

Bettendorf, Timo; Heinlein, Reinhold

Article — Published Version

Connectedness between G10 currencies: Searching for the causal structure

International Journal of Finance & Economics

Provided in Cooperation with:

John Wiley & Sons

Suggested Citation: Bettendorf, Timo; Heinlein, Reinhold (2022) : Connectedness between G10 currencies: Searching for the causal structure, International Journal of Finance & Economics, ISSN 1099-1158, John Wiley & Sons, Ltd., Chichester, UK, Vol. 28, Iss. 4, pp. 3938-3959, <https://doi.org/10.1002/ijfe.2629>

This Version is available at:

<https://hdl.handle.net/10419/288116>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by/4.0/>

RESEARCH ARTICLE

Connectedness between G10 currencies: Searching for the causal structure

Timo Bettendorf¹  | Reinhold Heinlein²

¹DG Economics, Deutsche Bundesbank, Frankfurt am Main, Germany

²Bristol Business School, University of the West of England, Bristol, UK

Correspondence

Timo Bettendorf, DG-Economics, Deutsche Bundesbank, Frankfurt am Main, Germany.

Email: timo.bettendorf@bundesbank.de

Reinhold Heinlein, Bristol Business School, University of the West of England, Bristol, UK.

Email: reinhold.heinlein@uwe.ac.uk

Abstract

This paper presents a new approach for modelling the connectedness between asset returns. We adapt the measure of Diebold and Yilmaz, which is based on the forecast error variance decomposition of a VAR model. However, their connectedness measure hinges on critical assumptions with regard to the variance–covariance matrix of the error terms. We propose to use a more agnostic empirical approach, based on a machine learning algorithm, to identify the contemporaneous structure. In a Monte Carlo study, we compare the different connectedness measures and discuss their advantages and disadvantages. In an empirical application we analyse the connectedness between the G10 currencies. Our results suggest that the US dollar as well as the Norwegian krone are the most independent currencies in our sample. By contrast, the Swiss franc and New Zealand dollar have a negligible impact on other currencies. Moreover, a cluster analysis suggests that the currencies can be divided into three groups, which we classify as: commodity currencies, European currencies and safe haven/carry trade financing currencies.

KEYWORDS

connectedness, exchange rates, graph theory, networks

1 | INTRODUCTION

Triggered by the seminal work of Diebold and Yilmaz (2009, 2014), the measurement of spillover effects and connectedness between asset returns has gained popularity in the economic literature. Their approach, which is based on the forecast error variance decomposition (FEVD) of a VAR model, hinges on critical assumptions with regard to a recursive ordering of the variables (e.g., Cholesky). In this paper, we propose an alternative and more agnostic approach to modelling the connectedness between asset returns, which is based on a causal search algorithm that imposes no a priori recursive ordering. We compare its

properties with those of other identification measures using a Monte Carlo experiment and apply it to the G10 currencies.¹

Given the new procedure, our first goal is to estimate the network structure between nine currencies vis-à-vis an appropriate numéraire currency (i.e., pound sterling).¹ We focus on the network structure and not the dynamics of a connectedness measure, because we aim to understand the relationships between currencies. Such estimates provide important information for policy makers and practitioners. The network indicates the extent to which a certain currency or group of currencies is affected by domestic and foreign shocks. In this sense, it

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *International Journal of Finance & Economics* published by John Wiley & Sons Ltd.

helps to gain a better understanding of potential contagion. The second goal is to utilise spillover intensities in order to identify clusters which can be interpreted as currency blocs or groups of common influence factors such as target currencies for carry trades.

Research in this area has created several extensions to the original work by Diebold and Yilmaz (2009), who estimated the return and volatility spillovers with respect to global equity markets. Diebold and Yilmaz (2014) studied the connectedness of financial institutions during the financial crisis period. In Diebold and Yilmaz (2015), they also estimated the connectedness between returns of other asset classes such as bilateral exchange rates, for instance. The approach is based on the idea that the network—or the spillover effects—between asset returns can be estimated given a FEVD of a vector autoregressive (VAR) model. Diebold and Yilmaz (2009) suggest orthogonalising the VAR model residuals with the help of a Cholesky decomposition, while pointing to the problem that the results of such a factorisation depend on the ordering of the variables, as zero restrictions are imposed on the upper triangular contemporaneous matrix without any theoretical or statistical motivation (i.e., preventing contemporaneous spillover effects between certain variables). Given the arbitrariness with respect to the ordering of the variables, the estimated model may not capture the spillover effects correctly. Instead of choosing one specific ordering, Klößner and Wagner (2014) propose considering all possible variable permutations. They replicate the paper by Diebold and Yilmaz (2009) and show that given different permutations, differences in spillover intensity can be large. Their approach, however, is not only computationally intensive, but also induces a high degree of model uncertainty. In other words, Klößner and Wagner (2014) average on many misspecified models and one correct model, which is unknown.

Apart from Cholesky decompositions, the literature also employs generalised impulse response functions (see Pesaran & Shin, 1998) in order to obtain variance decompositions which are invariant to the ordering of the variables, see for example Diebold and Yilmaz (2012, 2014). Greenwood-Nimmo et al. (2016), for instance, use generalised variance decompositions to study exchange rate return and volatility connectedness. In a rolling-window approach, Greenwood-Nimmo et al. (2017) apply the generalised approach in order to analyse the change in European debt connectedness distributions over time. This approach, however, has two shortcomings: First, the explained shares of forecast error variance do not sum to unity. In order to avoid re-scaling the shares, Lanne and Nyberg (2016) propose an alternative generalised FEVD which yields shares summing up to unity by construction. However, shocks are not orthogonalised. Second,

and more importantly, the approach is unable to model contemporaneous causal linkages. The weaknesses of generalised variance decompositions have been pointed out by De Santis and Zimic (2018), who instead propose absolute magnitude restrictions to identify SVAR models. They show that generalised variance decompositions tend to overestimate connectedness.

An alternative and more agnostic approach is identification with the help of causal search algorithms from the machine learning literature. Such an approach for structural VAR models has been suggested by Swanson and Granger (1997). Demiralp and Hoover (2003) introduced the causal search methods for identification. These algorithms use information from the reduced VAR residuals in order to uncover the contemporaneous causal structure. Applications can be seen in Heinlein and Krolzig (2012) and Demiralp et al. (2014). We follow this literature on empirical identification and systematically analyse in a Monte Carlo experiment as well as in an application on returns of G10 currencies how the identification strategy impacts on the measures of connectedness. To the best of our knowledge, the only papers using an empirical identification strategy in the connectedness literature are Scida (2018) and Yang et al. (2021), but they do not systematically study the impact of this approach or compare the empirical identification with other identification methods. The machine learning approach is very appealing because it does not require any prior assumptions with regard to the contemporaneous causal structure between the variables. On the contrary, we derive with our data-driven approach a causal ordering, which can be evaluated and discussed.

Another important strand of literature in this context focuses on VAR model parameter reduction. With an increasing number of variables to be modelled, the number of coefficients to be estimated increases exponentially. This problem is often referred to as the *curse of dimensionality*. Demirer et al. (2018) use lasso-type dimension reduction methods combined with generalised variance decompositions in order to estimate the connectedness between 150 bank stocks. An even sparser approach is proposed by Barigozzi and Brownlees (2019), who use lasso-type reduction methods not only to shrink the VAR lag matrices but also to shrink the variance-covariance matrix. Our causal search algorithm delivers an over-identified model, reducing the number of coefficients to be estimated, and hence eases the issue of dimensionality.

This paper contributes in several ways to the existing literature. First, we propose an alternative identification strategy which detects causal linkages. As Demiralp and Hoover (2003) show, the empirical procedure is very effective in detecting the true causal connections among

different variables. Second, we analyse the performance of our algorithm with respect to the Diebold and Yilmaz (2014) measure of connectedness and show in a Monte Carlo experiment that our algorithm outperforms other approaches.² Third, we apply our algorithm to the G10 currencies and pay special attention to the choice of the numéraire currency. This choice is of particular importance because it can have strong effects on the estimates, as we will discuss later.

Our results suggest that the US dollar as well as the Norwegian krone are the most independent currencies in our sample. By contrast, the Swiss franc and New Zealand dollar have a negligible impact on other currencies. Moreover, a cluster analysis suggests that the currencies can be divided into three groups, which can be identified as: commodity currencies, European currencies, and safe haven/carry trade financing currencies. We show that following the Brexit referendum, the within cluster dispersion is very low.

2 | METHODOLOGY

For \mathbf{y}_t being a $K \times 1$ vector of endogenous variables, we consider a SVAR(1) as follows:

$$\mathbf{B}_0 \mathbf{y}_t = \mathbf{B} \mathbf{y}_{t-1} + \mathbf{w}_t, \quad (1)$$

where \mathbf{B} refers to the $K \times K$ coefficient matrix of the lagged vector of endogenous variables. \mathbf{B}_0 defines the $K \times K$ contemporaneous coefficient matrix. Uncorrelated structural shocks are denoted by $\mathbf{w}_t \sim NID(\mathbf{0}, \boldsymbol{\Sigma}_w)$. Note that the off-diagonal entries of $\boldsymbol{\Sigma}_w$ are 0. We follow the notation of Kilian and Lütkepohl (2017). For brevity, we work here with just one lag and no deterministic terms. For the estimation, however, a constant is included, and the lag order is chosen according to AIC.

The reduced form of this model can be written as follows:

$$\mathbf{y}_t = \mathbf{A} \mathbf{y}_{t-1} + \mathbf{u}_t, \quad (2)$$

with $\mathbf{A} = \mathbf{B}_0^{-1} \mathbf{B}$ and $\mathbf{u}_t = \mathbf{B}_0^{-1} \mathbf{w}_t$.

Traditionally, the contemporaneous matrix \mathbf{B}_0 is uncovered with the help of restrictions motivated by economic theory. For a VAR model of exchange rate returns, economic theory does not provide a unique causal structure that can be imposed on the contemporaneous matrix. However, we achieve (over-)identification using a graph theoretical causal search algorithm which finds contemporaneous causality in the reduced form residuals \mathbf{u}_t . The correct contemporaneous effects are an important factor in the computation of the FEVD and consequently

in the connectedness measure of Diebold and Yilmaz (2014).

2.1 | The PC causal search algorithm and its application to the identification of SVAR models

The PC algorithm belongs to the literature on graph-theoretic analysis of causal structures, see Pearl (2000) and Spirtes et al. (2001).³ A causal structure is represented by a graph with arrows from causes to caused variables. The algorithm uses the residual variance-covariance matrix of the reduced form model as an input to detect the causal structure of a system, a directed acyclical graph (DAG). The PC algorithm cannot necessarily determine the DAG uniquely, but only down to a Markov equivalence class of the DAG. All members of an equivalence class encode the same conditional independence information. By using the conditional independence information the algorithm can only determine the equivalence class, but not distinguish between members of a class. In this way, the algorithm finds some undirected edges. We will determine these undirected edges with the help of a bootstrap procedure, which we will explain in the following.

To find the DAG, the algorithm performs an elimination stage and an orientation stage. The elimination stage starts with a graph where all the variables are linked to each other with an undirected link. Then, links are removed based on unconditional and conditional correlation tests, with a tuning parameter α for Fisher's Z-statistic being used as a significance level. First, connections are removed between two variables, which are unconditionally uncorrelated. Then, connections are eliminated for variables which are uncorrelated conditional on other variables. Here, the correlation of a pair of variables is conditioned on every other variable individually, then on all possible pairs of variables, thereafter on all subsets of three variables and so on up to all possible subsets of conditioning. When there is no more link to be removed, the elimination stage is finished and the skeleton of the graph is identified.

In the orientation stage, triples of linked variables $A-B-C$ are analysed. Unshielded colliders (v-structures) $A \rightarrow B \leftarrow C$ can be determined when A and C are independent when conditioned on possible sets of variables, but dependent when conditioned also on B. The algorithm searches for unshielded colliders and directs the edges accordingly. Finally, some more links might be oriented on the basis of logic. Some directions of links would lead to new unshielded colliders or to cyclicity, hence they need to be directed the other way around.

Cyclicity, like $A \rightarrow B \rightarrow C \rightarrow A$, is not permitted, hence bi-directional links are likewise not possible. Demiralp et al. (2008) show that a bootstrap procedure is successful in directing the undirected edges. The residuals of the reduced form VAR are drawn randomly with replacement, and new dataset are generated, to which the PC algorithm is applied. The undirected edges are finally directed in the direction, which is prominent more often in the bootstrap runs. Sampling errors or latent variables can lead to conflicting information about edge directions. In these cases, the algorithm returns a bi-directed edge. We decide on the bi-directed edges via our bootstrap procedure. Thus, the bootstrap procedure decides on the undirected edges (Markov equivalence class) and on the edges with conflicting information.

If the final graph is a DAG, then it can be mapped in the contemporaneous matrix B_0 , and due to the acyclicity property of the DAG, the contemporaneous matrix can be written as an overidentified lower triangular matrix for some ordering of the variables. Hence the SVAR model is identified. If the final graph contains cyclicity, which might arise due to some conflicting information about certain v-structures, the order condition is fulfilled, but it will not be possible to write the SVAR model as an overidentified recursive form.⁴

It is not clear from the onset which alpha value should be chosen in the PC algorithm. With increasing alpha values the algorithm becomes more liberal and so chooses fewer zero restrictions. Following our Monte Carlo simulation and Demiralp et al. (2014), we choose an alpha value of 10% in our application.

2.2 | A connectedness measure using forecast error variance decompositions

We use the connectedness measure of Diebold and Yilmaz (2014). The approach is based on the computation of FEVD.⁵

The stationary SVAR model in Equation (1) can be written in an MA representation as

$$y_t = \sum_{i=0}^{\infty} \Phi_i u_{t-i} = \sum_{i=0}^{\infty} \Theta_i w_{t-i}, \tag{3}$$

where Φ_i are reduced-form impulse responses and Θ_i are the structural impulse responses with $\Theta_i \equiv \Phi_i B_0^{-1}$. The matrixes Φ_i can be retrieved recursively by computing $\Phi_0 = I_K$ and $\Phi_i = (B_0^{-1} B)^i$.

We compute a FEVD

$$d_{jk}^h = 100 \sum_{i=0}^{h-1} \left(e_j' \Theta_i e_k \right)^2 / \sum_{i=0}^{h-1} \sum_{k=1}^K \theta_{jk,i}^2, \tag{4}$$

where $\theta_{jk,i}$ are the jk th element of Θ_i and e_k is the k th column of I_K . The measure d_{jk}^h is the proportion of the h -step forecast error variance of variable j , accounted for by innovations from variable k . We multiply the fractions by 100 to obtain percentages. Following Diebold and Yilmaz (2014), the pairwise directional connectedness from k to j is defined as

$$C_{j \leftarrow k}^h = d_{jk}^h. \tag{5}$$

In general $C_{j \leftarrow k}^h \neq C_{k \leftarrow j}^h$, so there are $K^2 - K$ separate pairwise directional connectedness measures.

The measure of total connectedness can be defined as

$$C^h = \frac{1}{K} \sum_{j,k=1, j \neq k}^K d_{jk}^h. \tag{6}$$

In the following sections, we will compare this measure with other measures of connectedness. One of these alternative measures is the generalised forecast error variance decomposition (GFEVD). For the computation of a GFEVD we follow Lanne and Nyberg (2016)⁶

$$d_{jk,g}^h = 100 \frac{\sum_{i=0}^{h-1} \left(e_j' \Phi_i \Sigma_u e_k \sigma_{kk}^{-1/2} \right)^2}{\sum_{i=0}^{h-1} \sum_{k=1}^K \left(e_j' \Phi_i \Sigma_u e_k \sigma_{kk}^{-1/2} \right)^2}, \tag{7}$$

where σ_{kk} are the diagonal entries of Σ_u .

2.3 | The algorithm

We make use of the R software package ‘pcalg’ by Kalisch et al. (2012).⁷ Our proposed algorithm (see Algorithm 1) starts with the estimation of a reduced form VAR model where the lag order is determined by the Akaike information criterion (AIC). Then, we apply the PC algorithm (PC) to the reduced form residuals and test if the resulting graph is a DAG. If this is the case, we can proceed and determine the contemporaneous matrix (B_0) in accordance with the obtained DAG. Otherwise, we bootstrap the reduced form VAR 10,000 times, apply the PC algorithm in each run, and collect the 10,000 suggested graphs. Note that it is important to draw

ALGORITHM 1

```

1: procedure IDENTIFICATION
2:    $[\mathbf{u}_t, \mathbf{A}] \leftarrow \text{VAR}(\text{data}, p = \text{AIC})$ 
3:    $\text{graph} \leftarrow \text{PC algorithm}(\alpha, \mathbf{u}_t)$ 
4:   if  $\text{graph}$  is directed-acyclical-graph then
5:      $\text{DAG} \leftarrow \text{graph}$ 
6:   else
7:      $\text{DAG} \leftarrow \text{Bootstrap}(\mathbf{u}_t, \mathbf{A}, \alpha, \text{graph})$ 
8:   end if
9:    $\mathbf{B}_0 \leftarrow \text{DAG}$ 
10:   $\text{connectedness} \leftarrow \text{FEVD}[\text{SVAR}(\mathbf{B}_0, \mathbf{A})]$ 
11: end procedure

12: function BOOTSTRAP( $\mathbf{u}_t, \mathbf{A}, \alpha, p, \text{graph}$ )
13:  for  $\text{runs} \in \{1, 2, \dots, 10,000\}$  do
14:     $\text{artificial data} \leftarrow \text{Create artificial data}(\mathbf{u}_t, \mathbf{A})$ 
15:     $[\mathbf{u}_t^{\text{BS}}] \leftarrow \text{VAR}(\text{artificial data}, p)$ 
16:     $\text{bootstrap-graph}(\text{runs}) \leftarrow \text{PC algorithm}(\alpha, \mathbf{u}_t^{\text{BS}})$ 
17:  end for
18:   $\text{return} \leftarrow \text{direct undirected edges in graph according to bootstrap} - \text{graph}$ 
19: end function

```

vectors from the residuals in such a way that the correlation between the residuals is preserved. Afterwards, we modify the original graph in such a way that the undirected edges become directed according to the direction preferred by the bootstrap. Note that we consider the bootstrap only in order to decide on edges which were originally undirected or bi-directed. Having obtained a DAG, we may proceed with the specification of \mathbf{B}_0 (line 10 of Algorithm 1). Finally, the connectedness measure—or spillover matrix—can be derived from the estimated structural VAR model where the shocks are orthogonalised by \mathbf{B}_0 .

3 | THE EFFECTIVENESS OF THE PC ALGORITHM IN THE CONNECTEDNESS APPROACH: A MONTE CARLO STUDY

In this section, we evaluate the impact of different identification strategies on the measures of connectedness by performing a Monte Carlo experiment. We generate artificial data with the help of a known data generating process (DGP). Afterwards, we estimate the connectedness matrices for different identification strategies and benchmark them with the theoretical result of the known DGP. We compare the empirical identification with the generalised approach (see Lanne & Nyberg, 2016) and the

average-of-all-Cholesky-orderings (see Klößner & Wagner, 2014) approach. The empirical identification is performed with two different algorithms, the PC algorithm and the greedy equivalence search (GES) algorithm of Chickering (2002).⁸ Because the appropriate significance level α for the individual conditional independence tests of the PC algorithm is not clear, we use two conventional options: 5% and 10%.

The artificial data are generated recursively according to the SVAR(1):

$$\mathbf{y}_t = \mathbf{B}_0^{-1} \mathbf{B} \mathbf{y}_{t-1} + \mathbf{B}_0^{-1} \mathbf{w}_t, \quad (8)$$

with structural shocks $\mathbf{w}_t \sim NID(\mathbf{0}, \boldsymbol{\Sigma}_w)$ and $\mathbf{y}_0 = \mathbf{0}$. To eliminate dependence on the initial condition we discard the first 80% of the generated data in all cases. The lag matrix, \mathbf{B} , is $K \times K$ with random uniform coefficients between -0.05 and 0.05 . The residuals, \mathbf{w}_t , are drawn randomly from independent normal distributions with mean 0 and variance 1. For the contemporaneous matrix, \mathbf{B}_0 , we generate random directed acyclic graphs (DAGs) with a fixed expected number of neighbours. We use random Erdős-Rényi graphs for the DAGs, multiply the matrix entries by -1 and add an identity matrix. In this way we generate a sparse contemporaneous matrix with some negative off-diagonal entries between 0 and -1 .⁹ We perform this Monte Carlo study for $N = 100$ datasets in each MC experiment. The categories are: two different

TABLE 1 Monte Carlo simulation: Comparing connectedness measures for different identification strategies relative to the correct connectedness measures using 100 random DAGs dimension 8

$d = 1$	$T = 250$				$T = 2500$			
	\mathcal{C}	\mathcal{T}	\mathcal{S}	\mathcal{K}	\mathcal{C}	\mathcal{T}	\mathcal{S}	\mathcal{K}
avgChol	1.849	3.648	1.798	14.836	1.487	0.558	1.576	13.606
GFEVD	2.823	12.233	1.523	11.965	2.233	7.126	1.260	10.400
PCalg 5%	1.644	2.594	0.455	4.245	1.031	0.534	0.228	2.166
PCalg 10%	1.664	2.824	0.452	4.334	1.045	0.548	0.245	2.300
GES	1.421	2.731	0.444	4.243	1.017	0.405	0.238	2.290

$d = 3$	$T = 250$				$T = 2500$			
	\mathcal{C}	\mathcal{T}	\mathcal{S}	\mathcal{K}	\mathcal{C}	\mathcal{T}	\mathcal{S}	\mathcal{K}
avgChol	5.386	4.337	1.331	5.342	5.195	2.219	1.310	5.350
GFEVD	8.276	28.478	1.207	4.644	8.284	23.241	0.862	3.637
PCalg 5%	4.722	3.249	0.520	3.261	3.584	2.738	0.354	2.217
PCalg 10%	4.668	3.256	0.502	3.086	3.630	2.554	0.366	2.247
GES	5.221	3.983	0.411	2.449	4.143	1.811	0.363	2.208

$d = 5$	$T = 250$				$T = 2500$			
	\mathcal{C}	\mathcal{T}	\mathcal{S}	\mathcal{K}	\mathcal{C}	\mathcal{T}	\mathcal{S}	\mathcal{K}
avgChol	9.695	6.007	1.559	3.596	9.643	5.091	1.564	3.586
GFEVD	11.552	30.634	1.113	3.098	13.508	28.869	0.858	2.642
PCalg 5%	9.592	8.656	0.754	3.730	9.157	6.858	0.576	2.561
PCalg 10%	9.728	7.621	0.671	3.217	9.227	6.807	0.579	2.634
GES	10.914	6.249	0.442	1.865	10.628	5.452	0.415	1.762

Note: 100 random Erdős-Rényi graphs with eight nodes. d (1, 3, 5) corresponds to the expected number of neighbours per node, more precisely the expected sum of the in- and out-degree. Sample size 250/2500 observations. \mathcal{C} is the MAE of the off-diagonal entries of the connectedness matrix. \mathcal{T} is the MAE of the total connectedness. \mathcal{S} and \mathcal{K} are the MAEs of the skewness and kurtosis of the distribution of the off-diagonal entries of the connectedness matrix.

system dimensions ($K = 8/16$), three different levels of sparsity ($d = 1/3/5$) and two different sample lengths ($T = 250/2500$).¹⁰

The results are evaluated as follows. For each identification method, we compute four measures in terms of recovering the true connectedness matrix. For all four measures, we report the mean absolute error (MAE) of the estimated measure relative to the measure for the true connectedness matrix. The first measure, \mathcal{C} , is the MAE of the off-diagonal elements of the connectedness matrix:

$$\mathcal{C} = \frac{1}{N} \sum_{i=1}^N \frac{1}{K^2 - K} \sum_{j,k=1, j \neq k}^K |C_{j \leftarrow k, i}^h - C_{j \leftarrow k, i}^{h^*}|, \quad (9)$$

whereby the variables with a star are the true connectedness values. \mathcal{C} is an important measure, as it places a strong weight on the direction of the connectedness. The second measure, \mathcal{T} , is the MAE of the total connectedness:

$$\mathcal{T} = \frac{1}{N} \sum_{i=1}^N |C_i^h - C_i^{h^*}|. \quad (10)$$

Here, it is not so much the direction of the links that is evaluated, but rather whether the over-identifying zeros of the PC algorithm are appropriate. The third measure, \mathcal{S} , is the MAE of the skewness of the distribution of the off-diagonal entries of the connectedness matrix:

$$\mathcal{S} = \frac{1}{N} \sum_{i=1}^N \left| \text{Skew} \left\{ C_{j \leftarrow k, i}^h \right\}_{j,k=1 \dots K, j \neq k} - \text{Skew} \left\{ C_{j \leftarrow k, i}^{h^*} \right\}_{j,k=1 \dots K, j \neq k} \right|. \quad (11)$$

While the fourth measure, \mathcal{K} , is the MAE of the kurtosis of the distribution of the off-diagonal entries of the connectedness matrix:

TABLE 2 Monte Carlo simulation: Comparing connectedness measures for different identification strategies relative to the correct connectedness measures using 100 random DAGs dimension 16

<i>d</i> = 1	<i>T</i> = 250				<i>T</i> = 2500			
	<i>C</i>	<i>T</i>	<i>S</i>	<i>K</i>	<i>C</i>	<i>T</i>	<i>S</i>	<i>K</i>
avgChol	1.291	8.006	2.603	33.602	0.864	0.997	2.219	30.262
GFEVD	2.066	20.489	2.572	30.278	1.334	8.775	1.763	22.862
PCalg 5%	1.020	5.867	0.613	10.554	0.581	0.680	0.322	5.870
PCalg 10%	1.023	6.179	0.586	9.586	0.586	0.670	0.321	5.810
GES	0.995	6.202	0.602	9.510	0.497	0.641	0.251	4.384

<i>d</i> = 3	<i>T</i> = 250				<i>T</i> = 2500			
	<i>C</i>	<i>T</i>	<i>S</i>	<i>K</i>	<i>C</i>	<i>T</i>	<i>S</i>	<i>K</i>
avgChol	3.021	7.009	1.791	12.483	2.762	2.477	1.752	12.510
GFEVD	4.871	36.769	2.112	13.264	4.468	26.757	1.460	9.546
PCalg 5%	2.309	2.932	0.505	5.531	1.729	1.896	0.494	5.094
PCalg 10%	2.276	3.332	0.515	5.633	1.671	1.796	0.453	4.607
GES	2.455	5.582	0.430	4.572	1.881	1.850	0.417	4.272

<i>d</i> = 5	<i>T</i> = 250				<i>T</i> = 2500			
	<i>C</i>	<i>T</i>	<i>S</i>	<i>K</i>	<i>C</i>	<i>T</i>	<i>S</i>	<i>K</i>
avgChol	5.018	8.309	2.163	10.084	4.900	5.410	2.140	10.111
GFEVD	6.420	37.451	1.720	9.183	6.684	33.617	1.593	8.289
PCalg 5%	4.662	4.533	0.772	6.909	3.952	4.408	0.640	5.155
PCalg 10%	4.606	4.381	0.822	7.276	4.000	4.532	0.685	5.542
GES	5.262	7.762	0.648	4.871	4.854	5.255	0.511	3.802

Note: 100 random Erdős-Rényi graphs with eight nodes. *d* (1, 3, 5) corresponds to the expected number of neighbours per node, more precisely the expected sum of the in- and out-degree. Sample size 250/2500 observations. *C* is the MAE of the off-diagonal entries of the connectedness matrix. *T* is the MAE of the total connectedness. *S* and *K* are the MAEs of the skewness and kurtosis of the distribution of the off-diagonal entries of the connectedness matrix.

$$\mathcal{K} = \frac{1}{N} \sum_{i=1}^N \left| Kurt \left\{ C_{j \leftarrow k, i}^h \right\}_{j, k=1 \dots K, j \neq k} - Kurt \left\{ C_{j \leftarrow k, i}^{h^*} \right\}_{j, k=1 \dots K, j \neq k} \right|. \tag{12}$$

These distributional measures aim to evaluate whether the extreme values in the connectedness matrices of the identification strategies are comparable to the theoretical connectedness. The results of the Monte Carlo experiments are displayed in Tables 1 and 2.

The PC algorithm performs, in general, better than the average-of-all-Cholesky-orderings approach and the generalised approach. While the other two identification strategies usually split the connectedness between *j* and *k* in such a way that $C_{j \leftarrow k}^h \approx C_{k \leftarrow j}^h$, the causal search algorithm manages to find the true causal connectedness. Even when the PC algorithm might find incorrect directions for some links, overall the true causal structure is uncovered to a much higher degree, which can be seen in

TABLE 3 Monte Carlo simulation: comparing connectedness measures for different identification strategies relative to the correct connectedness measures using 100 random DAGs. Specification similar to the application: dimension = 9, *d* = 3.8, *T* = 2048

	<i>C</i>	<i>T</i>	<i>S</i>	<i>K</i>
avgChol	6.042	3.342	1.332	4.750
GFEVD	9.371	28.402	0.803	3.249
PCalg 5%	4.863	4.033	0.467	2.725
PCalg 10%	4.875	3.661	0.484	2.831
GES	5.410	2.939	0.463	2.645

Note: 100 random Erdős-Rényi graphs with nine nodes. *d* = 3.8 corresponds to the expected number of neighbours per node, more precisely the expected sum of the in- and out-degree. Sample size 2048 observations. *C* is the MAE of the off-diagonal entries of the connectedness matrix. *T* is the MAE of the total connectedness. *S* and *K* are the MAEs of the skewness and kurtosis of the distribution of the off-diagonal entries of the connectedness matrix.

the lower MAE values of our measure *C*. The generalised approach overestimates the total connectedness strongly in all cases, which can be seen in the high MAE in the measure *T*. This result is in line with findings by De

Santis and Zimic (2018). The average-of-all-Cholesky-orderings approach usually overestimates the total connectedness slightly. The PC algorithm partly overestimates and partly underestimates the total connectedness, depending on the sparsity of the network, performing well for many DGPs, but poorly in a small number of cases. The GES algorithm tends to overestimate the total connectedness more often. The empirical approaches are superior in estimating connectedness for models with sparse contemporaneous matrices. In general there is no clear favourite between a PC algorithm with a 5% and 10% significance level. The GES algorithm performs equally well as the PC algorithm and is particularly strong for sparse contemporaneous matrices.

To reinforce the relevance of the Monte Carlo simulation to our empirical application, we perform a simulation experiment where we mimic the settings of the application, see Table 3. In the application we have a dimension of 9, a level of sparsity of $d = 3.8$ and 2048 observations.¹¹ In the Monte Carlo experiment, the causal search PC algorithm performs strongly, achieving low MAEs, especially in category \mathcal{C} , which measures the direction of connectedness.

4 | AN APPLICATION TO EXCHANGE RATE DATA

The bilateral exchange rate can be interpreted as the relative price between two currencies. Here, it is defined as the foreign currency price of buying one unit of home currency (quantity quotation). A positive shock to the bilateral exchange rate in quantity quotation can thus be interpreted as a positive shock to the demand of the home currency or a negative shock to the demand of the foreign currency. These shocks can trigger movements in other exchange rates as well. The reasons behind the international effects are manifold. One could think of currency (basket) pegs or international substitution effects, for instance.

The aim of this exercise is to uncover the network of spillover effects between exchange rate returns.¹² We apply the proposed algorithm to the G10 bilateral euro exchange rates and cluster the exchange rates in order to uncover potential currency blocs afterwards. All bilateral exchange rates are downloaded in daily frequency from the ECB statistical data warehouse (SDW) and correspond to the ECB reference rates, representing the 14:15 CET fixing.¹³ The sample covers the period between January 2010 and December 2017. We start in 2010 in order to exclude potential effects arising from the 2008 financial crisis. As ECB reference rates are expressed in quantity quotation and quoted against the euro, we

transform the rates in such a way that the pound sterling becomes the numéraire. All transformed series enter our model in log differences. The reasoning behind changing the numéraire currency is discussed in the following section.

4.1 | Choice of the numéraire currency

The bilateral (or multilateral) nature of exchange rates poses a problem for researchers and practitioners. When regressing exchange rate returns on exchange rate returns, the correct choice of the numéraire currency (or basket) is crucial, because the numéraire can have substantial effects on the estimates. If both currencies were pegged to the numéraire, the regression coefficient would be zero, implying that despite the common peg no relationship would exist. This problem has been extensively discussed by the literature on currency baskets and blocks, for example by Frankel and Wei (2008), Frankel and Xie (2010), or Ohno (1999).

The US dollar, the most heavily traded currency, appears to be a good choice as the numéraire currency. But the afore mentioned statistical problems arise if currencies, pegged to the dollar enter the model. Despite the peg, these exchange rates—expressed in US dollar—would appear to be unconnected. More importantly, however, the connectedness of the US dollar could not be estimated if it served as the numéraire currency. This would eliminate important information, as several studies have pointed out that the US dollar shares certain properties with other currencies, for example its status as a safe haven currency (see Hossfeld & MacDonald, 2015). Other studies such as Frankel and Wei (2008) and Ohno (1999) relied on the Swiss franc as the numéraire currency. The Swiss franc seemed to be an appealing choice, because its trading volume is high and the currency was independent at that time. The Swiss franc lost this property when the Swiss National Bank introduced a minimum rate vis-à-vis the euro on 6 September 2011.¹⁴ When the numéraire currency is pegged to another currency in the sample, the exchange rate has no variance, which can be explained by other currencies. Apart from the numerical problems that arise from this, the series would not have any variance and should thus not be employed within the Diebold and Yilmaz (2014) approach.¹⁵ Deutsche Bundesbank (2019) provides numerical examples with respect to this issue.

The literature on basket weights proposes using a basket of different currencies as the numéraire. Frankel and Xie (2010) claim that monetary authorities are more likely to use a weighted average of currencies as a reference for possible interventions when the exchange rate

regime is a managed float or target zone and propose the IMF Special Drawing Right (SDR) basket as the numéraire. On the other hand, Aloosh and Bekaert (2021) propose the use of unweighted baskets in a somewhat different setting. Nevertheless, a closed basket as the numéraire cannot be employed in the Diebold and Yilmaz (2014) approach, because the first currency basket is nothing but a linear combination of the other $N - 1$ currency baskets. A VAR model cannot be estimated for such a dataset. Hence, even when baskets are used, one currency needs to be excluded.

Therefore, we have to find a suitable numéraire currency which is freely floating, not an anchor currency and preferably not, or at least only to a small extent, part of currency baskets which central banks peg their currency to. A currency which sufficiently fulfils these criteria is the pound sterling (see IMF, 2016). The pound sterling is not as heavily traded as the US dollar. But for this approach, prices and not quantities matter. Arbitrage is supposed to eliminate any discrepancy in prices within one period (here: 14:15 CET to 14:15 CET on the following trading day). Overall, the pound sterling appears to be independent and relatively uncorrelated with other currencies (see Aloosh & Bekaert, 2021; Hossfeld & MacDonald, 2015). The yen is also a relatively independent currency. But our study will show that the yen shares certain properties with other currencies such as the US dollar and the Swiss franc.

4.2 | Uncovering the causal structure

We follow the proposed algorithm and estimate the VAR model Equation (2) for the rate of change in nine G10 exchange rates, taking the pound sterling as the numéraire currency. The number of observations is 2048, and the lag length of $p = 1$ is chosen according to the Akaike Information Criterion (AIC).¹⁶ Before applying the PC algorithm, we investigate the correlation structure of the residuals \mathbf{u}_t (see Table 4).

The table reveals relatively strong correlations between the residuals of commodity currencies such as the Australian dollar, Canadian dollar, New Zealand dollar and Norwegian krone. Particularly striking is the strong correlation between the Australian dollar and the New Zealand dollar ($\rho_{AUD,NZD}^u = 0.76$). The correlation between the Swedish krona and euro residuals ($\rho_{SEK,EUR}^u = 0.75$) is similarly strong. The Swedish krona residuals, however, are also correlated with the Norwegian krone residuals, thus linking the euro to the commodity currencies. The euro itself is also correlated with the Swiss franc. The correlation between the Japanese yen and the US dollar ($\rho_{JPY,USD}^u = 0.66$) is somewhat uncoupled. These

correlations provide us with a first insight into international exchange rate connectedness. They suggest the existence of a commodity currency bloc which is connected with the euro, and a strong relationship between the Japanese yen and US dollar. These correlations, however, do not reflect any type of causality. We only learn that certain relationships may exist.

In order to estimate the causal structure we apply the PC algorithm with an α size of 0.1 to the residuals. The algorithm yields the adjacency matrix

$$F = \begin{array}{c} \begin{array}{cccccccc} \text{AUD} & \text{CAD} & \text{CHF} & \text{EUR} & \text{NOK} & \text{NZD} & \text{SEK} & \text{USD} & \text{JPY} \end{array} \\ \left[\begin{array}{cccccccc} 1 & 1 & 0 & 0 & 1 & 0 & \underline{1} & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & \underline{1} & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & \underline{1} & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & \underline{1} & 0 & 0 \\ 1 & 1 & \underline{1} & 0 & 0 & 1 & 0 & 0 & 1 \\ \underline{1} & 0 & 0 & 0 & \underline{1} & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \underline{1} & 0 & 0 & 0 & 1 & \underline{1} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \underline{1} & 1 \end{array} \right] \begin{array}{l} \text{AUD} \\ \text{CAD} \\ \text{CHF} \\ \text{EUR} \\ \text{NOK} \\ \text{NZD} \\ \text{SEK} \\ \text{USD} \\ \text{JPY} \end{array} \end{array}$$

where 1 indicates that the currency of the related column is causing the currency of the corresponding row. By contrast, 0 indicates that the currency of the related column is not causing the currency of the corresponding row. Hence, if $F[i,j] \neq F[j,i]$ (with $i \neq j$) the algorithm was able to direct the link between the two currencies: if $F[i,j] = F[j,i] = 0$, then there is no contemporaneous relationship between the two currencies and an over-identification of the variance covariance matrix is possible. A problem occurs, however, if $F[i,j] = F[j,i] = 1$ (see underlined numbers). In that particular case, the PC algorithm was unable to direct the link between the two currency pairs (Markov equivalence class). Note, undirected edges do not indicate causation in both directions. Undirected edges have been found for the pairs *AUD—SEK*, *CHF—NZD*, *EUR—USD*, *NOK—SEK* and *USD—JPY*. As explained earlier, we apply the bootstrap approach of Demiralp et al. (2008) in order to direct the undirected edges.¹⁷

More specifically, we bootstrap the VAR, let the PC algorithm determine the causal structure between the residuals, and save the output of the algorithm for each of the 10,000 runs. The results are reported in Table 5. Of particular interest are the entries $F[7,1]$ (*AUD—SEK*), $F[6,3]$ (*CHF—NZD*), $F[8,4]$ (*EUR—USD*), $F[7,5]$ (*NOK—SEK*), and $F[9,8]$ (*JPY—USD*).

The table shows for the currency pair *AUD—SEK* that $AUD \rightarrow SEK$ is preferred over $AUD \leftarrow SEK$ by 35.68% vs. 1.17% of the draws. The results are slightly less clear for *CHF—NZD*, as only 9.23% of the draws prefer $CHF \rightarrow NZD$. According to the bootstrap, the edge *EUR—USD* should be directed such that $EUR \leftarrow USD$. Lastly, we direct the edges *NOK—SEK* and *JPY—USD* such that $NOK \rightarrow SEK$ and $JPY \leftarrow USD$.

TABLE 4 Correlation between residuals (u_t)

	AUD	CAD	CHF	EUR	NOK	NZD	SEK	USD	JPY
AUD	1.00								
CAD	0.66	1.00							
CHF	0.31	0.33	1.00						
EUR	0.44	0.42	0.61	1.00					
NOK	0.53	0.50	0.40	0.64	1.00				
NZD	0.76	0.57	0.33	0.43	0.48	1.00			
SEK	0.51	0.45	0.44	0.75	0.73	0.46	1.00		
USD	0.34	0.53	0.39	0.43	0.24	0.31	0.30	1.00	
JPY	0.31	0.36	0.48	0.45	0.22	0.33	0.28	0.64	1.00

Note: The table shows the cross-correlation between reduced form VAR residuals. Correlation coefficients >0.5 are marked in bold.

Given the bootstrap information, the matrix F can now be updated and written as

$$F_{BS} = \begin{bmatrix} \text{AUD} & \text{CAD} & \text{CHF} & \text{EUR} & \text{NOK} & \text{NZD} & \text{SEK} & \text{USD} & \text{JPY} \\ \begin{bmatrix} 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} & \begin{matrix} \text{AUD} \\ \text{CAD} \\ \text{CHF} \\ \text{EUR} \\ \text{NOK} \\ \text{NZD} \\ \text{SEK} \\ \text{USD} \\ \text{JPY} \end{matrix} \end{bmatrix}$$

This matrix no longer exhibits any cycles or undirected edges and can now be used to (over-)identify the variance-covariance matrix of the VAR model. Each 1-entry represents a coefficient which has to be estimated. Figure 1 presents the adjacency matrix F_{BS} as a graph.

For a better understanding of the dynamics, we reorder the adjacency matrix into the over-identified matrix B_0 in such a way that we recover its recursive form.¹⁸

$$B_0 = \begin{bmatrix} \text{USD} & \text{JPY} & \text{NOK} & \text{CAD} & \text{AUD} & \text{SEK} & \text{EUR} & \text{CHF} & \text{NZD} \\ \begin{bmatrix} 1 & & & & & & & & \\ b_{21} & 1 & & & & & & & \\ 0 & 0 & 1 & & & & & & \\ b_{41} & 0 & b_{43} & 1 & & & & & \\ 0 & 0 & b_{53} & b_{54} & 1 & & & & \\ 0 & 0 & b_{63} & 0 & b_{65} & 1 & & & \\ b_{71} & b_{72} & b_{73} & 0 & 0 & b_{76} & 1 & & \\ 0 & b_{82} & 0 & 0 & 0 & 0 & b_{87} & 1 & \\ 0 & b_{92} & 0 & b_{94} & b_{95} & 0 & 0 & b_{98} & 1 \end{bmatrix} & \begin{matrix} \text{USD} \\ \text{JPY} \\ \text{NOK} \\ \text{CAD} \\ \text{AUD} \\ \text{SEK} \\ \text{EUR} \\ \text{CHF} \\ \text{NZD} \end{matrix} \end{bmatrix}$$

This enables us to uncover and interpret the causal structure between the structural shocks. Note here that the ordering is entirely determined by the data-driven causal-search algorithm (PC). We observe that the USD is

ordered first, suggesting that it is the most independent currency in the sample. Shocks to the US dollar affect the Japanese yen, the Canadian dollar and the euro contemporaneously, while no foreign shock has contemporaneous effects on the US dollar. The New Zealand dollar is ordered last. Thus, shocks to this currency have no contemporaneous effects on any other currency in the sample. However, it is affected by shocks to the Japanese yen, the Canadian dollar, the Australian dollar and the Swiss franc. Interestingly, these are all currencies which are related to carry trades. The commodity currencies (AUD, CAD) are often referred to as carry trade target currencies, while the other two (CHF, JPY) are used by market participants for carry trade funding (see Ferreira Filipe & Suominen, 2013; Hossfeld & MacDonald, 2015). The Norwegian krone appears to be another important currency. It is ordered third, but the PC algorithm suggests that shocks to the US dollar and the Japanese yen, which is ordered second, have no contemporaneous effects on the krone. However, shocks to the krone affect other commodity currencies (CAD, AUD) and geographical neighbours (EUR, SEK) contemporaneously. Overall, we observe a causal structure which is not only closely related to the correlation of reduced form residuals (Table 4), but also economically plausible.

The coefficients in the contemporaneous matrix (B_0) can be obtained by re-estimating the (structural) VAR equation by equation, whereby the contemporaneous effects (according to B_0) are included in each equation. The estimates are presented in Appendix A.


The matrix B_0 is overidentified with 19 zero restrictions. When testing the 19 restrictions with a likelihood ratio test, the null hypothesis that all these 19 coefficients can be restricted to zero needs to be rejected with a test statistic of 268.66 and a p value of 0.000. When we return

TABLE 5 PC algorithm results in percent (bootstrap with 10,000 draws)

$F[i,j]$		PC	# of Undirected	# of Left	# of No edge	# of Right	# of Bi-directed
i	j						
2	1	-1	38.41	0.60	0.00	60.98	0.01
3	1	0	0.00	0.01	99.99	0.00	0.00
4	1	0	0.00	0.00	100.00	0.00	0.00
5	1	-1	43.65	12.72	2.28	41.28	0.07
6	1	1	46.13	22.10	0.00	31.59	0.18
7	1	2	62.26	35.68	0.84	1.17	0.05
8	1	0	11.00	0.07	87.40	1.53	0.00
9	1	0	0.00	0.00	99.95	0.00	0.05
3	2	0	0.00	0.00	100.00	0.00	0.00
4	2	0	0.00	0.00	100.00	0.00	0.00
5	2	-1	39.84	11.43	0.00	48.60	0.13
6	2	1	12.79	77.39	8.15	0.66	1.01
7	2	0	0.19	0.05	99.71	0.05	0.00
8	2	-1	46.67	1.15	0.00	51.29	0.89
9	2	0	2.15	0.22	97.28	0.35	0.00
4	3	-1	31.55	5.88	0.00	62.56	0.01
5	3	0	0.81	1.45	97.59	0.15	0.00
6	3	2	31.41	9.23	54.85	4.46	0.05
7	3	0	11.48	1.49	85.70	1.33	0.00
8	3	0	16.96	2.04	79.69	1.31	0.00
9	3	-1	32.84	4.85	0.00	62.31	0.00
5	4	-1	33.40	4.74	0.00	61.78	0.08
6	4	0	9.68	4.57	83.26	2.33	0.16
7	4	-1	55.26	10.49	0.00	34.24	0.01
8	4	2	57.21	8.06	18.79	15.81	0.13
9	4	-1	1.90	0.03	37.10	60.97	0.00
6	5	0	18.55	12.59	68.59	0.27	0.00
7	5	2	67.08	29.71	0.00	3.19	0.02
8	5	0	0.01	0.00	99.99	0.00	0.00
9	5	0	0.00	0.00	100.00	0.00	0.00
7	6	0	0.68	0.17	98.79	0.36	0.00
8	6	0	0.00	0.00	100.00	0.00	0.00
9	6	-1	6.54	4.32	15.80	60.98	12.36
8	7	0	0.03	0.01	99.95	0.01	0.00
9	7	0	0.89	0.03	98.77	0.31	0.00
9	8	2	75.20	17.25	0.00	7.55	0.00

Note: The table shows the decisions of the PC algorithm for each entry in the adjacency matrix $F[i,j]$. The PC column refers to the PC algorithm decision with respect to the VAR point estimate residuals. Here, '0' denotes 'no edge', '1' stands for 'right', '-1' for 'left' and '2' for an 'undirected' linkage. The columns for the percentages (abbreviated as #) of the decisions 'undirected', 'left', 'no edge', 'right' and 'bi-directed' refer to the bootstrap. Whenever the algorithm finds no direction using the point estimate residuals (i.e., where the PC column has the entry 2), the edge is directed according to the (maximum) percentage in the 'left' and 'right' columns.

FIGURE 1 Visualization of the adjacency matrix F_{BS} . This figure shows the contemporaneous causality structure which is used to orthogonalise the SVAR residuals (see matrix B_0). Note that this graph does not represent the connectedness matrix, which will be estimated in section 4.3.



Notes: This figure shows the contemporaneous causality structure which is used to orthogonalise the SVAR residuals (see matrix B_0). Note that this graph does not represent the connectedness matrix, which will be estimated in section (4.3).

more coefficients to the contemporaneous matrix, beginning with the highest entries in the variance covariance matrix of the SVAR model, we find that already with 7 over-identifying restrictions the null needs to be rejected with a test statistic of 16.17 and a p value of 0.024. We attribute this finding to the fact that we are working with bilateral exchange rates. The contemporaneous correlations of the currency pairs are high in all cases, see Table 4, which might be due to a UK effect in all exchange rates. To further scrutinise this point we perform a robustness check at the end of section 4.3, where we compare the connectedness matrix for the PC algorithm approach with the connectedness matrix of a version of the SVAR model without over-identifying restrictions on the contemporaneous matrix but the same ordering.

4.3 | Connectedness

This section shows the connectedness between exchange rates for different identification methods. First, we use a simple Cholesky decomposition as in the seminal paper by Diebold and Yilmaz (2009). As we are completely agnostic with regard to causality, the ordering of the variables is random. Table 6 presents the FEVD (i.e., the connectedness). The entries represent the shares of forecast error variance (in percent) of the variables in rows, which are explained by shocks to the variables in columns. Hence, rows add up to 100. For instance, the estimates suggest that 16.8% of US dollar forecast error variance is explained by shocks to the Canadian dollar, while only

0.3% of Canadian dollar forecast error variance is explained by shocks to the US dollar. This result is surprising, because the United States is usually considered as a large and less dependent economy. Accordingly, one would expect causality to point from USD to CAD rather than the other way round. Later, we will see that a different causal ordering yields completely different results from those obtained by the PC algorithm. But also the measures of total connectedness show a surprising picture. The row 'IN' represents the contribution of international shocks to the forecast error variance of the variables in rows (i.e., the row sum minus the idiosyncratic contribution). This measure is referred to as in-connectedness. Particularly the Australian dollar, which is ordered first, is almost entirely driven by its own shocks. Then again, the out-connectedness of the Australian dollar (OUT; i.e., the column sum of contributions subtracted by the idiosyncratic component) is extraordinarily high compared with other currencies. Overall, we observe that the shares on the lower diagonal are substantially higher than those on the upper diagonal, reflecting the lower diagonal structure of the Cholesky factorisation which has been applied. Hence, a Cholesky decomposition has to be applied with caution.

One way to circumvent this problem is to apply a generalised FEVD. This decomposition is derived from generalised impulse response functions which were originally proposed by Pesaran and Shin (1998). The intuition behind this approach is that every variable is treated as it would be ordered first in a Cholesky decomposition. Hence, any variable can have contemporaneous effects on any variable in the system. The

TABLE 6 Connectedness: Cholesky (random ordering)

	AUD	CAD	CHF	EUR	NOK	NZD	SEK	USD	JPY	IN
AUD	99.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.2
CAD	43.8	55.5	0.0	0.0	0.1	0.1	0.0	0.3	0.1	44.5
CHF	9.9	2.7	86.9	0.0	0.1	0.0	0.3	0.0	0.0	13.1
EUR	18.8	3.2	22.0	55.4	0.0	0.0	0.1	0.2	0.3	44.6
NOK	28.5	3.6	4.5	13.3	49.5	0.0	0.2	0.0	0.4	50.5
NZD	58.1	0.6	0.8	0.4	0.0	39.9	0.0	0.0	0.2	60.1
SEK	25.9	2.3	7.2	24.9	7.0	0.0	32.2	0.1	0.4	67.8
USD	12.0	16.8	5.4	1.7	3.2	0.0	0.1	60.9	0.0	39.1
JPY	9.4	4.1	14.0	1.8	2.9	0.4	0.2	16.9	50.3	49.7
OUT	206.5	33.4	53.8	42.2	13.4	0.5	0.9	17.4	1.5	$C^{10} = 41.1$

Note: The table shows the (10 periods ahead) forecast error variance decomposition of the SVAR model which is identified by a Cholesky decomposition with random ordering. The column 'IN' corresponds to the row sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by [international] shocks). The column 'OUT' corresponds to the column sum of the non-diagonal variance shares (i.e., the total share of variance, which is explained by the corresponding column variable). C^{10} refers to the measure of total connectedness (see section 2.2).

problem, however, is that this approach does not orthogonalise the shocks, implying that forecast error variance shares do not sum to unity for any given variable.¹⁹ Instead of normalizing the variance shares in such a way that they sum to unity (see Greenwood-Nimmo et al., 2016, for instance), we proceed as outlined in section 2.2 and apply an alternative form of the generalised FEVD where variance shares sum to unity by construction (see Lanne & Nyberg, 2016). This decision is based on the results of Chan-Lau (2017), who finds that the generalised variance decomposition by Lanne and Nyberg (2016) performs better than the one by Pesaran and Shin (1998).

The results of the GFEVD measure are presented in Table 7. Now, we observe that the differences between in- and out-connectedness have decreased for all currency pairs. For many pairs, the degrees have even become roughly equivalent: 3.9% of the Japanese yen's forecast error variance is explained by the Norwegian krone and vice versa. Overall, the linkages are qualitatively similar to those obtained from the Cholesky decomposition with the difference that the information with regard to the direction of causality has been lost.

The same applies to the results based on the fastSOM algorithm by Klößner and Wagner (2014), which represent the average variance shares over all possible Cholesky permutations (Table 8). Here, the differences between in- and out-connectedness are even smaller than in the previous case. We learn, for instance, that the Australian dollar and the New Zealand dollar are connected, but the differences between the variance shares ($AUD \rightarrow NZD$ [18.6%] and $AUD \leftarrow NZD$ [17.6%]) are not strong enough to make a statement about the direction.

Additionally, when studying all possible permutations of orderings, the maximum and minimum levels of connectedness per country pair can be computed, see Tables B1 and B2. It is obvious that the connectedness between currency pairs depends crucially on an arbitrary ordering of the variables in the model.

The ambiguity diminishes drastically when our proposed algorithm is applied (Table 9). Clear causal patterns appear, which help in understanding the network topology of G10 exchange rates and the direction of edges, in particular. The estimates suggest that the US dollar and the Norwegian krone are both important drivers of international exchange rate fluctuations. Their in-connectedness is very low (0.7 and 1.4, respectively), suggesting that they are barely affected by shocks to foreign currencies. However, their out-connectedness is relatively high (91.1 and 152.4, respectively). Shocks to the US dollar explain 21.2% of Canadian dollar and 41.2% of Japanese yen forecast error variance. Note that these are clearly not bi-directed linkages. Our estimates suggest that the causality is directed from the US dollar to both currencies. The Norwegian krone appears to explain high shares of the Australian dollar, Canadian dollar, Swiss franc, euro, New Zealand dollar and Swedish krona forecast error variance. Hence, there seem to be strong ties between commodity currencies and European currencies. The strong effects of the Norwegian krone are somewhat surprising. It is possible that the Norwegian krone reflects influences of other currencies from oil exporting countries.

Also other studies find a high influence of the Norwegian krone, for example Le et al. (2018) use measures like right-hand eigenvector centrality, Harmonic closeness centrality, out-degree and out-strength and detect a high

TABLE 7 Connectedness: GFEVD (Lanne & Nyberg, 2016)

	AUD	CAD	CHF	EUR	NOK	NZD	SEK	USD	JPY	IN
AUD	33.7	11.4	3.8	3.9	9.3	21.6	8.0	2.8	5.5	66.3
CAD	17.2	28.9	4.7	4.1	9.6	13.3	7.1	7.5	7.7	71.1
CHF	4.6	3.8	44.1	8.9	7.6	5.0	8.1	4.3	13.5	55.9
EUR	7.4	5.2	13.8	20.6	13.6	7.1	17.4	4.3	10.6	79.4
NOK	11.7	7.3	6.5	9.2	32.9	9.3	17.8	1.5	3.9	67.1
NZD	20.8	8.7	4.4	4.0	7.9	38.7	6.8	2.3	6.3	61.3
SEK	10.3	6.0	7.4	12.1	17.4	8.2	30.9	2.2	5.4	69.1
USD	6.0	10.4	7.7	5.2	4.4	4.7	4.0	31.3	26.2	68.7
JPY	4.1	4.1	9.8	4.8	3.9	4.5	3.4	11.2	54.0	46.0
OUT	82.2	56.9	58.0	52.2	73.7	73.7	72.6	36.2	79.3	$C^{10} = 65.0$

Note: The table shows the (10 periods ahead) forecast error variance decomposition of the SVAR model which is identified by using generalised impulse responses. The column 'IN' corresponds to the row sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by [international] shocks). The column 'OUT' corresponds to the column sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by the corresponding column variable). C^{10} refers to the measure of total connectedness (see section 2.2).

TABLE 8 Connectedness: fastSOM

	AUD	CAD	CHF	EUR	NOK	NZD	SEK	USD	JPY	IN
AUD	53.5	11.2	1.3	2.8	5.5	17.6	4.7	1.9	1.5	46.5
CAD	11.4	59.3	1.6	2.5	5.0	6.9	3.2	7.8	2.3	40.7
CHF	1.3	1.5	69.7	10.1	3.1	1.7	3.7	3.0	5.9	30.3
EUR	2.8	2.6	8.7	51.0	9.3	2.8	15.3	3.3	4.1	49.0
NOK	5.8	4.8	2.8	9.7	56.3	4.0	14.6	0.9	1.1	43.7
NZD	18.6	7.1	1.7	2.9	4.1	58.9	3.4	1.4	1.9	41.1
SEK	4.8	3.3	3.4	15.7	14.1	3.4	52.4	1.3	1.6	47.6
USD	2.1	8.4	3.0	3.8	1.0	1.5	1.3	65.1	13.8	34.9
JPY	1.4	2.3	5.8	4.5	1.0	1.8	1.2	14.2	67.8	32.2
OUT	48.2	41.3	28.4	52.0	43.0	39.7	47.6	33.7	32.1	$C^{10} = 40.7$

Note: The table shows the (10 periods ahead) forecast error variance decomposition of the SVAR model as an average of all possible Cholesky orderings. The column 'IN' corresponds to the row sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by [international] shocks). The column 'OUT' corresponds to the column sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by the corresponding column variable). C^{10} refers to the measure of total connectedness (see section 2.2).

influence of the Norwegian krone in a set of 34 currencies. Wen and Wang (2020) study volatility connectedness, finding a high connectedness of the Norwegian krone to the Swedish krona, but not to the US dollar, confirming our results. Other studies on currency connectedness, like Wan and He (2021) measure dynamic connectedness of currencies in G7 countries, finding the US dollar to be the highest transmitter of spillovers to other currencies. While Huynh et al. (2020) study the role of trade policy uncertainty for currency connectedness, Kühl (2018) focuses on the role of macroeconomic fundamentals and non-fundamental factors in comovements of currencies.

We also find evidence suggesting that the Australian dollar causes the New Zealand dollar, which was not obvious when other measures of connectedness were applied. The same accounts for the $USD \rightarrow JPY$ relationship. Another interesting edge is the $EUR \rightarrow CHF$ link, which reflects the minimum exchange rate of CHF 1.20 per EUR.²⁰

As a robustness analysis, we estimate the spillover matrix of Table 9 without over-identifying restrictions on the contemporaneous matrix. More specifically, we use one of the possible causal orderings which is consistent with the findings of the PC algorithm and estimate the structural VAR using a Cholesky decomposition. The

TABLE 9 Connectedness: PC algorithm ($\alpha = 0.1$)

	AUD	CAD	CHF	EUR	NOK	NZD	SEK	USD	JPY	IN
AUD	53.0	16.8	0.1	0.0	24.1	0.0	0.0	5.9	0.1	47.0
CAD	0.1	61.6	0.1	0.1	16.8	0.0	0.0	21.2	0.1	38.4
CHF	0.1	0.0	61.9	9.1	8.4	0.0	3.5	8.9	8.0	38.1
EUR	0.5	0.3	0.1	39.6	36.1	0.0	13.8	6.9	2.7	60.4
NOK	0.4	0.1	0.0	0.0	98.6	0.0	0.2	0.2	0.5	1.4
NZD	24.3	11.2	0.4	0.0	14.5	42.3	0.0	6.6	0.6	57.7
SEK	1.6	0.6	0.0	0.0	52.0	0.0	45.1	0.3	0.5	54.9
USD	0.2	0.1	0.0	0.1	0.2	0.0	0.0	99.3	0.0	0.7
JPY	0.0	0.0	0.0	0.0	0.3	0.0	0.1	41.2	58.4	41.6
OUT	27.1	29.2	0.7	9.4	152.4	0.1	17.6	91.1	12.5	$C^{10} = 37.8$

Note: The table shows the (10 periods ahead) forecast error variance decomposition of the SVAR model, which is identified by the PC algorithm. The column 'IN' corresponds to the row sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by [international] shocks). The column 'OUT' corresponds to the column sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by the corresponding column variable). C^{10} refers to the measure of total connectedness (see section 2.2).

obtained spillover matrix (see Table C1) shows a very similar pattern when compared to Table 9.²¹ The total connectedness increases slightly because of missing over-identifying restrictions.²² However, the relative importance of the shocks remains the same, therefore the results are qualitatively similar. For instance, shocks to the euro explain 8.1% (9.1% with over-identifying restrictions) of Swiss franc forecast error variance and only 0.3% (0.0% with over-identifying restrictions) of New Zealand forecast error variance, although the coefficient b_{97} is now unrestricted. It is also surprising to see that the measure of total connectedness, $C^{10} = 41.0$ is very similar to the Cholesky application with random ordering (here, $C^{10} = 41.1$). This finding supports Diebold and Yilmaz (2014), who argue that the system-wide summary measure (C^h) is often robust to the Cholesky ordering. However, this does not change the fact that the result is due to a random ordering and thus (potentially) a result of misspecification.

In summary, the PC algorithm provides us with well-defined directed edges, which enable us to unveil a directed network of exchange rates. This gives the PC algorithm a clear advantage over the other presented methods. However, the PC algorithm is computationally also the most intense procedure, particularly because of the time-consuming bootstrap.²³

4.4 | Cluster analysis


In this section, we exploit the connectedness between exchange rates (displayed in Table 9) in order to divide the network into clusters (also known as communities or modules). A cluster is characterised by a high number of

edges between nodes within the cluster, relative to the number of edges to nodes outside the cluster. In this sense, we visualise the previously estimated connectedness and identify groups of exchange rates with a relatively high intra-group connectedness. These groups can be interpreted as currency blocs. The currencies of a bloc are likely to move in tandem, which is important information for policy makers and the management of currency risk. Note that our definition of a currency bloc is more general than the definition by Fischer (2016), for instance.

The quality of the partitioning of a whole network, which can consist of as many clusters as nodes, is thus often expressed by a measure, depending on the differences between the numbers of edges within clusters and the numbers of edges that would exist if it were a random network model. Hence, positive values indicate the existence of clusters. This measure is referred to as *modularity*. For a detailed explanation, we refer the reader to Blondel et al. (2008), who propose a popular algorithm (hereafter: Louvain algorithm) which detects the *best* clustering by maximizing modularity. Initially, the algorithm assigns each node to a single cluster. In a second step, the algorithm moves nodes to new clusters if gains in modularity can be achieved until no additional gain can be achieved (see Blondel et al., 2008). One drawback of the Louvain algorithm is that it is designed for undirected networks. Consequently, it is not feasible given the causal structure of our network. Dugué and Perez (2015) provide a solution to this common problem. They modify the Louvain algorithm in such a way that it allows for *directed modularity* as defined by Leicht and Newman (2008).


Using the Directed Louvain algorithm by Dugué and Perez (2015), we aim to partition the network in Table 9.

FIGURE 2 Partition according to the directed Louvain algorithm. This figure shows the exchange rates clustered according to the spillover matrix of the PC algorithm (Table 9). Cluster 1 (dark grey): AUD, CAD, NZD; Cluster 2 (grey): EUR, NOK, SEK; Cluster 3 (light grey): CHF, USD, JPY. Causation propagates clockwise.



Notes: This figure shows the exchange rates clustered according to the spillover matrix of the PC algorithm (Table 9). Cluster 1 (dark grey): AUD, CAD, NZD; Cluster 2 (grey): EUR, NOK, SEK; Cluster 3 (light grey): CHF, USD, JPY. Causation propagates clockwise.

FIGURE 3 Exchange rate movements following the Brexit referendum. This figure shows pound sterling exchange rate movements following the Brexit referendum (pound sterling in quantity quotation). Exchange rates are marked according to their corresponding cluster. Cluster 1 (solid): AUD, CAD, NZD; Cluster 2 (dashed): EUR, NOK, SEK; Cluster 3 (dotted): CHF, USD, JPY. Source: ECB.



Notes: This figure shows pound sterling exchange rate movements following the Brexit referendum (pound sterling in quantity quotation). Exchange rates are marked according to their corresponding cluster. Cluster 1 (solid): AUD, CAD, NZD; Cluster 2 (dashed): EUR, NOK, SEK; Cluster 3 (dotted): CHF, USD, JPY. Source: ECB.

We observe in Figure 2 that the algorithm classifies the G10 currencies into three different clusters. The first cluster contains the AUD, CAD and NZD, which are often referred to as commodity currencies. Another common property is that investments in these countries provide

the investor with a relatively high yield. The NOK is often also referred to as a commodity currency, but it is part of the second cluster. In addition to the NOK, this cluster also contains the EUR as well as the SEK and thus European currencies only. The third cluster contains the

CHF, USD and JPY. These currencies are often referred to as safe haven and/or carry funding currencies (see Ferreira Filipe & Suominen, 2013; Hossfeld & MacDonald, 2015). Thus, the latter group has the tendency to appreciate in times of financial stress, either because investors are seeking a safe haven for their investments or due to the unwinding of carry trades.

4.5 | Empirical assessment of the clustering: Brexit referendum

In this section, we assess the quality of the partition suggested by the Directed Louvain algorithm. To this end we normalise all exchange rates before the referendum on the UK's membership of the EU, the result of which surprised many market participants, and discuss their movements during the trading days following the referendum (see Figure 3). The Brexit referendum is an appealing example, as it is a shock to the numéraire, which affects all other currencies. Following the referendum, the pound sterling depreciated against all currencies in our sample. Figure 3 shows the movements of the pound sterling exchange rates against all currencies in quantity quotation (cross rates of ECB reference rates). In order to simplify the interpretation, exchange rates have been normalised to 100 on the day of the referendum (23 June 2016 is day 0).²⁴ The similarity of movements within clusters is striking. Currencies within the first (solid lines) and second (dashed lines) cluster, in particular, move closely in tandem. Only currencies in the third cluster (dotted lines) display a slightly larger dispersion. Nor is it surprising that the pound sterling depreciates strongly against the third cluster, which reflects safe haven and carry funding currencies. These are supposed to appreciate in times of financial stress. Additionally, it is expected that the European currencies appreciate the least of the three clusters, because the uncertainty surrounding Brexit means uncertainty for the European Monetary Union. The Swiss franc, which has been found by our procedure to belong to the cluster of safe havens, moves in the case of the Brexit experiment more closely in line with European currencies.

Overall, we observe that exchange rate movements follow a very similar pattern, but we also see that the dispersion within clusters is strikingly low.

5 | CONCLUSIONS

The literature on connectedness between exchange rates has so far ignored a potential causal structure. Research along the lines of Diebold and Yilmaz (2014) is based on

a FEVD in a VAR framework. The difficulty in this context is the identification of the variance-covariance matrix in order to orthogonalise the shocks. We show that a Cholesky decomposition, which is frequently used, can lead to arbitrary results, as the outcome depends heavily on the ordering of the variables. A generalised FEVD is independent of the ordering of the variables, but it is unable to detect causality between the shocks. The same applies when all possible orderings of variables are considered (see Klößner & Wagner, 2014).

We address this problem by employing a causal search algorithm from the machine learning literature, which is able to find causality in contemporaneous data. This approach is then applied to the G10 currencies, whereby nine currencies are modelled vis-à-vis the pound sterling as the numéraire currency. Our results suggest that the US dollar and the Norwegian krone are the most independent currencies in our sample. Shocks to these currencies affect a large set of other currencies. We also observe that connectedness between commodity currencies and those that are often referred to as safe haven and/or carry funding currencies is particularly high.

Using a clustering algorithm, we identify three currency clusters which confirm the previous findings. The first cluster contains commodity currencies such as the AUD, CAD and NZD. The second cluster comprises the European currencies EUR, NOK and SEK. Finally, the third cluster contains the CHF, USD and JPY—currencies, which are often referred to as safe haven or carry funding currencies. They have the tendency to appreciate in times of financial stress.

In an additional exercise, we evaluate the movements of currencies with respect to their clusters following the Brexit referendum. We observe that the dispersion of exchange rate movements within clusters is indeed relatively low, particularly for the first and second clusters. The third cluster shows the strongest appreciation against the pound sterling following the referendum. The Swiss franc, however, appears to move more closely in line with other European currencies (second cluster).

Overall, these estimates provide important information for policy makers and practitioners, as they shed light on potential co-movements between certain exchange rates.

ACKNOWLEDGEMENTS

The authors are grateful to Joscha Beckmann, Geert Bekaert, Christoph Fischer, Ulrich Grosch, Kevin D. Hoover, M. Hashem Pesaran and an unknown referee, as well as participants of the Bundesbank research seminar, the ICMAIF 2019, Rethymno, 25th CEF conference, Ottawa, the 6th IAAE conference, Nicosia, and the Swiss Society of Economics and Statistics conference 2021, Zürich, for helpful comments and suggestions. The views

expressed in this paper are those of the authors and do not necessarily coincide with the views of the Deutsche Bundesbank or the Eurosystem.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in the ECB Statistical Data Warehouse (SDW) at https://urldefense.com/v3/__https://sdw.ecb.europa.eu/

ORCID

Timo Bettendorf  <https://orcid.org/0000-0002-6375-055X>

ENDNOTES

¹ Note that the G10 currencies refer to the 10 most heavily traded currencies and not to the Group of Ten countries.

¹ In order to measure spillover effects using variance decompositions, the numéraire currency should neither be an anchor currency nor be pegged to another currency. Returns of fixed exchange rates have no volatility, implying that variance decompositions would be meaningless. Moreover, we interpret an exchange rate as an asset price and focus in our empirical analysis on exchange rate changes. Within this strand of literature, exchange rate changes (log differences) are referred to as exchange rate returns.

² Previous studies have focussed on the ability of the PC algorithm, a causal search algorithm which will be explained in the following sections, to detect the correct causal structure. We, however, focus on the ability of different approaches to detect the correct degree of network connectedness.

³ PC stands for the initials of its inventors, Peter Spirtes and Clark Glymour.

⁴ On these grounds, we estimate our SVAR models equation by equation using OLS.

⁵ For an introduction to FEVDs, see Lütkepohl (2005).

⁶ Chan-Lau (2017) studies the advantages of the Lanne and Nyberg (2016) approach in a connectedness application. Note that this is not the same GFEVD approach as in Diebold and Yilmaz (2014). Nevertheless, it is also subject to the shortcoming that it fails in modelling contemporaneous causality.

⁷ The chosen settings are: (conservative = TRUE, solve.confl = TRUE, u2pd = c('relaxed')). By choosing the conservative rule instead of the retry option, the algorithm produces a fully order-independent output; see Colombo and Maathuis (2014).

⁸ Note that the PC algorithm is a constraint-based approach, while the GES algorithm is a score-based method. Constraint-based approaches work with conditional independence tests. Score-based approaches assign scores to particular graph structures based on the data fit, for example using scoring metrics like the BIC score, which we use here.

⁹ Given positive correlations between variables in applications to most markets, it is a reasonable assumption to focus on negative entries in the contemporaneous matrix.

¹⁰ In contrast to our application, we perform 100 bootstrap runs in the Monte Carlo experiment.

¹¹ We identified 17 links in our application. Hence, the sum of in- and out-degree is 34. Thirty-four divided by 9 is 3.8.

¹² As our focus is on the identification of the contemporaneous causal structure and its impact on the connectedness between returns, we are not interested in dynamic total connectedness, which is a standard procedure in this literature.

¹³ Note that all exchange rates are fixed at the same time. Hence, trading times do not overlap.

¹⁴ The minimum rate was abandoned on 13 January 2015, which caused the Swiss franc to appreciate strongly against several major currencies.

¹⁵ Note that even if the numéraire is an independent currency, hard pegs among other currencies in the sample cause collinearity.

¹⁶ When performing Ljung-Box tests on the residuals of the VAR (1), the null hypothesis of no autocorrelation cannot be rejected in each of the equations under the 5% significance level. Hence, a lag length of one in the VAR is adequate to model the multivariate dynamics of the system.

¹⁷ Bi-directed edges have not been found in our application. Where a bi-directed edge is detected in the bootstrap of the application, these edges are displayed under 'bi-directed' in Table 5. Overall, bi-directed edges are not prominent in our application, no bi-directed edge has been found on the original data, and only one edge showed a relevant occurrence of bi-directed outcomes in the bootstrap. The absence of bi-directed edges indicates that the system of exchange rates is contemporaneously self-contained. Latent variables, which might affect several exchange rates, seem to enter only with a lag.

¹⁸ Note that due to the over-identifying restrictions, this ordering is not unique. For example, the Japanese yen could also be ordered behind the Norwegian krone and the following other currencies, but it needs to be before the euro. Hence, some other orderings would also be consistent with the output of the causal search algorithm. However, the computed connectedness measures in the following are not influenced by our choice of a recursive ordering.

¹⁹ We observe that the sum of variance shares is close to unity, but not exactly unity.

²⁰ The minimum exchange rate was introduced by the Swiss National Bank on 6 September 2011 and abandoned on 15 January 2015. It served as a key monetary policy instrument.

²¹ Note that in Table C1 the currencies follow the recursive ordering.

²² Given the Cholesky decomposition, we no longer have exact zero entries in the lower diagonal of the contemporaneous matrix. The non-zero coefficients translate into minor changes in the variance shares.

²³ However, we found that working with a small number of bootstrap runs usually leads to the same result as with a large number of runs. For example, for our application we ran 100 experiments with bootstraps with 100 runs each and found that 86 experiments lead to the same contemporaneous matrix as a bootstrap with 10,000 runs, while in 14 experiments one of the edges is found to be directed differently.

²⁴ Note that the fixing of ECB reference rates takes place at 14:15 CET—before the results of the referendum were published.

REFERENCES

- Aloosh, A., & Bekaert, G. (2021). Currency factors. *Management Science* 0(0). <https://doi.org/10.1287/mnsc.2021.4023>
- Barigozzi, M., & Brownles, C. (2019). Nets: Network estimation for time series. *Journal of Applied Econometrics*, 34(3), 347–364.
- Blondel, V., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, P10008, 1–12.
- Chan-Lau, M. J. A. (2017). *Variance decomposition networks: Potential pitfalls and a simple solution*. Working paper 17/107. International Monetary Fund.
- Chickering, D. M. (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov), 507–554.
- Colombo, D., & Maathuis, M. H. (2014). Order-independent constraint-based causal structure learning. *The Journal of Machine Learning Research*, 15(1), 3741–3782.
- De Santis, R. A., & Zimic, S. (2018). Spillovers among sovereign debt markets: Identification through absolute magnitude restrictions. *Journal of Applied Econometrics*, 33(5), 727–747.
- Demiralp, S., & Hoover, K. D. (2003). Searching for the causal structure of a vector autoregression. *Oxford Bulletin of Economics and Statistics*, 65, 745–767.
- Demiralp, S., Hoover, K. D., & Perez, S. J. (2008). A bootstrap method for identifying and evaluating a structural vector autoregression. *Oxford Bulletin of Economics and Statistics*, 70(4), 509–533.
- Demiralp, S., Hoover, K. D., & Perez, S. J. (2014). Still puzzling: Evaluating the price puzzle in an empirically identified structural vector autoregression. *Empirical Economics*, 46(2), 701–731.
- Demirer, M., Diebold, F. X., Liu, L., & Yilmaz, K. (2018). Estimating global bank network connectedness. *Journal of Applied Econometrics*, 33(1), 1–15.
- Deutsche Bundesbank. (2019). Parallels in the exchange rate movements of major currencies. *Monthly Report*, July 2019, 19–37.
- Diebold, F. X., & Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *The Economic Journal*, 119(534), 158–171.
- Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting*, 28(1), 57–66.
- Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1), 119–134.
- Diebold, F. X., & Yilmaz, K. (2015). *Financial and macroeconomic connectedness: A network approach to measurement and monitoring*. Oxford University Press.
- Dugué, N., & Perez, A. (2015). *Directed louvain: Maximizing modularity in directed networks* (PhD thesis). Université d'Orléans.
- Ferreira Filipe, S., & Suominen, M. (2013). *Currency carry trades and funding risk*. AFA 2014 Philadelphia Meetings Paper.
- Fischer, C. (2016). Determining global currency bloc equilibria: An empirical strategy based on estimates of anchor currency choice. *Journal of International Money and Finance*, 64, 214–238.
- Frankel, F., & Wei, S.-J. (2008). Estimation of de facto exchange rate regimes: Synthesis of the techniques for inferring flexibility and basket weights. *IMF Staff Papers*, 55(3), 384–416.
- Frankel, J., & Xie, D. (2010). Estimation of de facto flexibility parameter and basket weights in evolving exchange rate regimes. *American Economic Review*, 100, 568–572.
- Greenwood-Nimmo, M., Nguyen, V. H., & Rafferty, B. (2016). Risk and return spillovers among the G10 currencies. *Journal of Financial Markets*, 31, 43–62.
- Greenwood-Nimmo, M., Nguyen, V. H., & Shin, Y. (2017). *What's mine is yours: Sovereign risk transmission during the European debt crisis*. Working Paper 17/17, Melbourne Institute.
- Heinlein, R., & Krolzig, H.-M. (2012). Effects of monetary policy on the US Dollar/UK pound exchange rate. Is there a 'delayed overshooting puzzle'? *Review of International Economics*, 20, 443–467.
- Hossfeld, O., & MacDonald, R. (2015). Carry funding and safe haven currencies: A threshold regression approach. *Journal of International Money and Finance*, 59, 185–202.
- Huynh, T. L. D., Nasir, M. A., & Nguyen, D. K. (2020). Spillovers and connectedness in foreign exchange markets: The role of trade policy uncertainty. *The Quarterly Review of Economics and Finance*. <https://doi.org/10.1016/j.qref.2020.09.001>
- IMF. (2016). *Annual report on exchange arrangements and exchange restrictions 2016*. Technical Report, International Monetary Fund.
- Kalisch, M., Mächler, M., Colombo, D., Maathuis, M. H., & Bühlmann, P. (2012). Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, 47(11), 1–26.
- Kilian, L., & Lütkepohl, H. (2017). *Structural vector autoregressive analysis*. Cambridge University Press.
- Klößner, S., & Wagner, S. (2014). Exploring all VAR orderings for calculating spillovers? Yes, we can!—A note on Diebold and Yilmaz (2009). *Journal of Applied Econometrics*, 29(1), 172–179.
- Kühl, M. (2018). Excess comovements between the euro/US dollar and pound sterling/US dollar exchange rates. *Applied Economics*, 50(34–35), 3664–3685.
- Lanne, M., & Nyberg, H. (2016). Generalized forecast error variance decomposition for linear and nonlinear multivariate models. *Oxford Bulletin of Economics and Statistics*, 78(4), 595–603.
- Le, T., Martin, F., & Nguyen, D. (2018). Dynamic connectedness of global currencies: A conditional Granger-causality approach. Working Paper hal-01806733.
- Leicht, E. A., & Newman, M. E. (2008). Community structure in directed networks. *Physical Review Letters*, 100(11), 118703.
- Lütkepohl, H. (2005). *New introduction to multiple time series analysis*. Springer.
- Ohno, K. (1999). *Exchange rate management in developing Asia*. Working paper 1. Asian Development Bank Institute.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Pesaran, H. M., & Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58(1), 17–29.
- Scida, D. (2018). Structural VAR and financial networks: A minimum distance approach to spatial modeling. SSRN. <https://doi.org/10.2139/ssrn.2860866>
- Spirites, P., Glymour, C., & Scheines, R. (2001). *Causation, prediction, and search* (2nd ed.). MIT Press.
- Swanson, N. R., & Granger, C. W. (1997). Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions. *Journal of the American Statistical Association*, 92(437), 357–367.

- Wan, Y., & He, S. (2021). Dynamic connectedness of currencies in G7 countries: A Bayesian time-varying approach. *Finance Research Letters*, *41*, 101896.
- Wen, T., & Wang, G.-J. (2020). Volatility connectedness in global foreign exchange markets. *Journal of Multinational Financial Management*, *54*, 100617.
- Yang, J., Tong, M., & Yu, Z. (2021). Housing market spillovers through the lens of transaction volume: A new spillover index approach. *Journal of Empirical Finance*, *64*, 351–378.

How to cite this article: Bettendorf, T., & Heinlein, R. (2023). Connectedness between G10 currencies: Searching for the causal structure. *International Journal of Finance & Economics*, *28*(4), 3938–3959. <https://doi.org/10.1002/ijfe.2629>

APPENDIX A: ESTIMATED MATRICES A

$$B_0 = \begin{bmatrix} \text{AUD} & \text{CAD} & \text{CHF} & \text{EUR} & \text{NOK} & \text{NZD} & \text{SEK} & \text{USD} & \text{JPY} \\ \mathbf{1.0000} & \mathbf{-0.6059} & 0.0000 & 0.0000 & \mathbf{-0.2865} & 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & \mathbf{1.0000} & 0.0000 & 0.0000 & \mathbf{-0.3645} & 0.0000 & 0.0000 & \mathbf{-0.4582} & 0.0000 \\ 0.0000 & 0.0000 & \mathbf{1.0000} & \mathbf{-0.6591} & 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{-0.2328} \\ 0.0000 & 0.0000 & 0.0000 & \mathbf{1.0000} & \mathbf{-0.1597} & 0.0000 & \mathbf{-0.4512} & \mathbf{-0.0956} & \mathbf{-0.1258} \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{1.0000} & 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ \mathbf{-0.7157} & \mathbf{-0.0903} & \mathbf{-0.0692} & 0.0000 & 0.0000 & \mathbf{1.0000} & 0.0000 & 0.0000 & \mathbf{-0.0542} \\ \mathbf{-0.1559} & 0.0000 & 0.0000 & 0.0000 & \mathbf{-0.6207} & 0.0000 & \mathbf{1.0000} & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{1.0000} & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{-0.9102} & \mathbf{1.0000} \end{bmatrix} \begin{matrix} \text{AUD} \\ \text{CAD} \\ \text{CHF} \\ \text{EUR} \\ \text{NOK} \\ \text{NZD} \\ \text{SEK} \\ \text{USD} \\ \text{JPY} \end{matrix}$$

which can be inverted and scaled by the standard deviations of the residuals to.

$$\begin{bmatrix} \text{AUD} & \text{CAD} & \text{CHF} & \text{EUR} & \text{NOK} & \text{NZD} & \text{SEK} & \text{USD} & \text{JPY} \\ \mathbf{0.0047} & \mathbf{0.0026} & 0.0000 & 0.0000 & \mathbf{0.0032} & 0.0000 & 0.0000 & \mathbf{0.0015} & 0.0000 \\ 0.0000 & \mathbf{0.0044} & 0.0000 & 0.0000 & \mathbf{0.0023} & 0.0000 & 0.0000 & \mathbf{0.0026} & 0.0000 \\ \mathbf{0.0002} & \mathbf{0.0001} & \mathbf{0.0053} & \mathbf{0.0020} & \mathbf{0.0020} & 0.0000 & \mathbf{0.0012} & \mathbf{0.0020} & \mathbf{0.0019} \\ \mathbf{0.0003} & \mathbf{0.0002} & 0.0000 & \mathbf{0.0031} & \mathbf{0.0030} & 0.0000 & \mathbf{0.0018} & \mathbf{0.0013} & \mathbf{0.0008} \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{0.0062} & 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ \mathbf{0.0034} & \mathbf{0.0023} & \mathbf{0.0004} & \mathbf{0.0001} & \mathbf{0.0026} & \mathbf{0.0045} & \mathbf{0.0001} & \mathbf{0.0018} & \mathbf{0.0005} \\ \mathbf{0.0007} & \mathbf{0.0004} & 0.0000 & 0.0000 & \mathbf{0.0044} & 0.0000 & \mathbf{0.0041} & \mathbf{0.0002} & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{0.0056} & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & \mathbf{0.0051} & \mathbf{0.0060} \end{bmatrix} \begin{matrix} \text{AUD} \\ \text{CAD} \\ \text{CHF} \\ \text{EUR} \\ \text{NOK} \\ \text{NZD} \\ \text{SEK} \\ \text{USD} \\ \text{JPY} \end{matrix}$$

APPENDIX B: MAXIMUM AND MINIMUM CONNECTEDNESS B

TABLE B1 Maximum connectedness considering all permutations of recursive orderings

	AUD	CAD	CHF	EUR	NOK	NZD	SEK	USD	JPY
AUD	99.8	44.0	9.9	18.9	28.4	58.1	26.0	11.8	9.5
CAD	43.8	99.1	11.0	17.7	24.7	31.9	20.2	28.4	12.7
CHF	9.9	11.0	99.5	36.5	15.9	11.2	18.8	15.4	23.0
EUR	18.8	17.7	36.5	99.3	40.6	18.7	55.7	18.3	20.2
NOK	28.5	24.5	15.8	40.3	98.6	22.6	52.5	5.6	5.2
NZD	58.1	32.0	11.2	18.8	22.6	99.7	20.7	9.5	10.8
SEK	25.9	20.1	18.9	55.6	52.7	20.6	99.0	9.0	8.5
USD	12.0	28.4	15.4	18.2	5.7	9.6	9.1	99.4	41.1
JPY	9.4	12.7	23.0	19.9	4.9	10.7	8.1	41.2	99.6

Note: The table shows (10 periods ahead) forecast error variance decomposition values of SVAR models which are identified by a Cholesky decomposition. The entries of the matrix show the maximum entries which can be achieved with a Cholesky decomposition approach considering all possible orderings.

TABLE B2 Minimum connectedness considering all permutations of recursive orderings

	AUD	CAD	CHF	EUR	NOK	NZD	SEK	USD	JPY
AUD	32.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1
CAD	0.0	42.3	0.0	0.0	0.1	0.0	0.0	0.0	0.1
CHF	0.0	0.0	57.1	0.0	0.0	0.0	0.1	0.0	0.0
EUR	0.0	0.0	0.0	30.7	0.0	0.0	0.0	0.0	0.2
NOK	0.3	0.0	0.0	0.0	38.9	0.0	0.0	0.0	0.2
NZD	0.0	0.0	0.0	0.0	0.0	39.4	0.0	0.0	0.1
SEK	0.0	0.0	0.0	0.0	0.0	0.0	32.1	0.0	0.3
USD	0.0	0.0	0.0	0.0	0.1	0.0	0.0	45.6	0.0
JPY	0.0	0.0	0.0	0.0	0.2	0.0	0.0	0.0	50.3

Note: The table shows (10 periods ahead) forecast error variance decomposition values of SVAR models which are identified by a Cholesky decomposition. The entries of the matrix show the minimum entries which can be achieved with a Cholesky decomposition approach considering all possible orderings.

APPENDIX C: ROBUSTNESS ANALYSIS C

TABLE C1 Connectedness: Cholesky (ordering according to PC algorithm)

	USD	JPY	NOK	CAD	AUD	SEK	EUR	CHF	NZD	IN
USD	99.4	0.0	0.2	0.1	0.2	0.0	0.1	0.0	0.0	0.6
JPY	41.2	58.4	0.3	0.0	0.0	0.1	0.0	0.0	0.0	41.6
NOK	5.6	1.2	92.5	0.1	0.4	0.2	0.0	0.0	0.0	7.5
CAD	28.4	0.1	14.6	56.5	0.1	0.0	0.1	0.1	0.0	43.5
AUD	11.8	1.3	20.9	16.2	49.6	0.0	0.0	0.1	0.0	50.4
SEK	9.0	1.9	44.4	0.3	1.1	43.2	0.0	0.0	0.0	56.8
EUR	18.3	5.2	28.6	0.2	0.1	12.3	35.1	0.1	0.0	64.9
CHF	15.4	8.9	8.5	0.0	0.0	2.0	8.1	57.2	0.0	42.8
NZD	9.5	3.0	16.1	10.6	21.0	0.0	0.3	0.1	39.4	60.6
OUT	139.2	21.6	133.6	27.5	22.9	14.7	8.6	0.4	0.1	$C^{10} = 41.0$

Note: The table shows the (10 periods ahead) forecast error variance decomposition of the SVAR model which is identified by a Cholesky decomposition determined by the PC algorithm. The column 'IN' corresponds to the row sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by [international] shocks). The column 'OUT' corresponds to the column sum of the non-diagonal variance shares (i.e., the total share of variance which is explained by the corresponding column variable). C^{10} refers to the measure of total connectedness (see section 2.2).