South African Reserve Bank
Working Paper Series
WP/21/15

Risk and Return Spillovers in a Global Model of the Foreign Exchange Network

Matthew Greenwood-Nimmo, Daan Steenkamp and Rossouw van Jaarsveld

Authorised for distribution by Konstantin Makrelov

4 August 2021

South African Reserve Bank

South African Reserve Bank Working Papers are written by staff members of the South African Reserve Bank and, on occasion, by consultants under the auspices of the South African Reserve Bank. The papers deal with topical issues and describe preliminary research findings, and develop new analytical or empirical approaches in their analyses. They are solely intended to elicit comments and stimulate debate.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the South African Reserve Bank or South African Reserve Bank policy. While every precaution is taken to ensure the accuracy of information, the South African Reserve Bank shall not be liable to any person for inaccurate information, omissions or opinions contained herein.

South African Reserve Bank Working Papers are externally refereed.

Information on South African Reserve Bank Working Papers can be found at
https://www.resbank.co.za/en/home/publications/Papers/working-papers

Enquiries relating to the Working Paper Series can be addressed to:

Head: Economic Research Department
South African Reserve Bank
P O Box 427
Pretoria 0001

Tel. +27 12 313 3911

Risk and Return Spillovers in a Global Model of the Foreign Exchange Network

Matthew Greenwood-Nimmo*, Daan Steenkamp†and Rossouw van Jaarsveld‡

4 August 2021

Abstract

We developed a network model to capture the dynamic interactions among foreign exchange returns and realised risk measures among 20 developed and emerging market currencies, including the rand (ZAR). We demonstrate how this framework can be used to assess the sensitivity of a given currency to shocks from other currencies and to provide narratives contextualising currency movements, focusing on the ZAR. We show that variations in the risk-return profile of the USDZAR correlate with variations in the risk-return profile of many other currencies, and that this is especially notable with respect to emerging market currencies. We interpret this as evidence of the ZAR's role as a bellwether emerging market currency. We show that the model is able to highlight risk transmission channels in a timely manner during foreign exchange flash crashes and periods of heightened financial market uncertainty.

JEL classification: F31, G01, G15

Keywords: Foreign exchange markets; Higher-order moment risk; Realised moments; Network modelling; Spillovers.

*Department of Economics, University of Melbourne and Centre for Applied Macroeconomic Analysis, Australian National University. Email: matthew.greenwood@unimelb.edu.au.

†Corresponding author. South African Reserve Bank, PO Box 427, Pretoria, South Africa, 0001. Email: daan.steenkamp@resbank.co.za.

‡SARB. Email: Rossouw.VanJaarsveld.co.za.

1. Introduction[1]

There has long been concern among policymakers about the relatively high volatility of the South African rand (ZAR) (see Loewald 2021). This reflects concern that exchange rate volatility imposes costs on the economy by interfering with investment and consumption decisions, raising hedging costs and potentially adversely affecting the development of South Africa's exporting sector.

Despite the extensive policy debate about the relative volatility of the ZAR, there has been relatively little research into the nature and drivers of the ZAR's volatility.[2] Likewise, the depreciating trend of the ZAR has drawn attention to directional risks, yet there is an absence of research into the skewness of the ZAR (i.e. the implied risk that the ZAR could move in a particular direction) or its kurtosis (i.e. the implied risk of very large ZAR changes). Given the interconnectedness of global foreign exchange (FX) markets and the unconventional monetary policy measures (such as quantitative easing) implemented by major economy central banks recently, there has also been a lot of policy interest in measuring spillovers between different markets, but few papers consider samples that include many emerging market (EM) currencies.

In this paper, we characterise the dynamics of the ZAR, as well as those of a large number of EM and developed market (DM) currency pairs and then assess the relationship between ZAR dynamics and those of other currencies. To this end, we construct a database of log-returns and higher-order realised moments for the ZAR, 7 other EM currencies and 12 DM currencies over the period 12 July 2006 to 10 February 2021, inclusive. Next, following the precedent in the literature (e.g. Greenwood-Nimmo et al., 2016, 2019b), we develop and estimate an empirical network model to measure the intensity of bilateral FX risk and return spillovers among currencies and to capture the dynamics of these spillovers. We adopt a modified version of the framework for network analysis put forth by Diebold and Yilmaz (2009, 2014), in which the structure of the network of bilateral spillover effects is estimated via a decomposition of the variance of the forecast errors obtained from a reduced form vector autoregression (VAR). By analysing the structure of the network on a rolling-sample basis, it is possible to track the evolution of spillovers over time.

Our empirical network model enjoys several benefits relative to competing techniques for the analysis of spillovers. First, by virtue of our consideration of the higher-order moments of the exchange rate distribution, our model achieves greater generality than popular ARCH-type models of FX volatility transmission (e.g. Engle et al., 1990, and the many studies that have adopted similar methods). Second, unlike many existing studies that develop separate models for each variable group under consideration (e.g. Cai et al., 2008; Diebold and Yilmaz, 2009),

[1] The computer programs required to replicate the results in this paper using R version 4.0.2 are available from the authors on request.

[2] Exceptions include Farrell et al. (2012), who show that inflation surprises have high-frequency impacts on the USDZAR, Hassan (2015), who suggests that there are relationships between macroeconomic fundamentals and ZAR volatility, and Arezki et al. (2014), who find that gold price volatility can explain the ZAR's volatility.

our use of a shrinkage and selection estimator allows us to develop a single high-dimensional model for multiple variable groups simultaneously that accounts for spillovers between different variable groups (e.g. spillovers between skewness and returns). Furthermore, relative to alternative techniques for network analysis based on Granger causality (e.g. Billio et al., 2012), our technique provides a means to capture not just the direction but also the strength of pairwise spillovers and to account not just for lagged dependence but also for contemporaneous dependence across currencies.

The spillover statistics that we obtain can be used to shed light on how changes in a currency's risk-return profile transmit through global FX markets and to assess the degree to which changes in the risk-return profile of a given currency reflect global and local conditions. Our analysis focuses primarily on spillovers to and from the ZAR and provides new insights into the sensitivity of the ZAR to shocks arising in FX markets around the world, as well as the informational role played by the ZAR in global FX markets. We show that our model is able to highlight risk transmission channels in a timely manner during FX flash crashes and periods of heightened financial market uncertainty. In addition, we show that variations in the risk-return profile of the USDZAR correlate with variations in the risk-return profile of many other currencies, and that this is especially notable with respect to EM currencies. We interpret this as evidence of the ZAR's role as a bellwether EM currency. Our results suggest that global conditions have an important impact on the ZAR, which would likely limit the effectiveness of interventions aimed at dampening currency volatility. Our results do, however, suggest that the ZAR also reacts strongly and rapidly to South Africa-specific shocks, such as bouts of political instability. This is consistent with the exchange rate playing the role of a shock absorber in the economy, as argued by Soobyah and Steenkamp (2019) and Loewald (2021).

2. The dataset

We construct risk and return measures at daily frequency for a group of $M = 20$ currency pairs, with each expressed in units of foreign currency per US dollar (USD).[3] Our sample contains 12 DM currencies, as follows: the Australian dollar (USDAUD), the Canadian dollar (USDCAD), the Swiss franc (USDCHF), the euro (USDEUR), the British pound (USDGBP), the Hong Kong dollar (USDHKD), the Japanese yen (USDJPY), the Korean won (USDKRW), the New Zealand dollar (USDNZD), the Norwegian krone (USDNOK), the Swedish krona (USDSEK), and the Singaporean dollar (USDSGD). In addition, our sample includes the following 8 EM currencies: the Brazilian real (USDBRL), the Chinese renminbi (USDCNY)[4], the Indian rupee (USDINR),

[3] Maintaining consistency in the quoting convention across currencies facilitates comparisons of higher-order risk measures across currencies.

[4] China maintains a managed floating currency framework. The renminbi is traded in two markets: while the offshore renminbi (USDCNH) is largely based on market trading, the onshore renminbi (USDCNY) is regulated and only allowed to trade within a 2% range from a given value that is set each day at 9.15 am. The framework has evolved over time. Between July 2005 and July 2008, the currency referenced a basket of currencies and a fixed trading band; between July 2008 and June 2010, the CNY was fixed to the USD; after July 2010 the managed float was reinstated and the trading band gradually expanded (from 0.3% between January 1994 and May 2007, then 0.5% to April 2012, then 1% to March 2014 and thereafter it was set at its current level of 2%).

the Mexican peso (USDMXN), the Polish złoti (USDPLN), the Russian ruble (USDRUB), the Turkish lira (USDTRY), and the South African rand (USDZAR).[5] Our sample spans the period 12 July 2006 to 10 February 2021, inclusive.

For each currency, we characterise the distribution of exchange rate returns at daily frequency. While it is possible to obtain data on the returns and a variety of higher-moment risk measures for each currency from data providers such as Bloomberg, we elect to construct our own measures using high-frequency intraday data from Refinitiv. The benefit of this approach is that it affords us precise control at all stages of data cleaning and moment construction.

We begin by constructing mid-rates as the midpoint of the bid and ask prices at intraday five-minute intervals. This results in $n = 288$ intraday observations. We then use these intraday mid-rates to compute the first four realised moments of the returns distribution at daily frequency: the log-return, realised variance, skewness and kurtosis. The use of a five-minute sampling frequency has been shown to balance asymptotic considerations in the estimation of the realised moments against the adverse impact of microstructure noise. Following the conventions described in Andersen et al. (2003), we consider a trading day to start at 21:05 GMT and end at 21:00 GMT. We remove weekends from our sample, as well as a selection of fixed and floating public holidays in the US.[6] The daily log-return for currency i on day t is obtained as follows:

$$r_{it} = \sum_{j=1}^{n} r_{it,j}$$

where $r_{it,j} = \log\left(p_{it,j} / p_{it,j-1}\right)$ is the jth period-to-period intraday log-return, with $p_{it,j}$ denoting the jth intraday mid-rate. The simplest estimator of the realised variance of returns for the ith currency is obtained as the sum of squared intraday returns:

$$UV_{it} = \sum_{j=1}^{n} r_{it,j}^2$$

It is common to scale the realised variance by the number of trading days per year, N, to obtain

The fixing mechanism has gradually changed as well to shift to a more market-determined exchange rate. The current regime, implemented in August 2018, bases the daily midpoint rate on the close of the CNY against a basket of currencies on the previous day and quotes from inter-bank dealers, along with a countercyclical adjustment that gives authorities flexibility to affect the level of the exchange rate. Unfortunately the CNH is only available from February 2011 in our dataset, so we are constrained to working with the CNY in order to include the global financial crisis in our sample. However, the CNY and CNH are highly correlated over our sample period and share similar dynamics.

[5] Our classification of DMs and EMs is consistent with the distinction between 'advanced economies' and 'emerging market and developing economies' adopted by the International Monetary Fund (2020).

[6] We discuss our data cleaning routine and compare our risk and return measures against corresponding series obtained from Bloomberg in our Data Appendix. In general, our measures are very similar to those of Bloomberg. The principal differences are observed in the higher-order realised moments, with our estimates exhibiting fewer large outliers than their Bloomberg equivalents. This reflects differences in our approach to data cleaning relative to Bloomberg, as well as our use of a five-minute sampling frequency as opposed to the 30-minute sampling frequency used by Bloomberg.

the following annualised variance measure:

$$V_{it} = N \cdot UV_{it}$$

The realised variance captures the dispersion of returns, which is a measure of uncertainty. Realised skewness measures the degree of asymmetry of the returns distribution and can be estimated as follows:

$$Q_{it} = \sqrt{n} \sum_{j=1}^{n} r_{it,j}^{3} \Big/ \left(UV_{it}\right)^{3/2}$$

Under our quoting convention, $Q_{it} > 0$ implies that there is a tendency towards depreciation of currency i against the USD, while $Q_{it} < 0$ reveals a tendency towards appreciation against the USD. Realised kurtosis can be constructed as follows:

$$K_{it} = n \sum_{j=1}^{n} r_{it,j}^{4} \Big/ \left(UV_{it}\right)^{2}$$

The realised kurtosis measures the mass in the tails of the return distribution; larger values of K_{it} indicate a higher probability of extreme exchange rate returns.

Our Data Appendix provides descriptive statistics for the log-returns and realised moments for all currency pairs. To account for the fact that the realised moments display considerably greater persistence than the exchange rate log-returns, we follow Greenwood-Nimmo et al. (2016) and isolate the innovations in each of the higher-order moments as the residual from a first order autoregression. For example, the realised variance innovation for currency i in period t, v_{it}, is estimated as:

$$\hat{v}_{it} = V_{it} - \hat{a}_i + \hat{b}_i V_{it-1} \tag{1}$$

where $\hat{a}_i$ and $\hat{b}_i$ are the estimated intercept and autoregressive parameters, respectively. The set of innovations obtained in this way for the ith currency is denoted by v_{it}, q_{it} and k_{it}, with the hat symbols suppressed to avoid notational clutter.

3. Measuring connectedness via error variance decomposition

Diebold and Yilmaz (2009) were the first to show that a decomposition of the FEV from a VAR model can be interpreted as a weighted directed network. Their technique has subsequently been generalised by Diebold and Yilmaz (2014) to achieve invariance to the order of the variables in the VAR model and by Greenwood-Nimmo et al. (2021) to facilitate block aggregation of the estimated network. We adopt the aggregation-robust framework of Greenwood-Nimmo et al., because it allows us to switch freely between different aggregation routines in order to study different aspects of FX connectedness.

Consider a set of currencies indexed by $i = 1, 2, \ldots, M$ observed at daily frequency over time

periods indexed by $t = 1, 2, \ldots, T$. For each currency, we observe the daily log-return on the spot exchange rate, expressed in foreign currency units per US dollar, r_{it}. In addition, for each currency, we observe the innovation to the realised variance, skewness and kurtosis of returns, denoted by v_{it}, q_{it} and k_{it}, respectively. The 4×1 vector $y_{it} = (r_{it}, v_{it}, q_{it}, k_{it})'$ collects the endogenous variables for the ith currency. The currency-specific variables can be collected into a vector of global variables to obtain $y_t = (y_{1t}, y_{2t}, \ldots, y_{Mt})'$. The dimension of y_t is $d \times 1$, where $d = 4M$.

The first step in constructing our network statistics is to estimate a VAR model that captures the cross-sectional and dynamic interactions among the variables in y_t. VARs are typically estimated using standard techniques such as ordinary least squares (OLS). However, in a high-dimensional setting such as ours, the overfitting problem associated with unrestricted estimation by OLS can be severe. Consequently, we follow Nicholson et al. (2017) and introduce sparsity into the VAR parameter matrices by implementing the least absolute shrinkage and selection operator (LASSO) proposed by Tibshirani (1996).[7]

Nicholson et al. (2017) propose a unified framework for the estimation of large sparse VARs called VARX-L. The VARX-L framework accommodates many alternative sparsity patterns, including the basic setting in which predictors are selected in an unstructured manner, as well as several structured settings in which predictors are selected in groups (e.g. by lag). We estimate what Nicholson et al. refer to as a *basic VARX-L model*, in which the LASSO L_1 penalty term is applied in order to zero out individual elements of the VAR parameter matrices without imposing any structure on the sparsity pattern. A notable benefit of the unstructured approach relative to structured alternatives is its computational tractability, which arises from the separability of the objective function.

Prior to estimation by the LASSO, it is necessary to standardise the elements of y_t by centring each series and dividing by its standard deviation. This ensures that the magnitude of the penalty term is unaffected by differences in the scale of the variables in y_t. The standardised data is denoted $\check{y}_t$. The pth order reduced form VAR in $\check{y}_t$ is given by:

$$\check{y}_t = \mu + \sum_{j=1}^{p} A_j \check{y}_{t-j} + u_t \tag{2}$$

where μ is a $d \times 1$ vector of intercepts, the A_j's are $d \times d$ matrices of unknown autoregressive coefficients and u_t is a $d \times 1$ vector of reduced form errors with mean zero and positive definite covariance matrix $\mathbf{\Sigma}$.[8] We estimate (2) by the LASSO, allowing for the penalty parameters to be chosen optimally for each equation (as opposed to the selection of a single penalty parameter for the system as a whole). For the ith equation of (2), abstracting from deterministic terms

[7] Shrinkage and selection estimators have been used to address the overfitting problem in the financial connectedness literature by Demirer et al. (2017), Greenwood-Nimmo et al. (2019a), Bostanci and Yilmaz (2020) and Greenwood-Nimmo and Tarassow (2020), among others.

[8] In practice, we find that a lag order of $p = 1$ is sufficient to capture the autocorrelation structure in the data, which decreases rapidly beyond lag 1 for all variables in the system.

and working with a lag order of 1, the LASSO solves the following minimisation problem:

$$\widehat{a}_i = \arg\min_a \left(\sum_{t=1}^{T} \left(\check{y}_{it} - \sum_{j=1}^{d} a_{ij} \check{y}_{j,t-1} \right)^2 + \lambda_i \sum_{j=1}^{d} \left| a_{ij} \right| \right) \tag{3}$$

where a_i denotes the ith row of the VAR(1) parameter matrix and where the equation-specific penalty parameter $\lambda_i > 0$. We select λ_i following the procedure described in Nicholson et al. (2017), using 10 gridpoints and a grid depth of 50.

The VAR model (2) has the following Wold representation:

$$\check{y}_t = \alpha + \sum_{j=0}^{\infty} B_j u_{t-j} \tag{4}$$

where the B_j's are obtained recursively as $B_j = A_1 B_{j-1} + A_2 B_{j-2} + \ldots$, for $j = 1, 2, \ldots$ with $B_0 = I_d$ and $B_j = 0$ for $j < 0$. Pesaran and Shin (1998) show that the h-steps-ahead generalised forecast error variance decomposition (GVD) is given by:

$$\vartheta_{i \leftarrow j}^{(h)} = \frac{\sigma_{jj}^{-1} \sum_{\ell=0}^{h} \left(\varepsilon_i' B_\ell \Sigma \varepsilon_j \right)^2}{\sum_{\ell=0}^{h} \varepsilon_i' B_\ell \Sigma B_\ell' \varepsilon_i} \tag{5}$$

for $i, j = 1, \ldots, d$, where σ_{jj} is the jth diagonal element of Σ and ε_i is a $d \times 1$ selection vector, the ith element of which is 1 with zeros elsewhere. $\vartheta_{i \leftarrow j}^{(h)}$ expresses the proportion of the h-steps-ahead forecast error variance (FEV) of variable i that can be attributed to shocks in the equation for variable j. The GVD is order-invariant, unlike alternative variance decompositions based on triangular decompositions of the residual covariance matrix. However, because the GVD is based on non-orthogonal disturbances, the sum of the estimated FEV shares may exceed 1. For this reason, following Greenwood-Nimmo et al. (2021), we apply the following normalisation:

$$\theta_{i \leftarrow j}^{(h)} = \frac{\vartheta_{i \leftarrow j}^{(h)}}{d \sum_{j=1}^{d} \vartheta_{i \leftarrow j}^{(h)}} \tag{6}$$

which restores a proportional interpretation of the FEV shares. Specifically, $\theta_{i \leftarrow j}^{(h)}$ measures the proportion of the *systemwide* FEV at horizon h accounted for by the jth variable's contribution to the FEV of the ith variable.[9]

The h-steps-ahead $d \times d$ connectedness matrix that summarises the d^2 interactions among

[9] Note that this normalisation differs from the normalisation proposed by Diebold and Yilmaz (2014) due to the presence of the d in the denominator. The normalisation proposed by Greenwood-Nimmo et al. (2021) facilitates aggregation of the spillover statistics obtained from the network model.

the d variables in $\check{y}_t$ is as follows:

$$C^{(h)} = \begin{bmatrix} \theta^{(h)}_{1\leftarrow 1} & \theta^{(h)}_{1\leftarrow 2} & \cdots & \theta^{(h)}_{1\leftarrow d} \\ \theta^{(h)}_{2\leftarrow 1} & \theta^{(h)}_{2\leftarrow 2} & \cdots & \theta^{(h)}_{2\leftarrow d} \\ \vdots & \vdots & \ddots & \vdots \\ \theta^{(h)}_{d\leftarrow 1} & \theta^{(h)}_{d\leftarrow 2} & \cdots & \theta^{(h)}_{d\leftarrow d} \end{bmatrix} \tag{7}$$

where each row of $C^{(h)}$ sums to $1/d$ and the matrix grand sum is 1. The ith diagonal entry of $C^{(h)}$, $\theta^{(h)}_{i\leftarrow i}$, measures the unilateral spillover (or loop) from variable i onto itself. We refer to this as the own variance share, $O^{(h)}_{i\leftarrow i}$. Meanwhile, the (i,j)th off-diagonal element of $C^{(h)}$ measures the spillover from variable j to variable i. The spillover from the rest of the system to variable i is:

$$F^{(h)}_{i\leftarrow\bullet} = \sum_{j=1, j\neq i}^{d} \theta^{(h)}_{i\leftarrow j} \tag{8}$$

while the spillover from the ith variable to the rest of the system is given by:

$$T^{(h)}_{\bullet\leftarrow i} = \sum_{j=1, j\neq i}^{d} \theta^{(h)}_{j\leftarrow i} \tag{9}$$

The total and net h-steps-ahead connectedness of variable i, $A^{(h)}_i$ and $N^{(h)}_i$, are defined as follows:

$$A^{(h)}_i = T^{(h)}_{\bullet\leftarrow i} + F^{(h)}_{i\leftarrow\bullet} \quad \text{and} \quad N^{(h)}_i = T^{(h)}_{\bullet\leftarrow i} - F^{(h)}_{i\leftarrow\bullet} \tag{10}$$

The spillover index introduced by Diebold and Yilmaz (2009) can be defined as:

$$S^{(h)} = \sum_{i=1}^{d} F^{(h)}_{i\leftarrow\bullet} \equiv \sum_{i=1}^{d} T^{(h)}_{\bullet\leftarrow i} \tag{11}$$

The spillover measures defined above focus on pairwise spillovers among individual variables. Following Greenwood-Nimmo et al. (2021), spillovers among variable groups can be evaluated by means of block aggregation of the connectedness matrix (7). For illustrative purposes, suppose that we wish to evaluate the connectedness among the M currencies in the model in a combined manner that encompasses all four variables relating to each currency. First, note that we may write:

$$C^{(h)} = \begin{bmatrix} B_{1\leftarrow 1} & B_{1\leftarrow 2} & \cdots & B_{1\leftarrow M} \\ B_{2\leftarrow 1} & B_{2\leftarrow 2} & \cdots & B_{2\leftarrow M} \\ \vdots & \vdots & \ddots & \vdots \\ B_{M\leftarrow 1} & B_{M\leftarrow 2} & \cdots & B_{M\leftarrow M} \end{bmatrix} \tag{12}$$

where:

$$B_{i\leftarrow j} = \begin{bmatrix} \theta^{(h)}_{r_i\leftarrow r_j} & \theta^{(h)}_{r_i\leftarrow v_j} & \theta^{(h)}_{r_i\leftarrow q_j} & \theta^{(h)}_{r_i\leftarrow k_j} \\ \theta^{(h)}_{v_i\leftarrow r_j} & \theta^{(h)}_{v_i\leftarrow v_j} & \theta^{(h)}_{v_i\leftarrow q_j} & \theta^{(h)}_{v_i\leftarrow k_j} \\ \theta^{(h)}_{q_i\leftarrow r_j} & \theta^{(h)}_{q_i\leftarrow v_j} & \theta^{(h)}_{q_i\leftarrow q_j} & \theta^{(h)}_{q_i\leftarrow k_j} \\ \theta^{(h)}_{k_i\leftarrow r_j} & \theta^{(h)}_{k_i\leftarrow v_j} & \theta^{(h)}_{k_i\leftarrow q_j} & \theta^{(h)}_{k_i\leftarrow k_j} \end{bmatrix}$$

for $i, j = 1, 2, \ldots, M$, where the block $B_{i\leftarrow i}$ collects all of the local spillover effects for currency i (i.e. all spillovers occurring among the risk and return measures for currency i), while the block $B_{i\leftarrow j}$ collects all of the spillover effects from currency j to currency i (i.e. all of the spillovers affecting the risk and return measures for currency i that originate from the risk and return measures for currency j). The total within-currency spillover effect for currency i is as follows:

$$\mathscr{W}^{(h)}_{i\leftarrow i} = \iota_4' B_{i\leftarrow i} \iota_4 \tag{13}$$

where ι_4 is a 4×1 vector of ones. As such, $\mathscr{W}^{(h)}_{i\leftarrow i}$ measures the proportion of the h-steps-ahead systemwide FEV accounted for by spillovers among the variables in $\check{y}_{it}$.[10] Meanwhile, the pairwise spillover from currency j to currency i can be measured by:

$$\mathscr{F}^{(h)}_{i\leftarrow j} = \iota_4' B_{i\leftarrow j} \iota_4 \tag{14}$$

It is straightforward to construct the aggregate spillover measures for currency i with respect to the other $M-1$ currencies as follows:

$$\mathscr{F}^{(h)}_{i\leftarrow \bullet} = \sum_{j=1, j\neq i}^{M} \iota_4' B_{i\leftarrow j} \iota_4 \tag{15}$$

$$\mathscr{T}^{(h)}_{\bullet\leftarrow i} = \sum_{j=1, j\neq i}^{M} \iota_4' B_{j\leftarrow i} \iota_4 \tag{16}$$

Given $\mathscr{F}^{(h)}_{i\leftarrow \bullet}$ and $\mathscr{T}^{(h)}_{\bullet\leftarrow i}$, the total and net connectedness of each currency can be easily computed as:

$$\mathscr{A}^{(h)}_i = \mathscr{T}^{h}_{\bullet\leftarrow i} + \mathscr{F}^{(h)}_{i\leftarrow \bullet} \text{ and } \mathscr{N}^{(h)}_i = \mathscr{T}^{(h)}_{\bullet\leftarrow i} - \mathscr{F}^{(h)}_{i\leftarrow \bullet} \tag{17}$$

where $\mathscr{A}^{(h)}_i$ measures the overall strength of currency i's linkages with the other currencies in the system and $\mathscr{N}^{(h)}_i$ reflects the role of currency i in the system as a net transmitter (receiver) if $\mathscr{N}^{(h)}_i > 0$ (if $\mathscr{N}^{(h)}_i < 0$). The aggregate spillover index at the currency level is obtained as

[10] It is straightforward to decompose these currency-level connectedness measures into own-variable effects (e.g. return-to-return spillovers) and cross-variable effects (e.g. variance-to-return spillovers) if desired. To see this, consider the within-currency connectedness of the ith currency, $\mathscr{W}^{(h)}_{i\leftarrow i}$, which may be decomposed to separate the own-variable FEV contributions within currency i, $\mathscr{O}^{(h)}_{i\leftarrow i}$, from the cross-variable effects, $\mathscr{C}^{(h)}_{i\leftarrow i}$, as follows:

$$\mathscr{O}^{(h)}_{i\leftarrow i} = \text{trace}\left(B^{(h)}_{i\leftarrow i}\right) \text{ and } \mathscr{C}^{(h)}_{i\leftarrow i} = \mathscr{W}^{(h)}_{i\leftarrow i} - \mathscr{O}^{(h)}_{i\leftarrow i}$$

where $\mathscr{O}^{(h)}_{i\leftarrow i}$ is the total own-variable spillover effect within currency i and $\mathscr{C}^{(h)}_{i\leftarrow i}$ represents spillovers between the log-return, variance, skewness and kurtosis innovations for currency i.

follows:

$$\mathscr{S}^{(h)} = \sum_{i=1}^{M} \mathscr{F}_{i\leftarrow\bullet}^{(h)} \equiv \sum_{i=1}^{M} \mathscr{T}_{\bullet\leftarrow i}^{(h)} \tag{18}$$

As shown by Greenwood-Nimmo et al. (2021), the aggregation procedure described above can be modified to accommodate any desired block structure in the connectedness matrix, $C^{(h)}$.

4. Results

Before we construct the network, it is first necessary to select the forecast horizon used to construct the GVD, h, as well as the window length to be used in the rolling-sample exercises, w. Initial experimentation with horizons in the set $h = \{5, 10, 15\}$ reveals that the network is largely insensitive to the choice of forecast horizon. Therefore, we select $h = 10$ trading days, which is a common choice in the literature. In selecting the window length for use in sections 4.2 and 4.3, we must balance two important considerations: (i) a longer window will yield more precise estimates of the VAR parameters in each rolling sample; and (ii) a shorter window will allow for greater timeliness in our rolling-sample analysis. We proceed by comparing the network obtained using window lengths in the set $w = \{200, 250, 300\}$ trading days, as in Greenwood-Nimmo et al. (2016). In practice, we find that the spillover statistics obtained in each case are similar, so we proceed on the basis that $w = 250$ without loss of generality. A concise comparison of the network statistics obtained from all pairwise combinations of $h = \{5, 10, 15\}$ and $w = \{200, 250, 300\}$ may be found in our Data Appendix.

4.1 Full-sample network statistics

We begin by summarising the connectedness among all d variables in the system on a disaggregate basis in the form of a heatmap in Figure 1. To facilitate interpretation of the heatmap, the variables in the model are gathered into moment groups, such that all of the returns appear first, followed by all of the variance innovations and so on. This has no effect on the results, because the GVD is order-invariant – it is purely a presentational choice.

Several features of the network are easily seen. First, bilateral spillovers are the dominant force in the network, accounting for 73.4% of systemwide FEV. The importance of spillover effects reflects the aggregate behaviour of FX investors, whose portfolio management decisions take account of differences in the risk-return profiles across currencies.

Second, although bilateral spillovers dominate the system on aggregate, the strongest individual spillovers lie on the prime diagonal of the connectedness matrix, indicating a substantial role of idiosyncratic (variable-specific) information.

Third, if one imagines partitioning the connectedness matrix conformably into 4^2 blocks of dimension $\frac{1}{4}d \times \frac{1}{4}d$, then the four blocks on the prime diagonal contain the majority of strong

bilateral spillovers in the system. This indicates that spillovers within moment groups (e.g. return-to-return spillovers) are generally stronger than spillovers between moment groups (e.g. variance-to-return spillovers). On a superficial level, this finding provides some support for the common practice of modelling return and risk spillovers separately in the empirical network literature (for an early example of this practice in the context of equity market spillovers, see Diebold and Yilmaz, 2009). However, developing separate network models in this way amounts to the assumption that the spillover effects contained in the off-diagonal blocks of the connectedness matrix are negligible. This is clearly not the case in Figure 1, where strong spillovers arise between moment groups, especially between returns and skewness innovations, as well as between variance and kurtosis innovations. The diagonal patterns seen in the top right and bottom left quadrants of the connectedness matrix indicate that these cross-moment spillovers are most pronounced within currencies (e.g. between the USDZAR return and the USDZAR skewness innovation), but the presence of off-diagonal blocks in the connectedness matrix indicates that cross-moment, cross-currency spillovers cannot be considered negligible. Hence, there is value in modelling risk and return spillovers jointly. A key contribution of our paper is to include a large number of EM and DM currency pairs in the same network. We show that this is important: spillovers between these markets are meaningful, as expected given the interconnectedness of global FX markets.

The block structure of Figure 1 indicates that there are strong spillovers between currencies. To develop intuition for the behaviour of each currency within the FX network, we summarise the connectedness among all M currencies over the full sample in Figure 2. This is obtained by block aggregation of the disaggregate connectedness matrix shown in Figure 1 following the procedure devised by Greenwood-Nimmo et al. (2021). For convenience, the currencies are ordered into four informal groups – non-commodity DMs, commodity DMs, commodity EMs and non-commodity EMs. Each currency exhibits a strong within-currency effect, although cross-currency spillovers account for 55.8% of systemwide FEV at the 10-days-ahead forecast horizon. The within-currency effect is larger than the sum of inward spillovers from the system for 4/8 EMs but for just 3/12 DMs, which suggests that DM currencies are more integrated within the FX network than EM currencies on average, perhaps due to their greater average liquidity. Several blocks can be discerned in Figure 2, with a particularly notable block of heavily traded European currencies (EUR, GBP and CHF), a broader block of European currencies (EUR, GBP, CHF, SEK, NOK and PLN) and a faint block of commodity currencies (AUD, CAD, NOK, NZD and ZAR). The absence of well-defined Asian or EM blocks is noteworthy. This might reflect the management of several of the EM currencies in our sample or a lack of clear EM or Asian blocks in FX market liquidity.[11]

[11] Using bid-ask spreads for the same currency pairs that we study, Olds et al. (2021) document a lack of co-movement in Asian or EM FX liquidity, although they find that common variation in major exchange rates is associated with common variation in FX market liquidity for the full sample of currency pairs. Understanding whether there is a linkage to FX trading volumes and FX spillovers and whether EM block correlation has been strengthening over time with the proliferation of automation platforms that distribute risks across EM currency blocks represents an interesting area for future research.

Figure 1: Disaggregate connectedness over the full sample

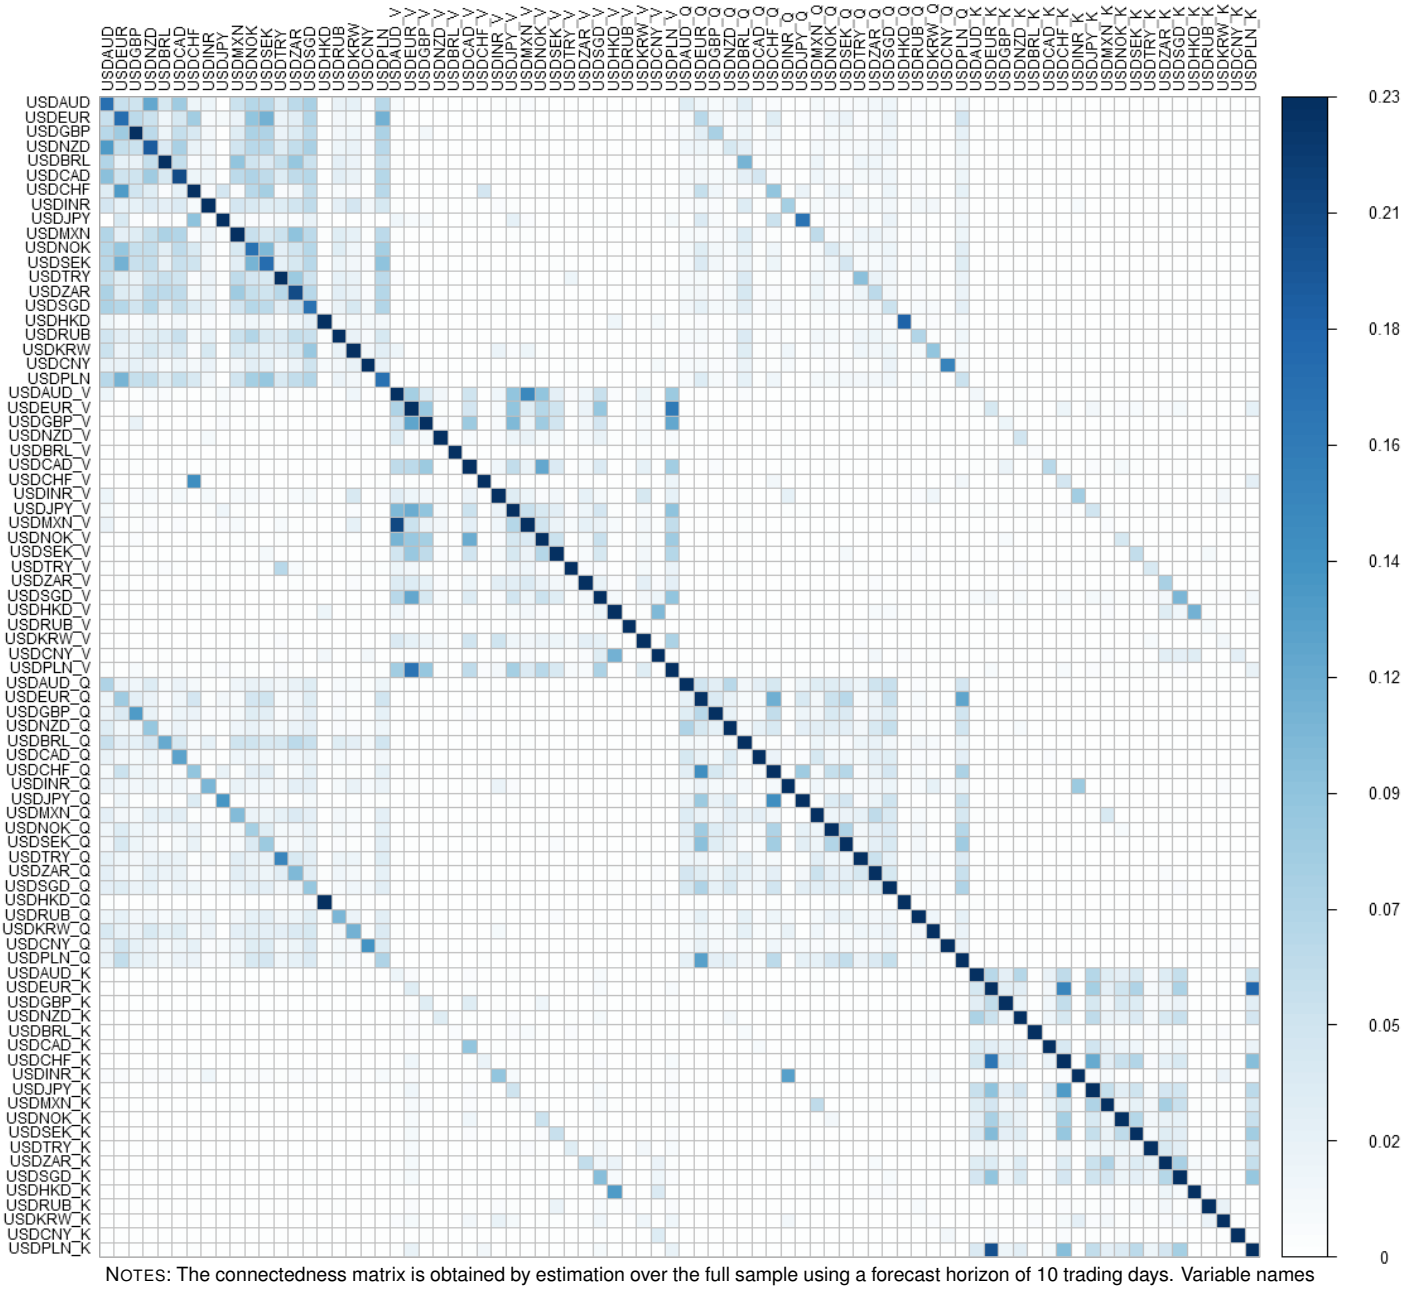

NOTES: The connectedness matrix is obtained by estimation over the full sample using a forecast horizon of 10 trading days. Variable names are abbreviated as follows: for currency XXX, USDXXX is the log-return, USDXXX_V is the realised variance innovation, USDXXX_Q is the realised skewness innovation and USDXXX_K is the realised kurtosis innovation. The depth of shading is proportional to the strength of the spillover effect.

To further illuminate the role of each currency in the network, in Figure 3 we plot the position of each currency in dependence-influence space. Following Greenwood-Nimmo et al. (2021), we define the dependence of the ith currency as $\mathscr{D}_i^{(h)} = \mathscr{F}_{i\leftarrow\bullet}^{(h)} / (\mathscr{F}_{i\leftarrow\bullet}^{(h)} + \mathscr{W}_{i\leftarrow i}^{(h)})$ and the influence of the ith currency as $\mathscr{I}_i^{(h)} = \mathscr{N}_i^{(h)} / (\mathscr{T}_{\bullet\leftarrow i}^{(h)} + \mathscr{F}_{i\leftarrow\bullet}^{(h)})$. Consequently, the dependence index takes values in the unit interval, with smaller (larger) values indicating that the ith currency is less (more) dependent on external conditions. Meanwhile, the influence index takes values in the interval $[-1, 1]$, with values close to 1 signifying a highly influential currency that strongly influences external conditions and values close to -1 signifying a currency that is strongly influenced by external conditions.

Figure 2: Connectedness among currencies over the full sample

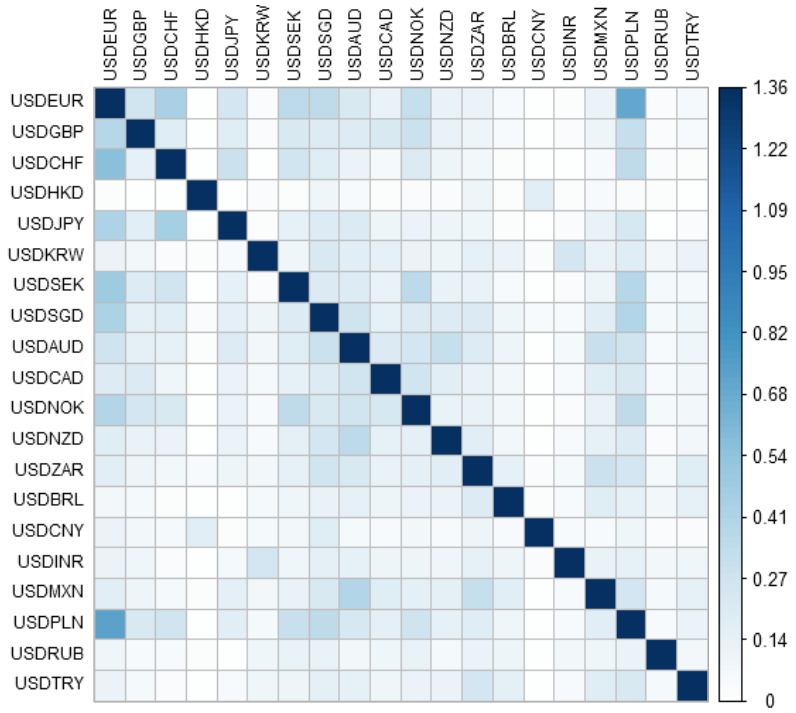

NOTES: The currency connectedness matrix is obtained by block aggregation of the 10-days-ahead connectedness matrix evaluated over the full sample. The depth of shading is proportional to the strength of the spillover effect.

First, consider the cluster of currencies at the upper right extreme of Figure 3: USDEUR, USDNOK and USDSEK. All three exhibit positive influence, which suggests that they play an important informational role in the global FX market. This is not surprising, as all three are members of the G10 group of highly liquid and heavily traded global currencies. They also exhibit strong external dependence, which at least partly reflects the strong bilateral spillovers among the European currencies documented in Figure 2.

At the other extreme, the USDCNY and the USDHKD record the lowest external dependence and influence among the currencies in our sample. For both currencies, their managed adjustment limits their fluctuations to a narrow range, and therefore affects the extent of their external dependence. In addition, the low influence of the yuan may result in part from our use of the onshore USDCNY exchange rate, as opposed to the offshore USDCNH exchange rate.[12]

The remaining currencies in our sample can be roughly partitioned into two groups. First, there is a group of DM currencies and commodity currencies for which the dependence index is typically greater than 0.5 and the influence index is typically close to zero. The other group

[12] The particular behaviour of the USDHKD in our model may also be affected by Hong Kong's unique role in yuan trading. Hong Kong is the principal location for conversion between the onshore CNY and offshore CNH. Offshore demand for Chinese equities is an important driver of demand for the HKD, because the Shanghai Connect and Shenzen Connect programs allow overseas investors to trade eligible shares on the Shanghai and Shenzen stock exchanges subject to a daily quota. Prior to the establishment of this programme, offshore trading in mainland Chinese stocks was largely prohibited. These trades are settled in CNY, with many offshore investors converting from a foreign currency to HKD or CNH and then to CNY. This implies that a significant source of demand for the HKD is driven by demand for the onshore CNY, generating spillovers between the two currencies (note that the bilateral spillovers between USDHDK and USDCNY are among the strongest spillovers affecting either currency in Figure 2).

Figure 3: Influence and dependence of currencies over the full sample

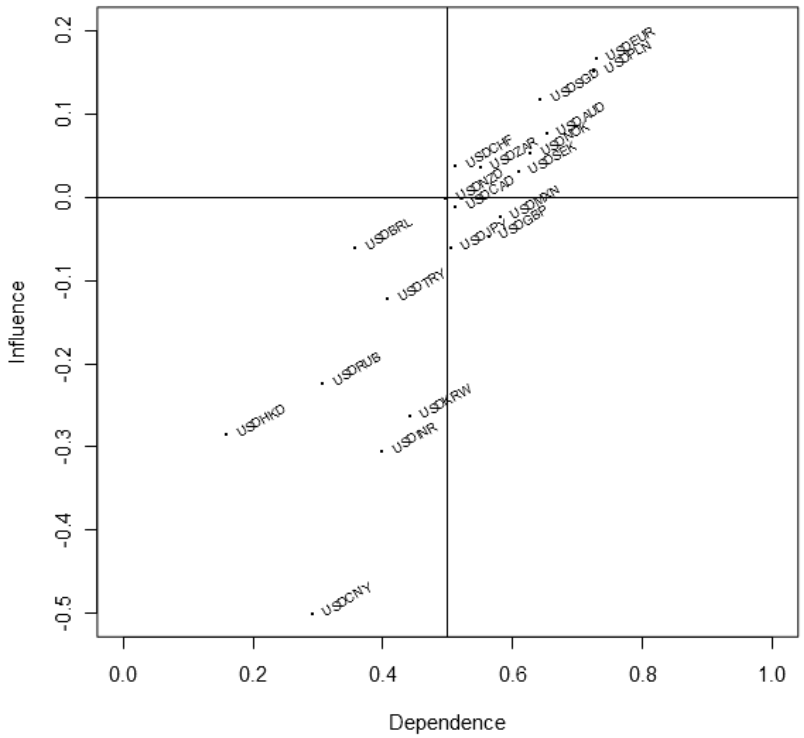

NOTES: The dependence of the ith currency is computed as $\mathscr{D}_i^{(h)} = \mathscr{F}_{i\leftarrow\bullet}^{(h)} / (\mathscr{F}_{i\leftarrow\bullet}^{(h)} + \mathscr{W}_{i\leftarrow i}^{(h)})$, while the influence of the ith currency is computed as $\mathscr{I}_i^{(h)} = \mathscr{N}_i^{(h)} / (\mathscr{T}_{\bullet\leftarrow i}^{(h)} + \mathscr{F}_{i\leftarrow\bullet}^{(h)})$. Both are constructed using a forecast horizon of 10 trading days over the full sample.

contains mostly EM currencies, many with dependence indices less than 0.5 and all with negative influence indices, often substantially so, which suggests that many EMs play a relatively weak informational role in the FX network.

The behaviour of the ZAR is interesting. Aside from the PLN, which occupies a unique position among the group of EM currencies due to Poland's membership of the European Union, the ZAR and the MXN are the only EM currencies to record a positive influence score. Over the full sample, on average, the ZAR is a modest net transmitter of spillovers to other currencies, with a net spillover effect of 0.426%. However, separating ZAR spillovers with respect to DM and EM currencies yields an interesting finding – the net spillover from the ZAR to the remaining EM currencies in the sample is 0.404%, compared with a net spillover of just 0.022% relative to the 12 DM currencies in the sample. This suggests that the ZAR may play an important role in EM portfolios. Market analysts regularly point to the role of the ZAR as a bellwether EM currency for investors but, as far as we are aware, this is the first academic study to provide evidence of this relationship in terms of risk and return spillovers. The ZAR's role as a bellwether could be attributed to the relatively sophisticated financial markets in South Africa and the relatively high FX liquidity of the ZAR relative to other EMs.

4.2 Rolling-sample network statistics

Full-sample network analysis provides a static impression of the average structure of the FX network over our sample. To investigate the evolution of the network over time, we now turn to rolling-sample analysis. Figure 4 reports the 10-days-ahead disaggregate and currency-level spillover indices evaluated over 250 trading day rolling samples. The two spillover indices share similar dynamics – both show notable time variation, with elevated spillover activity in periods of stress, including the global financial crisis, the European debt crisis, the wake of the Brexit referendum and during the COVID-19 pandemic. The finding that currency markets become more strongly interconnected during periods of stress is consistent with the results in Greenwood-Nimmo et al. (2016) and suggests more frequent or more substantial portfolio rebalancing by FX investors during periods of elevated uncertainty. Furthermore, a careful inspection of Figure 4 reveals that aggregate spillovers tend to rise faster than they fall, which suggests that currency market spillovers may respond asymmetrically to good and bad news.

Figure 4: Connectedness among currencies over 250-day rolling samples

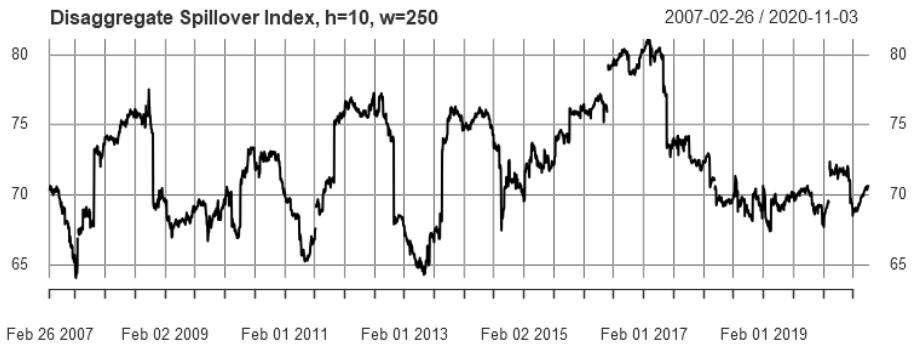

(a) Disaggregate spillover index

(b) Currency spillover index

NOTES: Panel (a) reports the rolling-sample spillover index computed from the disaggregate connectedness matrix, while panel (b) reports the corresponding spillover index computed from the currency connectedness matrix. In both cases, the forecast horizon is 10 trading days and the window length is 250 trading days. The dates reported on the horizontal axis relate to the end date of each rolling-sample. The unit of measurement is %. Nine out of 3 413 rolling samples yield unstable solves of the VAR model and are therefore excluded from the analysis, which results in some discontinuities in the time series plots.

A broadly similar pattern of time variation characterises the inward and outward spillovers between the ZAR and the DM/EM currency groups (Figure 5). However, while the ZAR acts as a net source of spillovers with respect to the DMs only in periods of stress (such as during the global financial crisis), it is a significant source of net spillovers with respect to the EMs over

much of our sample period. We interpret the latter finding as further evidence of the ZAR's role as a bellwether EM, which leads it to play an important informational role among EM currency portfolios in general. This same phenomenon may explain why we observe positive net spillovers from the ZAR to DMs in periods of stress, which often generate reversals in capital flows between EMs and DMs due to rising risk-off sentiment.

Figure 5: Average connectedness of the ZAR with DM/EM currencies

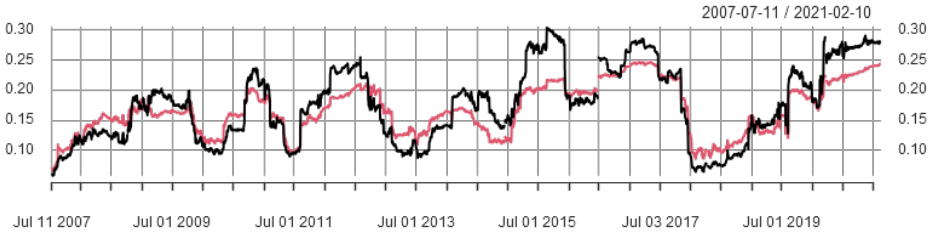

(a) ZAR vs. DM currencies (average per currency)

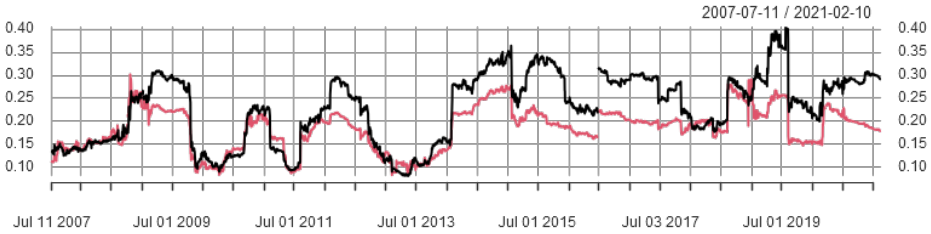

(b) ZAR vs. EM currencies (average per currency)

NOTES: Panel (a) reports the ZAR's average outward (black) and inward (red) spillovers with respect to DM currencies over 250-day rolling samples. Panel (b) repeats the same exercise with respect to all EM currencies apart from the ZAR. In both cases, the forecast horizon is 10 trading days and the dates reported on the horizontal axis relate to the end date of each rolling sample. The unit of measurement is %. Nine out of 3 413 rolling samples yield unstable solves of the VAR model and are therefore excluded from the analysis, which results in some discontinuities in the time series plots.

4.3 Using the FX network for narrative analysis

In addition to providing a framework for monitoring FX risk transmission, our model can be used to contextualise abrupt movements in exchange rates. In this section, we demonstrate that changes in the structure of the FX network occur at the time of known risk events for the ZAR and we show that these changes can shed light on the key sources of risk affecting the ZAR. We consider three classes of event: (i) idiosyncratic domestic events; (ii) idiosyncratic foreign events; and (iii) global events.

4.3.1 Idiosyncratic domestic events

We consider the following selection of domestic political events: (i) the dismissal of Nhlanhla Nene as Finance Minister on 9 December 2015; (ii) the appointment of Pravin Gordhan as Nene's replacement on 14 December 2015; (iii) the revelation that Gordhan would be required to report to the Directorate for Priority Crime Investigation to answer questions regarding his work at the South African Revenue Service on 23 August 2016; and (iv) the dismissal of Gordhan as Finance Minister on 30 March 2017.[13]

[13] Because Gordhan was dismissed as part of a late-night cabinet reshuffle on Thursday, 30 March 2017, we treat the event as though it occurred on Friday, 31 March 2017. This adjustment is important because we define a

To illustrate our method, in Figure 6 we provide a detailed visualisation of the change in the FX network at the time of Nene's dismissal. The figure reveals that the dismissal of Nene was associated with an abrupt increase in the within-currency effect for South Africa, which is met by a reduction in inward spillovers to the ZAR from many currencies. In addition, outward spillovers from the ZAR to almost all other currencies drop substantially at this time. These phenomena collectively indicate that the dismissal of Nene, which came as a shock to many market participants, raised political uncertainty and generated a marked increase in South African idiosyncratic risk. It is also natural to see that the positive net spillover from the ZAR to the other currencies – both EM and DM – reported in Figure 5 dropped substantially at this time. An important implication of this finding is that the ZAR's role as a bellwether EM currency diminishes in the face of increasing South African idiosyncratic risk.

Figure 6: Change in spillovers among currencies when Nene was dismissed

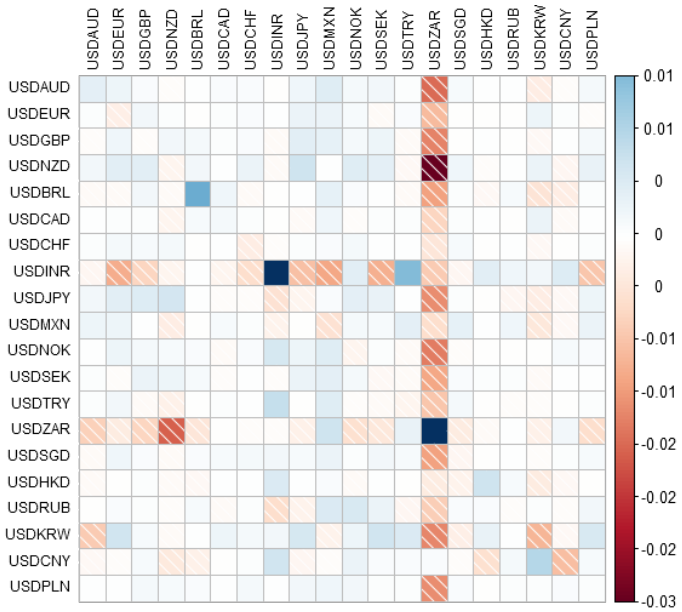

NOTES: The $\{i, j\}$th cell of the heatmap reports the change in the value of the $\{i, j\}$th element of the 10-days-ahead spillover matrix between the 250-day rolling-sample ending on the trading day prior to Nene's dismissal and the 250-day rolling-sample ending on the day of Nene's dismissal. The depth of shading is proportional to the magnitude of the change. The unit of measurement is %.

Given that our focus in this section is on spillovers involving the ZAR, it is not necessary to report the change in spillover activity among all pairwise combinations of countries, as in Figure 6. Figure 7(a) offers a more concise summary of the change in spillover activity at the time of Nene's dismissal that focuses exclusively on spillovers to/from the ZAR. The stark increase in the within-currency effect coupled with the reduction in external connectedness of the ZAR is easily seen.

The next event that we consider is the appointment of Gordhan as Nene's replacement. Figure 7(b) reveals a large drop in the within-currency effect for the ZAR, reflecting a reduction in South African idiosyncratic risk as markets were calmed by news of a safe choice to replace Nene. Note, however, that the reduction in the within-currency effect (-0.024% of systemwide FEV) is not sufficient to offset the increase at the time of Nene's dismissal (+0.135%).[14] The

trading day as starting at 21:05 GMT.

[14] This is consistent with contemporary commentary, in which Gordhan's appointment was typically viewed

Figure 7: Change in spillovers to/from the ZAR on domestic event days

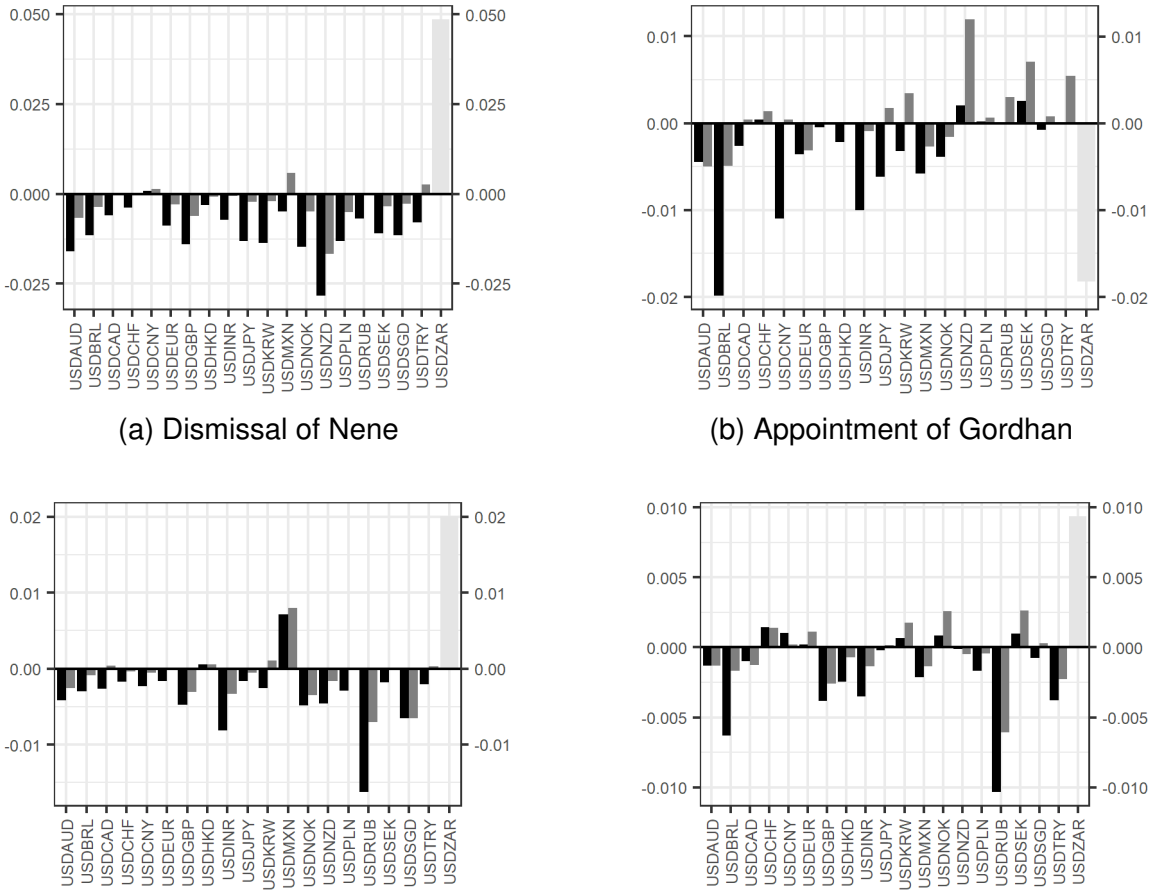

(a) Dismissal of Nene

(b) Appointment of Gordhan

(c) Questioning of Gordhan

(d) Dismissal of Gordhan

NOTES: The figure summarises the change in the within (pale grey), inward (dark grey) and outward (black) spillovers for the ZAR between the 250-day rolling-sample ending on the trading day prior to the specified event and the 250-day rolling-sample ending on the day of the specified event. The forecast horizon is 10 trading days and the unit of measurement is %.

reduction in the within-currency effect is offset by an increase in spillover onto the ZAR from several other currencies, most notably the CHF, INR, NOK and TRY. The figure also reveals a general reduction in outward spillovers from the ZAR to other currencies, with particularly large changes with respect to the AUD and NZD, as well as the EUR, MXN and SGD. This is consistent with the results reported in Figure 5, which shows that outward spillovers from the ZAR to other currencies fall sharply at the time of Nene's dismissal and remain low over several rolling samples.

Next, we consider the change in the FX network at the time of the escalating controversy surrounding Gordhan's involvement with an anti-crime unit during his tenure at the South African Revenue Service. Figure 7(c) shows that South African idiosyncratic risk spiked when news broke that Gordhan would be questioned by the Directorate for Priority Crime Investigation, echoing the impact of Nene's dismissal. The dismissal of Gordhan as part of a cabinet reshuffle on 30 March 2017 is also reflected in a similar way in the FX network, as shown in Figure 7(d).

favourably but not as an immediate resolution to the political crisis brought on by Nene's dismissal. See, for example, https://www.reuters.com/article/us-safrica-gordhan-idUSKBN0TW0T320151213.

Collectively, these results suggest that periods of heightened political uncertainty in South Africa drive up idiosyncratic risk and weaken the ZAR's connectedness with other currencies, as investors come to perceive that changes in the risk-return profile of the ZAR reflect idiosyncratic factors to a greater degree than marketwide trends. This is also consistent with Reserve Bank modelling of the ZAR, which identifies the ZAR's depreciation around the time of 'Nenegate' as being driven by South Africa-specific risk (see South African Reserve Bank, 2019).[15] A corollary to this observation is that if the ZAR only acts as a bellwether EM currency during periods of moderate South African idiosyncratic risk, then its role in EM portfolios is likely to diminish if South African fundamentals continue to weaken.

4.3.2 Idiosyncratic foreign events

We now consider two foreign events that spilled over strongly to the ZAR. First, the Chinese stock market slump of 11 January 2016 caused risk-off sentiment that spread to the ZAR via trading in the Asian session. Figures 8(a) and (c) reveal that the HKD played a key role in transmitting the shock, which is a natural finding given the unique role played by the HKD in converting between the onshore Yuan and foreign currencies. With the exception of the HKD, bilateral spillovers between the ZAR and most other currencies fell sharply at this time, as many EM investors unwound their ZAR positions.

Figures 8(b) and (d) relate to the sharp depreciation of the TRY on 13 August 2018 in response to a combination of economic and geopolitical tensions, including an announcement by President Trump of his intention to double US tariffs on Turkish steel and aluminium imports. The collapse of the TRY generated fears of contagion among EMs, with the ZAR being viewed as particularly vulnerable due to weak economic fundamentals and a growing debt stock. The figure reveals that the increased spillover from the TRY to the ZAR is met almost one-to-one by a reduction in the within-currency effect for the ZAR, which indicates a strong influence of TRY spillovers on the risk-return profile of the ZAR at this time. This is consistent with the ZAR acting as a proxy for EMs: the ZAR often reacts to and magnifies broader EM currency developments when these shocks are common to other EM currencies (South African Reserve Bank, 2019).[16] Some South Africa-specific shocks that did not give rise to EM-wide financial contagion such as the dismissal of Minister Nene did not appear to generate large FX spillovers.

4.3.3 Global events

We close by analysing an event with global ramifications for the FX market. Figure 9(a) shows the change in the FX network over the period of the Brexit referendum. The plot reveals a surge in outward spillovers from the GBP, coupled with a widespread rise in bilateral connect-

[15] There is also evidence of a weakening of the ZAR's role as a gauge of global risk appetite with the heightened perception of South Africa-specific credit risk. See https://www.economist.com/finance-and-economics/2013/02/09/vix-inhaler-no-more for details.

[16] This is consistent with the finding in Olds et al. (2021) of a strong correlation between ZAR and TRY liquidity.

Figure 8: Change in spillovers to/from the ZAR on foreign event days

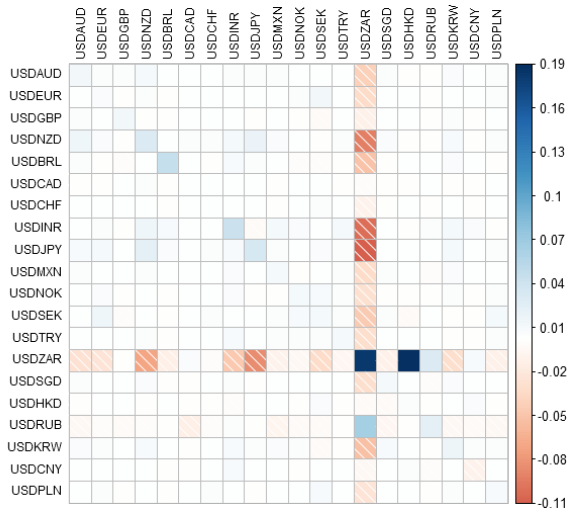

(a) Chinese stock market slump

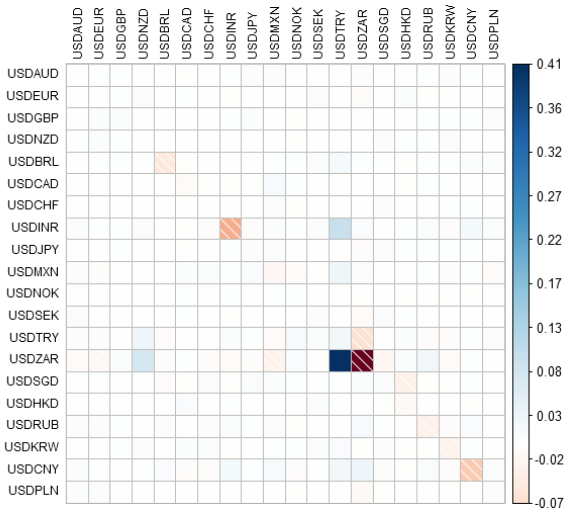

(b) Turkish lira slump

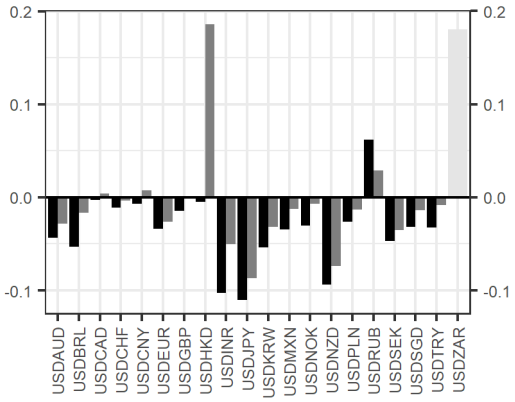

(c) Chinese stock market slump

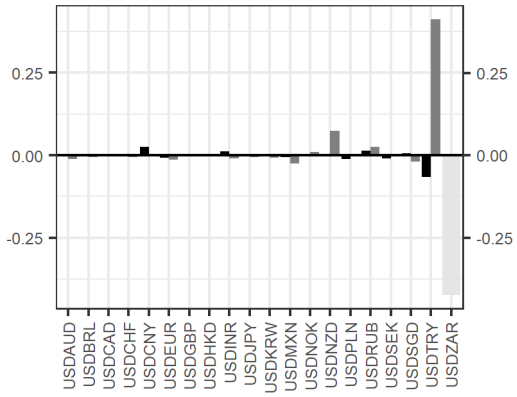

(d) Turkish lira Slump

NOTES: Panels (a) and (b) show heatmaps, the $\{i, j\}$th cell of the which reports the change in the value of the $\{i, j\}$th element of the spillover matrix between the 250-day rolling sample ending on the day of the specified event relative to the 250-day rolling sample ending on the trading day prior to the specified event. The depth of shading is proportional to the magnitude of the change. Panels (c) and (d) summarise the change in the within (pale grey), inward (dark grey) and outward (black) spillovers for the ZAR between the 250-day rolling-sample ending on the day of the specified event relative to the 250-day rolling-sample ending on the trading day prior to the specified event. The forecast horizon is 10 trading days and the unit of measurement is % in all panels.

edness throughout the network. Within-currency effects fall substantially for almost all currencies, reflecting the global impact of the shock. Figure 9(b) focuses on the change in ZAR spillovers observed over the period of the Brexit referendum and reveals a massive drop in the within-currency effect (the largest of any event that we consider) coupled with a near-universal increase in both inward and outward bilateral spillovers. As in the previous sub-sections, this demonstrates how our framework can be used to assess the sensitivity of currencies to shocks arising in other FX markets and to provide narratives contextualising currency movements.

Figure 9: Change in spillovers to/from the ZAR over the period of the Brexit referendum

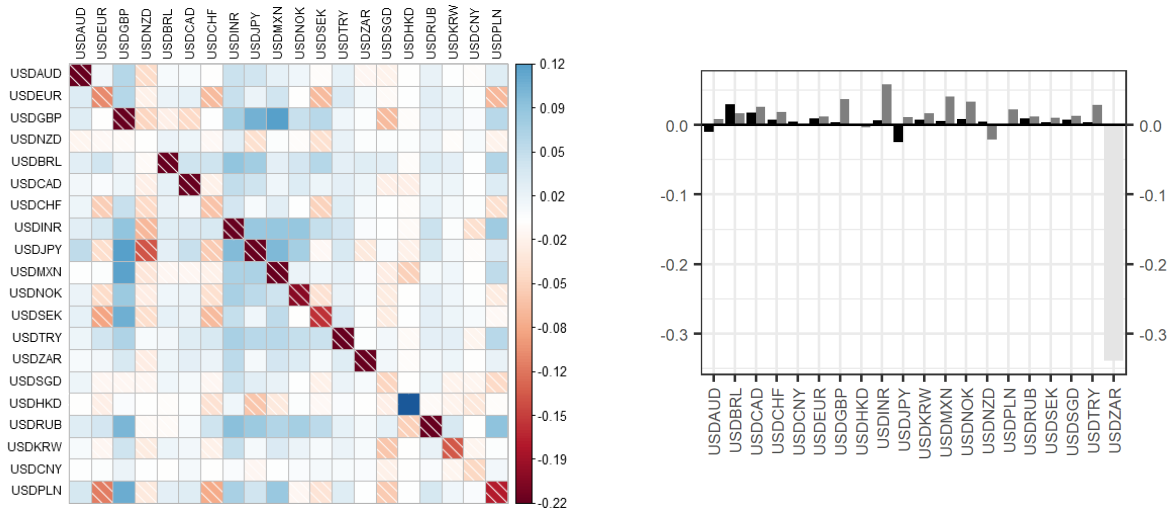

(a) Systemwide connectedness (b) ZAR connectedness

NOTES: Panel (a) shows a heatmap, the $\{i, j\}$th cell of the which reports the change in the value of the $\{i, j\}$th element of the spillover matrix between the 250-day rolling-sample ending on the trading day after the Brexit referendum relative to the 250-day rolling-sample ending on the trading day before the Brexit referendum. The depth of shading is proportional to the magnitude of the change. Panels (c) and (d) summarise the change in the within (pale grey), inward (dark grey) and outward (black) spillovers for the ZAR between the 250-day rolling-sample ending on the trading day after the Brexit referendum relative to the 250-day rolling-sample ending on the trading day before the Brexit referendum. The forecast horizon is 10 trading days and the unit of measurement is % in both panels. Note that the rolling-sample ending on the day of the Brexit referendum yields and unstable solve, which reflects the magnitude of the shock to global currency markets.

5. Conclusion

We examined the dynamic interactions among FX returns and innovations to realised measures of risk for a group of 20 major currencies. We computed daily log-returns and realised measures of volatility, skewness and kurtosis for 20 major currencies quoted against the USD. We developed a network model that captures the dynamic interactions among the daily log-returns and realised risk measures and showed how it can be used to study the transmission of risk between currency markets.

We showed that the model is able to highlight risk transmission channels in a timely manner during FX flash crashes and periods of heightened financial market uncertainty. Focusing on the ZAR, we showed that variations in the risk-return profile of the USDZAR correlate with variations in the risk-return profile of many other currencies, and that this is especially true of the EM currencies in our sample. We interpret this as evidence of the ZAR's role as a bellwether EM currency, at least in times when South African idiosyncratic risk is not elevated.

Our results indicate that global conditions have an important impact on the ZAR. This shows potentially limited effectiveness of interventions aimed at dampening currency volatility. Our results do, however, suggest that the ZAR also reacts strongly and rapidly to South Africa-specific shocks, such as bouts of political instability. This is consistent with the exchange rate playing a shock absorber role in the economy (as argued by Soobyah and Steenkamp 2019 and Loewald 2021).

Our approach could be used in future research into the high frequency dynamics of the ZAR (such as understanding movements of the ZAR around Monetary Policy Commitee decisions)

or to develop macroeconomic network models to study the cross-border propagation of inflation and growth shocks.

References

Andersen, T G, Bollerslev, T, Diebold, F X and Labys, P. 2003. 'Modeling and forecasting realized volatility'. *Econometrica* 71: 579–625.

Arezki, R, Dumitrescu, E, Freytag, A and Quintyn, M. 2014. 'Commodity prices and exchange rate volatility: Lessons from south africa's capital account liberalization'. *Emerging Markets Review* 19: 96–105.
URL https://ideas.repec.org/a/eee/ememar/v19y2014icp96-105.html

Billio, M, Getmansky, M, Lo, A and Pelizzon, L. 2012. 'Econometric measures of connectedness and systemic risk in the finance and insurance sectors'. *Journal of Financial Economics* 104: 535–559.

Bostanci, G and Yilmaz, K. 2020. 'How connected is the global sovereign credit risk network?' *Journal of Banking and Finance* .

Cai, F, Howorka, E and Wongswan, J. 2008. 'Informational linkage across trading regions: Evidence from foreign exchange markets'. *Journal of International Money and Finance* 27: 1212–1243.

Demirer, M, Diebold, F X, Liu, L and Yilmaz, K. 2017. 'Estimating global bank network connectedness'. *Journal of Applied Econometrics* 33: 1–15.

Diebold, F and Yilmaz, K. 2009. 'Measuring financial asset return and volatility spillovers, with application to global equity markets'. *Economic Journal* 119: 158–171.

Diebold, F and Yilmaz, K. 2014. 'On the network topology of variance decompositions: measuring the connectedness of financial firms'. *Journal of Econometrics* 182: 119–134.

Engle, R F, Ito, T and Lin, W. 1990. 'Meteor showers or heat waves? heteroskedastic intra-daily volatility in the foreign exchange market'. *Econometrica* 58: 525–524.

Farrell, G, Hassan, S and Viegi, N. 2012. 'The high-frequency response of the rand-dollar rate to inflation surprises', *Working Papers 5028*, South African Reserve Bank.
URL https://ideas.repec.org/p/rbz/wpaper/5028.html

Greenwood-Nimmo, M J, Huang, J and Nguyen, V H. 2019a. 'Financial sector bailouts, sovereign bailouts and the transfer of credit risk'. *Journal of Financial Markets* 42: 121–142.

Greenwood-Nimmo, M J, Nguyen, V H and Rafferty, B. 2016. 'Risk and return spillovers among the G10 currencies'. *Journal of Financial Markets* 31: 43–62.

Greenwood-Nimmo, M J, Nguyen, V H and Shin, Y. 2019b. 'Quantifying informational linkages in a global model of currency spot market', in *Advances in Applied Financial Econometrics*, edited by J Chevalier, S Goutte, D Guerreiro, S Saglio and B Sanhaji, Routledge, London.

Greenwood-Nimmo, M J, Nguyen, V H and Shin, Y. 2021. 'Measuring the connectedness of the global economy'. *International Journal of Forecasting* 37: 899–919.

Greenwood-Nimmo, M J and Tarassow, A. 2020. 'Bootstrap-based probabilistic analysis of spillover scenarios in macroeconomic and financial networks', Mimeo: University of Melbourne.

Hassan, S. 2015. 'Speculative flows, exchange rate volatility and monetary policy: the south african experience', *Working Papers WP/15/02*, South African Reserve Bank.

International Monetary Fund. 2020. *World Economic Outlook: a long and difficult ascent*, International Monetary Fund, Washington D.C.

Loewald, C. 2021. 'Macro works: a decision-tree approach to exchange rate policy', *Working Paper 21-10*, South African Reserve Bank.

Nicholson, W B, Matteson, D S and Bien, J. 2017. 'Varx-l: structured regularization for large vector autoregressions with exogenous variables'. *International Journal of Forecasting* 33: 627–651.

Olds, T, Steenkamp, D and van Jaarsveld, R. 2021. 'The impact of global FX liquidity on the rand'. *SARB Economic Note* 21-02.

Pesaran, M and Shin, Y. 1998. 'Generalized impulse response analysis in linear multivariate models'. *Economics Letters* 58: 17–29.

Soobyah, L and Steenkamp, D. 2019. 'The role of the rand as a shock absorber', *Working Paper 19-02*, South African Reserve Bank.

South African Reserve Bank. 2019. 'Monetary policy review (october, box 4)' .

Tibshirani, R J. 1996. 'Regression shrinkage and selection via the lasso'. *Journal of the Royal Statistical Society, Series B* 58: 267–288.

Appendices

A Robustness to window and horizon selection

Figures 10 and 11 indicate that neither the choice of window length for the rolling-sample analysis nor the choice of forecast horizon used to construct the GVD exerts a dominant influence on our results.

Figure 10: Robustness to variations in the rolling-window and forecast horizon

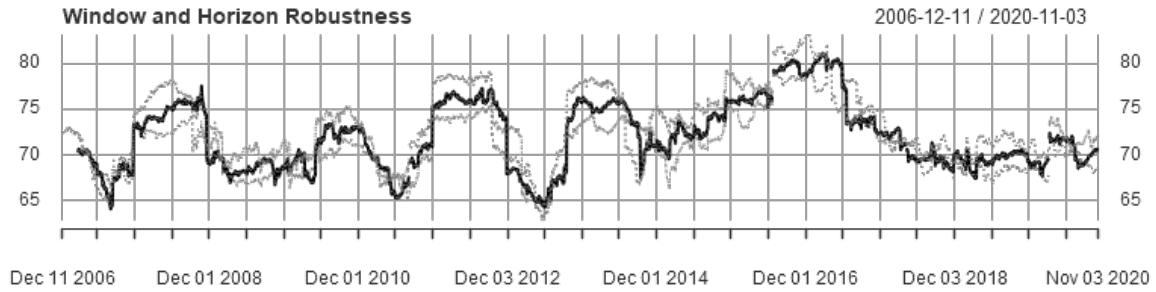

NOTES: The figure plots the disaggregate spillover indices obtained using all pairwise combinations of $w = \{200, 250, 300\}$ and $h = \{5, 10, 15\}$ trading days. Results for our baseline setting with $w = 250$ and $h = 10$ trading days are shown in black, while results using the other combinations of parameters are shown in grey. The unit of measurement is %.

Figure 11: Pairwise commmon sample correlation among spillover indices obtained using different rolling-windows and forecast horizons

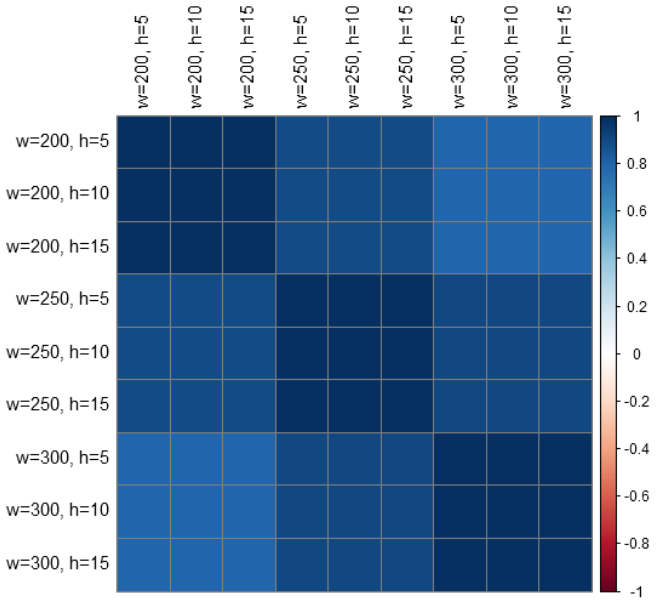

NOTES: The figure shows the pairwise common sample correlation between the disaggregate spillover indices obtained using all pairwise combinations of $w = \{200, 250, 300\}$ and $h = \{5, 10, 15\}$ in the form of a heatmap.

Figure 10 shows that the disaggregate spillover indices obtained using all pairwise combinations of $w = \{200, 250, 300\}$ and $h = \{5, 10, 15\}$ trading days exhibit similar dynamics and share common peaks and troughs. To the naked eye, Figure 10 appears to show only three

lines because, for a given window length, the choice of forecast horizon has a negligible effect on the spillover index, meaning that the three spillover indices for a given window length are almost indistinguishable from one another. This is reflected in a clear block structure in Figure 11, as only changes in the window length result in any notable variation in pairwise correlation among the spillover indices. Nonetheless, all of the correlations reported in Figure 11 are strongly positive: to three decimal places, the lowest pairwise correlation is 0.845, the median is 0.927 and the highest is 1.000 (excluding diagonal elements).

Additional robustness tests based on aggregated spillover indices are available from the authors on request.