for Applied Macroeconomic Analysis

Incorporating Relevant Multivariate Information for Characterizing Half-Life with an Application to Purchasing Power Parity

CAMA Working Paper 47/2017
July 2017

Benjamin Wong

Reserve Bank of New Zealand and
Centre for Applied Macroeconomic Analysis, ANU

Abstract

Half-lives are summary measures of persistence, and are usually characterized from impulse response functions (IRFs) of univariate time series models. Two issues which occur with half-life characterization in multivariate time series are IRFs become conditional on specific shocks and are often also not uniquely identified. I introduce an approach for characterizing the half-life in multivariate time series models which circumvents both issues. An empirical application suggests the half-life of the real exchange rate estimated from multivariate models is generally longer relative to univariate models.

Keywords

Half-Life, Purchasing Power Parity, Multivariate Information

JEL Classification

C31, F41

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Incorporating Relevant Multivariate Information for Characterizing Half-Life with an Application to Purchasing Power Parity

Benjamin Wong*

Reserve Bank of New Zealand and Centre for Applied Macroeconomic
Analysis, The Australian National University[†]

July 17, 2017

Abstract

Half-lives are summary measures of persistence, and are usually characterized from impulse response functions (IRFs) of univariate time series models. Two issues which occur with half-life characterization in multivariate time series are IRFs become conditional on specific shocks and are often also not uniquely identified. I introduce an approach for characterizing the half-life in multivariate time series models which circumvents both issues. An empirical application suggests the half-life of the real exchange rate estimated from multivariate models is generally longer relative to univariate models.

JEL Classification: C31, F41

Keywords: Half-Life, Purchasing Power Parity, Multivariate Information

*The views do not necessarily reflect those of the Reserve Bank of New Zealand. I thank Yu-Chin Chen, Charles Engel, James Hamilton, Punnoose Jacob, Güneş Kamber, Leo Krippner, John McDermott, James Morley, Luis Uzeda, Tugrul Vehbi, and James Yetman and seminar participants at the Bank for International Settlements Asian Office, and participants of the RBNZ Morning Metrics Workshop for helpful comments and discussion. Any remaining errors are mine.

[†]Reserve Bank of New Zealand. 2 The Terrace, Wellington 6011, New Zealand. Email: benjamin.wong@rbnz.govt.nz Tel: +64 44713957

1 Introduction

Persistence profiles are a key facet in learning about the dynamic adjustment of macroeconomic variables. The real exchange rate literature stands out for its substantial attention to the characterization of persistence through the calculation of half-lives. The half-life is the time it takes for half the effect of a shock to dissipate. It is a summary measure of persistence which quantifies the speed of mean reversion, where the estimation of the half-life of the real exchange rates is used to test how quickly Purchasing Power Parity (PPP) is restored. Characterization of the half-life of the real exchange rate occurs almost exclusively through the lens of univariate models, mostly through estimating short-order autoregressive (AR) processes (e.g. Cheung and Lai, 2000; Kilian and Zha, 2002; Rossi, 2005; Chortareas and Kapetanios, 2013; Lo and Morley, 2015). A broad consensus is that the half-life of the real exchange rate is roughly between three to five years, an estimate which is acknowledged as “puzzlingly high” (see, e.g. Rogoff, 1996).

The contribution of this paper is to present methods to characterize the half-life within a multivariate framework, thus expanding the tools available for measuring persistence. There are theoretical and consistency reasons for why one might characterize persistence in multivariate frameworks despite the widespread practice of half-life estimation within univariate environments. The proliferation of tools such as Vector Autoregressions (VAR) and Dynamic Stochastic General Equilibrium (DSGE) models at least suggests that macroeconomists regard joint determination of variables as a useful way to think about macroeconomic fluctuations. Even so, it appears common practice to write down general equilibrium models to match univariate time series properties of the real exchange rate, often in the form of autocorrelations or the univariate impulse response functions (e.g. Steinsson, 2008; Carvalho and Nechio, 2011). This represents a form of cognitive dissonance, given it is unclear how to reconcile a multivariate data generating process with empirical evidence of real exchange rate persistence from univariate processes. At the same time, models like Eichenbaum and Evans (1995) and Ng (2003) study the persistence of the real exchange rate within SVARs which are *conditional* upon different identified shocks. It is once again not clear how one should reconcile *conditional* characterization of persistence from say, an identified monetary shock from an SVAR, to the persistence profile from univariate models, which may be closer to unconditional measures. Therefore, at a minimum, the development of tools to characterize persistence in multivariate environments is a step in aiding reconciliation of theory with data. To the extent that theory or the assumed underlying data generating process is often from the perspective of a multivariate process like a VAR or DSGE model, empirical validation through multivariate models is likely to provide a more consistent basis to makes comparisons.

Characterizing half-life through multivariate models is however less straightforward because discussion inevitably center about the impulse response function (IRF) (e.g Che-

ung and Lai, 2000; Chortareas and Kapetanios, 2013). The focus on IRFs in multivariate environments such as an N variable VAR is less helpful as one can identify IRFs for up to N shocks. Not only are there multiple IRFs in multivariate models, these IRFs are also conditional. Persistence is inherently a summary of the autocovariance function (ACF). While the distinction between the IRF and ACF can be easily ignored in univariate AR models, this can cause confusion if we consider multivariate models. In particular, viewing the problem through the IRF means multiple IRFs in multivariate models may be an impediment to calculating half-life, but a single variable within a multivariate system still only possess a *single* ACF. Therefore, focussing one's attention on some characterization of the ACF, rather than the IRF, greatly clarifies how one can measure the half-life in multivariate environments.

I present an empirical example by calculating the half-life of the real exchange rate of the G7, Australia and New Zealand relative to the U.S. In general, multivariate information increases the estimated half-life of the real exchange rate. A key reason for this empirical finding is because multivariate information contains information about real exchange rate fluctuations, and this is confirmed by Granger causality tests.

A Monte Carlo exercise utilizing a standard two country DSGE model as a data generating process lends support to the empirical results. In particular, the results of the Monte Carlo experiment suggest the omission of multivariate information is a possible underlying cause of underestimating the half-life for the univariate model, consistent with the empirical finding. Even allowing for the usual finite sample issues with estimating real exchange rate half-lives, the Monte Carlo experiment suggests incorporating relevant multivariate information results in *less* biased estimates of the half-life of the real exchange rate. A large sample Monte Carlo simulation suggests that the approach I am proposing is also much closer to uncovering the true persistence in population.

The rest of this paper is organized as follows. In Section 2, I will first describe the proposed approach. Section 3 presents an empirical application estimating the half-life of the real exchange rate for the G7, Australia and New Zealand. Section 4 presents a Monte Carlo simulation using a standard two country DSGE model to lend some context to the empirical results. Section 5 briefly notes how the ideas developed can be easily extended into a fully identified VAR. The final section provides some concluding remarks.

2 A Multivariate Approach to Characterizing Half-Life

Let y_t be a vector of N variables. Suppressing the constant without any loss of generality, a Vector Autoregression (VAR) of lag order p can be written as

$$y_t = A_1 y_{t-1} + A_2 y_{t-2} + \dots + A_p y_{t-p} + \nu_t, \quad (1)$$

where ν_t is a vector of mean zero innovations, $\mathbb{E}[\nu_t' \nu_t] = \Sigma_\nu$ and $\mathbb{E}[\nu_t' \nu_{t-i}] = 0 \forall i > 0$. Given the half-life literature typically makes the argument from the perspective of an impulse response function (IRF), it is more convenient for the exposition to work with the moving average (MA) representation. Assuming the roots of the VAR are invertible, the Wold representation theorem states that the VAR in Equation (1) has a infinite MA representation

$$y_t = B_0 \nu_t + B_1 \nu_{t-1} + B_2 \nu_{t-2} + \dots \quad (2)$$

Equation (2) is generally not sufficient to pin down the IRF because of correlation across ν_t . To get independent variation, let ϵ_t is a vector of N independent innovations with covariance matrix I_N . Suppose there is a matrix C where $\Sigma_\nu = CC'$, then $\nu_t = C\epsilon_t$. Substituting into Equation (2),

$$y_t = B_0 C \epsilon_t + B_1 C \epsilon_{t-1} + B_2 C \epsilon_{t-2} + \dots \quad (3)$$

Let $\phi(i)$ be the IRF at horizon i for a one unit shock to ϵ_t , the IRF at horizon i can be written as

$$\phi(i) = B_i C. \quad (4)$$

If one can identify ϵ_t , then one can calculate the IRF, and all other characterizations of IRF such as half-lives in typical set ups. The problem arises that while C is uniquely pinned down in the univariate case, it is generally not in the multivariate case. To see this, let Q be an $N \times N$ orthonormal matrix where $QQ' = I_N$, and $\tilde{C} = CQ$. It is then trivial to verify $\Sigma_\nu = \tilde{C}\tilde{C}'$. So long as Q is not the identity matrix, both C and \tilde{C} are valid, yet different, orthogonalizations to obtain ϵ_t . In the univariate case when $N = 1$, $Q = 1$ and thus C is uniquely pinned down (i.e $C = \Sigma_\nu^{\frac{1}{2}}$). Therefore, the problem only arises because C is generally not unique in the multivariate setting, or when $N > 1$. In the multivariate case, because C is not uniquely pinned down, the likelihood function cannot uniquely pin down the IRF. On the other hand, in the univariate case, because C is uniquely identified, knowledge of the likelihood is sufficient to pin down the IRF, and other quantities of interest, such as half-lives.

Yet, the uniqueness of the orthogonalization is not the only concern. Even if there is agreement to how to identify ϵ_t , each variable has N sets of IRFs, one for each of the N shocks. In other words, each of the IRFs is conditional on each of the N shocks, and so any characterizations of the IRF like half-lives, are similarly conditional. Once again, when $N = 1$, this is not an issue as the unconditional IRF coincides with the conditional one given there is only a single set of IRF. But when $N > 1$, it is not obvious which set of the IRFs best represents the half-life of the variable of interest. At times, the half-life conditional of different shocks in a multivariate framework is of interest (e.g. Ng, 2003). In other situations, this is less satisfactory as it would appear overly restrictive to only be unable to calculate the unconditional half-life. Therefore, even comparing half-life profiles across univariate and multivariate specifications is not natural given the conditionality under which these profiles are derived do not necessarily coincide.

From this perspective, it is hardly surprising that univariate testing of half-life is so prevalent in the wider literature given such tasks are much easier to handle conceptually in the univariate, relative to a multivariate, framework. To the extent that modern macroeconomic theory write down general equilibrium models which emphasize interconnection and joint determination of many variables makes univariate testing less satisfactory and somewhat inconsistent.

It turns out that while the IRF in a multivariate system is not invariant to how a researcher orthogonalizes Σ_ν , the variance of the IRF is invariant. This can be easily verified by taking the variance of both the left and right hand side of the Wold form in Equation (3),

$$\begin{aligned} \text{Var}(y_t) &= B_0 C \epsilon_t \epsilon_t' C' B_0' + B_1 C \epsilon_{t-1} \epsilon_{t-1}' C' B_1' + B_2 C \epsilon_{t-2} \epsilon_{t-2}' C' B_2' + \dots \\ &= B_0 \Sigma_\nu B_0' + B_1 \Sigma_\nu B_1' + B_2 \Sigma_\nu B_2' + \dots \end{aligned} \quad (5)$$

This result suggests that a summary persistence measure, such as the half-life, derived from working off the variance will be invariant to any arbitrary orthogonalization of Σ_ν , because one can entirely sidestep the issue of orthogonalising Σ_ν . Taking the square of (4), we can write the square of the IRF as

$$\phi(i)^2 = B_i \Sigma_\nu B_i'. \quad (6)$$

The cumulative *squared* IRF at horizon i , $\Phi(i)$, can be expressed as

$$\begin{aligned} \Phi(i) &= B_0 \Sigma_\nu B_0' + B_1 \Sigma_\nu B_1' + B_2 \Sigma_\nu B_2' + \dots + B_i \Sigma_\nu B_i' \\ &= \phi(0)^2 + \phi(1)^2 + \phi(2)^2 + \dots + \phi(i)^2 \end{aligned} \quad (7)$$

Note that as the horizon approaches infinity, the expression for the cumulative sum of the squares of the IRFs, $\lim_{i \rightarrow \infty} \Phi(i)$, is the unconditional covariance matrix, and coincides with the expression in Equation (5). So long as the real exchange rate is covariance stationary, the expression is summable (i.e. $\sum_{i=0}^{\infty} \phi(i)^2 < \infty$). The above discussion suggests working with the variance may yield a possible characterization of the half-life in a multivariate framework. A related approach suggested by Chortareas and Kapetanios (2013) adopts an alternative definition as the half-life as being the time it takes for half the total volatility to subside. Let the real exchange rate be the k^{th} variable in the vector y_t and e_i be a selector vector of zeros apart from 1 as its i^{th} element. Adopting an alternative definition of half-life as the time taken for half the total volatility to subside, the cumulative volatility (CuVo) half-life, denoted h_{CuVo}^* , will satisfy

$$e_k \Phi(h_{CuVo}^*) e_k' = \frac{\lim_{i \rightarrow \infty} e_k \Phi(i) e_k'}{2}.$$

To allow for non-integer values for h_{CuVo}^* , we can instead evaluate¹

$$\int_0^{h_{CuVo}^*} e_k \phi(i)^2 e_k' di = \int_{h_{CuVo}^*}^{\infty} e_k \phi(i)^2 e_k' di. \quad (8)$$

The CuVo half-life sidesteps the issue of identifying shocks and only uses information contained in the likelihood function (i.e. B and Σ_ν) without any auxiliary identification assumptions.² Notice that as we focus attention on a single element, namely the cumulative volatility of the k^{th} variable, the issue of *which* IRF to focus on is sidestepped as well. Therefore, the CuVo half-life provides an unconditional measure of persistence of any variable of interest within the multivariate system.

Relation to Existing Work

One interpretation is that the CuVo half-life is a multivariate extension and generalization of the cumulative volatility approach in AR models suggested by Chortareas and Kapetanios (2013). In univariate environment, the CuVo half-life is *exactly* as suggested by Chortareas and Kapetanios (2013). The link with Chortareas and Kapetanios is that the object of interest should be characterizing an object closer to the autocovariance function (ACF), rather than the IRF. The discussion by Chortareas and Kapetanios occurs

¹I follow Chortareas and Kapetanios (2013) to use linear interpolation between integer values of i to work out the non-integer values of the half-life.

²A reader contrasting with the earlier univariate literature may wonder if it is possible to characterize the persistence of the real exchange rate by considering the eigenvalues of the companion matrix (e.g., see Rossi, 2005). While the eigenvalues entirely characterizes the persistence of the variable in question in a univariate system, this is not necessarily true in the multivariate system. The eigenvalues in the multivariate system are influenced by all the variables in the system, and the largest eigenvalue may *not* be associated with the variable of interest. For sake of argument, suppose in a bivariate VAR, one variable is more persistent, while the other exhibits more white noise behavior. In this example, we learn little about the persistence of the less persistent variable by characterizing the largest eigenvalue.

through viewing the half-life through the IRF, like much of the half-life literature. This distinction is not crucial because Chortareas and Kapetanios (2013) work within univariate AR frameworks. However, recognizing the distinction between the ACF and the IRF is important in allowing the generalization of Chortareas and Kapetanios (2013) results into a multivariate setting. In particular, given the structure presented in Equation (1) and Equation (2), the i^{th} order ACF of the k^{th} variable, which we denote $\vartheta_k(i)$ is

$$\vartheta_k(i) = e_k B_i e_k' V.$$

where V is the unconditional variance of the variable of interest.³ To see the distinction between a characterization of the ACF or the IRF, recall from Equation (4) that the IRF, $\phi(i) = B_i C$. Therefore, we can rewrite the ACF as

$$\begin{aligned} \vartheta_k(i) &= e_k B_i C C^{-1} e_k' V \\ &= e_k \phi(i) C^{-1} e_k' V \end{aligned}$$

In the univariate case, recall $C = \Sigma_\nu^{1/2}$, which is a scalar. Therefore, in the univariate case, the distinction between $\vartheta_k(i)$ and $\phi(i)$ only differs by a constant scaling factor, $V^{-1} \Sigma_\nu^{1/2}$. This is why the distinction between characterizing the IRF or ACF in a univariate environment is a matter of semantics. As should be clear by now, in the multivariate case, the identification of C matters, and therefore the distinction between the ACF and the IRF matters. The CuVo half-life, has a tighter link to the ACF through its evaluation of sequences of $\vartheta_k(i)^2 V^{-1} \Sigma_\nu(k, k)$, where $\Sigma_\nu(k, k)$ is the (k, k) element of Σ_ν .⁴

Recognizing that characterizing the half-life is a summary statement about the ACF entirely sidesteps the issues about IRFs in multivariate models. While an N variable multivariate model has N sets of IRFs, it still only has a *single* ACF per variable. The CuVo half-life is analogous to evaluating the squared ACF, normalized by the inverse of the variance of the variable of interest. Therefore, a subtle shift of the discussion from the IRF to the ACF greatly clarifies the object of interest.

³We can calculate V if we rewrite the VAR in companion form, $y_t = F y_{t-1} + \zeta_t$, where $F = \begin{bmatrix} A_1 & A_2 & \dots & A_p \\ I & 0 & 0 & 0 \\ 0 & \ddots & \vdots & 0 \\ 0 & \dots & I & 0 \end{bmatrix}$ and $Q = \begin{bmatrix} \Sigma_\nu & 0 \\ 0 & 0 \end{bmatrix}$, $vec(R) = [I - F \otimes F]^{-1} vec(Q)$, $V = e_k diag(R) e_k'$.

⁴It would seem equally valid to redefine the half-life as the cumulative standard deviation, given standard deviations are in the units of the variable of interest and still retain the link to the ACF. The resulting cumulative standard deviation half-life will be a factor of $\sqrt{\frac{1}{2}}$ or approximately seven-tenths of the CuVo half-life. I will not pursue this further given such a scaling can be trivially implemented, as well as the cumulative volatility has tighter links to existing work by Cogley, Primiceri, and Sargent (2010) and Chortareas and Kapetanios (2013). Moreover, an adjustment which I will describe later, further detailed in Footnote 5, will make this distinction between cumulative standard deviation and cumulative volatility a second order issue. I thank James Hamilton for pointing out this possibility.

The CuVo half-life is also related to work on measuring predictability or forecastability (e.g., see Diebold and Kilian, 2001; Cogley, Primiceri, and Sargent, 2010). In particular, let i and j be forecast horizons where $j \gg i$. The idea is that if a variable is more persistent, then the degree of predictability at forecast horizon i will be more dissimilar to the predictability at horizon j . Using notation that I have already introduced, an example is Cogley, Primiceri, and Sargent's R^2 persistence statistic after i periods, which we can view as a special case of Diebold and Kilian's predictability statistic when $j = \infty$ and variable k is stationary,

$$R^2(i) = 1 - \frac{e_k \Phi(i) e_k'}{e_k \Phi(\infty) e_k'}$$

where $R^2(i)$ measures the share of the forecastable variation in variable k after i periods. $R^2(\infty) = 0$ by construction since in the infinitely long future, none of the variation of variable k is forecastable. The CuVo half-life serves as an alternative reframing of these predictability and forecastability measures where we calculate how large should i be before the statistic attains a value of one-half or $R^2(h_{CuVo}^*) = 0.5$. This reframing to a half-life measure provides an interpretation where we relate the ideas of the degree of persistence of a variable of interest to degree of predictability.

Uniqueness of the CuVo Half-Life

The definition of the traditional half-life is the time it takes for half the effect of a shock has decayed. Let $\hat{\rho}_{AR1}$ is an estimated coefficient from an AR(1) regression of the real exchange rate, it is well known that the traditional half-life, h_{THL}^* can be calculated

$$h_{THL}^* = \frac{\log(0.5)}{\log(\hat{\rho}_{AR1})}. \quad (9)$$

Calculation of traditional half-lives from an AR(1) holds intuitive appeal because decay is monotonic, therefore solution to the half-life is unique. Even within univariate environments, the calculation of half-life using higher order AR(p) models poses conceptual challenges because decay is not guaranteed to be monotonic, and empirical analysis often does reveal non-monotonic decay. Non-monotonicity compromises the uniqueness of the half-life, which partly explains the appeal of working with an AR(1). Some examples of such non-monotonic decay are hump-shaped IRFs (see Cheung and Lai, 2000) and also multiple solutions as the IRF attains half the value of an initial shock multiple times. Therefore, the traditional measure of half-life has a surfeit of issues to resolve once we move away from an AR(1), and these are issues that are inherent even before one considers the issue of how to incorporate multivariate information. Part of Chortareas and Kapetanios's motivation of suggesting working with the cumulative volatility is precisely to work with a monotonic object, and thus be able to calculate a unique half-life.

Naturally, being also a multivariate generalization of their approach, the CuVo half-life also similarly obviates the issues related to non-monotonicity. This can be verified by observing Equation (7). Given that the diagonals of the matrices are variances and so strictly non-negative, $e_k\Phi(i)e'_k$ is monotonically increasing in i , and therefore solution to the CuVo half-life is unique.

An AR(1) Normalization

The CuVo half-life, however, does require a definitional change of the half-life before one can compare the CuVo half-life with the wider literature. While there may be compelling reasons to adopt a definitional shift given the traditional half-life struggles to deal with non-monotonic processes as previously discussed, there is no common ground to do comparisons with the wider literature if there is no shift in the definition of the traditional half-life. This is especially true if one is interested in comparing against the three to five year consensus described by Rogoff (1996), where the vast majority of estimates of this consensus occurs almost exclusively through the definition of the traditional half-life.

I thus propose a possible normalization of the CuVo half-life into a summary statistic which can be made comparable to the traditional half-life, especially if one is not prepared to compromise on a definitional change. The starting point of such a normalization is to first recognize that there is at least agreement on how to characterize a traditional half-life from an AR(1), which as per Equation (9), is the familiar $h_{THL}^* = \log(0.5)/\log(\hat{\rho}_{AR1})$. It is also not difficult to verify that for a given $\hat{\rho}_{AR1}$, the traditional half-life is longer than the CuVo half-life.

One suggestion to normalize the CuVo half-life is to recognize for a given AR(1) regression with parameter $\hat{\rho}$, this process has a unique CuVo half-life estimate due to the monotonicity of the cumulative variance. Denote this estimate as $h_{CuVo}^*(\hat{\rho})$. Suppose we wanted to adjust a CuVo half-life estimate from an alternative model. This model may be a AR process with richer dynamics or a multivariate process. Denote this $h_{CuVo}^*(\hat{\phi})$. Given $h_{CuVo}^*(\hat{\phi})$, we can solve for $\hat{\rho}$ where

$$h_{CuVo}^*(\hat{\phi}) = h_{CuVo}^*(\hat{\rho}). \quad (10)$$

We can then use $\hat{\rho}$ by appealing to the traditional half-life estimate from an AR(1) where we can convert the CuVo half-life to a cumulative volatility adjusted (CuVoA) half-life where

$$h_{CuVoA}^* = \frac{\log(0.5)}{\log(\hat{\rho})}. \quad (11)$$

The CuVoA half-life adjustment “converts” the CuVo half-life estimate into “units” of the traditional half-life. This will facilitate comparison with the wider empirical litera-

ture, especially is one was not prepared to accept the definitional change of the cumulative volatility. The CuVoA half-life also retains all the properties of uniqueness and monotonicity of the CuVo half-life given we are mapping relative to an AR(1) process.⁵

3 An Application to PPP

For the empirical application, I estimate half-lives of the real exchange rate for the G7 countries, Australia and New Zealand. All real exchange rates are expressed relative to the U.S. dollar. They are constructed as the ratio of the U.S. CPI relative to the domestic CPI adjusted for the domestic to U.S dollar exchange rate. I consider three additional variables; relative inflation, relative interest rates, and relative real GDP growth rates to the U.S.⁶

Generally, most of the countries have samples that starts in the 1960s, with the earliest in 1960Q2 and all ending in 2017Q1.⁷ Whenever possible, I try to use the longest possible sample. While this poses few issues with GDP and CPI data, there is often not a long continuous span of time series for interest rates. I generally use a short term treasury bill that has a long history, and if needed, use a relationship with an alternative interest rate to backdate the sample as far as possible. Most of the CPI and interest rate data are from the International Financial Statistics database. The GDP data is generally from the World Economic Outlook database. I leave the specifics of data sources, coverage, and construction to the online appendix.

For each of the countries, I estimate both univariate AR models and four variable VARs to characterize the half-life of the real exchange rate using the CuVo half-life procedure which I introduced, as well as the CuVoA half-life. I will produce two sets of results. In one set, I use the Akaike Information Criterion (AIC) to select the lag order. For the AR models, the AIC chooses two lags for every country except Japan, which it chooses one lag. For the four variable VAR specifications, the AIC chooses one lag for every country. The selection of one lag for the VAR models is of concern as there is evidence that criteria like AIC may pick too low a lag order with misspecified models. Therefore, I report another set of results where I will uniformly use four lags for all the AR and VAR models, with the choice of four informed by the general rule of thumb for quarterly data.⁸ I do not

⁵ The CuVoA half-life also solves a slight issue of the CuVo measure described in Footnote 4, namely whether to construct the persistence measure based on the variance or standard deviation. Normalizing relative to a traditional half-life's AR(1) means that whether one uses the cumulative standard deviation or cumulative volatility, both map into the same CuVoA half-life as implied by a traditional half-life AR(1), given all three produce unique solutions relative to their respective definitions of a half-life.

⁶While it is less clear whether to use relative variables, or just the raw variable for each country, the use of relative variables for the empirical analysis decreases the number of variables in the VAR from seven to four, thus preserving precious degrees of freedom within the VAR system.

⁷The sample start dates are as follows. Australia: 1960Q2, Canada: 1962Q1, France: 1964Q1, Germany: 1970Q2, Italy: 1960Q2, Japan: 1960Q2, New Zealand: 1964Q1, and United Kingdom: 1960Q2.

⁸Wiriyawit and Wong (2016) and Kamber, Morley, and Wong (2017) show Monte Carlo evidence

conduct unit root tests given the real exchange rate needs to be stationary for the half-life to exist. Perhaps the most important issue uncovered by the prior literature, but relevant in the current context, is that unit root tests are based entirely on a univariate AR(p) test specification. It is not entirely clear how to reconcile conclusions from univariate tests to the multivariate specifications for which I will be conducting the empirical application. Moreover, it is known that there are a multitude of issues with regards to reconciling unit root tests with the existence of a finite half-life (see, e.g. Engel, 2000). I will therefore regard the exchange rate as mean-reverting for the empirical analysis.

Bias-Correction It is known that OLS estimates from autoregressive process are downward biased, and this bias becomes more serious the closer the process is to non-stationarity. This has emerged as a particularly important issue for half-life estimation because most estimated real exchange rate dynamics are close to being non-stationary, and so these biases have a quantifiable impact on half-life estimates. While there are methods to correct for the bias in AR models (e.g. Andrews and Chen, 1994), and some variants of such bias adjustments have been applied in the estimation of half-lives (e.g. Murray and Papell, 2002; Cashin and McDermott, 2003), it is not entirely clear how one should proceed in a multivariate settings. In my empirical application, I will utilize the “bootstrap-after-bootstrap” procedure introduced by Kilian (1998) in order to correct for possible biases and construct confidence intervals for the half-life estimates. In my application, I will utilize the first bootstrap to approximate the bias and use this approximation to subsequently bias-correct the half-life estimates.⁹ The second bootstrap will be applied in order to construct confidence intervals. To make the confidence intervals and bias-corrected estimates comparable between the univariate and multivariate models, I apply the same bootstrapping and bias-correction procedure to both the AR and VAR models.

Empirical Results

CuVo Half-life Estimates Figures 1 and 2 present the estimated half-lives, with the lag order chosen by AIC and four lags respectively. The estimated half-life from the OLS coefficient is indicated with a cross, and the bias corrected estimate is indicated with a solid circle. The whiskers indicate the width of the 90% confidence interval. If any of the point estimates, or the upper whiskers are not presented, this indicates that the estimated half-life is too large to be plotted on the chosen scale, and should be

where the true model is not in the set choice models due to issues like possible model misspecification or moving average terms, lag order test may revert to choosing a very low lag order AR or VAR model. Choosing rule of thumb lag orders, such as four for quarterly data, may then sometimes be a superior alternative to using information criterion.

⁹The online appendix contains more specific details of the bootstrapping algorithm.

interpreted as being close to infinity, or non-mean reverting.¹⁰ There is a large degree of heterogeneity of half-life estimates across all the countries in the sample. Canada and Japan stand out as having very long estimated real exchange rate half-life, though this is less surprising given similar findings have previously been reported elsewhere (e.g., see Obstfeld and Rogoff, 2001; Kilian and Zha, 2002; Murray and Papell, 2002). We can note a number of similarities. In general, the estimated CuVo half-life from using a rule of thumb four lags appears to be slightly shorter than their counterpart estimate using AIC, which is always estimated with a lag order shorter than four. This suggest higher order dynamics appear important in explaining mean reversion behaviour of the real exchange rate. A striking observation, and perhaps the most important result within the context of this paper, is that most of the estimated CuVo half-lives, especially after applying bias-correction, are longer for the multivariate VAR model relative to the univariate AR model. The confidence interval also appears to be wider for estimates from multivariate models relative to the univariate AR models. The uncertainty around these half-life estimates, while extremely wide, is not surprising and entirely consistent with the wider literature (see Murray and Papell, 2002; Rossi, 2005, for specific examples). The key observation to note is that relevant multivariate information appears to increase the estimated half-life relative to a univariate model.

CuVoA Half-life Estimates We turn to estimates of the CuVoA half-life. As described earlier, the purpose of implementing the adjustment is mainly to reconcile with the definition of the half-life in the wider literature through a normalization to the AR(1) of a traditional half-life. A key reconciliation is with the three to five years “remarkable consensus” of the wider real exchange rate literature (see Rogoff, 1996) without necessarily invoking a change of definition that the CuVo half-life entails. Figures 3 and 4 presents these results. Given the CuVoA half-life is just a transformation of the CuVo half-life, it should come as no surprise that Canada and Japan once again stand out for very long half-lives. It is also unsurprising that multivariate information leads to a larger CuVoA half-life estimate relative to a univariate model, given these once again reflects the results we already know from the CuVo half-life. With the four lags specification, the multivariate half-lives with the exception of Japan and Canada, while larger than their univariate counterpart, are still largely within the three to five year consensus. This is not true from the models which consider AIC lag selection. Once again, excluding Canada and Japan, four of the six remaining countries have a multivariate half-life estimate which is slightly

¹⁰To preserve the readability of the Figure, I use a cut-off of 25 years if the confidence interval and point estimates cannot be presented on the same scale. Note this means we regard even finite half-life like hundreds or thousands of years as being non-mean reverting. To give some idea about the magnitudes involved, an AR(1) model will require an AR parameter of 0.993 to yield a traditional half-life of 25 years with quarterly data. With the CuVo half-life, an AR(1) model will need a parameter of 0.997 to correspond with a half-life of 25 years.

higher than the consensus, as they are over five years, but under six years. However, multivariate information also substantially increases the estimated CuVoA half-life estimate for Australia and Italy, suggesting a half-life of about thirteen and eight years respectively. In general, once we consider multivariate information without a definitional shift, reconciliation with the “remarkable consensus” is more mixed. With a rule of thumb four lags, the results appear more consistent with this consensus range, though this consensus appears to dissipate once we consider lag selection using AIC.

Multivariate Information and Reconciliation to the “Remarkable Consensus”

It is important to stress that whether one uses the CuVo or CuVoA half-life estimate does not alter the empirical results that multivariate information is important. The extent of which whether multivariate information matters for this “remarkable consensus” may be definitional. If one can accept that we can compare the CuVo half-life with the wider literature, then a general conclusion is multivariate information often increases the half-life estimate into the consensus range, undoing Chortareas and Kapetanios’ conclusion of a lower half-life based on a definition change as they only estimated univariate models. If one can only accept a comparison on the basis that the traditional half-life definition is the right one, then it appears that the estimated CuVoA half-lives might be higher than the consensus range, depending on model specification. Regardless of one’s persuasion, it is clear that conclusions regarding half-life estimates may be altered based on multivariate information, and ideas developed within this paper may be an important step to reconciling these results.

The Role of Bias-Correction To better understand the role of bias-correction on persistence, we can study the largest eigenvalue within the estimated AR and VAR system, with the caveat that the largest eigenvalue may not necessarily be driven by the real exchange rate, and so such analysis is merely indicative. Figure 5 plots the largest eigenvalue from the estimated models in the empirical analysis. Models estimated using four lags appear to have less persistence, given the modulus of the dominant eigenvalue appears smaller, consistent with our previous analysis that higher order model are more likely to evidence quicker mean reversion. While bias-correction will by construction result in a larger estimated half-life, we observe, especially in specifications which use AIC to conduct lag order selection, a number of instances that the *bias-uncorrected* multivariate model evidences more persistence than the *bias-corrected* univariate counterpart. On some level, this suggests that even bias-correcting the univariate AR model may not be sufficient on its own in producing more persistence relative to considering relevant multivariate information. More generally, the analysis of the largest eigenvalue suggest that the univariate AR models, perhaps through omission of relevant multivariate information, are much less persistent than the VAR models, and these difference in persistence are further

exacerbated when we correct for the bias in the least squares estimates. These results are fully consistent with the previous empirical analysis.

Granger causality Given the preceding exercise suggests multivariate information as being important for the estimation of the half-life, we delve deeper into the content of the multivariate information. A straightforward test for whether multivariate information is relevant is to conduct a Granger causality test. In particular, if variables Granger cause the real exchange rate, then we can conclude that the data supports that the multivariate information is relevant. A related consideration is Granger causality is only a sufficient, but not necessary, condition for whether a variable has relevant informational content to forecast the real exchange rate. For example, a variable may not Granger cause the real exchange rate, but Granger causes a variable which Granger causes the real exchange rate. Such a variable in question still has relevant multivariate information for forecasting the real exchange rate and this has been variously described as Granger causal prior (e.g., see Jarociński and Maćkowiak, 2017), or “indirect” Granger causality (see, e.g. Dufour and Renault, 1998), and is a much more challenging condition to establish and test for relative to Granger causality. I thus test conduct standard Granger causality test, as well as the long horizon Granger causality tests introduced by Dufour, Pelletier, and Renault (2006), with the latter used as an attempt to understand if such indirect and complex forms of relevant multivariate information exist.¹¹ Table 1 presents the results of the Granger causality tests. Once again, the lag order is selected either through AIC or by applying a uniform rule of thumb of four lags. If the null of a variable not Granger causing the real exchange rate can be rejected, the appropriate entry in the Table is marked. When we consider the standard Granger causality tests, it appears apart from Canada and Italy, at least one variable does Granger cause the real exchange rate. When we apply the test of Dufour, Pelletier, and Renault (2006) to test for “indirect” causality, the evidence for whether the multivariate information is relevant is much stronger. In particular, when we consider the four lag test specification, we often find two out of the three variables as being relevant multivariate information for the real exchange rate. On the basis of the Granger causality tests, it would suggest that the longer estimated half-lives from the multivariate models are because there is genuine multivariate information which has been incorporated into the half-life characterization. The Granger causality tests provide strong evidence that the considered multivariate information is empirically important. Multivariate information should therefore be part of the characterization of the persistence of the real exchange rate because they are empirically relevant.

¹¹The test by Dufour, Pelletier, and Renault (2006) are largely standard Granger causality test, but instead of regressing against y_t , regress against y_{t+h} to establish if there is Granger causality at horizon h . I test for this form of “long horizon” Granger causality from $h = 1, 2 \dots 6$, and conclude there is “indirect” causality if the null can be rejected for at least one of the tested horizons.

4 Monte Carlo Exercise

So far, the paper has presented methods of estimating half-life in multivariate systems. The empirical application of PPP for a group of G7 countries, Australia and New Zealand generally provides support that it is important to consider multivariate information when studying the persistence of the real exchange rate. I now conduct a Monte Carlo exercise to better understand the empirical results.

To conduct the Monte Carlo exercise, I choose a standard two country open economy Dynamic Stochastic General Equilibrium (DSGE) model as introduced by Lubik and Schorfheide (2006) as the Data Generating Process (DGP). The model featuring optimizing households, monopolistically competitive firms and a monetary and fiscal authority in each country. The DGP is parametrized using the posterior mode of the estimated model which is estimated using Euro Area and U.S. data on the sample used by Lubik and Schorfheide (2006).¹² The choice DSGE model is natural given many of its features are ubiquitous in a variety of open economy DSGE models (e.g. Bergin, 2006; Jacob and Peersman, 2013). Moreover, given that Lubik and Schorfheide (2006) estimate their model on Euro Area-U.S. data, the parameterization at least provides simulated data which match observed time series dynamics. Model details are as per Lubik and Schorfheide (2006) (pages 318-326), with the exact parametrization of the DGP in the online appendix.

For each Monte Carlo sample of simulated data, I generate seven observables from the DSGE model; domestic and foreign output growth, domestic and foreign inflation rates, domestic and foreign interest rates, and the real exchange rate. I then fit both an AR model on the real exchange rate and a VAR on the seven variables to study the role of multivariate information. I use the linearized solution of the DSGE model as the DGP, so that the underlying DGP can be well approximated by a linear model.¹³ An important issue to note is that the underlying DGP has a VARMA representation as its reduced form. Therefore, while the AR model is misspecified due to omitted multivariate information, the estimated VAR may also be misspecified because it does not feature MA terms. Nonetheless, it is known that the omitted state variables may simply result in a re-parameterized model of fewer variables, where a potential ARMA structure or long order AR model may be able to approximate the underlying DGP (e.g., see Wallis, 1977; Kapetanios, Pagan, and Scott, 2007). This point suggests that we should not immediately condemn the AR model to failure, given the theoretical possibility exists that a very high order AR model may be able to approximate the multivariate DGP.

¹²I used the Dynare code on Frank Schorfheide's website to estimate the model and subsequently simulated the model parameterized using the posterior mode. I thank both Frank Schorfheide and Thomas Lubik for making the code available.

¹³Ahmad, Lo, and Mykhaylova (2013) show nonlinear dynamics can be particularly striking from the perspective of nonlinear tests when they generate data from a DSGE model from it higher order, and nonlinear, solution. Given the focus on the paper to first propose a solution to measure persistence within multivariate linear systems, I will not pursue the idea of nonlinearity.

I report the estimated CuVo half-life. The choice between reporting the CuVo half-life or CuVoA half-life is less crucial within the Monte Carlo experiment as we are only investigating whether the fitted models can capture the persistence of the real exchange rate within the DGP. In particular, if one can pin down the underlying autocovariance structure, or the ACF, in the DGP, then one will estimate the correct CuVo half-life as per the underlying DGP, as this is just a summary statistic of the DGP's ACF.

I first simulate a long Monte Carlo sample with 200,000 observations in order to study the asymptotic properties of the persistence estimates and these results are presented in the top panel of Figure 6. Because degrees of freedom issues are irrelevant in such a sample, and possible VARMA dynamics may be approximated with an arbitrarily high AR lag order, I thus estimate both the VAR and AR model with 100 lags.¹⁴ The CuVo half-life using the multivariate VAR model is close to being unbiased, falling short of the DGP CuVo half-life by only about one period (i.e. one quarter with quarterly data). The AR models, by omitting multivariate information, results in an extremely downward biased estimate of CuVo half-life, with the estimated CuVo half-life less than half of that of the DGP's. On the basis of the large sample Monte Carlo, I conclude that even a very high order univariate model is sufficient to offset the omission of relevant multivariate information in population. Therefore, measures of persistence from univariate models will systematically underestimate the persistence of the real exchange rate given an underlying multivariate DGP.

To understand small sample issues, I also simulate 2000 Monte Carlo samples, each with 250 observations. These results are presented in the second and third panel. The CuVo half-life has a distribution which is possibly of an unknown form. Because the least we know is that the distribution of the CuVo half-life should be extremely skewed, taking the mean across Monte Carlo draws is not a meaningful way to understand the small sample distribution. Therefore, to get an idea of the possible central tendency of the estimated half-life in small samples, I present two summary statistics. The first calculates the median estimated half-life across the 2000 Monte Carlo samples and is presented in the middle panel. The second approach calculates the mean across the the interquartile range, and thus attempts to understand the distribution without the skewed tail. While neither will be entirely suitable statistics for the task at hand, they at least provide some information to understand the small sample properties. I choose the lag order of the AR and VAR model similar to how I conducted the empirical exercise. I either choose four lags as a rule of thumb for quarterly data, or use the AIC, of which

¹⁴Given weak identification issues of MA terms are an important issue on its own and would warrant a separate research agenda, I will not pursue studying MA terms in the Monte Carlo exercise. Note that if one could sidestep all the weak identification issues of fitting MA terms, then one can estimate a model with MA terms and apply the ideas introduced in this paper to estimate half-lives with multivariate information, as long as the estimated VARMA is stationary and has a vector moving average representation.

the results are presented with p in the parenthesis. Once again, we observe the CuVo half-life, estimated using multivariate models is less biased. Interestingly, without bias-correction, both multivariate and univariate models are almost as severely downward biased in terms of their CuVo half-life. The bias-correction helps push the multivariate models towards the true CuVo half-life. It is not surprising that bias-correction helps to produce less biased persistence measures, given these corrections are meant to solve small sample issues. It is however somewhat more surprising that, at least within the context of the Monte Carlo exercise, the correction only has a tangible effect of producing less bias with the multivariate models relative to the univariate ones.

Overall, I make two key observations based on the Monte Carlo exercise. First, in a multivariate DGP environment, multivariate models produce less biased measures of real exchange rate persistence relative to univariate models. This result is consistent with the empirical analysis where we find more persistent real exchange rates with multivariate models. One interpretation consistent with the Monte Carlo exercise and the empirical analysis is that the lower half-lives estimated from univariate models could thus be attributed to omitting relevant multivariate information. Therefore, the practical implications are clear. Unless one believed the real exchange rate is generated as a univariate process and not through some general equilibrium or multivariate structure, then one should consider a multivariate approach when estimating the half-life of the real exchange rate. Omitting relevant multivariate information can result in a large downward biases in estimated persistence resulting from the omitted information. Second, the lower real exchange rate persistence estimated from univariate models is not a small sample issue. Our large sample Monte Carlo exercise suggests that the downward bias in persistence estimated through univariate models does not vanish asymptotically. On the other hand, it appears that at least in an asymptotic sense, the multivariate model can eventually uncover the true real exchange rate persistence.

5 Application to SVAR with Identified Shocks

While the identification of particular shocks is unnecessary for the calculation of the multivariate half-life measures, the ideas within this paper can be trivially extended if one wishes to identify particular shocks within a multivariate framework. To illustrate such an application, I will present an empirical example of a model proposed by Ng (2003), which is a five variable SVAR which has five identified shocks. The interested reader is referred to the relevant paper for the underlying motivation of the restrictions which underpin the identification scheme.

The model has two countries, a large country, A, and a small open economy, B. The model has inflation rates and GDP growth of both countries, as well as the real exchange

rate between the two countries.¹⁵ Using notation introduced much earlier in the article in Equation(3), she proposes the following identification scheme for the five variable VAR.

$$\begin{aligned}
 y_t &= B_0 C \epsilon_t + B_1 C \epsilon_{t-1} + B_2 C \epsilon_{t-2} + \dots \\
 \begin{pmatrix} \pi_t^A \\ \pi_t^B \\ g_t^A \\ g_t^B \\ q_t \end{pmatrix} &= B_0 \begin{pmatrix} C_{11} & 0 & 0 & 0 & C_{15} \\ 0 & C_{22} & 0 & 0 & C_{25} \\ C_{31} & 0 & C_{33} & 0 & C_{35} \\ C_{41} & C_{42} & 0 & C_{44} & C_{45} \\ C_{51} & C_{52} & C_{53} & C_{54} & 0 \end{pmatrix} \begin{pmatrix} \epsilon_{0t}^A \\ \epsilon_{0t}^B \\ \epsilon_{1t}^A \\ \epsilon_{1t}^B \\ \epsilon_t^G \end{pmatrix} + B_1 C \epsilon_{t-1} + B_2 C \epsilon_{t-2} + \dots
 \end{aligned}$$

where the superscripts A and B refer to country A and B, and π_t , g_t and q_t are inflation, GDP growth and the real exchange rate respectively. The second line of the exposition expands on the C matrix, which identify the ϵ_t vector of structural shocks. ϵ_{0t} are country specific shocks, while ϵ_{1t} are sticky price shocks. The key difference which distinguishes sticky price shocks and the country specific shocks is that the latter has an instantaneous effect on prices, whereas the sticky price shocks only have a delayed effect on prices. Finally, ϵ_t^G represent global shocks. The idea behind Ng's identification scheme is probably best interpreted as a semi-structural attempt to understand the differences between real and nominal shocks, which are analogous to the country specific and the sticky price shocks respectively.

The idea of sticky price shocks motivates the restrictions $C_{13} = C_{24} = 0$. As a global shock affects both countries, the identification scheme assumes that the global shock has no impact on the real exchange rate, given it is a relative price, motivating $C_{55} = 0$. Finally, all the other zero restrictions reflect the fact that shocks that originate in the foreign country do not have an instantaneous impact on the domestic country. This in principle will place 11 restrictions on the model, which is one more than needed to just identify the VAR. Ng thus proposes to leave C_{41} free, arguing the with country A being assumed to be larger than country B, country specific shocks in country A can have an instantaneous impact on country B. The matrix C encompasses 10 identifying zero restrictions, and thus satisfies the right number of restrictions for the structural model to be just identified from the reduced form VAR.

Certainly, it can never be stressed enough that the credibility of any SVAR analysis depends on how tenable are the underlying identifying restrictions. My objective is neither to support nor object to the tenability of the current identification scheme, but instead aim to use a prior published piece to serve as an empirical example. Therefore, how seriously one should take the results of any empirical result presented within the SVAR

¹⁵Ng's analysis uses monthly data, and thus uses industrial production as a proxy of output. I will use GDP growth and also quarterly data to be somewhat consistent with my previous empirical analysis.

analysis is always subject to the same caveats about the plausibility of any identification scheme.

For the empirical example, I will present results from the Canadian-U.S. and U.K-U.S. real exchange rate. Country A is thus the U.S. for both models, with country B being either Canada or the U.K. The VAR is once again estimated with four lags.¹⁶ We first calculate the CuVo half-life of the real exchange rate, conditional on the five identified shocks. Reusing the previously introduced notation, let $\phi(i)$ be the IRF at the horizon i and $\Gamma_{m,n}(i)$ be the cumulative squared IRF of the m^{th} variable to the n^{th} shock. Therefore,

$$\Gamma_{m,n}(i) = \sum_{k=0}^i e_n \phi(k)^2 e'_m.$$

Then the conditional CuVo half-life of variable m to shock n , denoted $h_{CuVo|e_t^n}^*$ will satisfy

$$\Gamma_{m,n}(h_{CuVo|e_t^n}^*) = \frac{\lim_{i \rightarrow \infty} \Gamma_{m,n}(i)}{2}.$$

If one wishes to work with the CuVoA half-life, it is also straightforward to convert the conditional CuVo half-life into a conditional CuVoA half-life, by once again applying Equations (10) and (11).

The top panels of Figure 7 presents the estimated half-life, conditional on each of the five identified shocks. I will present the conditional CuVoA half-life estimates, based on the bias-corrected VAR coefficients. The horizontal dotted line marks out the unconditional CuVoA half-life, as calculated from the procedure proposed earlier in the paper. The unconditional CuVoA half-lives are about sixteen and four and a half years respectively for the Canada-U.S. and the U.K-U.S. real exchange rate.¹⁷ We observe that for both countries, the global shock has the longest CuVoA half-life estimate. In fact, for both the Canadian-U.S., and U.K.-U.S. models, the persistence of the global shock is by far the most persistent shock, with a CuVoA half-life estimate that is nearly twice that of the next most persistent shock. Intuitively, the unconditional persistence should be a weighted average of the conditional persistence across all shocks, with the weights depending on the relative importance of the shocks. To get an idea of the importance of the shocks, the bottom panel presents the forecast error variance decompositions of the five shocks. I present the 20 quarters ahead and the asymptotic forecast error variance decomposition. While the Canadian sticky price shock is the most important shock for the Canada-U.S. real exchange rate, the U.S. country specific shock is the most important shock for the U.K.-U.S. real exchange rate. Both of these shocks explain over 50% of

¹⁶Like the main empirical analysis, the Canadian-U.S. sample is from 1962Q1 to 2017Q1. The U.K. sample is 1960Q2 to 2017Q1.

¹⁷These unconditional CuVoA half-life estimates are similar to the analysis in Figure 4, but differ only to the extent that the information sets used in the VARs deviate slightly.

the variance of their respective real exchange rate. The global shock, while is the most persistent, is unimportant relative to the other four shocks. This explains why while the global shock is found to have the most persistent impact on the real exchange rate, the unconditional CuVoA half-life is not weighted towards the high persistence of the global shock, but instead towards its most important shock. In fact, for the U.K.-U.S. real exchange rate, the conditional CuVoA half-life to the U.S. country specific shock is nearly equal to the unconditional CuVoA half-life, reflecting the fact that it derives its unconditional persistence from its most important shock.

Overall, this section illustrates that the ideas proposed in this paper can be readily extended to VARs with identified shocks. Even if the approach of estimating the unconditional half-life in multivariate systems can occur without identifying any of the underlying shocks, one can always identify shocks, and then reconcile the conditional persistence to the unconditional persistence. Therefore, the ideas presented in this paper do not in any way restrict, and probably expand, the type of analysis which one can perform even if one chooses to work with a fully or partially identified VAR.

6 Conclusion

Measuring persistence is an integral part of summarizing data dynamics. Half-life characterization is a widely-cited summary measure of persistence, and has commanded wide attention within the real exchange rate literature. The goal of this paper is to propose an approach to characterize half-life in multivariate models. Modern macroeconomic models are often multivariate and general equilibrium in nature, and thus provides a natural role for tools to measure persistence in multivariate environments. To the extent that these models are often written down to explicitly match the observed data dynamics suggest such exercises form has an integral role in matching theory with data.

An empirical application of real exchange rates for G7 countries, Australia and New Zealand suggest that accounting for multivariate information is empirically important. The empirical findings suggest half-lives estimated from multivariate models are usually longer than its univariate counterpart. A Monte Carlo simulation using a DSGE model is also consistent with the empirical results, where the omission of multivariate information can play a crucial role in severely underestimating the persistence of the real exchange rate.

While the empirical exercise provides compelling evidence that considering relevant multivariate information matters in measuring the persistence of the real exchange rate, it is a broader empirical question whether there might be other sources of multivariate information which may also matter, but have been neglected in the empirical exercise. The ideas introduced in this paper can serve as a good starting point for such analysis.

References

- AHMAD, Y., M. C. LO, AND O. MYKHAYLOVA (2013): “Causes of nonlinearities in low-order models of the real exchange rate,” *Journal of International Economics*, 91(1), 128 – 141.
- ANDREWS, D. W. K., AND H.-Y. CHEN (1994): “Approximately Median-Unbiased Estimation of Autoregressive Models,” *Journal of Business & Economic Statistics*, 12(2), 187–204.
- BERGIN, P. R. (2006): “How well can the New Open Economy Macroeconomics explain the exchange rate and current account?,” *Journal of International Money and Finance*, 25(5), 675–701.
- CARVALHO, C., AND F. NECHIO (2011): “Aggregation and the PPP Puzzle in a Sticky-Price Model,” *American Economic Review*, 101(6), 2391–2424.
- CASHIN, P., AND C. J. MCDERMOTT (2003): “An Unbiased Appraisal of Purchasing Power Parity,” *IMF Staff Papers*, 50(3), 1.
- CHEUNG, Y.-W., AND K. S. LAI (2000): “On the purchasing power parity puzzle,” *Journal of International Economics*, 52(2), 321–330.
- CHORTAREAS, G., AND G. KAPETANIOS (2013): “How Puzzling is the PPP Puzzle? An Alternative Half-life Measure of Convergence to PPP,” *Journal of Applied Econometrics*, 28(3), 435–457.
- COGLEY, T., G. E. PRIMICERI, AND T. J. SARGENT (2010): “Inflation-Gap Persistence in the U.S.,” *American Economic Journal: Macroeconomics*, 2(1), 43–69.
- DIEBOLD, F. X., AND L. KILIAN (2001): “Measuring predictability: theory and macroeconomic applications,” *Journal of Applied Econometrics*, 16(6), 657–669.
- DUFOUR, J.-M., D. PELLETIER, AND E. RENAULT (2006): “Short run and long run causality in time series: inference,” *Journal of Econometrics*, 132(2), 337–362.
- DUFOUR, J.-M., AND E. RENAULT (1998): “Short Run and Long Run Causality in Time Series: Theory,” *Econometrica*, 66(5), 1099–1126.
- EICHENBAUM, M., AND C. L. EVANS (1995): “Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates,” *The Quarterly Journal of Economics*, 110(4), 975–1009.
- ENGEL, C. (2000): “Long-run PPP may not hold after all,” *Journal of International Economics*, 51(2), 243–273.
- JACOB, P., AND G. PEERSMAN (2013): “Dissecting the dynamics of the US trade balance in an estimated equilibrium model,” *Journal of International Economics*, 90(2), 302–315.
- JAROCIŃSKI, M., AND B. MAĆKOWIAK (2017): “Granger-causal-priority and choice of variables in vector autoregressions,” *The Review of Economics and Statistics*, 99(2), 319–329.
- KAMBER, G., J. MORLEY, AND B. WONG (2017): “Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter,” *Review of Economics and Statistics*, forthcoming.
- KAPETANIOS, G., A. PAGAN, AND A. SCOTT (2007): “Making a match: Combining theory and evidence in policy-oriented macroeconomic modeling,” *Journal of Econometrics*, 136(2), 565–594.


- KILIAN, L. (1998): “Small-Sample Confidence Intervals For Impulse Response Functions,” *The Review of Economics and Statistics*, 80(2), 218–230.
- KILIAN, L., AND T. ZHA (2002): “Quantifying the uncertainty about the half-life of deviations from PPP,” *Journal of Applied Econometrics*, 17(2), 107–125.
- LO, M. C., AND J. MORLEY (2015): “Bayesian analysis of nonlinear exchange rate dynamics and the purchasing power parity persistence puzzle,” *Journal of International Money and Finance*, 51(C), 285–302.
- LUBIK, T., AND F. SCHORFHEIDE (2006): “A Bayesian Look at the New Open Economy Macroeconomics,” in *NBER Macroeconomics Annual 2005, Volume 20*, NBER Chapters, pp. 313–382. National Bureau of Economic Research, Inc.
- MURRAY, C. J., AND D. H. PAPELL (2002): “The purchasing power parity persistence paradigm,” *Journal of International Economics*, 56(1), 1–19.
- NG, S. (2003): “Can sticky prices account for the variations and persistence in real exchange rates?,” *Journal of International Money and Finance*, 22(1), 65–85.
- OBSTFELD, M., AND K. ROGOFF (2001): “The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?,” in *NBER Macroeconomics Annual 2000, Volume 15*, NBER Chapters, pp. 339–412. National Bureau of Economic Research, Inc.
- ROGOFF, K. (1996): “The Purchasing Power Parity Puzzle,” *Journal of Economic Literature*, 34(2), 647–668.
- ROSSI, B. (2005): “Confidence Intervals for Half-Life Deviations From Purchasing Power Parity,” *Journal of Business & Economic Statistics*, 23, 432–442.
- STEINSSON, J. (2008): “The Dynamic Behavior of the Real Exchange Rate in Sticky Price Models,” *American Economic Review*, 98(1), 519–533.
- WALLIS, K. F. (1977): “Multiple Time Series Analysis and the Final Form of Econometric Models,” *Econometrica*, 45(6), 1481–97.
- WIRIYAWIT, V., AND B. WONG (2016): “Structural VARs, deterministic and stochastic trends: How much detrending matters for shock identification,” *Studies in Nonlinear Dynamics & Econometrics*, 20(2), 141–157.

Table 1: Granger Causality Tests

		Lag Order of Test Equation			
		AIC		Four Lags	
		Granger	“Indirect”	Granger	“Indirect”
		Causality	Causality	Causality	Causality
Australia	Relative Inflation		X	X	X
	Relative GDP Growth				X
	Relative Interest Rates		X		X
Canada	Relative Inflation		X		X
	Relative GDP Growth				X
	Relative Interest Rates				
France	Relative Inflation	X	X		X
	Relative GDP Growth				
	Relative Interest Rates	X	X		X
Germany	Relative Inflation				
	Relative GDP Growth				X
	Relative Interest Rates	X	X		X
Italy	Relative Inflation		X		X
	Relative GDP Growth				
	Relative Interest Rates				X
Japan	Relative Inflation			X	X
	Relative GDP Growth				
	Relative Interest Rates	X	X		X
New Zealand	Relative Inflation		X		
	Relative GDP Growth				
	Relative Interest Rates	X	X	X	X
United Kingdom	Relative Inflation			X	X
	Relative GDP Growth	X	X		
	Relative Interest Rates	X	X	X	X


*Notes: All variables are relative to the U.S. Test equation includes lags of all the variables and the lags of the real exchange rate all regressed on the real exchange rate. The test for Granger causality is a standard test where the covariates are regressed on the current real exchange rate. The test for “indirect” causality is a test of the same covariate against the h step ahead real exchange rate, where $h = 1, 2 \dots 6$ (see Dufour, Pelletier, and Renault, 2006). The lag order is either selected using AIC or set at four lags. An **X** in the Granger causality column indicates that the variable Granger causes the real exchange rate at 10% level of significance. An **X** in under the “indirect” causality column indicates the variables Granger cause the real exchange for at least one of the horizons, $h = 1, 2 \dots 6$, at the 10% level of significance.*

Figure 1: Estimated Cumulative Volatility (CuVo) half-life, Lag Order Selected by AIC



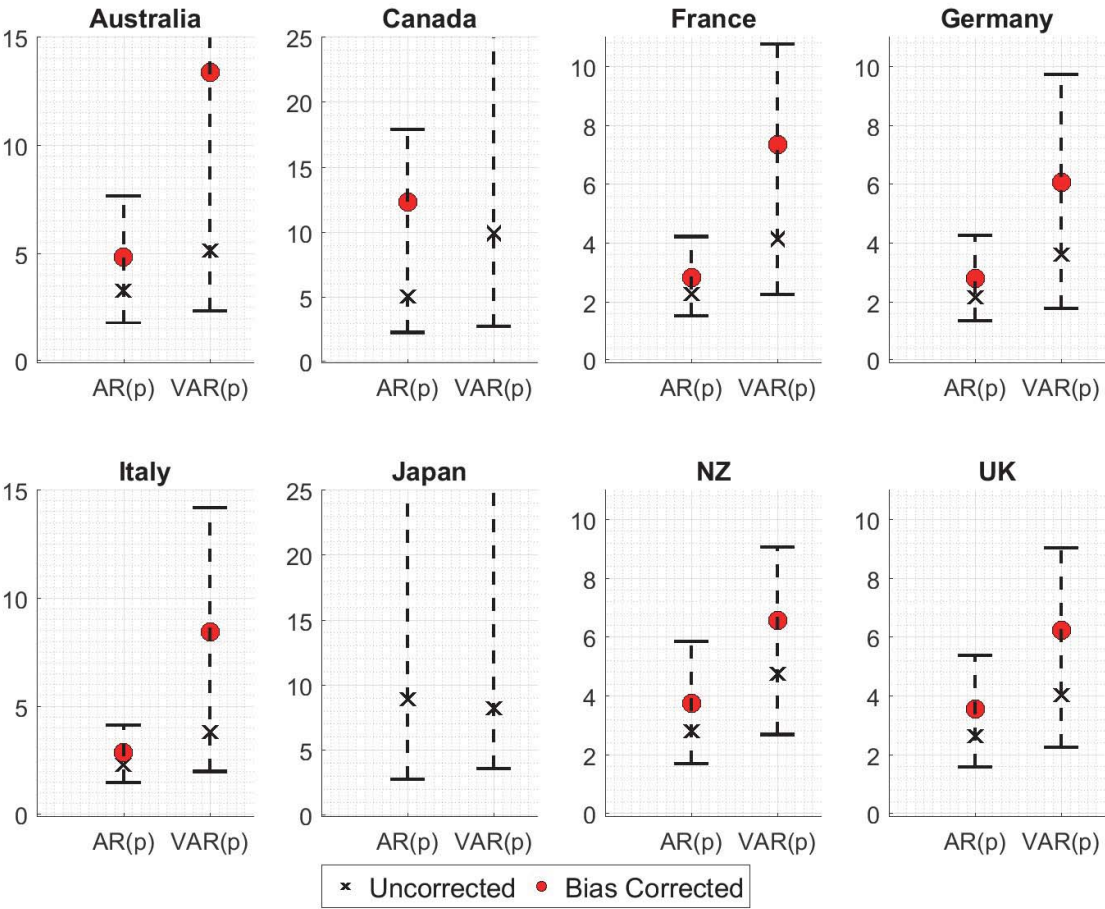
Notes: CuVo Half-life expressed in years. The whiskers represents the bounds of the 90% confidence intervals of the estimated CuVo Half-life. Any unrepresented bound or point estimate should be interpreted as infinitely and non-mean reverting.

Figure 2: Estimated Cumulative Volatility (CuVo) half-life, Lag Order Set to Four Lags




Notes: CuVo Half-life expressed in years. The whiskers represents the bounds of the 90% confidence intervals of the estimated CuVo Half-life. Any unrepresented bound or point estimate should be interpreted as infinitely and non-mean reverting.

Figure 3: Estimated Cumulative Volatility Adjusted (CuVoA) half-life, Lag Order Selected by AIC



Notes: CuVoA Half-life expressed in years. The whiskers represents the bounds of the 90% confidence intervals of the estimated CuVo Half-life. Any unrepresented bound or point estimate should be interpreted as infinitely and non-mean reverting.

Figure 4: Estimated Cumulative Volatility Adjusted (CuVoA) half-life, Lag Order Set to Four Lags



Notes: CuVoA Half-life expressed in years. The whiskers represents the bounds of the 90% confidence intervals of the estimated CuVo Half-life. Any unrepresented bound or point estimate should be interpreted as infinitely and non-mean reverting.

Figure 5: Modulus of Largest Eigenvalue of Companion Matrix

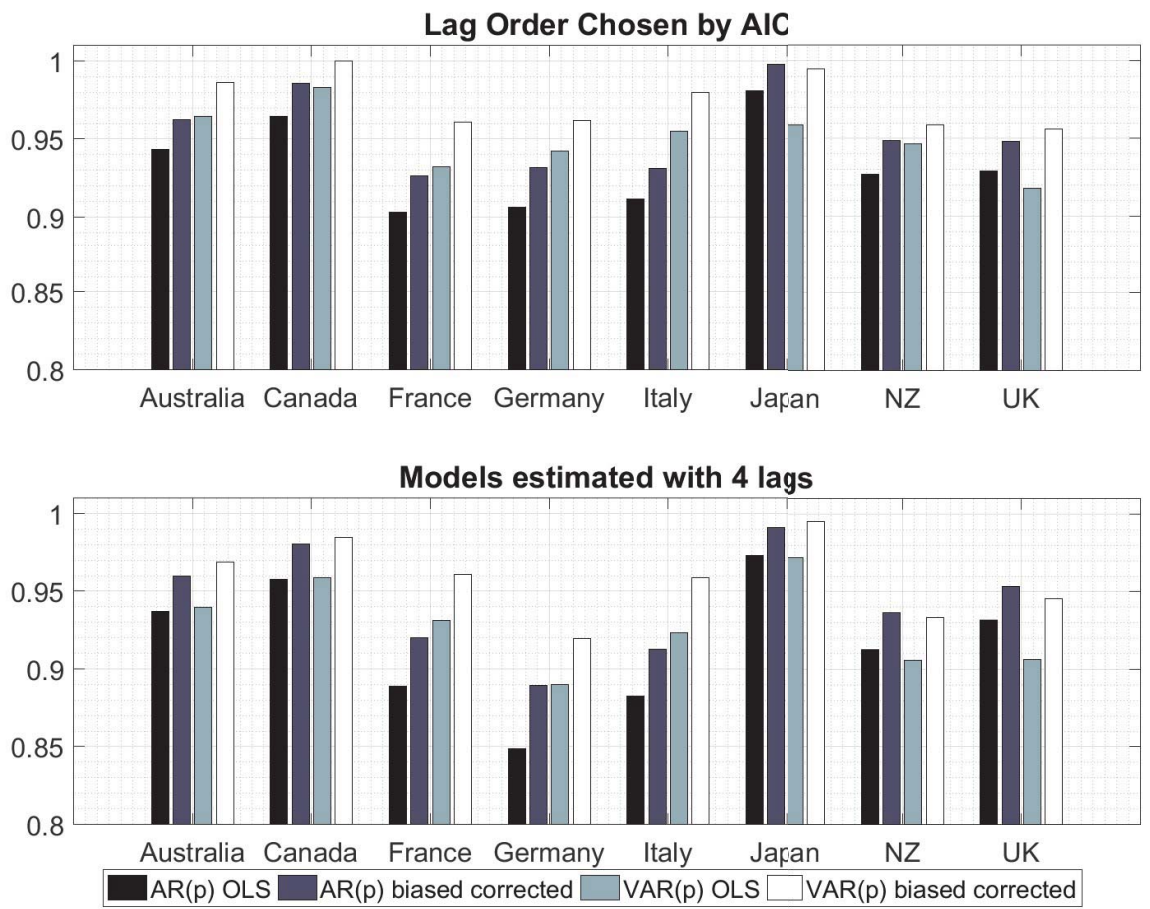
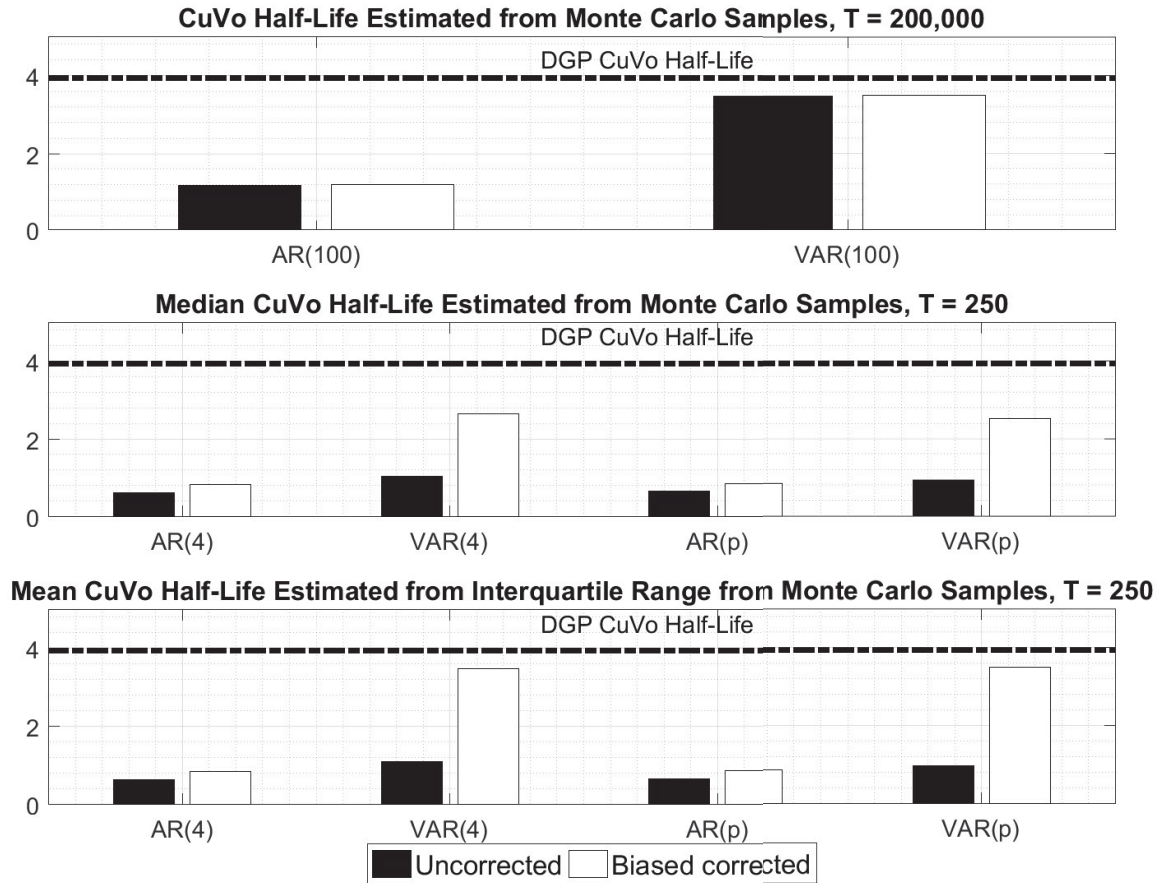
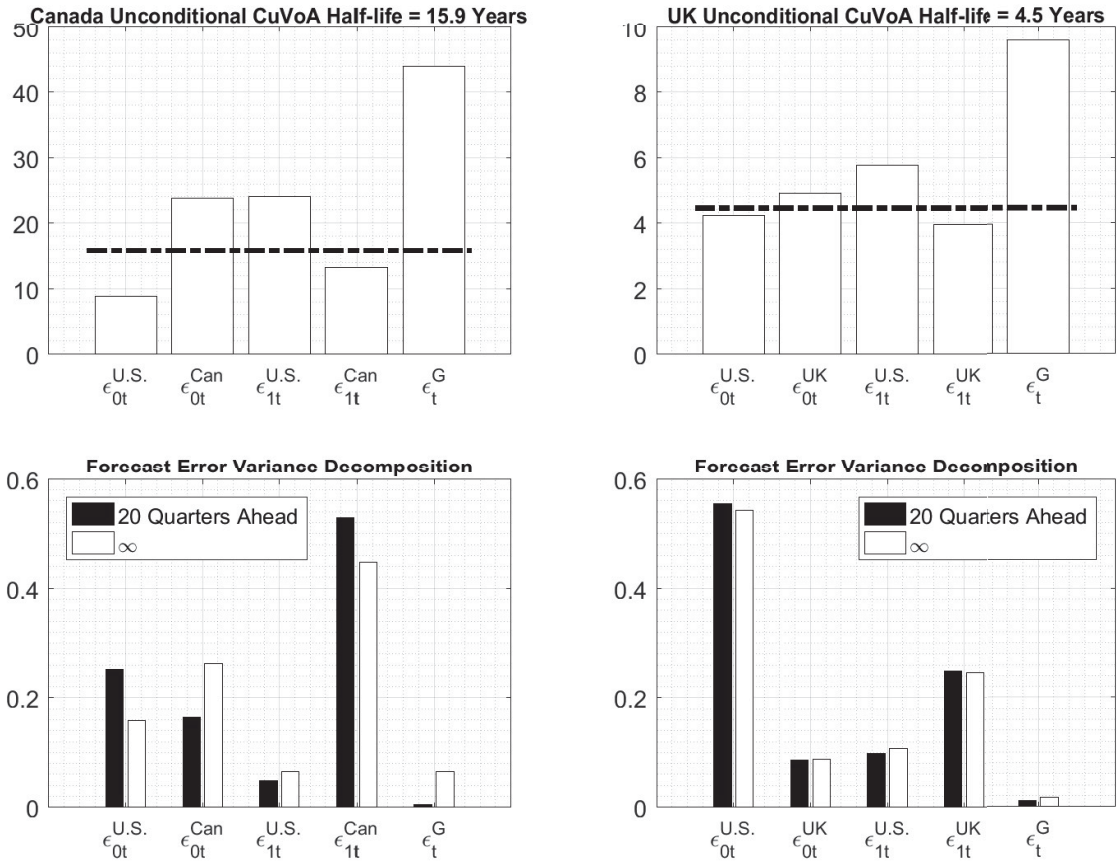


Figure 6: Results from Monte Carlo Experiment



Notes: CuVo half-life estimates are in years. Underlying DGP is a two country open economy DSGE model by Lubik and Schorfheide (2006). The dotted line represents the CuVo half-life from the underlying DGP. The number in parenthesis indicates the number of lags in the AR or VAR model. p in the parenthesis indicates the lag order is chosen by AIC. The bar represents the half-life estimates from the various exercises. The VAR model is estimated with the seven observables generated from the DGP, with the bar representing the estimated half-life from the real exchange rate. AR represents half-life estimates of the real exchange rate from an AR model. Bias-correction is done via a bootstrap as per the first step bootstrap described by Kilian (1998). Summary statistics from the $T = 250$ Monte Carlo experiments are done from 2000 Monte Carlo samples.

Figure 7: Conditional CuVoA Half-lives and Variance Decomposition of the Canadian-U.S. and U.K.-U.S. Real Exchange Rate



Notes: The left and right panels present results of the Canada-U.S. and U.K.-U.S. real exchange rate to the various identified shocks. The top panels present the conditional CuVoA half-life estimate, while the bottom panels present the forecast error variance decomposition. ϵ_{0t} and ϵ_{1t} refer to country specific and sticky price shocks respectively. The superscript refers to the countries U.S., Canada and the U.K. ϵ_t^G refers to global shocks.

Online Appendix to “*Incorporating Relevant
Multivariate Information for Characterizing
Half-Life with an Application to Purchasing Power
Parity*”

Benjamin Wong*

Reserve Bank of New Zealand and Centre for Applied Macroeconomic
Analysis, The Australian National University

July 17, 2017

*The views do not necessarily reflect those of the Reserve Bank of New Zealand. Reserve Bank of New Zealand, 2 The Terrace, Wellington 6011, New Zealand. Email: benjamin.wong@rbnz.govt.nz Tel: +64 44713957.

A1 Raw Data Sources

All the data are from the OECD Economic Outlook and the IMF's International Finance Statistics (IFS). The data is accessed through Haver Analytics. Haver Analytics is a subscription based system which aggregates various databases. For Germany, the data is sourced through Haver Analytics' G10 database because it combines pre and post unification data, and so allows for a longer span of data for Germany to enter the empirical analysis.

CPI

IFS Canada, France, Italy, Japan, Australia and New Zealand

Bureau of Labor Statistics U.S.

Bundesbank with adjustment by Haver Germany

Office for National Statistics, Retail Price Index United Kingdom

Nominal Exchange Rate Relative to the U.S.

IFS Australia, Canada, Japan, New Zealand and United Kingdom

Legacy Exchange Rates, Haver Analytics G10 Databank with Source Data from Federal Reserve Bank Germany, Italy and France

GDP

OECD Outlook, Real GDP (volume) U.S., Canada, France, Italy, Japan, United Kingdom, Australia, New Zealand

Deutche Bundesbank, Accessed through Haver Analytics G10 Databank Germany

Short Term Interest Rate

Whenever possible, a 3 month bank bill is taken to be the domestic interest rate. To get the longest possible span of interest rate data, another long span of interest rate data is usually taken to backcast using a linear regression. While I do try as much as possible to use another short term interest rate, sometimes a long rate is used for the backcasting. In principle, given interest rates tend to co-move, this should provide enough information

to construct a long span of data for short term interest rates. Table A1 describes the main interest rate series used, and if backcasted, the series used to backcast using a linear regression. All the data is from IFS, except Germany, which is from the G10 databank of Haver Analytics.

Table A1: Interest Rates Data Sources

U.S.
3-Month Treasury Bill
Canada
3-Month Treasury Bill Yield
France
Treasury Bill Rate: 13 weeks
Backcast Using: Money Market Opening Rate: Day-to-Day Loans Against Private Bills
Germany
Base Rate (formerly discount rate)
Italy
3-12 Month Weighted New Treasury Bill Rate (%)
Backcast Using: 9-10 Year Government Bond Yield (%)
Japan
Financing Bill Rate (% Per Annum)
United Kingdom
91-Day Treasury Bill Tender Rate
Backcast Using: London Clearing Banks: Instant Access Deposits
Australia
Money Market Rate: Short-Term, Weighted Average of Loans Outstanding
Backcast Using: 10-Year Government Nonrebate Bond Yield (%)
New Zealand
3-Month Treasury Bill Tender Rate
Backcast Using: 5+ Year Government Bond Yield to Maturity

A2 Bias Correction and Construction of Confidence Intervals

Almost all of this discussion draws on and repeats the algorithm section of Kilian (1998, page 220). Let

$$y_t = \hat{c} + \hat{A}_1 y_{t-1} + \hat{A}_2 y_{t-2} + \dots + \hat{A}_p y_{t-p} + \hat{v}_t \quad (\text{A.1})$$

be the estimated VAR model where y represents the vector of variable in the VAR, and \hat{c} , \hat{v}_t and \hat{A}_i , $i \in \{1, 2, \dots, p\}$ are respectively the matrices or vectors of the constants, residuals and estimated VAR coefficients obtained by fitting the VAR equation using least squares. Let T be the length of the sample. I then apply the following steps.

1. Simulate 1000 datasets, each of length T , centring the DGP from Equation (A.1) and sampling with replacement from the empirical distribution of the residuals, \hat{v}_t . Let A^* be the mean of the VAR coefficients across the 1000 simulated datasets. Approximate the bias $\Psi = A^* - \hat{A}$.
2. Calculate the modulus of the largest root of the companion matrix. Denote this as $m(\hat{A})$. If $m(\hat{A}) \geq 1$, set $\tilde{A} = \hat{A}$. Otherwise, set $\delta_1 = 1$ and $\delta_{i+1} = \delta_i - 0.01$. Let $\bar{A}(i) = \tilde{A} - \Psi \Pi_{j=1}^i \delta_j$. Iterate on $i = 1, 2, \dots$ until $m(\bar{A}(i)) \leq 1$, then set $\tilde{A} = \bar{A}(i)$. This step retains the stationarity of the bias-adjusted DGP as one can get arbitrarily close to the unit circle. Use \tilde{A} to construct the bias adjusted half-life estimate.¹
3. Replace \hat{A} with \tilde{A} in Equation (A.1), and generate 2000 bootstrap samples. Use the 2000 bootstrapped samples to calculate the confidence interval of the half-life estimate.

¹In the empirical exercise, the estimated VAR coefficients, \hat{A} , always imply stationarity. Therefore, despite not testing for unit roots, this implies that the half-life exists, at least under the condition of stationarity of the estimated AR or VAR coefficients in the empirical exercise.

A3 Parameterization of Monte Carlo DGP

A DSGE model is used as the Data Generating Process (DGP) for the Monte Carlo exercise. The model is a two country New Keynesian DSGE model used by Lubik and Schorfheide (2006). The reader is referred to the relevant pages of their paper for details of the model. The structure of the model can be found on pages 318-326 of their paper. The model has an underlying structure of what is now considered a standard two country open economy DSGE model. I use the same priors as Lubik and Schorfheide (2006), with the model estimated on their dataset for Euro-U.S. data. The Dynare code to estimate their model and the dataset can be found on Frank Schorfheide's website. The DGP is parametrized using the posterior mode. In Table A2, I report the parameterization of the parameters in the DGP which I had subsequently used to generate data. The notation used in Table A2 correspond with Lubik and Schorfheide's notation for ease of cross referencing.

Table A2: Parameterization of Lubik and Schorfheide (2006) two country model used for DGP

θ_H	Fraction of home producers adjusting prices to steady state inflation	0.62
θ_F	Fraction of domestic importers adjusting prices to steady state inflation	0.42
θ_H^*	Fraction of foreign importers adjusting prices to steady state inflation	0.90
θ_F^*	Fraction of foreign producers adjusting prices to steady state inflation	0.59
τ	Coefficient of relative risk aversion	3.95
h	Habit persistence	0.48
α	Import share	0.19
η	Intratemporal substitution elasticity between home and foreign consumption goods	0.29
ψ_1	Coefficient on inflation on home monetary policy rule	1.53
ψ_2	Coefficient on output growth on home monetary policy rule	0.61
ψ_3	Coefficient on nominal exchange rate on home monetary policy rule	0.03
ψ_1^*	Coefficient on inflation on foreign monetary policy rule	1.55
ψ_2^*	Coefficient on output growth on foreign monetary policy rule	0.70
ψ_3^*	Coefficient on nominal exchange rate on foreign monetary policy rule	0.02
ρ_A	AR(1) coefficient on the home technology shock	0.81
ρ_R	AR(1) coefficient on for interest rate smoothing in home monetary policy rule	0.74
ρ_G	AR(1) coefficient on home government spending	0.91
ρ_A^*	AR(1) coefficient on the foreign technology shock	0.89
ρ_R^*	AR(1) coefficient on for interest rate smoothing in foreign monetary policy rule	0.78
ρ_G^*	AR(1) coefficient on foreign government spending	0.91
ρ_Z	AR(1) coefficient on worldwide technology shock	0.51
$r^{(A)}$	Steady state real interest rate	0.00
γ	Steady state growth rate	0.41
$\pi^{(A)}$	Annualized steady state inflation	6.65
σ_A	Standard deviation of the home technology shock	1.55
σ_G	Standard deviation of home government spending shock	0.50
σ_R	Standard deviation of home monetary policy shock	0.18
σ_A^*	Standard deviation of the foreign technology shock	1.16
σ_G^*	Standard deviation of foreign government spending shock	0.55
σ_R^*	Standard deviation of foreign monetary policy shock	0.17
σ_Z	Standard deviation of worldwide technology shock	0.34
σ_E	Standard deviation of PPP shock	4.42

References

- KILIAN, L. (1998): "Small-Sample Confidence Intervals For Impulse Response Functions," *The Review of Economics and Statistics*, 80(2), 218–230.
- LUBIK, T., AND F. SCHORFHEIDE (2006): "A Bayesian Look at the New Open Economy Macroeconomics," in *NBER Macroeconomics Annual 2005, Volume 20*, NBER Chapters, pp. 313–382. National Bureau of Economic Research, Inc.