FACULTY OF BUSINESS & ECONOMICS

THE UNIVERSITY OF MELBOURNE

Melbourne Institute Working Paper Series

Working Paper No. 17/14

Quantifying Informational Linkages in a Global Model of Currency Spot Markets

Matthew Greenwood-Nimmo, Viet Hoang Nguyen and Yongcheol Shin

MELBOURNE INSTITUTE®
of Applied Economic and Social Research

Quantifying Informational Linkages in a Global Model of Currency Spot Markets*

Matthew Greenwood-Nimmo[†], Viet Hoang Nguyen[‡] and Yongcheol Shin[§]

[†] Department of Economics, The University of Melbourne

[‡] Melbourne Institute of Applied Economic and Social Research, The University of Melbourne

[§] Department of Economic and Related Studies, University of York

Melbourne Institute Working Paper No. 17/14

ISSN 1328-4991 (Print)

ISSN 1447-5863 (Online)

ISBN 978-0-7340-4356-6

July 2014

* We are especially indebted to Heather Anderson, Tony Garratt, Guay Lim, Adrian Pagan and Kevin Reilly for their many helpful comments and suggestions and to Martin Evans for sharing the dataset. This paper has benefited greatly from the thoughtful discussion of participants at the Australian Conference of Economists (Canberra, July 2011) and of seminar participants at the Melbourne Institute. Any remaining errors or omissions are the sole responsibility of the authors. The Gauss programmes used in estimation are available on request or from <www.greenwoodeconomics.com>. Corresponding author: Dr Nguyen, telephone: +61 3 90353621; email <vietn@unimelb.edu.au>.

Melbourne Institute of Applied Economic and Social Research

The University of Melbourne

Victoria 3010 Australia

***Telephone* (03) 8344 2100**

***Fax* (03) 8344 2111**

***Email* melb-inst@unimelb.edu.au**

***WWW Address* http://www.melbourneinstitute.com**

Abstract

We develop a global vector autoregressive model to study the transmission of information between currency spot markets. Our model accounts for both simultaneous and dynamic interactions between exchange rates and order flows using historical data from the Reuters Dealing 2000–1 platform for the period May–August 1996. By analysing the network topography of the system, we find that currency markets are intricately linked and that the Deutsche Mark and the Yen exert a leading influence over the European currencies. Furthermore, using a novel technique we find that the Yen and Sterling act as safe haven currencies.

JEL classification: C32, C51, F31, G15

Keywords: Exchange rates, order flows, global VAR, connectedness and spillovers, safe haven currency

1 Introduction

It is well established that information is transmitted between financial markets through a variety of complex channels. Among the most influential papers in the field is that of Engle, Ito, and Lin (1990), which likens the transmission of volatility between markets to a *meteor shower* which gives rise to volatility clustering across markets. Subsequent studies of volatility transmission in a similar vein include Fung and Patterson (1999), Melvin and Melvin (2003) and Cai, Howorka, and Wongswan (2008). The considerable weight of evidence of volatility transmission across markets is highly suggestive of underlying informational linkages between markets. However, the study of volatility transmission alone cannot shed light on the mechanisms by which trading information in one market affects prices in another.

The microstructure literature has opened new avenues to model these informational linkages directly. This literature exploits information about the trading process to develop exchange rate models with superior in- and out-of-sample performance to the class of macro exchange rate models. In this framework, the order flow (defined as the number of buyer-initiated orders less the number of seller-initiated orders) plays an important role as it conveys a great deal of information about current fundamentals as well as traders' private information or beliefs about future fundamentals. Early empirical results presented in Evans and Lyons (2002b) indicate that the order flow can explain approximately 60% of variations in the DEM/USD market and 40% in the JPY/USD market. Subsequent research by Evans and Lyons (2008) and Love and Payne (2008) demonstrates that approximately one-third of public information or macro news is impounded into prices via the order flow.

In a seminal contribution, Evans and Lyons (2002a) develop a static multi-currency portfolio shifts model which explains how trading activity in one market can influence price determination in other markets. The principal limitation of the Evans and Lyons approach is that multi-currency linkages are addressed only in a static setting under the strong assumption of continuous market equilibrium. However, it is widely documented that the collective action of imperfectly informed noise traders may frequently cause the market to under- or over-react with respect to fundamental news, while various risks in the financial markets may also limit the arbitrage of rational traders (e.g. DeLong, Shleifer, Summers, and Waldmann, 1990; Barberis, Shleifer, and Vishny, 1998; Abreu and Brunnermeier, 2002). As a result, we may observe persistent mispricing and significant episodes of market disequilibrium.

We contribute to the literature on currency market linkages by deriving a dynamic error correction representation of the Evans and Lyons (2002a) multi-currency portfolio shifts model. Our model explicitly addresses the concern raised by Sager and Taylor (2006) that dealers may find themselves holding overnight imbalances as a result of mispricing. We then develop a well-specified empirical analogue of our dynamic multi-currency portfolio shifts model using the global vector autoregressive (GVAR) framework advanced by Pesaran, Schuermann, and Weiner (2004) and Dees, di Mauro, Pesaran, and Smith (2007). The GVAR model is built by combining market-specific VAR models using so-called 'link matrices', which contain granular weights used in the construction of cross-sectional averages of the market-specific data.

The GVAR framework is uniquely well suited to the analysis of informational linkages across markets. By virtue of its vector autoregressive form combined with its ability to exploit the panel structure of the price and trading data across markets, the GVAR model explicitly accounts for the dynamic interactions between prices and order flows simultaneously and across all of the currency markets in the system. Specifically, the GVAR model captures the interactions across currency markets through two distinct channels: (i) the price and order flow in market i depend directly on the contemporaneous and lagged values of the weighted average order flow from the other markets in the system, $j \neq i$; and (ii) the weak correlation of shocks across different currency markets. This structure allows for an explicit treatment of the effect of trading information in one market on the prices and order flows observed in other markets.

Our GVAR model significantly extends the frontier demarcated by the existing empirical literature in a number of directions. Unlike the majority of the dynamic literature which deduces information about the underlying linkages among markets via the analysis of volatility spillovers (e.g. Fung and Patterson, 1999; Melvin and Melvin, 2003), we directly model the dynamic linkages between prices and order flows in a framework which is firmly rooted in the theoretical work of Evans and Lyons (2002b). Of the existing dynamic models, ours is most closely related to that of Cai et al. (2008), which studies the transmission of information between five different trading centres (Asia Pacific, the Asia–Europe overlap, Europe, the Europe–America overlap, and America). They estimate separate VAR models for a variety of different informational proxies including both volatility and the order flow. In this way, they find that local information (i.e. within-region) tends to dominate spillovers between regions in terms of economic significance. However, our model differs from theirs in two crucial respects. Firstly, since Cai et al. estimate separate VAR models for each of their informational proxies, their approach cannot illuminate the linkage between price and order flow, which is central to the portfolio shifts model of Evans and Lyons (2002b). Secondly, and on a closely related note, Cai et al.'s analysis is based on the trading of the same currency pair (e.g. trading the Dollar versus the Yen) across different regional trading centres. As such, their model cannot address spillovers *between* currency pairs, which is essential if one wishes to capture the rebalancing of investor portfolios in response to shocks and to uncover how trading in one currency may affect prices and order flows in another.

Our framework is also considerably more general than the static models employed by Evans and Lyons (2002a) and more recently by Daníelsson, Luo, and Payne (2012). Both of these studies analyse a separate single equation model for each market in their respective datasets, where the price of currency i is assumed to depend on both the order flow for currency i and the order flows for other currencies $j \neq i$. At an elementary level, unlike the dynamic models discussed above, static models must inherently abstract from the possibility that prices may not fully adjust to reflect new information within the trading day. More subtly, however, the use of a single-equation framework implies that the trading information from all markets — including the domestic market — is exogenous, which reduces the scope for contemporaneous interactions among markets. For example, such a model cannot directly account for the possibility that foreign order flows may induce local order flows which, in turn, may affect prices. By contrast, the GVAR model distinguishes between within-market information (which is treated as endogenous) and between-market information (which is treated as weakly exogenous). Weak exogeneity testing strongly supports the validity of this distinction in our dataset.

Thus, the GVAR model enjoys a number of methodological advantages relative to the approaches that have adopted in the existing literature. In particular, it allows the modeller to account for both the *direct association* between prices and order flows as well as the *dynamic interactions* among prices and order flows in all currency markets. At the time of writing, we are unaware of any existing studies that have addressed both of these issues. Furthermore, the innovative structure of the GVAR model permits consistent estimation of very large models using reasonably small datasets.

Since the GVAR model is essentially a large VAR model, it can be readily inverted into its vector moving average form, based on which it is straightforward to apply all of the standard tools of dynamic analysis including impulse response functions and forecast error variance decomposition. Furthermore, we demonstrate that the connectedness methodology developed by Diebold and Yilmaz (2009, 2014) can be applied to explore the network topology of a simple GVAR model such as ours without difficulty. Using these tools, and adopting the terminology of Engle et al. (1990), we evaluate the importance of heatwave (within-market or local) effects and meteor shower (between-market or foreign) effects for each of the markets in our sample. Moreover, by comparing the strength of the meteor shower from market j to market i against the heatwave effect in market i, we develop a novel data-driven method to identify leader-recipient relationships between pairs of markets. Finally, by studying the trading patterns across all markets following local currency selling pressure in the i-th market, we are able to identify safe haven currencies with respect to currency i in a data-driven manner.[1]

Our dataset is drawn from the classic order flow dataset of Evans and Lyons (2002a, 2002b) which has subsequently underpinned notable papers by Cao, Evans, and Lyons (2006), Evans and Lyons (2008) and Sager and Taylor (2008), among others. Our dataset contains daily observations over the four-month period May–August 1996 on both the bilateral spot exchange rate against the Dollar and the order flow in eight currency markets: the Deutsche Mark (DEM), the British Pound (GBP), the Japanese Yen (JPY), the Swiss Franc (CHF), the French Franc (FRF), the Belgian Franc (BEF), the Italian Lira (ITL) and the Dutch Guilder (NLG). By working with Evans and Lyons' seminal dataset, we achieve direct comparability of our results against a large and influential body of existing order flow literature, which therefore provides us with a natural benchmark against which to gauge our findings.

Our GVAR model allows us to map out the complex network of interlinkages among the currency markets in the Evans and Lyons dataset in considerable detail for the first time. Our results unambiguously confirm the central contention of the order flow literature – order flow is an important determinant of the price. Furthermore, in keeping with the literature on information transmission across markets, we find that the order flow in market j exerts a non-negligible influence on the trading behaviours and prices observed in other markets, particularly if market j is a large, heavily-traded and liquid market (Daníelsson et al., 2012).

Our results can be broken down into five principal findings. Firstly, we find that error correction is not instantaneous and that disequilibrium can persist beyond the day in which a shock occurs. Static models are inherently unable to accommodate such behaviour and therefore omit an important aspect of information integration within and between currency markets. Hence, we concur with Sager and Taylor (2006) that future theoretical work should strive to account for overnight portfolio imbalances and that empirical research should accommodate

[1] Subsection 2.4 provides a detailed definition of the safe haven currency hypothesis. See Kaul and Sapp (2006) and Ranaldo and Soderlind (2007) as good examples of the safe haven currency literature.

partial adjustment.

Secondly, in contrast to Cai et al. (2008), we find that while heatwave effects are important, they are by no means the dominant factor influencing global currency markets. At the one-day ahead horizon, heatwave effects account for between 13% and 80% of the forecast error variance of prices across markets with the remainder due to meteor showers. However, the equivalent range at the five-days ahead horizon is 12–70%, with six out of eight markets recording a heatwave effect of less than 35%. Hence, our results are strongly consistent with Evans and Lyons (2002a) and Daníelsson et al. (2012), both of which conclude that informational linkages between markets play an important role in price determination.

Thirdly, we find that the smaller markets in our sample (BEF, FRF, ITL, NLG) behave quite differently than the deeper and more liquid markets (DEM, GBP, JPY, CHF). Specifically, among the smaller markets we find that the influence of within-market order flows is much weaker and the role of meteor showers much greater. Furthermore, the meteor shower effects arising from the smaller markets are generally rather weak. This suggests that the larger markets exert a dominant influence in the global foreign exchange market. This leads directly to our fourth main result – the meteor showers from the Deutsche Mark and the Yen are so strong that they exceed the heatwave effect in the markets for the Belgian Franc, the French Franc, the Swiss Franc and the Dutch Guilder. We therefore conclude that there exists a leader-recipient relationship between the Deutsche Mark/Yen and each of these four markets. Interestingly, while the influence of shocks to the Deutsche Mark on other currencies is immediate and profound, the meteor shower effects from the Yen are relatively muted initially before they gradually intensify. This perhaps reflects delays in the transmission of news regarding fundamentals from Japan to the market in London given the time zone difference as well as the available communications technologies in 1996.

Finally, by innovative use of impulse response analysis, we are able to identify safe haven currencies with respect to each market in our sample. We consider that currency j is a safe haven for currency i if buying pressure for currency j intensifies following an adverse shock to currency i which is associated with excess sales of currency i. This mechanism is consistent with investors rebalancing their portfolios in favour of the safe haven currency. Therefore, for each of the eight markets in our sample, we investigate the effect of local currency selling pressure and identify the safe haven as that market which experiences the strongest cumulative local currency buying pressure over the next five trading days. On this basis, we find that the Dollar and the Yen are the main safe havens for the European currencies while the Dollar and Sterling are safe havens for the Yen.

This paper proceeds in 6 sections as follows. Section 2 provides a brief summary of the multi-currency portfolio shifts (MPS) model of Evans and Lyons (2002a) and develops an empirical framework within which to estimate the MPS model in a dynamic setting with a potentially large number of markets. Section 3 reviews the properties of the dataset and presents the core estimation results for the GVAR model. Section 4 analyses the connectedness of the global system using the methodology developed by Diebold and Yilmaz (2009, 2014) in order to measure the strength of meteor shower effects and to identify leader-recipient relationships among the markets. Section 5 analyses impulse response functions in order to identify market-specific safe haven currencies. Finally, Section 6 concludes.

2 A Global Model of Currency Spot Markets

2.1 The Multi-Currency Portfolio Shifts (MPS) Model

The MPS model extends the single-market portfolio shifts model advanced by Evans and Lyons (2002a) by explicitly incorporating the linkages among currency markets. Evans and Lyons (2002b) examine a pure dealership-type exchange economy with T trading days and $K+1$ assets (currencies), where one asset is riskless (i.e. its gross returns are normalised to 1) and the other K risky assets have stochastic payoffs. Agents buy and sell the riskless asset for risky assets so there are K asset markets for K risky assets. Hence, buying pressure for the i-th risky asset signals decreasing risk exposure of that asset. The MPS model is given by:

$$\Delta \boldsymbol{P}_t = \Delta \boldsymbol{R}_t + \boldsymbol{\Lambda} \Delta \boldsymbol{Q}_t, \tag{2.1}$$

where $\Delta \boldsymbol{P}_t = (\Delta P_{1t}, ..., \Delta P_{Kt})'$ is a $K \times 1$ vector of price changes in which P_{it} denotes the price of asset i relative to the riskless asset, $\Delta \boldsymbol{R}_t = (\Delta R_{1t}, ..., \Delta R_{Kt})'$ is a $K \times 1$ vector of payoff increments for the risky assets, $\Delta \boldsymbol{Q}_t = (\Delta Q_{1t}, ..., \Delta Q_{Kt})'$ is a $K \times 1$ vector of order flows and $\boldsymbol{\Lambda}$ is a $K \times K$ matrix capturing the price impacts of order flows. In the general case in which $\boldsymbol{\Lambda}$ is non-diagonal, the price change of asset i is jointly determined by the trading in market i as well as the trading in the remaining $K-1$ markets. The MPS model therefore explicitly accounts for information integration among currency markets, as the linkage between the price in one market and the order flows in other markets is directly modelled.

The MPS model (2.1) is based on the assumption that all market participants are rational and thus all markets clear at the end of each trading day. In this setting, the long-run relationship between exchange rates (prices) and order flows can be expressed as follows:

$$\boldsymbol{P}_t = \boldsymbol{R}_t + \boldsymbol{\Lambda} \boldsymbol{Q}_t, \tag{2.2}$$

where $\boldsymbol{P}_t$, $\boldsymbol{R}_t$ and $\boldsymbol{Q}_t$ are the cumulative representations of $\Delta \boldsymbol{P}_t$, $\Delta \boldsymbol{R}_t$ and $\Delta \boldsymbol{Q}_t$, respectively. However, a number of studies have documented pervasive evidence of persistent mispricing in financial markets due to the presence of noise traders and the limits of arbitrage (e.g. DeLong et al., 1990; Barberis et al., 1998; Abreu and Brunnermeier, 2002)[2]. Under the Evans and Lyons (2002b) portfolio shifts framework, persistent mispricing implies that the market may not clear at the end of each trading day. Therefore, Sager and Taylor (2006) stress that micro exchange rate models which correspond closely to the real world should allow for the possibility that dealers hold overnight imbalances, as one cannot simply assume that customers always absorb dealers' daily inventories. Where the market-clearing assumption cannot be maintained, then (2.2) can be generalised to:

$$\boldsymbol{P}_t = \boldsymbol{R}_t + \boldsymbol{\Lambda} \boldsymbol{Q}_t + \boldsymbol{\xi}_t, \tag{2.3}$$

where $\boldsymbol{\xi}_t = (\xi_{1t}, ..., \xi_{Kt})'$ is a $K \times 1$ vector of error correction terms which represent the aggregate inventory imbalances of the dealers in the K markets. Therefore, if mispricing persists beyond a trading day, then (2.1) is misspecified due to the omission of the error-correction term, $\boldsymbol{\xi}_t$.

[2] For example, DeLong et al. (1990) argue that persistent mispricing arises because of noise trader risk which deters arbitrageurs from taking large positions, while Abreu and Brunnermeier (2002) discuss the importance of the synchronisation problem which arises due to the dispersion of opinions among arbitrageurs.

2.2 Country-Specific Modelling

Under the assumption that the public information increments, $\boldsymbol{R}_t$, are directly impounded into prices and have no effect on order flows, the relationship between price and order flow in currency market i — that is, the i-th equation of (2.3) — can be written as follows:

$$P_{it} = \lambda_i Q_{it} + \boldsymbol{\lambda}_i^{*\prime} \boldsymbol{Q}_{it}^* + \xi_{it}, \tag{2.4}$$

where λ_i is the long-run (equilibrium) impact of the order flow in the i-th market on the price of asset i, $\mathbf{Q}_{it}^* = (Q_{1t}, ..., Q_{i-1,t}, Q_{i+1,t}, ..., Q_{Kt})'$ is a $(K-1)\times 1$ vector recording the order flows in the remaining $K-1$ currency markets, $\boldsymbol{\lambda}_i^* = (\lambda_1, ..., \lambda_{i-1}, \lambda_{i+1}, ..., \lambda_K)$ is a $(K-1)\times 1$ vector of parameters capturing the long-run price impacts of these $K-1$ foreign order flows on the price of asset i and ξ_{it} is the aggregated inventory imbalance of the dealers in market i.

Estimating a dynamic form of (2.3) as a vector autoregression is infeasible with available datasets due to the curse of dimensionality. With K markets in which we observe both the price and the order flow, the dimension of the p-th order VAR model would be $2K(2Kp+1)$. To overcome this issue, we estimate a set of K market-specific VAR models, each of which does not include the $K-1$ foreign order flows directly but rather a weighted average of these foreign order flows. The aggregate foreign order flow with respect to market i is defined as $Q_{it}^* = \sum_{j=1}^{K} w_{ij} Q_{jt}$, where $w_{ij} \geq 0$ are the set of granular weights which satisfy $\sum_{j=1}^{K} w_{ij} = 1$ and $w_{ii} = 0$. The weight assigned to each market depends on its relative trading volume within the global currency market, with greater weight attached to more heavily traded markets which are likely to play a greater role in information dissemination. Specifically, the w_{ij}'s are calculated as follows:

$$w_{ij} = \frac{V_j}{\sum_{j=1}^{K} V_j - V_i} \text{ with } V_j = \sum_{t=1}^{T} \left|\Delta Q_{jt}\right| \text{ and } V_i = \sum_{t=1}^{T} \left|\Delta Q_{it}\right| \text{ for } i, j = 1, ..., K,$$

where ΔQ_{it} and ΔQ_{jt} are the order flows in the i-th and j-th markets, respectively. In other words, w_{ij} is the share of the aggregate excess orders (buying and selling) from market j over the sample period $(\sum_{t=1}^{T} \left|\Delta Q_{jt}\right|)$ in the total excess orders from all markets less market i $(\sum_{j=1}^{K} \sum_{t=1}^{T} \left|\Delta Q_{jt}\right| - \sum_{t=1}^{T} \left|\Delta Q_{it}\right|)$.[3] We therefore modify (2.4) as follows:

$$P_{it} = \lambda_i Q_{it} + \lambda_i^* Q_{it}^* + \xi_{it}, \tag{2.5}$$

where λ_i^* captures the price impact of the weighted-average foreign order flow. Embedding the long-run relationship, (2.5), into an otherwise unrestricted VAR(p) model and assuming that Q_{it}^* is weakly exogenous, we obtain the following market-specific VECM model:

$$\Delta \boldsymbol{X}_{it} = \boldsymbol{\Lambda}_i \Delta Q_{it}^* + \boldsymbol{\alpha}_i \boldsymbol{\beta}_i' \boldsymbol{Z}_{i,t-1} + \sum_{j=1}^{p-1} \boldsymbol{\Gamma}_{ij} \Delta \boldsymbol{Z}_{i,t-j} + \boldsymbol{\varepsilon}_{it}, \tag{2.6}$$

where $\boldsymbol{X}_{it} = (P_{it}, Q_{it})'$ is a 2×1 vector of endogenous variables containing the price and order flow for market i, $\boldsymbol{Z}_{it} = (\boldsymbol{X}_{it}', Q_{it}^*)'$, $\boldsymbol{\beta}_i = (1, \lambda_i, \lambda_i^*)$ is a 3×1 vector of cointegrating parameters,

[3] In principle, one can employ time-varying weights in the construction of the weakly exogenous variables. However, given the relatively short period spanned by our dataset (82 trading days) we do not pursue this option.

$\boldsymbol{\alpha}_i$ is the 2×1 vector of adjustment parameters, $\xi_{it} = \boldsymbol{\beta}_i' \boldsymbol{Z}_{it}$ and the dimensions of dynamic parameter matrices $\boldsymbol{\Lambda}_i$, $\boldsymbol{\Gamma}_{ij}$ are 2×1 and 2×3, respectively.

Using the market-specific model (2.6) we can directly account for the observation of Fung and Patterson (1999) that major and minor currencies play different informational roles, being either leaders or recipients. Note that (2.6) also allows for the presence of feedback trading through the parameters of the order flow equation and via free estimation of the covariance matrix which allows for the possibility that price and order flow shocks may be correlated. This is an important feature of the model in light of the compelling evidence of feedback trading in financial markets (e.g. Hasbrouck, 1991; Cohen and Shin, 2003) and the observation of Daníelsson and Love (2006) that unless data are sampled at the highest possible frequency then contemporaneous feedback trading may be observed.

2.3 The Global Model

The inclusion of the foreign order flow variables in the market-specific VAR models explicitly accommodates linkages among the currency markets. With the K estimated market-specific VAR model in hand, it is straightforward to combine them into the Global VAR model. First, note that (2.6) can be rewritten equivalently as:

$$\boldsymbol{A}_{i0}\boldsymbol{Z}_{it} = \boldsymbol{A}_{i1}\boldsymbol{Z}_{i,t-1} + \boldsymbol{A}_{i2}\boldsymbol{Z}_{i,t-2} + ... + \boldsymbol{A}_{i,p}\boldsymbol{Z}_{i,t-p} + \boldsymbol{\varepsilon}_{it}, \tag{2.7}$$

where $\boldsymbol{A}_{i0} = (\boldsymbol{I}_2, -\boldsymbol{\Lambda}_i)$, $\boldsymbol{A}_{i1} = \boldsymbol{A}_{i0} + \boldsymbol{\alpha}_i\boldsymbol{\beta}_i' + \boldsymbol{\Gamma}_{i1}$, $\boldsymbol{A}_{ij} = \boldsymbol{\Gamma}_{i,j+1} - \boldsymbol{\Gamma}_{i,j}$ for $j = 2, ..., p-1$, and $\boldsymbol{A}_{ip} = -\boldsymbol{\Gamma}_{ip}$. Now, we define the $m \times 1$ vector of global variables $\boldsymbol{X}_t = (\boldsymbol{X}'_{1t}, ..., \boldsymbol{X}'_{Kt})'$, where $\boldsymbol{X}_{it} = (P_{it}, Q_{it})'$ and $m = 2K$. It is possible to express $\boldsymbol{Z}_{it}$ compactly as:

$$\boldsymbol{Z}_{it} = \boldsymbol{W}_i\boldsymbol{X}_t, \; i = 1, 2, ..., K, \tag{2.8}$$

where $\boldsymbol{W}_i$ is the $(2+1) \times 2K$ link matrix. Careful construction of the link matrices is critical in the development of the GVAR model. We construct the $\boldsymbol{W}_i$'s as follows:

$$\underset{3\times16}{\boldsymbol{W}_i} = \begin{pmatrix} \boldsymbol{R}_{i1} & \boldsymbol{R}_{i2} & \boldsymbol{R}_{i3} & \cdots & \boldsymbol{R}_{iK} \\ \boldsymbol{W}_{i1} & \boldsymbol{W}_{i2} & \boldsymbol{W}_{i3} & \cdots & \boldsymbol{W}_{iK} \end{pmatrix}, \; i = 1, ..., K,$$

where:

$$\{\boldsymbol{R}_{ij}\}_{j=1}^K = \left\{ \begin{array}{ll} \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} & \text{if } j \neq i \\ \boldsymbol{I}_2 & \text{if } j = i \end{array} \right\}, \; \{\boldsymbol{W}_{ij}\}_{j=1}^K = \begin{bmatrix} 0 & w_{ij} \end{bmatrix}, \; i = 1, ..., K,$$

and w_{ij} is the weight assigned to market j in creating the weighted-average foreign order flow variable from the perspective of market i. Using (2.8) in (2.7) and stacking the results we obtain:

$$\boldsymbol{H}_0\boldsymbol{X}_t = \boldsymbol{H}_1\boldsymbol{X}_{t-1} + ... + \boldsymbol{H}_p\boldsymbol{X}_{t-p} + \boldsymbol{\varepsilon}_t, \tag{2.9}$$

where $\boldsymbol{\varepsilon}_t = (\boldsymbol{\varepsilon}'_{1t}, \boldsymbol{\varepsilon}'_{2t}, ..., \boldsymbol{\varepsilon}'_{Kt})'$ and $\boldsymbol{H}_j = \left(\boldsymbol{W}'_1\boldsymbol{A}'_{1j}, \boldsymbol{W}'_2\boldsymbol{A}'_{2j}, ..., \boldsymbol{W}'_K\boldsymbol{A}'_{Kj}\right)'$ for $j = 0, 1, ..., p$. The reduced-form GVAR is finally obtained as:

$$\boldsymbol{X}_t = \boldsymbol{G}_1\boldsymbol{X}_{t-1} + \boldsymbol{G}_2\boldsymbol{X}_{t-2} + ... + \boldsymbol{G}_p\boldsymbol{X}_{t-p} + \boldsymbol{\zeta}_t, \tag{2.10}$$

where $\boldsymbol{G}_j = \boldsymbol{H}_0^{-1}\boldsymbol{H}_j$, $j = 1, 2, ..., p$ and $\boldsymbol{\zeta}_t = \boldsymbol{H}_0^{-1}\boldsymbol{\varepsilon}_t$. Following the convention in the GVAR literature, we allow the market-specific shocks to be weakly correlated across markets such that $E\left(\boldsymbol{\varepsilon}_{it}\boldsymbol{\varepsilon}'_{jt}\right) = \boldsymbol{\Sigma}_{\varepsilon,ij}$ for $t = t'$ and 0 otherwise. This is an important feature of the global model as it accommodates contemporaneous feedback trading across markets since shocks to the equation for ΔP_{it} can have contemporaneous effects on ΔQ_{it} and ΔQ_{jt} for $j \neq i$.

The GVAR model in (2.10) links the K currency markets together in a coherent and flexible manner. The model accommodates informational linkages among currency markets through two distinct channels: (i) direct dependence of the domestic price and order flow on the weighted average foreign order flow and its lagged values; and (ii) non-zero contemporaneous dependence of shocks across markets. The first channel represents the direct impact of trading information in the $K-1$ foreign markets on the trading activity and price in market i. By contrast, the second channel accounts for the contemporaneous linkage of shocks across markets and is therefore of particular interest when one seeks to model spillover effects.

2.4 Dynamic Analysis of the Global Model

In the context of the multi-currency portfolio shifts framework, we are principally interested in evaluating the following:

(i) The *heatwave* effect, defined as the impact of a shock emanating from market i on the price and order flow in the same market (see Engle et al., 1990, for a discussion based on sequences of volatile trading days within a given market).

(ii) The *meteor shower* effect, defined as the impact of a shock emanating from market j on the price and order flow of another market $i \neq j$ (see Engle et al., 1990, on the transmission of volatility between markets). A meteor shower is simply a spillover between markets so we shall henceforth use these terms interchangeably.

(iii) The *leader–recipient* relationship. If the spillover from market j to the price in market i dominates the effect of both price and order flow shocks in market i on the price in market i, then market j is a leader and market i a recipient (see Fung and Patterson, 1999 for a related approach based on the error variance decomposition of a small VAR model).[4]

(iv) The *safe haven currency* hypothesis.[5] Currency j is considered a safe haven with respect to currency i if selling pressure for currency i results in buying pressure for currency $j \neq i$.

[4]Note that our concept of the leader-recipient relationship differs significantly from that of Sapp (2002), who examines price leadership from the institutional perspective by studying which institutions incorporate information into their quoted prices first.

[5]We note in passing that safe haven currency effects have been largely neglected in the literature, with only a few exceptions. Ranaldo and Soderlind (2007, p. 25) define a safe haven currency as one which "benefits from negative exposure to risky assets and appreciates when market risk and illiquidity increase". They find that the Swiss Franc, the Yen and, to some extent, the Euro appreciate against the US dollar when US equity prices decrease and when US bond prices and foreign exchange trading volatility increase, suggesting that they exhibit safe haven attributes. Meanwhile, Kaul and Sapp (2006) study safe haven trading in the EUR/USD spot and

That is, currency j is considered to have a lower risk exposure following a shock that leads to a depreciation of currency i.

Since the GVAR model is still a VAR model, all of the standard tools for the dynamic analysis of VAR models can be employed. While it is technically possible to structurally identify shocks in a GVAR model (particularly by means of sign restrictions which permit set identification; see, for example, Cashin et al., 2014), the large majority of existing studies have not attempted to do so due to the difficulties associated with achieving an uncontroversial structure in a multi-market setting. Similarly, reliance on a Wold causal identification scheme is acutely problematic in the setting of a GVAR model as it carries the implication that markets respond to one-another sequentially. Therefore, we follow the standard approach in the GVAR literature and employ generalised impulse response functions (GIRFs) and generalised forecast error variance decompositions (GFEVDs) which are based on non-orthogonalised shocks (Pesaran and Shin, 1998). This is a natural choice given that market participants trade at high frequency and so one would naturally expect to observe non-negligible correlation of shocks across markets when working with daily data (Daníelsson and Love, 2006).

It is natural to model the first three effects using the connectedness methodology advanced by Diebold and Yilmaz (2009, 2014). This technique involves the construction of a weighted directed network on the basis of forecast error variance decompositions. The proportion of the h-step-ahead forecast error variance of variable i explained by shocks originating within (outside) market i is a natural measure of the heatwave (meteor shower) effect. The GVAR model in (2.10) is first recast in its Wold representation as follows:

$$\boldsymbol{X}_t = \sum_{j=0}^{\infty} \boldsymbol{B}_j \boldsymbol{\zeta}_{t-j}, \tag{2.11}$$

where the $\boldsymbol{B}_j$'s are evaluated recursively as:

$$\boldsymbol{B}_j = \boldsymbol{G}_1\boldsymbol{B}_{j-1} + \boldsymbol{G}_2\boldsymbol{B}_{j-2} + \cdots + \boldsymbol{G}_{p-1}\boldsymbol{B}_{j-p+1},\ j = 1, 2, ...,\ \text{with } \boldsymbol{B}_0 = \boldsymbol{I}_m,\ \boldsymbol{B}_j = \boldsymbol{0} \text{ for } j < 0.$$

Following Pesaran and Shin (1998), the h-step ahead GFEVD is written as follows:

$$\varphi_{j\leftarrow i}^{(h)} = \frac{\sigma_{\varepsilon,ii}^{-1} \sum_{\ell=0}^{h-1} \left(\mathbf{e}_j' \boldsymbol{B}_\ell \boldsymbol{H}_0^{-1} \boldsymbol{\Sigma}_\varepsilon \mathbf{e}_i\right)^2}{\sum_{\ell=0}^{h-1} \mathbf{e}_j' \boldsymbol{B}_\ell \boldsymbol{\Sigma}_\zeta \boldsymbol{B}_\ell' \mathbf{e}_j} \tag{2.12}$$

for $i, j = 1, ..., m$ where $\boldsymbol{e}_i$ is an $m \times 1$ selection vector whose i-th element is unity with zeros elsewhere and $\boldsymbol{e}_j$ is an $m \times 1$ selection vector whose j-th element is unity with zeros elsewhere. Note that $\varphi_{j\leftarrow i}^{(h)}$ denotes the contribution of a one-standard error shock to variable i to the h-step forecast error variance (FEV) of variable j. Due to the non-orthogonal structure of $\boldsymbol{\Sigma}_\varepsilon$, the forecast error variance contributions need not sum to 100% across i. Diebold and Yilmaz (2014)

forward markets around Y2K and find that investors channeled their funds into the Dollar in the months preceding December 1999 and well into January 2000, suggesting that the Dollar was considered the principal safe haven currency as Y2K concerns grew.

therefore normalise $\varphi^{(h)}_{j\leftarrow i}$ as follows:

$$\phi^{(h)}_{j\leftarrow i} = 100\left(\frac{\varphi^{(h)}_{j\leftarrow i}}{\sum_{i=1}^{m}\varphi^{(h)}_{j\leftarrow i}}\right),$$

such that $\sum_{i=1}^{m}\phi^{(h)}_{j\leftarrow i} = 100$ and $\sum_{j,i=1}^{m}\phi^{(h)}_{j\leftarrow i} = 100m$. To demonstrate the computation of the spillover and heatwave effects via the Diebold–Yilmaz method, consider the simplest possible setting with two markets ($K = 2$) and hence $m = 2K = 4$ variables in total, in the order of $P_{1,t}$, $Q_{1,t}$, $P_{2,t}$, $Q_{2,t}$. As such, at any forecast horizon h, the forecast error variance contributions can be cross-tabulated as follows:

$$\mathbb{C}^{(h)} = \begin{bmatrix} \phi^{(h)}_{P_1\leftarrow P_1} & \phi^{(h)}_{P_1\leftarrow Q_1} & \phi^{(h)}_{P_1\leftarrow P_2} & \phi^{(h)}_{P_1\leftarrow Q_2} \\ \phi^{(h)}_{Q_1\leftarrow P_1} & \phi^{(h)}_{Q_1\leftarrow Q_1} & \phi^{(h)}_{Q_1\leftarrow P_2} & \phi^{(h)}_{Q_1\leftarrow Q_2} \\ \phi^{(h)}_{P_2\leftarrow P_1} & \phi^{(h)}_{P_2\leftarrow Q_1} & \phi^{(h)}_{P_2\leftarrow P_2} & \phi^{(h)}_{P_2\leftarrow Q_2} \\ \phi^{(h)}_{Q_2\leftarrow P_1} & \phi^{(h)}_{Q_2\leftarrow Q_1} & \phi^{(h)}_{Q_2\leftarrow P_2} & \phi^{(h)}_{Q_2\leftarrow Q_2} \end{bmatrix} \tag{2.13}$$

Note that (2.13) is structured such that the total h-step ahead FEV of the i-th variable in the system is decomposed into the elements of the i-th row of $\mathbb{C}^{(h)}$. Meanwhile, the contributions of the i-th variable to the h-step ahead FEV of all variables in the system are contained in the i-th column of $\mathbb{C}^{(h)}$. Therefore, based on (2.13) and following Diebold and Yilmaz (2009, 2014), the aggregate h-step ahead heatwave and spillover indices, denoted $H^{(h)}$ and $S^{(h)}$, respectively, can be computed as follows:

$$H^{(h)} = \frac{1}{m}\text{trace}\left(\mathbb{C}^{(h)}\right) \quad \text{and} \quad S^{(h)} = \frac{1}{m}\left(\mathbf{e}'\mathbb{C}^{(h)}\mathbf{e} - \text{trace}\left(\mathbb{C}^{(h)}\right)\right), \tag{2.14}$$

where $\mathbf{e}$ is an $m \times 1$ vector of ones. By construction, $H^{(h)} + S^{(h)} = 100$. Furthermore, with some simple algebra we are able to evaluate heatwave and meteor shower effects at the *market level* as opposed to the variable level (via $\mathbb{C}^{(h)}$) or the systemwide aggregate level (via $H^{(h)}$ and $S^{(h)}$). To see this, note that:

$$\begin{aligned} H^{(h)}_{1\leftarrow 1} &= \tfrac{1}{K}\left(\phi^{(h)}_{P_1\leftarrow P_1} + \phi^{(h)}_{P_1\leftarrow Q_1} + \phi^{(h)}_{Q_1\leftarrow P_1} + \phi^{(h)}_{Q_1\leftarrow Q_1}\right), \\ H^{(h)}_{2\leftarrow 2} &= \tfrac{1}{K}\left(\phi^{(h)}_{P_2\leftarrow P_2} + \phi^{(h)}_{P_2\leftarrow Q_2} + \phi^{(h)}_{Q_2\leftarrow P_2} + \phi^{(h)}_{Q_2\leftarrow Q_2}\right), \\ S^{(h)}_{1\leftarrow 2} &= \tfrac{1}{K}\left(\phi^{(h)}_{P_1\leftarrow P_2} + \phi^{(h)}_{P_1\leftarrow Q_2} + \phi^{(h)}_{Q_1\leftarrow P_2} + \phi^{(h)}_{Q_1\leftarrow Q_2}\right), \\ S^{(h)}_{2\leftarrow 1} &= \tfrac{1}{K}\left(\phi^{(h)}_{P_2\leftarrow P_1} + \phi^{(h)}_{P_2\leftarrow Q_1} + \phi^{(h)}_{Q_2\leftarrow P_1} + \phi^{(h)}_{Q_2\leftarrow Q_1}\right), \end{aligned} \tag{2.15}$$

where $H^{(h)}_{1\leftarrow 1}$ ($H^{(h)}_{2\leftarrow 2}$) denotes the heatwave effect within market 1 (market 2) and $S^{(h)}_{1\leftarrow 2}$ ($S^{(h)}_{2\leftarrow 1}$) denotes the meteor shower effect from market 2 to market 1 (from market 1 to market 2). Note that $H^{(h)}_{1\leftarrow 1} + S^{(h)}_{1\leftarrow 2} = 100$ and $H^{(h)}_{2\leftarrow 2} + S^{(h)}_{2\leftarrow 1} = 100$ by construction. With these definitions in place, the net connectedness between market 1 and market 2 can be defined as follows:

$$N^{(h)}_1 = S^{(h)}_{2\leftarrow 1} - S^{(h)}_{1\leftarrow 2} \quad \text{and} \quad N^{(h)}_2 = S^{(h)}_{1\leftarrow 2} - S^{(h)}_{2\leftarrow 1}. \tag{2.16}$$

Note that in this simple case with two markets, $N^{(h)}_1 = -N^{(h)}_2$ by construction. In a more general setting with K markets, the sign of $N^{(h)}_i$ indicates whether market i is a net transmitter of shocks to the system ($N_i^{(h)} > 0$) or a net receiver of shocks from the system ($N_i^{(h)} < 0$). Using these market-level connectedness measures, we can examine the relative importance of within-market (heatwave) and cross-market (meteor shower/spillover) information in explaining the trading activity and price movements in each currency market. By extension, the case where one is interested in spillovers between *groups of markets* (e.g. major versus minor markets) is conceptually analogous and computationally straightforward.

Finally, we can analyse the safe-haven currency hypothesis by means of generalised impulse response analysis. Recall that currency j is a safe haven with respect to currency i if an adverse shock to currency i leads investors to move away from currency i in favour of currency j. That is, currency j is a safe haven with respect to currency i if selling pressure for currency i results in buying pressure for currency j. Therefore, by analysing the directional impact of local currency selling pressure in one currency market on trading activities and prices in other markets, we can observe how investors switch between currencies in the wake of a shock and, therefore, which currencies act as safe havens.

3 Empirical Results

3.1 The Dataset

We employ the dataset used by Evans and Lyons (2002a), which contains daily data on direct interdealer trading for eight-two trading days over the four month period from May 1 to August 31, 1996. The data is drawn from the Reuters D2000-1 platform and covers trades in the following eight currency spot markets against the US Dollar: the German Mark (DEM), the British Pound (GBP), the Japanese Yen (JPY), the Swiss Franc (CHF), the French Franc (FRF), the Belgian Franc (BEF), the Italian Lira (ITL) and the Dutch Guilder (NLG).[6] In each case, trades are defined such that investors buy and sell the US Dollar for other currencies. Thus, in order to convert holdings of Japanese Yen into Sterling, for example, an investor must first sell the Yen to buy the Dollar and then sell the Dollar to buy the Pound.

Table 1 presents basic summary statistics for each market. P_{it} denotes the natural logarithm of the spot exchange rate of the nominated currency against the Dollar (local currency per US Dollar). ΔQ_{it} denotes the order flow in the i-th market in thousands of orders, measured as the difference between positive (Dollar buying) and negative (Dollar selling) orders; hence, for the i-th market, $\Delta Q_{it} > 0$ indicates local currency selling (i.e. Dollar buying) pressure and $\Delta Q_{it} < 0$ indicates local currency buying (Dollar selling) pressure. In all cases, unit root tests indicate that P_{it} and Q_{it} are difference stationary; results are available on request.

— Insert Table 1 here —

The eight markets that we consider divide naturally into two groups according to the volume of total excess trades, $\sum |\Delta Q_{it}|$. The more heavily traded group is composed of the DEM, JPY, CHF and GBP markets, in descending order of trade volume. These may be considered the major markets in our sample and the remainder (FRF, ITL, BEF, NLG) may be viewed as

[6] Additional trading data for the Danish Kroner was excluded due to two missing observations (May 27$^{\text{th}}$ and July 12$^{\text{th}}$, 1996). This is not a significant omission as this market is neither particularly deep nor liquid during our sample period.

the minor markets. The sum of the excess trades in the major markets during our sample period exceeds that for the minor markets by a factor of five, with the respective totals standing at 23,350 and 4,686. As explained above, we employ the total excess trades in each market to compute the weighting matrix required to construct the market-specific weakly exogenous foreign order flow variables (the Q^*_{it}'s). The resulting weight matrix is recorded in Table 2. The weights reveal the importance of the four major markets identified above and of the DEM and JPY markets in particular. The importance of the Deutsche Mark reflects its importance as an investment currency while the weight attached to the Yen is a manifestation of the Yen carry trade to a large degree (Gagnon and Chaboud, 2007).

Insert Table 2 here

Over the time period spanned by our sample, the Yen depreciated mildly while all of the other currencies appreciated. This pattern is clearly visible in Figures 1 (a)-(h), which plot both the log exchange rate and the cumulative order flow for each country. In all cases except the relatively illiquid BEF market, large movements in cumulative order flow are associated with large movements in the exchange rate, in keeping with the predictions of the MPS model. In five markets (DEM, JPY, CHF, FRF and NLG), the price and local order flow series show relatively close comovement which is reflected in the strong positive correlations reported in the seventh column of Table 1. This may be viewed as a normal situation as excess local currency selling orders lead to a depreciation of the local currency. Meanwhile, we observe a negative association between price and local order flow in the GBP, BEF and ITL markets where the respective currencies gradually appreciate against the Dollar in spite of a prolonged sequence of excess local currency selling orders in these markets.

— Insert Figure 1 here —

Two further features of the data are apparent in Figure 1. Firstly, there are notable similarities in the trading patterns and particularly in the prices of the continental European currencies. On an informal level, this is suggestive of either a common risk exposure of the European currencies or of a leader-recipient relationship among them. By contrast, the trading activities and price movements observed in the JPY market differ substantially from the European currency markets. To a lesser extent, this is also true of the market for GBP. This raises the possibility that traders may be able to exploit the Yen and/or Sterling as safe havens relative to the continental European currencies. These initial conjectures will be formally scrutinised below.

3.2 VECM Estimation Results

The Johansen Trace and Maximum Eigenvalue tests provide mixed support for cointegration. Specifically, there is evidence of one cointegrating relationship in the CHF, BEF, NLG and GBP markets but not in the other cases (test results are available on request). A similar result is recorded by Boyer and van Norden (2006), who test for market-specific cointegration between the exchange rate and cumulative order flow in a VAR framework using the same dataset we employ here. The authors conclude that failure to resoundingly detect a cointegrating relationship does not imply the absence of such a relationship. In practice, testing for market-specific cointegration will inevitably be difficult in the current context due to the relatively low power of traditional cointegration tests in small samples. To achieve higher power, we therefore make use of the panel structure of our dataset which allows us to exploit both the time-variation and the cross-section variation in our sample. Table 3 records the results of the panel cointegration tests devised by Westerlund (2007), which provide overwhelming evidence of cointegration between P_{it}, Q_{it} and Q^*_{it} at the panel level.

— Insert Table 3 here —

Although rejecting the null hypothesis of no cointegration in a panel does not confirm the existence of cointegration in each cross-section unit, we will proceed on the basis that there exists a single cointegrating vector between the three variables P_{it}, Q_{it} and Q^*_{it} in each market. It is straightforward to verify whether this assumption yields a dynamically stable error correction model by inspecting the persistence profile of the cointegrating vector with respect to a systemwide shock, as in Figure 2. The Figure plots the persistent profiles alongside 90% bootstrap intervals. Throughout this paper, we employ the non-parametric sieve bootstrap detailed in Dees et al. (2007, especially Supplement A) based on 10,000 stable iterations. From the Figure, we see that the persistence profile converges to equilibrium swiftly in each case, with the effect of the shock typically dying away to zero after 4 to 8 days. Furthermore, with the exception of the market for the Swiss Franc where we observe some mild noise, adjustment is smooth. As stressed by Pesaran and Shin (1996), the absence of excess persistence and the smooth adjustment reflected in the persistence profiles indicates that the long-run relationships embedded in the market-specific models are valid and well-specified.

— Insert Figure 2 here —

We estimate each market-specific VECM including a constant and a restricted time trend. The i-th market-specific model includes p_i lags of $\boldsymbol{Z}_{it} = (\boldsymbol{X}'_{it}, Q^*_{it})'$, where p_i is chosen from the set $(2, 3, 4, 5)$ using the Akaike Information Criterion as documented in Table 4. The market specific models achieve a very good fit to the data relative to the class of macro exchange rate models, with $\overline{R}^2$s in the range 10% to 80%. Furthermore, specification tests reveal that the models are dynamically stable and well-specified. Detailed results are available on request.

— Insert Table 4 here —

The long-run parameters recorded in Table 4 indicate the expected positive relationship between the exchange rate and the local order flow, Q_{it}, in the majority of markets. That is, our results show that selling the i-th currency in favour of the Dollar causes the i-th currency to depreciate against the Dollar. The notable exception is the BEF market, where we find a negative long-run association. This is not unexpected given the negative correlation between exchange rate and local order flow in the BEF market observed in the data and documented in Table 1. Furthermore, it is consistent with Evans and Lyons (2002a), who establish a comparable result and suggest that the market for the Belgian Franc is a recipient and, therefore, that the price of the Belgian Franc is not predominantly determined by local (within-market) trading information. This observation is certainly borne out by the coefficient on the foreign order flow, Q^*_i, which is large and significant for the BEF market.

Looking at the long-run coefficients on Q_i^* across all eight market-specific models, we see a considerable degree of heterogeneity in their magnitude, indicating varying degrees of sensitivity to conditions in foreign markets. In general, our coefficient estimates reveal that the long-run price impact of both domestic and foreign order flows is stronger in the smaller and less liquid markets where a given volume of trading conveys more information. This effect has been well documented in the microstructure literature (Kyle, 1985). Setting the scale differences aside, the sign of the coefficient on foreign order flow is positive in the large majority of cases, indicating that aggregate Dollar-buying pressure in foreign currency markets is typically associated with a depreciation of the local currency against the Dollar. This is an intuitively reasonable result, as an aggregate increase in demand for the Dollar will tend to cause it to appreciate against a basket of currencies.

The error correction coefficients reported in Table 4 convey a great deal of information about the dynamic adjustment of both prices and order flows. With the exception of the market for the Dutch Guilder, the error correction coefficients in the price equations are negative, indicating that disequilibrium errors are corrected in subsequent trading periods. The market for the Guilder is the smallest and least liquid in our sample so it is likely that a significant and negative error correction term would be uncovered with more data. The speed of error correction is relatively slow in all models, which is consistent with the prevalence of persistent mispricing discussed in subsection 2.1. Nevertheless, we observe relatively rapid convergence of the impulse response functions and forecast error variance decompositions employed below due to the joint influence of the error correction terms and the coefficients on lagged first differences in the market-specific VECMs. Finally, while the error correction parameters in the order flow equations may take either positive or negative values in our model, we find that they are typically negative, indicating that feedback trading will tend to correct deviations from equilibrium. Interestingly, the error correction terms in the major markets are considerably larger than in the minor markets, indicating that traders respond to disequilibrium more rapidly in these more liquid markets.

4 Market Connectedness and Spillover Analysis

Figure 3 presents the systemwide heatwave and spillover indices defined in (2.14) over horizons $h = 1, 2, \ldots, 10$. At the systemwide level, heatwave effects are initially marginally dominant, accounting for just over 50% of the 1-step-ahead FEV with spillovers accounting for the remainder. However, as the horizon increases the heatwave effects gradually diminish, reflecting the rapid integration of within-market information via the trading process. As a result, after one trading week meteor shower effects account for more than 60% of the systemwide FEV. This simple exercise highlights the importance of informational linkages among global currency markets and provides the backdrop against which the remainder of our analysis is conducted.

— Insert Figure 3 here —

Table 5 presents the one-step ahead 16×16 connectedness matrix, $\mathbb{C}^{(1)}$, in the form of a heat map. Note that the elements of the 2×2 blocks lying on the prime diagonal collect the within-market (or local) information. For each of these 2×2 blocks, the elements on the diagonal measure heatwave effects while the off-diagonal elements capture the linkage between price and local order flow. Meanwhile the off-diagonal blocks relate to meteor shower effects across markets.

— Insert Table 5 here —

Although heatwave effects are largely dominant in the aggregate at $h = 1$ (as noted above), there is nevertheless some significant variation. In particular, looking at the diagonal blocks of Table 5, we find an interesting distinction between the major markets (DEM, GBP, JPY and CHF) and the smaller markets. Specifically, the heatwave effects are generally weaker among the smaller markets and the average spillover from Q_{it} to P_{it} at $h = 1$ is considerably weaker in this group. This implies that prices in the smaller markets are more sensitive to foreign order flows than prices in the larger markets. For example, the meteor shower effect from the DEM exchange rate to the BEF exchange rate (26.4%) significantly exceeds the heatwave effect of BEF exchange rate (12.4%). A similar phenomenon is observed for both the FRF and NLG markets. Meanwhile, as expected, the meteor shower effects from the minor markets to the major markets are negligible in all cases.

In light of the apparent differences in the behaviour of the major and minor markets, Figure 4 presents an aggregated representation of the spillovers between these two groups over ten trading days. The contrast between the two groups is stark. Firstly, note that heatwave effects and spillovers among the major markets account for almost 90% of the FEV in these markets at $h = 1$ and for approximately 75% at $h = 10$. This indicates that trading activity and prices in the major markets are largely unaffected by conditions in the minor markets. Meanwhile, spillovers from the major markets to the minor markets account for approximately 50% of the FEV of the minor markets on average across all horizons. As a result, the major markets record strong positive net connectedness with respect to the system as a whole, reflecting the importance of shocks in the major markets for global currency trading.

— Insert Figure 4 here —

It is important to note, however, that price and order flow innovations in the DEM market are also key factors influencing the price of the Swiss Franc at $h = 1$. Similarly, we observe considerable spillovers from the Yen to the Swiss Franc. This implies that significant meteor shower effects are not felt only by the smaller markets but also occur among the major markets. Fung and Patterson (1999) also find a close linkage among currency futures markets for DEM, GBP, JPY and CHF. In particular, they find that volatilities in the DEM, GBP and JPY markets explain up to 90% of the volatility in the CHF market, contributing around 30% each. Collectively, the importance of shocks to the Deutsche Mark at the one-day ahead horizon is suggestive of a leader-recipient relationship between the Deutsche Mark and the other continental European currencies, an observation to which we will shortly return. However, we also observe non-negligible spillovers from the CHF, FRF, BEF and NLG exchange rates to the DEM exchange rate, which suggests that the European markets may also share some common risk exposure.

Table 6 presents a similar heat map at the five-days ahead horizon. While we still observe non-negligible heatwave effects, they are considerably weaker than in the one-day ahead case. By contrast, the meteor shower effects are much stronger. The key difference between Tables 5 and 6 lies in the relative magnitude of meteor shower effects coming from the DEM and JPY markets. At the one-day ahead horizon, the Deutsche Mark was by far the dominant currency. However, at $h = 5$ this dominance is challenged by the Yen. This shift arises because spillovers from the Deutsche Mark tend to subside as the horizon increases while spillovers from the Yen intensify at longer horizons. This suggests that while information from the market for the Deutsche Mark is immediately impounded into the trading activity and prices of the other markets, information from the Yen market is impounded more gradually.

— Insert Table 6 here —

The different informational roles of the Deutsche Mark and the Japanese Yen markets can be seen clearly in Table 7, which identifies the foreign market that exerts the strongest leadership effect with respect to each market in the system. To measure the strength of the leadership effect experienced by the i-th market at the h-day ahead horizon, we compute the following:

$$L_{i \leftarrow j}^{(h)} = \left(\phi_{P_i \leftarrow P_j}^{(h)} + \phi_{P_i \leftarrow Q_j}^{(h)} \right) - \left\{ \phi_{P_i \leftarrow P_i}^{(h)} + \phi_{P_i \leftarrow Q_i}^{(h)} \right\}$$

for $j = 1, \ldots, K$, $j \neq i$. The terms in rounded parentheses measure the meteor shower effect from market j to the price in market i while the terms in braces measure the effect of local information on the price in market i. Hence, if $L_{i \leftarrow j}^{(h)} > 0$ then market j leads market i at horizon h. For each market i, Table 7 records the leader as the market j for which $L_{i \leftarrow j}^{(h)} > 0$ is maximised. The growing importance of shocks to the Yen at longer horizons is readily apparent.

— Insert Table 7 here —

Careful consideration of the heat maps in Tables 5 and 6 provides one final insight. The price and order flow in the markets for both the Japanese Yen and the British Pound are largely determined by their domestic innovations. This suggests that these two markets may be insulated to some extent from shocks in the continental European markets. As such, it is likely that investors may view them as safe havens following adverse shocks to the European markets. Formal evaluation of the safe haven hypothesis requires directional information from impulse response analysis and will be conducted in the next section.

5 Impulse Response Analysis

To further examine the linkages among currency markets, we conduct two counterfactual exercises based on generalised impulse response analysis with respect to adverse shocks in the order flow equations of selected market-specific models. The order flow is often used in the micro exchange rate literature to proxy for publicly unavailable information which will be impounded into the exchange rate through the trading process. Hence, our counterfactual exercises trace the information integration process within and between currency markets given the arrival of news in a designated market. Significantly, and unlike the body of existing research, our GVAR model allows us to observe not only the price movement across markets following a given shock but also the order flows that drive this movement. Hence, our analysis will yield detailed insights into the dynamic interactions between the markets and thereby provide new insights into the leader-recipient and safe haven hypotheses in a manner that was hitherto infeasible.

Specifically, we analyse impulse responses following a positive shock to the order flow equation in the DEM market and in the JPY market. These shocks represent local currency selling pressure in the respective markets; that is, traders are selling either the Deutsche Mark or the Yen to buy the Dollar. They may then liquidate their Dollar position in favour of another currency. Our focus on shocks to the Deutsche Mark and the Yen is natural in light of the preceding analysis, which has shown that these markets are not only the most liquid in our sample but that they also exert a dominant effect over many other markets in the system. Furthermore, it is consistent with the focus of the seminal paper of Evans and Lyons (2002b), which studies the importance of order flow information as a determinant of the exchange rate in the Deutsche Mark/Dollar and Yen/Dollar markets using data of the same vintage that we employ here.

5.1 Deutsche Mark Selling Pressure in the DEM/USD Market

Figure 5 reports the first difference of the impulse responses following a shock of one standard deviation in the DEM order flow equation, which amounts to 110.79 excess Deutsche Mark selling orders in the DEM market. By working with the first difference of the impulse response function, we are able to frame our discussion naturally in terms of order flows and price changes as opposed to cumulative order flows and price levels. The figure reveals that Deutsche Mark selling pressure exerts strong heatwave effects. Furthermore, trading activity in the DEM market persists beyond the day in which the shock occurs and leads to a significant depreciation of the Mark *vis-à-vis* the Dollar.

Importantly, we also observe strong meteor-shower effects in the sense that an adverse shock in the DEM market causes portfolio shifts involving a number of other currencies, including both major and minor markets. This is a reflection of the intricate network of connections among markets identified using connectedness analysis in the preceding section. In particular, we see that Deutsche Mark selling pressure is associated with considerable local currency selling pressure in the GBP, CHF and FRF markets. This indicates that traders may consider the Deutsche Mark, the Pound, the Swiss Franc and the French Franc to share similar risk exposures. Furthermore, we see that trading activity in these markets also continues on the day after the shock occurs, indicating that traders' portfolio shifts occur in a gradual fashion, in part due to the limits of arbitrage and informational asymmetries.

— Insert Figure 5 here —

A further meteor shower effect is of particular interest, namely that from the DEM market to the JPY market. In response to local currency selling pressure in the DEM market, the JPY market experiences considerable Yen buying pressure as some investors convert positions from the Deutsche Mark to the Yen. This, in turn, leads to a marked appreciation of the Yen. Recall that currency j is a safe haven from the perspective of currency i if selling pressure for currency i triggers buying pressure for currency j. Therefore, our analysis reveals that investors regard the

Yen as a safe haven currency relative to the Mark.[7] This finding is consistent with Evans and Lyons (2002a), who find that Deutsche Mark selling pressure leads to an appreciation (albeit insignificant) of the Yen against the Dollar.

Finally, we observe either negligible or insignificant trading responses in the BEF, ITL and NLG markets. Nevertheless, each of these currencies experiences a significant depreciation against the Dollar, with a pattern which closely mimics that described by the Deutsche Mark. This is highly suggestive of a leader-recipient effect among the European currency markets — in keeping with our connectedness analysis — in which the Deutsche Mark is the dominant currency.

5.2 Yen Selling Pressure in the JPY/USD Market

We now consider a selling pressure shock of one standard deviation in the JPY market, which equates to 94.62 excess Yen selling orders. Recall that interest rates in Japan were set close to zero from the end of 1995, creating a significant negative interest rate differential against most other major economies. As of 1996Q3, the interbank rate in Japan was 0.46%, relative to 5.31% in the US, 3.2% in Germany, 5.79% in the UK and 1.89% in Switzerland.[8] It is well known that these interest rate differentials fuelled a sizeable Yen carry trade (Gagnon and Chaboud, 2007; Burnside, Eichenbaum, and Rebelo, 2007; Brunnermeier, Nagel, and Pedersen, 2009). Hence, over our sample period (May 1996–August 1996), a Yen selling pressure shock in the JPY/USD market mimics the conversion of borrowed funds denominated in Yen into investments denominated in Dollars or other currencies.

In keeping with the results of connectedness analysis presented above, Figure 6 shows that Yen selling pressure in the JPY market exerts a strong heatwave effect which is somewhat persistent, with significant feedback trading activity continuing for two days. The Japanese shock also exerts notable meteor shower effects, resulting in modest local currency selling pressure in most other currency markets. The notable exception is the GBP market which experiences pronounced Sterling buying pressure. This is consistent with the carry trade mechanism discussed above as Sterling offers the largest yield among any of the major markets in our sample. However, since the volume of excess Yen selling orders amounts to 113.39 at the five day horizon while the volume of excess Sterling buying orders is just 10.39 over the same period, it is clear that the Yen–Sterling carry trade only accounts for a small proportion of the portfolio shift in the wake of the shock. Much of the remainder is likely to be invested in Dollar-denominated assets to exploit the Yen–Dollar carry trade. Furthermore, since an adverse shock to the Yen causes both Dollar and Sterling buying pressure, it follows that the latter are safe havens with respect to the Yen.

— Insert Figure 6 here —

The shock causes some apparent movement in the Sterling exchange rate, although it is marginally insignificant at the 90% level. A careful examination of the trading data in the GBP

[7] Our analysis indicates that the Dollar is also a safe haven with respect to the Deutsche Mark given that our dataset contains information on bilateral exchange rates *vis-à-vis* the Dollar and that local currency selling pressure in the i-th market is equivalent to Dollar buying pressure in that market.

[8] Data are sourced from the 'Interbank Money Rate' series in the IMF's International Financial Statistics.

market (plotted in Figure 1(b)) suggests that the price and cumulative order flow in this market do not always move in the same direction over our sample period. This is also apparent from Table 1, which reveals that the pairwise correlations between P_{it}, Q_{it} and Q^*_{it} are considerably weaker than in most other markets. Evans and Lyons (2002a) also document that order flows in the JPY market have a negligible impact on the Pound. Nevertheless, the shock generates an immediate and significant depreciation of all the remaining currencies against the Dollar as it creates simultaneous local currency selling pressures (i.e. Dollar buying pressures) in these markets. This confirms our earlier finding that the Yen acts as a leader with respect to many European markets, particularly at longer horizons. A similar although somewhat narrower result is reported by Evans and Lyons (2002a). They find that selling pressure in the JPY/USD market is associated with a significant depreciation of not just the Yen but also of the Swiss Franc. The ability of our GVAR model to detect a greater number of significant responses than the Evans and Lyons model rests in its construction, which takes full account of informational linkages across markets in a dynamic setting in which all variables are globally endogenous.

5.3 Identifying Save Haven Currencies

Our analysis of the GIRFs above confirms that currency spot markets are intricately linked. Overall, we find that both the Deutsche Mark and the Yen play a major informational role in the global system, leading other markets. We also see that the impulse responses tend to converge to zero between one and two days after the initial shock, which suggests that cross-market information is integrated into prices with a modest lag. Finally, by studying the direction of meteor shower effects — which is often absent from existing studies — we are able to identify safe haven currencies with respect to market-specific shocks. In this way, we find that the Yen and the Dollar are safe havens with respect to the Deutsche Mark while the Pound and the Dollar are safe havens for Yen investors.

By extending our analysis to consider adverse shocks to each market in our sample, we are able to identify safe havens for every market. Table 8 identifies as a safe haven the market which experiences the largest cumulative negative order flow (i.e. local currency buying pressure) in the five trading days following an adverse shock to a chosen market. A striking pattern emerges among the major markets – the Yen is an important safe haven for each of the major markets while Sterling provides a safe haven for the Yen, as seen earlier. Turning to the smaller markets, the pattern is less clear although it is interesting to note that in all cases the safe haven identified using our method is one of the major currencies. Furthermore, across all eight markets, it is clear that Yen and Sterling (as well as the Dollar) are the dominant safe havens for the continental European currencies, reflecting their importance for investors wishing to manage their exposure to risks arising from European shocks through diversification.

— Insert Table 8 here —

6 Concluding Remarks

We develop a global model of currency spot markets using the Global VAR framework originated by Pesaran et al. (2004) and Dees et al. (2007). Our model represents a significant extension of the multi-currency portfolio shifts model of Evans and Lyons (2002a). As an error correction model, our GVAR explicitly accounts for persistent mispricing and for the possibility that dealers may hold overnight imbalances. Furthermore, as a dynamic model, it is able to trace the complex interactions between prices and order flows over time. Finally, by virtue of its dynamic panel structure, our model can analyse the interactions among variables both within a given market and also across markets. Hence, our approach can shed light on how trading activity in a chosen market may affect the trading activity and price in all other markets in the system in both the short- and the long-run. Therefore, our model offers a singularly rich framework to explore the intricate network of relationships underlying the global currency market. Such a level of detail has not been achieved by existing studies to date.

Working with the classic order flow dataset of Evans and Lyons (2002a, 2002b), we map out a complex network of interlinkages among the eight currency markets in our sample. Our results can be broken down into five principal findings. Firstly, our results indicate that adjustment to equilibrium following a shock is not necessarily completed within a single trading day, underscoring the necessity for future applied and theoretical work to account for overnight imbalances and persistent mispricing (Sager and Taylor, 2006). Secondly, we find that while heatwave effects are important, they are by no means the dominant factor influencing global currency markets. Meteor shower effects account for more than 60% of the forecast error variance in the system at the five-days ahead horizon. Third, we find that the minor markets (BEF, FRF, ITL, NLG) behave quite differently than the major markets (DEM, GBP, JPY, CHF) in the sense that the price in the minor markets is considerably more sensitive to meteor showers effects. Fourth, we find that the meteor showers from the Deutsche Mark and the Yen are so strong that they exert a leadership effect with respect to the minor markets as well as the Swiss Franc. Interestingly, while the influence of shocks to the Deutsche Mark on other currencies is immediate and profound, meteor showers from the Yen are initially muted but gradually intensify. This perhaps reflects delays in the transmission of news regarding fundamentals from Japan to the physical marketplace in London given the time zone difference as well as the available communications technologies in 1996. Finally, by innovative use of impulse response analysis, we identify safe haven currencies with respect to each market in our sample. The Dollar and the Yen are the main safe havens for the European currencies while the Dollar and Sterling are safe havens for the Yen.

Our work draws attention to the intricate linkages that exist between currency markets. As such, it offers a rich vein for continuing research which may, in time, shed light on some of the celebrated paradoxes of exchange rate research. Work toward a more complete understanding of these linkages is not only important for improving exchange rate modelling itself but is also likely to yield major gains in forecasting performance.

References

Abreu, D., & Brunnermeier, M. K. (2002). Synchronization Risk and Delayed Arbitrage. *Journal of Financial Economics*, *66*, 341 - 360.

Barberis, N., Shleifer, A., & Vishny, R. (1998). A Model of Investor Sentiment. *Journal of Financial Economics*, *49*, 307-343.

Boyer, M. M., & van Norden, S. (2006). Exchange Rates and Order Flow in the Long Run. *Finance Research Letters*, *3*, 235-243.

Brunnermeier, M. K., Nagel, S., & Pedersen, L. H. (2009). *Carry Trades and Currency Crashes* (NBER Working Paper 14473). National Bureau of Economic Research.

Burnside, C., Eichenbaum, M., & Rebelo, S. (2007). The Returns to Currency Speculation in Emerging Markets. *American Economic Review*, *97*, 333-338.

Cai, F., Howorka, E., & Wongswan, J. (2008). Informational Linkage Across Trading Regions: Evidence from Foreign Exchange Markets. *Journal of International Money and Finance*, *27*, 1212-1243.

Cao, H. H., Evans, M., & Lyons, R. K. (2006). Inventory Information. *The Journal of Business*, *79*, 325-364.

Cashin, P., Mohaddes, K., Raissi, M., & Raissi, M. (2014). The Differential Effects of Oil Demand and Supply Shocks on the Global Economy. *Energy Economics*, *in press*.

Cohen, B. H., & Shin, H. S. (2003). *Positive Feedback Trading under Stress: Evidence from the US Treasury Securities Market* (BIS Working Paper No. 122). Bank for International Settlements.

Daníelsson, J., & Love, R. (2006). Feedback Trading. *International Journal of Finance & Economics*, *11*, 35-53.

Daníelsson, J., Luo, J., & Payne, R. (2012). Exchange Rate Determination and Inter-Market Order Flow Effects. *European Journal of Finance*, *18*, 823-840.

Dees, S., di Mauro, F., Pesaran, M. H., & Smith, L. V. (2007). Exploring the International Linkages of the Euro Area: A Global VAR Analysis. *Journal of Applied Econometrics*, *22*, 1-38.

DeLong, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy*, *98*, 703-738.

Diebold, F., & Yilmaz, K. (2009). Measuring Financial Asset Return and Volatility Spillovers, with Application to Global Equity Markets. *Economic Journal*, *119*, 158-171.

Diebold, F., & Yilmaz, K. (2014). On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms. *Journal of Econometrics*, *in press*.

Engle, R. F., Ito, T., & Lin, W.-L. (1990). Meteor Showers or Heat Waves? Heteroskedastic Intra-daily Volatility in the Foreign Exchange Market. *Econometrica*, *58*, 525-524.

Evans, M. D. D., & Lyons, R. K. (2002a). Informational Integration and FX Trading. *Journal of International Money and Finance*, *21*, 807-831.

Evans, M. D. D., & Lyons, R. K. (2002b). Order Flow and Exchange Rate Dynamics. *Journal of Political Economy*, *110*, 170-180.

Evans, M. D. D., & Lyons, R. K. (2008). How is Macro News Transmitted to Exchange Rates? *Journal of Financial Economics*, *88*, 26-50.

Fung, H.-G., & Patterson, G. A. (1999). Volatility Linkage among Currency Futures Markets during US Trading and Non-trading Periods. *Journal of Multinational Financial Management*, *9*, 129-153.

Gagnon, J. E., & Chaboud, A. P. (2007). *What Can the Data Tell Us about Carry Trades in Japanese Yen?* (International Finance Discussion Papers 899). Board of Governors of the Federal Reserve System.

Hasbrouck, J. (1991). Measuring the Information Content of Stock Trades. *Journal of Finance*, *46*, 179-207.

Kaul, A., & Sapp, S. (2006). Y2K Fears and Safe Haven Trading of the U.S. Dollar. *Journal of International Money and Finance*, *25*, 760-779.

Kyle, A. S. (1985). Continuous Auctions and Insider Trading. *Econometrica*, *53*, 1315-35.

Love, R., & Payne, R. (2008). Macroeconomic News, Order Flows, and Exchange Rates. *Journal of Financial and Quantitative Analysis*, *43*, 467-488.

Melvin, M., & Melvin, B. P. (2003). The Global Transmission of Volatility in the Foreign Exchange Market. *The Review of Economics and Statistics*, *85*, 670-679.

Pesaran, M. H., Schuermann, T., & Weiner, S. M. (2004). Modeling Regional Interdependencies using a Global Error-Correcting Macroeconometric Model. *Journal of Business and Economic Statistics*, *22*, 129-162.

Pesaran, M. H., & Shin, Y. (1996). Cointegration and Speed of Convergence to Equilibrium. *Journal of Econometrics*, *71*, 117-143.

Pesaran, M. H., & Shin, Y. (1998). Generalized Impulse Response Analysis in Linear Multivariate Models. *Economics Letters*, *58*, 17-29.

Ranaldo, A., & Soderlind, P. (2007). *Safe Haven Currencies* (Working Papers). Swiss National Bank.

Sager, M. J., & Taylor, M. P. (2006). Under the Microscope: The Structure of the Foreign Exchange Market. *International Journal of Finance and Economics*, *11*, 81-95.

Sager, M. J., & Taylor, M. P. (2008). Commercially Available Order Flow Data and Exchange Rate Movements: Caveat Emptor. *Journal of Money, Credit and Banking*, *40*, 583-625.

Sapp, S. G. (2002). Price Leadership in the Spot Foreign Exchange Market. *Journal of Financial and Quantitative Analysis*, *37*, 425–448.

Westerlund, J. (2007). Testing for Error Correction in Panel Data. *Oxford Bulletin of Economics and Statistics*, *69*, 709-748.

Market	$\sum \Delta Q_{it}^{-}$	$\sum \Delta Q_{it}^{+}$	$\sum \Delta Q_{it}$	$\sum \lvert \Delta Q_{it} \rvert$	$\sum \Delta P_{it}$	$\rho_{(P_{it},Q_{it})}$	$\rho_{(P_{it},Q_{it}^{*})}$	$\rho_{(Q_{it},Q_{it}^{*})}$
DEM	-4.343	4.326	-0.017	8.669	-0.037	0.779	-0.528	-0.263
GBP	-1.082	2.162	1.080	3.244	-0.043	-0.501	-0.312	0.108
JPY	-1.954	4.907	2.953	6.861	0.033	0.769	-0.248	-0.567
CHF	-3.011	1.565	-1.446	4.576	-0.041	0.711	-0.256	-0.628
FRF	-0.919	0.894	-0.025	1.813	-0.022	0.889	0.041	0.218
BEF	-0.151	0.403	0.252	0.554	-0.035	-0.830	-0.016	0.531
ITL	-0.421	0.529	0.108	0.950	-0.035	-0.697	-0.349	0.427
NLG	-0.287	0.132	-0.155	0.419	-0.036	0.939	-0.009	-0.274

NOTES: Each exchange rate is bilateral *vis-à-vis* the US Dollar. For each market, $\sum \Delta Q_{it}^{-}$ (in thousands) denotes the total USD-selling orders over the sample period, $\sum \Delta Q_{it}^{+}$ (in thousands) the total USD-buying orders, $\sum \Delta \lvert Q_{it} \rvert$ the total trading orders, $\sum \Delta Q_{it}$ the net trading orders and $\sum \Delta P_{it}$ the aggregate price change over the sample period. Following Evans and Lyons (2002a), P_{it} is expressed in logarithmic form while Q_{it} and Q_{it}^{*} are cumulative order flows measured in thousands of orders. Q_{it}^{*} is defined as in subsection 2.2. $\rho_{(i,j)}$ denotes the correlation between $i, j \in (P_{it}, Q_{it}, Q_{it}^{*})$.

Table 1: Summary Trading Statistics

	DEM	GBP	JPY	CHF	FRF	BEF	ITL	NLG
DEM	0.000	0.176	0.373	0.248	0.098	0.030	0.052	0.023
GBP	0.364	0.000	0.288	0.192	0.076	0.023	0.040	0.018
JPY	0.429	0.160	0.000	0.226	0.090	0.027	0.047	0.021
CHF	0.385	0.144	0.305	0.000	0.081	0.025	0.042	0.019
FRF	0.343	0.128	0.271	0.181	0.000	0.022	0.038	0.017
BEF	0.327	0.122	0.259	0.172	0.068	0.000	0.036	0.016
ITL	0.332	0.124	0.263	0.175	0.069	0.021	0.000	0.016
NLG	0.325	0.122	0.257	0.172	0.068	0.021	0.036	0.000

NOTES: The (i,j)-th element of the weight matrix, w_{ij}, is calculated as $w_{ij} = \frac{V_j}{\sum_{j=1}^{K} V_j - V_i}$, for $i, j = 1, ..., K$, where $V_j = \sum_{t=1}^{T} \lvert \Delta Q_{jt} \rvert$, $V_i = \sum_{t=1}^{T} \lvert \Delta Q_{it} \rvert$ and ΔQ_{it} and ΔQ_{jt} are the order flows in the i-th and j-th markets, respectively. As such, w_{ij} is the share of the aggregate excess orders (buying and selling) from market j over the sample period ($\sum_{t=1}^{T} \lvert \Delta Q_{jt} \rvert$) in the total excess orders from all markets less market i ($\sum_{j=1}^{K} \sum_{t=1}^{T} \lvert \Delta Q_{jt} \rvert - \sum_{t=1}^{T} \lvert \Delta Q_{it} \rvert$).

Table 2: Matrix of Weights used to Construct Foreign Order Flows

Statistic	Value	z-value	p-value
G_τ	-3.15	-2.09	0.02
G_α	-18.14	-1.74	0.04
P_τ	-8.29	-2.01	0.02
P_α	-15.57	-2.13	0.02

NOTES: The Westerlund cointegration test is based on a panel error correction model (with an intercept and a time trend) for the null hypothesis that there is no cointegration between P_{it}, Q_{it} and Q^*_{it}. P_τ and P_α are the test statistics for the null of no cointegration against the alternative that the panel is cointegrated as a whole. G_τ and G_α are the test statistics for the null of no cointegration against the alternative that at least one unit in the panel is cointegrated.

Table 3: Westerlund Panel Cointegration Test Results

		DEM	GBP	JPY	CHF	FRF	BEF	ITL	NLG
(a) Lag Order	p_i	2	3	2	4	2	2	2	2
(b) LR Coefs	$P_{i,t-1}$	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	$Q_{i,t-1}$	-0.401	-1.818**	-2.452*	2.508	-6.495	14.644*	-2.845	-39.804**
	$Q^*_{i,t-1}$	-3.230**	0.483	-2.051	-6.413	-0.375	-4.639*	-0.935***	-0.098
(c) Adj. Coefs	$\xi^P_{i,t-1}$	-0.001*	-0.001*	-0.001**	-0.001*	0.000 ***	-0.001*	-0.001**	0.000
	$\xi^Q_{i,t-1}$	-0.042**	-0.009	-0.010	-0.015**	0.005***	-0.001	-0.003	-0.003*

NOTES: The lag order is selected using the Akaike Information Criterion. The reported values of the long-run (LR) coefficients are normalised on $P_{i,t-1}$ and are then multiplied by 100 to ease their interpretation (e.g. a dollar buying (positive) shock of 100 trades in the GBP/USD market causes a long-run Sterling depreciation against the US Dollar of 1.82%). Significance at the 10%, 5% and 1% levels is denoted by *, ** and ***.

Table 4: Key Parameters of the Market-Specific VECMs

	DEM P	DEM Q	GBP P	GBP Q	JPY P	JPY Q	CHF P	CHF Q	FRF P	FRF Q	BEF P	BEF Q	ITL P	ITL Q	NLG P	NLG Q
DEM P	33.9 (26.7,40.6)	18.3 (8.7,28.1)	1.2 (0.0,3.5)	0.8 (0.0,2.7)	4.4 (0.6,10.3)	3.8 (0.3,10.2)	8.6 (4.4,13.1)	1.9 (0.1,4.8)	8.7 (4.0,13.2)	1.0 (0.0,3.4)	7.3 (3.4,11.2)	0.4 (0.0,1.6)	1.6 (0.0,4.3)	0.6 (0.0,2.2)	6.9 (2.8,10.9)	0.4 (0.0,1.7)
DEM Q	30.4 (19.0,38.4)	58.8 (49.7,68.5)	0.7 (0.0,2.5)	1.3 (0.0,4.1)	0.7 (0.0,2.8)	1.6 (0.0,6.3)	0.5 (0.0,1.9)	1.1 (0.0,3.7)	0.8 (0.0,2.9)	0.7 (0.0,2.7)	0.6 (0.0,2.1)	0.5 (0.0,1.9)	0.5 (0.0,2.1)	0.7 (0.0,2.6)	0.6 (0.0,2.4)	0.4 (0.0,1.7)
GBP P	6.7 (1.5,13.2)	6.5 (1.1,13.7)	54.2 (43.1,65.5)	14.3 (5.5,23.7)	1.8 (0.0,6.2)	2.2 (0.0,7.4)	1.9 (0.0,6.0)	0.7 (0.0,2.7)	2.7 (0.0,7.4)	0.7 (0.0,2.8)	2.5 (0.0,6.8)	1.0 (0.0,3.6)	1.5 (0.0,5.2)	1.1 (0.0,3.9)	1.0 (0.0,3.4)	1.1 (0.0,4.0)
GBP Q	1.9 (0.0,5.5)	2.9 (0.1,7.2)	15.5 (6.8,24.2)	59.5 (48.6,71.4)	2.8 (0.1,7.6)	3.8 (0.2,8.9)	0.7 (0.0,2.8)	1.8 (0.0,5.7)	1.1 (0.0,3.8)	0.9 (0.0,3.4)	0.9 (0.0,3.2)	2.3 (0.0,6.1)	1.2 (0.0,4.5)	1.7 (0.0,5.5)	0.7 (0.0,2.5)	2.2 (0.0,7.0)
JPY P	0.7 (0.0,2.6)	0.9 (0.0,2.8)	4.0 (0.1,10.0)	2.1 (0.0,6.0)	54.7 (46.6,62.9)	25.6 (16.8,33.1)	3.2 (0.1,8.4)	0.7 (0.0,2.6)	0.9 (0.0,3.2)	0.5 (0.0,2.0)	0.6 (0.0,2.2)	2.3 (0.0,6.2)	0.8 (0.0,2.9)	1.2 (0.0,4.1)	1.0 (0.0,3.7)	0.7 (0.0,2.8)
JPY Q	2.6 (0.5,4.9)	3.2 (0.8,5.8)	0.8 (0.0,2.7)	3.4 (0.5,7.2)	22.5 (14.3,29.9)	53.2 (44.9,62.6)	0.7 (0.0,2.4)	1.6 (0.0,4.7)	1.9 (0.1,5.0)	1.1 (0.0,3.7)	2.5 (0.3,5.6)	2.2 (0.1,5.6)	0.5 (0.0,2.0)	0.7 (0.0,2.5)	2.6 (0.2,6.2)	0.6 (0.0,2.2)
CHF P	21.7 (14.1,29.0)	12.5 (4.0,22.2)	1.3 (0.0,4.1)	0.6 (0.0,2.1)	8.0 (2.0,15.6)	7.6 (1.9,15.8)	21.4 (14.9,28.6)	3.7 (0.7,7.9)	7.8 (3.3,12.4)	0.6 (0.0,2.1)	5.3 (1.5,9.4)	0.5 (0.0,1.8)	2.7 (0.2,6.1)	1.0 (0.0,3.3)	4.8 (1.1,8.8)	0.5 (0.0,1.9)
CHF Q	11.0 (3.3,19.1)	14.4 (5.1,24.5)	0.8 (0.0,2.9)	2.4 (0.0,7.5)	3.0 (0.0,9.0)	4.0 (0.1,11.8)	6.3 (1.0,13.2)	49.4 (35.9,62.9)	0.9 (0.0,3.1)	1.2 (0.0,4.5)	0.8 (0.0,2.9)	0.7 (0.0,2.7)	0.9 (0.0,3.4)	1.2 (0.0,4.4)	0.9 (0.0,3.1)	2.2 (0.0,6.9)
FRF P	23.6 (17.8,29.5)	14.6 (7.2,22.8)	1.8 (0.0,5.0)	1.2 (0.0,3.6)	3.3 (0.3,8.1)	3.5 (0.5,8.5)	8.3 (4.3,12.4)	1.5 (0.1,3.7)	18.9 (14.3,23.7)	2.5 (0.2,5.6)	7.1 (3.3,10.8)	0.5 (0.0,2.0)	3.5 (0.5,7.2)	1.2 (0.0,3.4)	7.7 (3.5,11.8)	0.8 (0.0,2.6)
FRF Q	12.4 (4.3,20.2)	15.6 (7.4,24.4)	0.6 (0.0,2.5)	1.6 (0.0,5.4)	2.8 (0.1,8.0)	4.0 (0.1,11.1)	0.8 (0.0,2.8)	0.9 (0.0,3.4)	5.7 (0.6,11.3)	46.6 (32.3,62.2)	1.5 (0.0,4.8)	0.8 (0.0,2.9)	1.5 (0.0,5.2)	1.0 (0.0,3.6)	3.2 (0.1,7.6)	1.1 (0.0,4.4)
BEF P	26.4 (18.8,33.4)	20.6 (11.6,29.8)	1.9 (0.1,5.1)	1.2 (0.0,3.8)	4.8 (0.7,10.9)	6.6 (1.4,14.3)	6.0 (2.5,9.9)	2.2 (0.2,5.3)	6.6 (2.3,11.1)	1.0 (0.0,3.4)	12.4 (7.8,17.1)	0.7 (0.0,2.5)	3.1 (0.3,7.0)	1.0 (0.0,3.2)	5.1 (1.3,9.2)	0.5 (0.0,1.8)
BEF Q	0.9 (0.0,3.4)	1.4 (0.0,5.2)	1.3 (0.0,4.5)	2.2 (0.0,6.1)	5.0 (0.3,11.1)	6.7 (1.4,13.2)	1.7 (0.0,5.5)	0.9 (0.0,3.5)	2.8 (0.1,6.9)	1.4 (0.0,4.7)	5.2 (0.6,10.8)	65.9 (51.9,79.9)	0.7 (0.0,2.6)	1.0 (0.0,3.7)	1.9 (0.0,5.5)	0.9 (0.0,3.4)
ITL P	9.2 (2.5,16.4)	7.3 (1.3,15.2)	1.4 (0.0,4.7)	0.6 (0.0,2.3)	3.2 (0.1,8.2)	5.1 (0.6,12.0)	4.6 (1.1,8.7)	1.3 (0.0,4.0)	5.9 (1.4,10.5)	1.1 (0.0,3.9)	5.6 (0.9,11.0)	0.5 (0.0,1.8)	37.4 (27.3,48.8)	11.3 (4.1,19.7)	4.9 (1.0,9.2)	0.6 (0.0,2.3)
ITL Q	2.9 (0.1,7.4)	2.0 (0.0,6.2)	1.4 (0.0,4.6)	2.2 (0.0,7.0)	4.4 (0.3,10.4)	6.1 (0.9,13.5)	2.0 (0.0,5.7)	1.9 (0.0,5.9)	2.4 (0.0,6.4)	1.0 (0.0,3.8)	1.6 (0.0,5.4)	0.7 (0.0,2.5)	15.1 (6.1,24.1)	51.8 (40.7,62.8)	1.8 (0.0,5.3)	2.6 (0.0,7.7)
NLG P	23.9 (17.7,30.1)	15.1 (7.4,23.6)	1.0 (0.0,3.1)	0.8 (0.0,2.8)	4.3 (0.6,9.7)	3.4 (0.4,8.6)	6.8 (3.4,10.4)	2.1 (0.2,4.9)	9.5 (5.1,13.6)	2.1 (0.1,5.1)	7.1 (3.5,10.7)	0.4 (0.0,1.6)	3.5 (0.7,6.8)	1.1 (0.0,3.2)	17.6 (13.0,22.4)	1.2 (0.0,3.5)
NLG Q	3.7 (0.1,9.6)	3.5 (0.1,9.2)	1.1 (0.0,3.8)	2.0 (0.0,6.6)	1.5 (0.0,5.8)	3.4 (0.0,9.7)	1.3 (0.0,4.3)	2.1 (0.0,6.9)	2.9 (0.0,8.0)	1.2 (0.0,4.7)	1.3 (0.0,4.4)	1.2 (0.0,4.4)	1.0 (0.0,3.8)	2.9 (0.0,8.2)	4.3 (0.1,10.6)	66.7 (51.8,80.1)

NOTES: The connectedness matrix is computed using normalised generalised forecast error variance decompositions following Diebold and Yilmaz (2014). Each row sums to 100% by construction. The depth of shading reflects the strength of the associated heatwave/spillover effect, with darker shading indicating a stronger effect. Values reported are empirical means with the corresponding 90% empirical confidence interval shown in rounded parentheses below. We employ a non-parametric sieve bootstrap based on 10,000 stable iterations. For more information about the sieve bootstrap procedure, see Dees et al. (2007), especially Supplement A.

Table 5: Bootstrapped Connectedness among Variables, One-Day Ahead

	DEM P	DEM Q	GBP P	GBP Q	JPY P	JPY Q	CHF P	CHF Q	FRF P	FRF Q	BEF P	BEF Q	ITL P	ITL Q	NLG P	NLG Q
DEM P	17.4 (9.5,29.3)	12.9 (3.1,32.8)	1.4 (0.1,3.8)	1.8 (0.1,6.3)	11.7 (1.8,25.3)	17.9 (2.5,33.6)	7.7 (3.3,13.4)	2.1 (0.3,6.2)	7.5 (3.3,12.4)	1.0 (0.1,3.0)	7.3 (2.9,12.4)	1.0 (0.1,3.0)	1.7 (0.3,3.9)	0.9 (0.1,2.7)	6.8 (2.6,11.9)	0.9 (0.1,2.9)
DEM Q	12.5 (4.5,26.8)	29.2 (7.1,57.0)	1.2 (0.1,3.7)	1.7 (0.1,5.9)	3.1 (0.2,10.9)	6.5 (0.2,20.0)	9.9 (1.4,17.7)	1.1 (0.1,3.6)	9.2 (1.1,16.3)	0.9 (0.0,3.1)	10.2 (1.2,17.1)	1.1 (0.0,3.6)	2.1 (0.2,5.3)	1.0 (0.0,3.2)	9.3 (1.0,16.3)	1.0 (0.0,3.3)
GBP P	6.3 (1.9,12.2)	5.6 (1.2,13.1)	32.1 (16.9,47.3)	23.6 (7.7,38.9)	2.8 (0.4,7.6)	4.7 (0.5,13.8)	4.8 (1.0,10.1)	1.2 (0.1,3.8)	4.1 (0.8,8.9)	1.0 (0.1,3.2)	4.1 (0.8,9.0)	0.9 (0.1,2.8)	1.8 (0.3,4.4)	1.3 (0.1,4.0)	4.0 (0.5,9.5)	1.9 (0.1,5.8)
GBP Q	1.9 (0.2,5.0)	2.7 (0.3,7.6)	15.3 (3.7,30.0)	57.0 (43.2,70.2)	2.5 (0.1,7.4)	3.1 (0.3,9.3)	2.0 (0.1,6.4)	1.8 (0.0,5.8)	1.6 (0.1,5.0)	0.9 (0.0,3.4)	1.6 (0.1,5.2)	1.9 (0.1,5.5)	1.9 (0.1,6.1)	1.7 (0.0,5.5)	1.6 (0.1,5.5)	2.5 (0.0,7.7)
JPY P	2.9 (0.3,7.3)	2.4 (0.1,7.4)	2.6 (0.2,6.9)	1.7 (0.1,4.9)	37.6 (21.1,52.6)	31.9 (16.4,41.6)	3.0 (0.6,7.4)	0.9 (0.0,3.2)	3.5 (0.4,8.6)	0.9 (0.0,2.8)	4.0 (0.5,9.3)	2.4 (0.2,5.9)	0.9 (0.1,2.9)	0.9 (0.0,3.1)	3.6 (0.4,9.1)	0.9 (0.0,3.0)
JPY Q	5.0 (1.0,9.5)	3.7 (0.5,7.8)	1.1 (0.0,3.6)	2.3 (0.2,5.6)	22.0 (7.5,36.2)	43.6 (33.9,53.5)	2.3 (0.1,6.9)	1.3 (0.0,4.1)	4.0 (0.5,8.9)	1.2 (0.0,3.9)	4.8 (0.9,9.9)	2.2 (0.1,5.5)	0.7 (0.0,2.6)	0.7 (0.0,2.4)	4.6 (0.5,10.4)	0.6 (0.0,2.2)
CHF P	11.0 (5.5,19.0)	9.7 (1.7,26.8)	1.4 (0.1,3.9)	1.6 (0.1,5.8)	12.3 (2.5,25.6)	19.6 (4.6,34.2)	11.9 (6.4,18.0)	2.7 (0.6,7.5)	8.0 (3.3,13.3)	1.0 (0.1,3.1)	8.1 (2.8,14.0)	1.2 (0.1,3.4)	2.1 (0.5,4.5)	1.0 (0.1,2.9)	7.5 (2.5,13.4)	1.0 (0.1,3.0)
CHF Q	6.6 (2.3,12.6)	6.2 (0.9,18.9)	1.4 (0.1,4.4)	2.4 (0.1,7.8)	4.5 (0.6,13.3)	7.6 (0.6,19.9)	10.6 (3.8,19.1)	24.3 (6.8,48.7)	9.8 (2.4,16.6)	1.8 (0.1,5.7)	9.3 (2.2,15.7)	1.0 (0.1,3.4)	2.4 (0.4,6.0)	1.0 (0.1,3.1)	8.2 (1.7,14.5)	2.8 (0.2,7.6)
FRF P	14.3 (7.7,23.3)	12.7 (3.3,29.0)	2.0 (0.2,5.5)	2.6 (0.1,8.5)	10.0 (1.4,22.5)	15.1 (1.8,31.4)	7.4 (3.3,13.1)	2.0 (0.2,6.1)	11.7 (6.1,18.4)	3.5 (0.4,9.6)	6.3 (2.5,11.6)	0.8 (0.1,2.6)	2.7 (0.6,5.8)	1.6 (0.1,4.6)	6.2 (2.6,11.0)	1.0 (0.1,3.0)
FRF Q	8.6 (2.6,18.9)	12.9 (2.5,31.2)	1.3 (0.1,4.6)	2.9 (0.1,9.5)	9.3 (0.8,23.7)	15.0 (1.1,34.1)	4.1 (0.6,11.3)	1.4 (0.1,5.0)	6.2 (1.6,13.7)	24.9 (8.0,46.3)	4.1 (0.5,11.2)	0.9 (0.0,3.1)	1.7 (0.2,4.8)	1.5 (0.1,4.9)	3.9 (0.8,9.5)	1.3 (0.1,4.5)
BEF P	12.6 (6.7,20.3)	12.2 (2.7,30.9)	1.7 (0.2,4.6)	2.6 (0.1,8.6)	8.7 (1.4,21.0)	16.3 (3.0,31.3)	8.9 (3.2,15.7)	2.1 (0.2,6.3)	8.7 (3.4,14.3)	1.1 (0.1,3.3)	10.2 (5.1,15.2)	1.4 (0.1,4.7)	2.6 (0.6,5.2)	1.0 (0.1,3.0)	8.7 (2.9,14.7)	1.3 (0.1,4.0)
BEF Q	3.6 (0.5,8.2)	1.6 (0.1,5.1)	1.4 (0.1,4.1)	1.3 (0.1,4.0)	5.4 (0.7,12.4)	9.3 (2.7,17.1)	4.9 (0.8,10.0)	0.8 (0.0,2.7)	7.0 (1.6,12.2)	1.4 (0.1,4.1)	11.5 (2.8,19.4)	43.0 (26.1,63.3)	1.4 (0.1,4.1)	0.9 (0.0,3.0)	5.8 (1.0,11.2)	0.7 (0.0,2.3)
ITL P	7.6 (2.8,13.1)	6.2 (1.1,16.4)	1.4 (0.1,4.0)	1.4 (0.1,4.6)	6.9 (1.0,16.1)	12.4 (2.7,23.7)	6.3 (2.2,11.7)	1.6 (0.1,4.9)	6.5 (2.4,11.1)	1.1 (0.1,3.1)	6.7 (2.4,11.6)	0.9 (0.1,2.8)	22.2 (10.8,35.9)	11.4 (3.1,23.5)	6.0 (2.1,10.7)	1.3 (0.1,3.9)
ITL Q	3.7 (0.7,8.5)	2.1 (0.2,6.3)	1.3 (0.1,4.3)	3.3 (0.1,10.0)	6.2 (0.7,14.9)	10.1 (1.6,20.4)	3.5 (0.7,8.5)	1.8 (0.1,5.9)	3.5 (0.7,8.4)	1.0 (0.0,3.3)	4.3 (0.6,10.5)	0.9 (0.0,2.9)	8.7 (2.3,20.7)	42.2 (22.3,59.6)	3.6 (0.6,8.9)	3.8 (0.2,9.9)
NLG P	14.9 (8.0,23.7)	12.7 (3.5,28.1)	1.5 (0.1,4.2)	1.9 (0.1,6.6)	11.0 (1.9,24.0)	13.7 (1.3,29.7)	6.5 (2.8,11.9)	2.4 (0.3,7.0)	7.1 (3.3,11.5)	1.8 (0.2,4.4)	5.7 (2.4,9.9)	1.0 (0.1,3.0)	2.8 (0.7,6.0)	1.4 (0.1,4.1)	12.4 (5.9,20.5)	3.3 (0.2,9.4)
NLG Q	7.2 (1.0,16.9)	8.2 (0.9,21.8)	1.1 (0.1,3.4)	1.6 (0.1,4.7)	7.5 (0.5,19.6)	11.8 (0.8,29.0)	2.4 (0.3,6.2)	1.6 (0.2,4.7)	4.2 (0.5,10.3)	1.2 (0.1,3.6)	2.4 (0.2,6.5)	1.2 (0.0,4.0)	1.7 (0.1,4.9)	1.5 (0.2,4.4)	8.7 (1.0,18.1)	37.6 (19.2,58.4)

NOTES: The connectedness matrix is computed using normalised generalised forecast error variance decompositions following Diebold and Yilmaz (2014). Each row sums to 100% by construction. The depth of shading reflects the strength of the associated heatwave/spillover effect, with darker shading indicating a stronger effect. Values reported are empirical means with the corresponding 90% empirical confidence interval shown in rounded parentheses below. We employ a non-parametric sieve bootstrap based on 10,000 stable iterations. For more information about the sieve bootstrap procedure, see Dees et al. (2007), especially Supplement A.

Table 6: Bootstrapped Connectedness among Variables, Five-Days Ahead

	One Day Ahead		Five Days Ahead	
Market	Leader	Magnitude	Leader	Magnitude
DEM	–	–	–	–
GBP	–	–	–	–
JPY	–	–	–	–
CHF	DEM	9.10	JPY	17.29
FRF	DEM	16.83	DEM	11.81
BEF	DEM	33.89	JPY	13.31
ITL	–	–	–	–
NLG	DEM	20.21	DEM	11.93

NOTES: For the i-th market, magnitude is defined as $L^{(h)}_{i\leftarrow j} = \left(\phi^{(h)}_{P_i\leftarrow P_j} + \phi^{(h)}_{P_i\leftarrow Q_j}\right) - \left\{\phi^{(h)}_{P_i\leftarrow P_i} + \phi^{(h)}_{P_i\leftarrow Q_i}\right\}$ for $j = 1,\ldots,K,\ j \neq i$. The terms in rounded parentheses measure the spillover from market j to the price in market i while the terms in braces measure the effect of local information (price and order flow in market i) on the price in market i. Hence, if $L^{(h)}_{i\leftarrow j} > 0$ then there is a leader-recipient relationship between markets i and j, where j is the leader. The market identified as the leader in each case is that market j for which $L^{(h)}_{i\leftarrow j}$ is maximised. Where no leader is identified then $L^{(h)}_{i\leftarrow j} \leq 0\ \forall j \in (1,\ldots,K),\ j \neq i$.

Table 7: Leader-Recipient Relationships

Shock Origin	Size of Shock	Safe Haven	5 Day Response
DEM	110.79	JPY	-31.17
GBP	44.15	JPY	-18.94
JPY	94.62	GBP	-10.39
CHF	41.53	JPY	-11.99
FRF	18.77	JPY	-16.93
BEF	7.15	GBP	-7.37
ITL	12.90	DEM	-11.72
NLG	5.57	CHF	-17.30

NOTES: In each case, the safe haven market is that market which experiences the strongest cumulative negative order flow in the five trading days following the initial adverse shock to the named originating market. The five day cumulative responses are recorded in the final column and the size of the original shock is reported in the second column. The unit of measurement is net orders.

Table 8: Safe Haven Currencies by Market

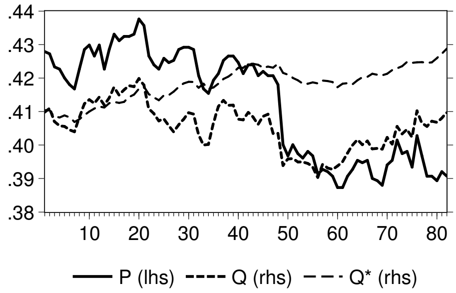

(a) DEM

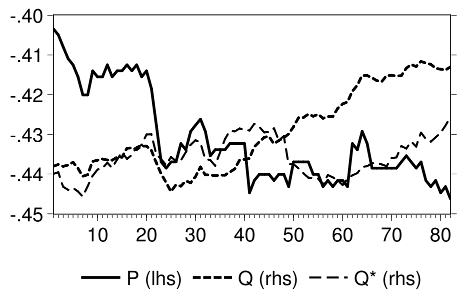

(b) GBP

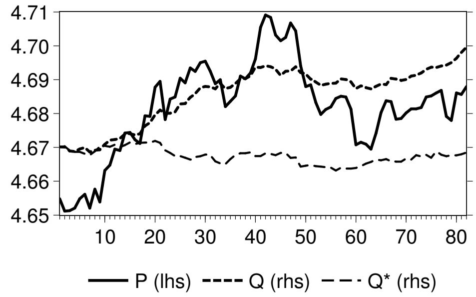

(c) JPY

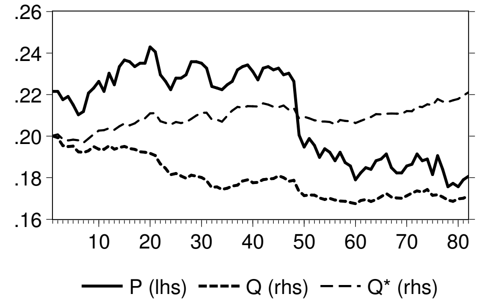

(d) CHF

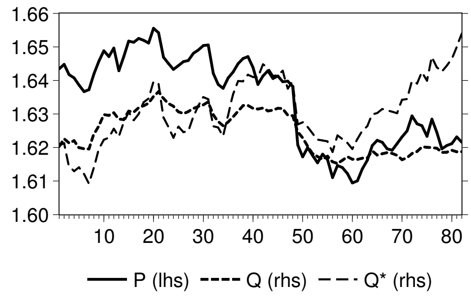

(e) FRF

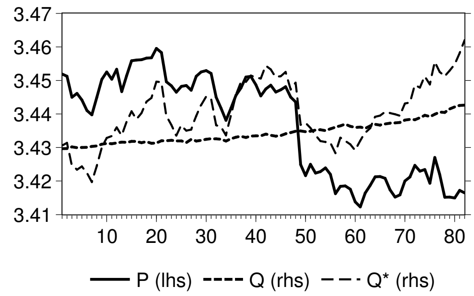

(f) BEF

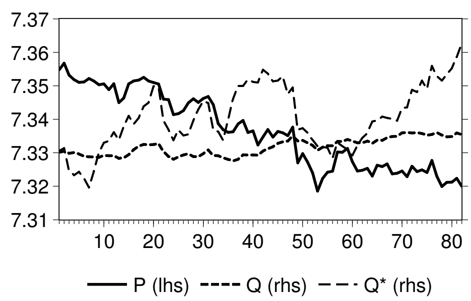

(g) ITL

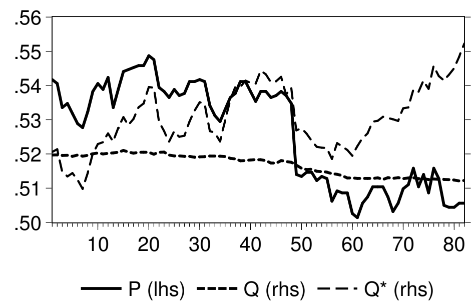

(h) NLG

NOTES: The labels on the horizontal time axis represent trading days 1 to 82. P_{it} denotes the natural logarithm of the exchange rate. Q_{it} denotes the cumulative local order flow in the named market, measured in thousands of trades. Q^*_{it} denotes the foreign cumulative order flow from the perspective of the named market, measured in thousands of trades.

Figure 1: Price and Cumulative Order Flow Data by Market

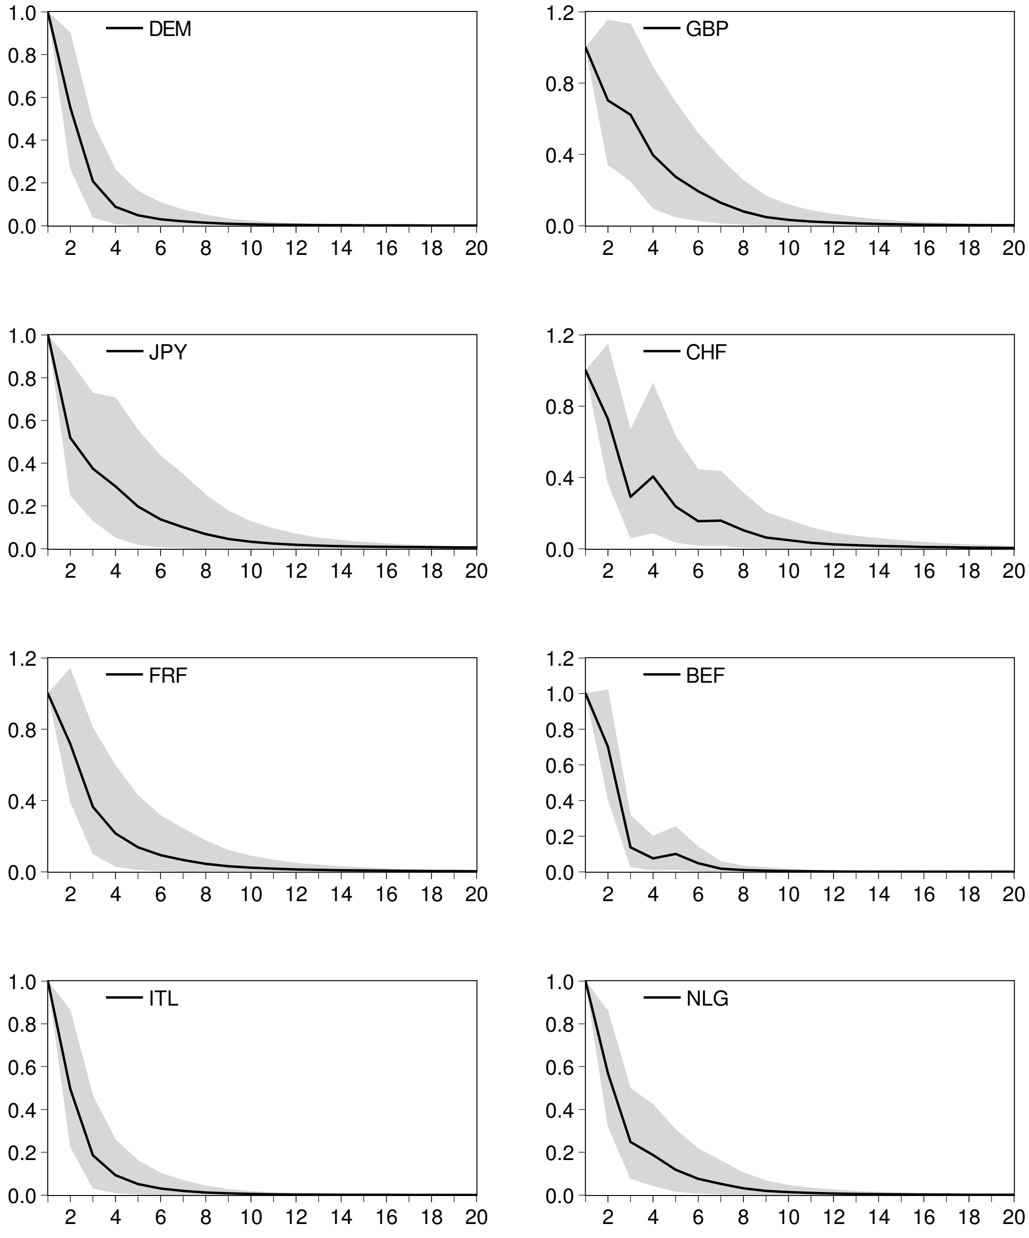

Notes: The horizontal axis records the horizon in trading days while the vertical axis records the response of the cointegrating vector to a systemwide shock. The response is normalised to unity on impact. In a stable and well-specified model, the persistence profiles will die away to zero smoothly and rapidly. See Pesaran and Shin (1996) for further details. The 90% bootstrap interval is shaded in each case.

Figure 2: Persistence Profiles of the Cointegrating Vectors

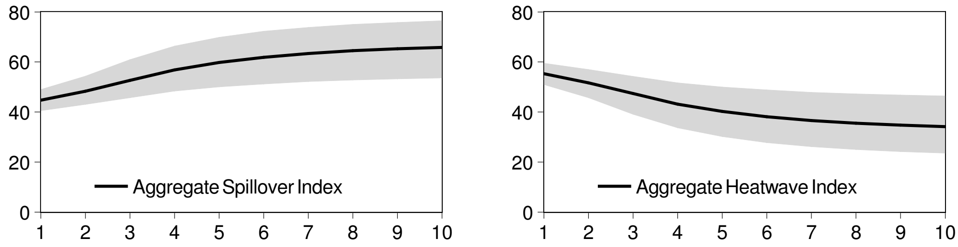

Notes: The horizontal axis records the horizon in trading days while the vertical axis is measured in percentage points. The aggregate heatwave and spillover indices, $H^{(h)}$ and $S^{(h)}$ respectively, are computed following equation (2.14) for horizons one to ten. Note that $H^{(h)} + S^{(h)} = 100 \ \forall h$. The 90% bootstrap interval is shaded in each case.

Figure 3: Aggregate Spillover and Heatwave Indices

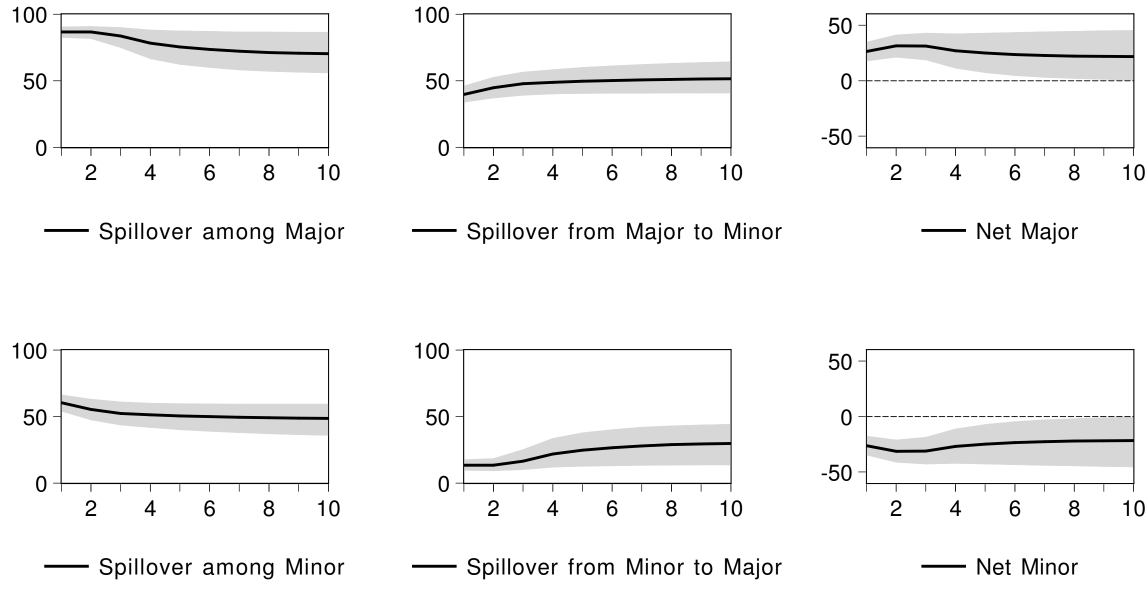

Notes: The horizontal axis records the horizon in trading days while the vertical axis is measured in percentage points. The major markets are DEM, GBP, JPY and CHF while the minor markets are FRF, BEF, ITL and NLG. The decomposed spillover indices reported here are computed following the method described in equation (2.15). In each case, the net connectedness is computed following equation (2.16). Note that the values of *net major* are equal in magnitude but of opposite sign to the values of *net minor* by construction. The 90% bootstrap interval is shaded in each case.

Figure 4: Spillover and Heatwave Indices for Major and Minor Markets

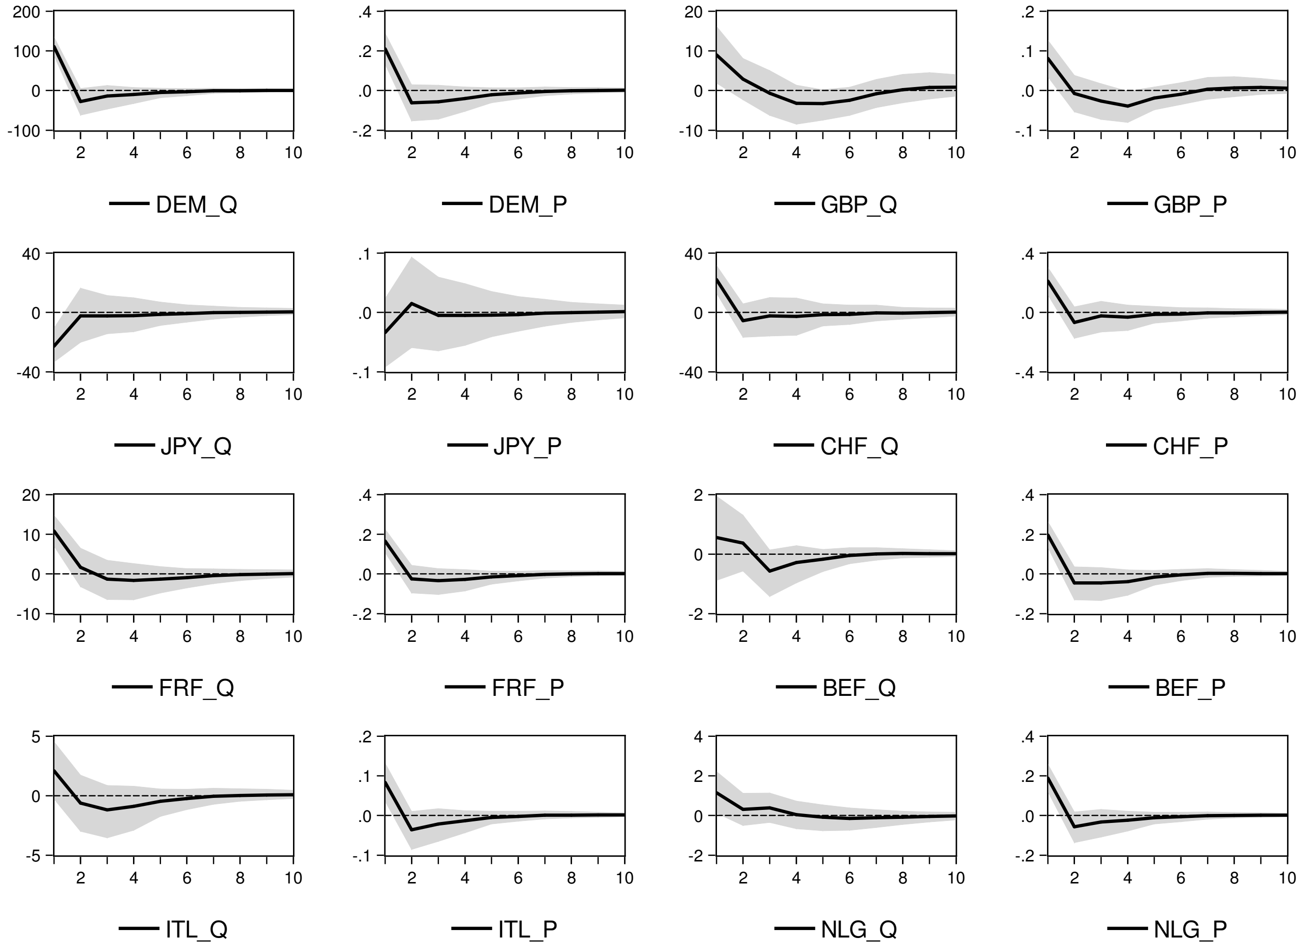

NOTES: the horizontal axis shows the horizon in trading days while the vertical axis measures the reponse of the named variable to the Deutsche Mark selling pressure in the DEM market. Price responses are measured in percent while the order flow responses are measured in trades. The 90% bootstrap interval is shaded in each case.

Figure 5: GIRFs with respect to Deutche Mark Selling Pressure

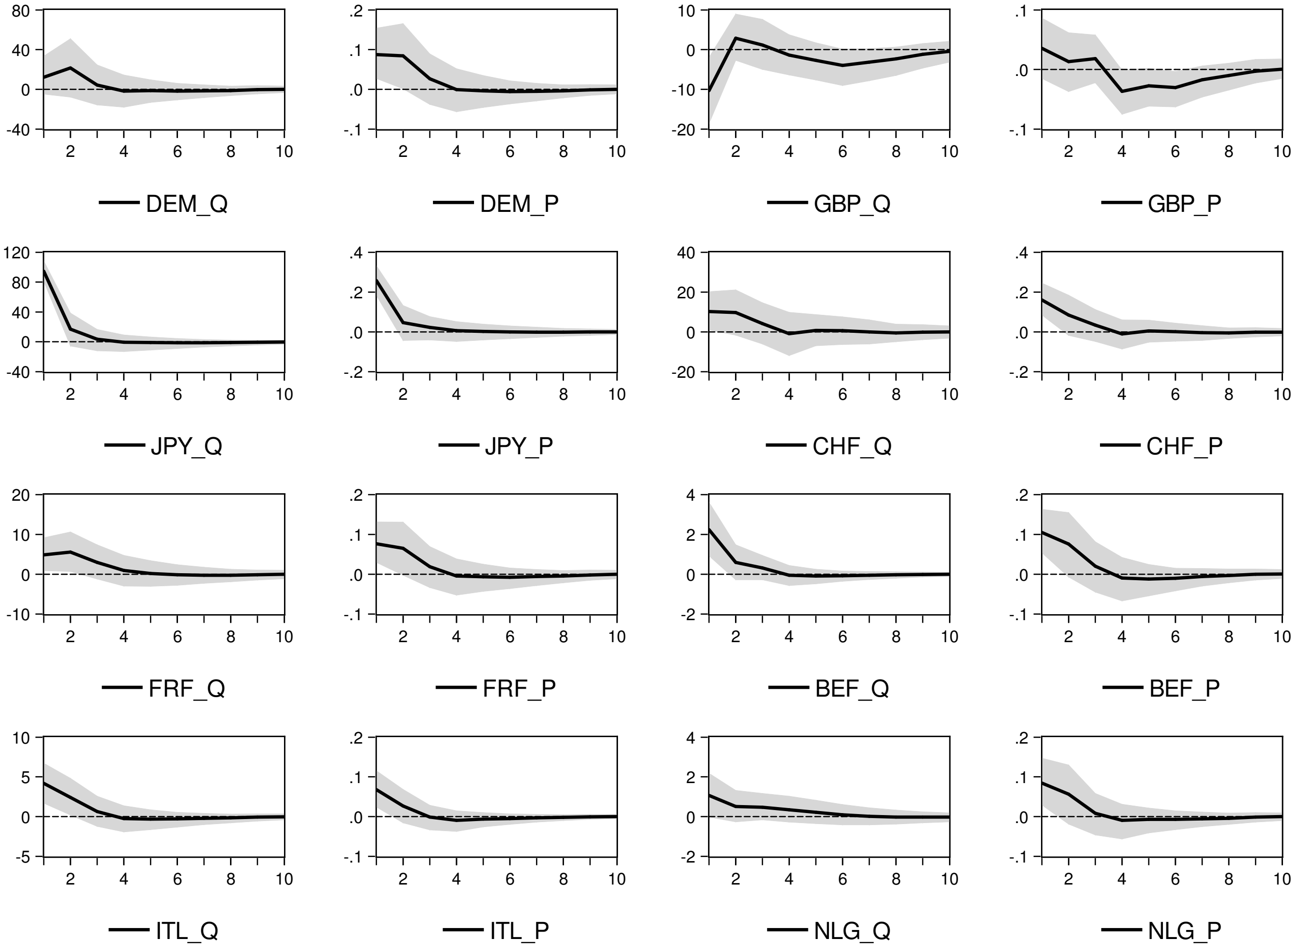

NOTES: the horizontal axis shows the horizon in trading days while the vertical axis measures the reponse of the named variable to the Yen selling pressure in the JPY market. Price responses are measured in percent while the order flow responses are measured in trades. The 90% bootstrap interval is shaded in each case.

Figure 6: GIRFs with respect to Yen Selling Pressure