



UNIVERSITÀ
DEGLI STUDI
FIRENZE

DISEI

DIPARTIMENTO DI SCIENZE
PER L'ECONOMIA E L'IMPRESA

WORKING PAPERS - ECONOMICS

Structural Interdependence of Price and Demand in a Model of the Foreign Exchange Market with Heterogeneous Speculators: Evidence from High-frequency Data

LEONARDO BARGIGLI, GIULIO CIFARELLI

WORKING PAPER N. 04/2020

*DISEI, Università degli Studi di Firenze
Via delle Pandette 9, 50127 Firenze (Italia) www.disei.unifi.it*

The findings, interpretations, and conclusions expressed in the working paper series are those of the authors alone. They do not represent the view of Dipartimento di Scienze per l'Economia e l'Impresa

Structural Interdependence of Price and Demand in a Model of the Foreign Exchange Market with Heterogeneous Speculators: Evidence from High-frequency Data

L. Bargigli¹ and G. Cifarelli²

^{1,2} Università di Firenze

March 20, 2020

Abstract

We assume that the variations of the exchange rate depend on the current net demand of the base currency as a consequence of market making, and that the current net demand of the base currency depends on current and past variations of the exchange rate as a consequence of how future price expectations are formed by bounded rational agents. We achieve identification supposing that the structural shocks of price variations and demand follow a GARCH process. Using high-frequency transaction data of the EUR/USD market in 2016, we show that the simultaneous effects of price on demand and viceversa are both significant and positive. Our estimates suggest that one important source of heterogeneity in demand might be missing from our model, since the structural errors are negatively correlated.

Keywords: Asset pricing model, heterogeneous beliefs, market making, foreign exchange market, SVAR-GARCH, high frequency data.

JEL codes: G12, D84, F31, C32, C55

1 Introduction

The market micro-structure approach provides powerful insights into the working of financial markets (Vitale, 2007; King *et al.*, 2013). This approach has gained popularity thanks to the availability of high frequency transaction data. Another popular stream of literature points instead to the possibility that market participants might deviate from the paradigm of full rationality. According to this family of models with heterogeneous agents (HAM), deviations of the price of assets from their fundamental value arise because the heterogeneous beliefs of market participants are myopic and evolve endogenously (Brock and Hommes, 1998). Although these two lines of research share the ambition to provide a more faithful representation of financial markets than the standard full-rationality-with-continuous-market-clearing setting, there are little contact points between these two approaches. In order to overcome this limitation, in this paper we present a simple model which incorporates two basic assumptions. The first one, coming from the market microstructure literature, is the dependency of price variations on the contemporaneous net order flow (Evans and Lyons, 2002). The second one, coming from the HAM literature, is the dependency of current demand on price variations.

The rationale for the first assumption relies on the behavior of a market maker who might find optimal to increase / decrease her bid and/or ask price following a positive / negative net order flow (Madhavan and Smidt, 1993). The rationale of the second assumption relies on the myopic nature of speculators who form their expectation of tomorrow's returns based on the returns observed to date. This general mechanism of expectation formation encompasses different typologies of speculators which are popular in the literature such as chartists, fundamentalists, contrarians etc. One shortcoming is that usually this approach assumes a delay, with which agents react to price signals, which is hard to justify. Indeed, bounded rationality is adopted as a good proxy for the actual behavior of market participants, but it is not very realistic to assume that the latter do not react as quickly as possible to price variations. The actual delay of real markets is so small that an instantaneous reaction, which is the assumption we adopt in this paper, appears to be much more realistic.

We further assume that net market demand is autocorrelated because of microstructural effects such as order-splitting or back-stop orders. Although the absence of significant linear correlations of price increments has been widely documented and is viewed as a proof that the market is efficient, autocorrelation cannot be ruled out at short time scales, which are those the market needs to react to new information. Moreover, there is evidence of

negative first-order autocorrelation of prices in high-frequency FX markets, which is explained as a microstructural effect (Zhou, 1996; Cont, 2001). Thus we also allow for the autocorrelation of price increments.

In order to take into account all these hypotheses, we need to refer to a bivariate unrestricted SVAR model in price variations Δp_t and net demand z_t . It's simple to show that this model must have time varying coefficients if the market participants are heterogeneous. Unfortunately, a model of this sort cannot be estimated with the currently available methods. However we can estimate the coefficients of a standard VAR model and then perform a stability analysis to verify the hypothesis of heterogeneity of market participants. In order to make both simultaneous coefficients identifiable, we adopt an approach based on heteroskedasticity (Rigobon, 2003; Sentana and Fiorentini, 2001). Equipped with these assumptions, we employ a dataset of high-frequency transactions on the EUR/USD market in 2016 provided by Nex data, in order to answer the two following questions: what is the simultaneous impact of price and demand on each other? Does heterogeneity affects the way market demand reacts to price variations and viceversa? The remaining of the paper is organized as follows. In Sec. 2 we present the theoretical model sketched above. In Sec. 3 we describe the dataset we employ for our empirical analysis. In Sec. 4 and in the Appendix we describe our identification and estimation strategy. In Sec. 5 we present the main results of the analysis. Finally, Sec. 6 concludes.

2 The model

We assume that the market maker is a profit maximizing monopolist who trades a zero yielding asset with a large number of different types of speculators whose participation to the market evolves endogenously. We assume that the market maker knows the optimal demand of each type of speculator and that she employs this information when she solves her optimization process. Then we suppose that the market is liquid enough to allow the market maker to adjust in advance her inventory at the current price, in order to match the projected orders of speculators. After this adjustment she announces the optimal price, taking into account a quadratic cost of inventory maintenance. Once the new price is revealed to speculators, the latter trade according to their optimal demand in such a way that, at the end of the period, the net variation of the inventory position of the market maker is zero. Finally, the market price is updated to the optimal price of the market maker.

The timeline of events in each period of the model is pictured in Fig. 1. We remark that the assumption we make on market liquidity is analogous to




Figure 1: Timeline of events occurring within a single period of the model.

the one made by Evans and Lyons (2002), who suppose that FX dealers trade on the wholesale market *before* the price adjustment on the retail market. It is also in line with evidence from the FX market and in particular with the practice of “hot potato” trading which allows FX dealers to profit from retail trading (King *et al.*, 2013). The current profit of the market maker is specified as follows:

$$\Pi_{d,t} = (p_t - p_{t-1}) z_t - \frac{\omega}{2} z_t^2 \quad (1)$$

where z_t is the net market demand and $\frac{\omega}{2} > 0$ measures the impact of the quadratic inventory cost z_t^2 on $\Pi_{d,t}$. We remark that, according to our hypotheses, the inventory cost is linked to z_t since at the end of each period the inventory is brought down to zero. In practice, it is the cost of holding an amount of risky asset equal to $z_t > 0$ until it is resold to the speculators¹.

The market maker maximizes $\Pi_{d,t}$ with respect to p_t . Taking into account the effect of the optimization variable on z_t the FOC reads:

$$z_t + (p_t - p_{t-1}) \frac{d}{dp_t} z_t - \omega z_t \frac{d}{dp_t} z_t = 0 \quad (2)$$

Setting $\gamma_{0,t} \equiv \frac{d}{dp_t} z_t$, we obtain

$$\Delta p_t = \left(\omega - \frac{1}{\gamma_{0,t}} \right) z_t \quad (3)$$

We follow the literature on asset pricing with heterogeneous speculators mentioned in Sec. 1 and exclude, for simplicity, a more complex forward looking setting, e.g. like the one of Madhavan and Smidt (1993). In the HAM setting speculators are myopic mean-variance maximizers, since their future wealth is uncertain. In practice, speculators of each type i maximize their expected risk adjusted profit in the next period:

¹ In the case in which $z_t < 0$, we might think that the risky asset must be exchanged against some other asset and that holding this other asset is costly too.

$$\max_{z_{i,t}} \left\{ E_{it} [\Pi_{i,t+1}] - \frac{1}{2D} V_{it} [\Pi_{i,t+1}] \right\} \quad (4)$$

where D^{-1} is a risk aversion parameter linked to the variance of future profit $V_{it} [\Pi_{i,t+1}]$. Indeed the profit of speculators at $t + 1$ is determined by their net demand at t and they ignore the future market price when taking their decision. The current profit of speculators of type i is written as follows:

$$\Pi_{i,t} = (p_t - p_{t-1}) z_{i,t-1} \quad (5)$$

where the right hand side of the equation above represents the profit obtained buying the amount $z_{i,t-1}$ at the price p_{t-1} and reselling the same amount at the price p_t .

Then we have:

$$E_{it} [\Pi_{i,t+1}] = (E_{it} [p_{t+1}] - p_t) z_{i,t} \quad (6)$$

$$V_{it} [\Pi_{i,t+1}] = V_{it} [p_{t+1}] z_{i,t}^2 \quad (7)$$

Taking into account Eqs. (6) and (7), supposing for simplicity that $V_{it} [p_{t+1}] = \sigma^2$ is constant across time and investor's types and finally letting this term be absorbed by D^{-1} , we may rewrite the objective as follows:

$$\max_{z_{i,t}} \left\{ E_{i,t} [\Delta p_{t+1}] z_{i,t} - \frac{z_{i,t}^2}{2D} \right\} \quad (8)$$

where $E_{i,t} [\Delta p_{t+1}]$ stands for the type i 's expectation of Δp_{t+1} .

We solve the FOC for $z_{i,t}$ to obtain the optimal demand of a generic speculator of type i :

$$z_{i,t} = D E_{i,t} [\Delta p_{t+1}] \quad (9)$$

Then the total market demand is

$$z_t = D \sum_{i=0}^{S-1} E_{i,t} [\Delta p_{t+1}] N_{i,t} \quad (10)$$

where $N_{i,t}$ is the number of speculators of type i at t and S is the number of different types of speculators on the market. We introduce the following specification for $E_{i,t} [\Delta p_{t+1}]$:

$$E_{i,t} [\Delta p_{t+1}] = \sum_{k=0}^{K-1} g_{i,k} \Delta p_{t-k} \quad (11)$$

Substituting (11) in (10) we obtain

$$z_t = D \sum_{h=0}^{S-1} N_{i,t} \sum_{k=0}^{K-1} g_{i,k} \Delta p_{t-k} \quad (12)$$

Changing the order of summation we obtain

$$z_t = \sum_{k=0}^{K-1} \Delta p_{t-k} \sum_{h=0}^{S-1} DN_{i,t} g_{i,k} \quad (13)$$

and defining

$$\gamma_{k,t} \equiv \sum_{i=0}^{S-1} DN_{i,t} g_{i,k} \quad (14)$$

we obtain that market demand is a time varying function of current and past prices:

$$z_t = \sum_{k=0}^{K-1} \gamma_{k,t} \Delta p_{t-k} \quad (15)$$

We see that, according to (15), $\frac{\partial z_t}{\partial p_t} = \gamma_{0,t}$ which is consistent with our previous definition in Eq. (3). Thus we end up with a system of 2 simultaneous equations in $(\Delta p_t, z_t)$ plus lagged values of price increments:

$$\begin{cases} \Delta p_t = \left(\omega - \frac{1}{\gamma_{0,t}} \right) z_t + \epsilon_{0,t} \\ z_t = \sum_{k=0}^{K-1} \gamma_{k,t} \Delta p_{t-k} + \epsilon_{1,t} \end{cases} \quad (16)$$

The bivariate random process $\epsilon_t = (\epsilon_{0,t}, \epsilon_{1,t})$ is added to take into account all the exogenous factors which impact price and demand. Thus we expect that this process has more structure than a simple i.i.d. white noise. In particular we allow for time varying conditional correlation and heteroskedasticity by assuming the following:

$$\epsilon_t = H_t^{\frac{1}{2}} v_t \quad (17)$$

where v_t is a bivariate i.i.d. vector with scaled first and second moments and $H_t^{\frac{1}{2}}$ is a positive definite matrix such that H_t is the conditional covariance matrix of ϵ_t .

The system (16) is equivalent to a restricted SVAR model with time varying coefficients:

$$\begin{aligned}
\begin{bmatrix} \Delta p_t & z_t \end{bmatrix} \begin{bmatrix} 1 & -\gamma_{0,t} \\ -\left(\omega - \frac{1}{\gamma_{0,t}}\right) & 1 \end{bmatrix} &= \\
= \begin{bmatrix} \Delta p_{t-1} & \Delta p_{t-2} & \dots & \Delta p_{t-K} \end{bmatrix} \begin{bmatrix} 0 & \gamma_{1,t} \\ 0 & \gamma_{2,t} \\ \dots & \dots \\ 0 & \gamma_{K-1,t} \end{bmatrix} + \begin{bmatrix} \epsilon_{0,t} & \epsilon_{1,t} \end{bmatrix} & \quad (18)
\end{aligned}$$

In the empirical estimation we prefer to lift all the restrictions on the r.h.s. since there is evidence that microstructural effects on the market lead to autocorrelation of both price variations and demand (see Sec. 1). So, setting $y'_t = (\Delta p_t, z_t)$, we end up with the following model:

$$A_{0,t}y_t = A_{1,t}y_{t-1} + \dots + A_{q,t}y_{t-q} + H_t^{\frac{1}{2}}v_t \quad (19)$$

where

$$A_{0,t} = \begin{bmatrix} 1 & -\left(\omega - \frac{1}{\gamma_{0,t}}\right) \\ -\gamma_{0,t} & 1 \end{bmatrix} \quad (20)$$

while all the coefficients in $A_{1,t}, \dots, A_{q,t}$ are unconstrained. Regarding the coefficients in $A_{0,t}$, the only ex ante hypothesis we can make is $\omega - \frac{1}{\gamma_{0,t}} > 0$, since we expect that market makers raise (lower) prices when there is a positive (negative) excess demand. We choose not to impose this sign restriction ex ante as a constraint and indeed, according to the estimation, it is confirmed by the data.

Finally, we observe that the elements of the matrices on the r.h.s. of (19) are a combination of time varying and constant coefficients. The time varying component is given by the coefficients $\gamma_{0,t}, \gamma_{1,t}, \dots, \gamma_{K-1,t}$ of the lagged values of Δp_t in the demand equation, which reflect the heterogeneous reaction of different types of speculators to price variations. The constant component is given by the coefficients of the lagged values of Δp_t and z_t in the price equation, and by the coefficients of the lagged values of z_t in the demand equation which, as we just explained, are introduced as microstructural effects whose magnitude and direction should not change over time.

3 Data description

We analyze in this paper tick-by-tick transaction data, recorded on the EBS FX Spot trading platform and provided by NEX data. For the purpose of our

analysis, the data are sampled with a 5 minutes frequency. More precisely, the following time series are investigated: the (last) Euro/ Dollar midpoint between bid and offer prices (p_t) expressed in USD cents; the total values in period t of the bid (BSZ_t) and offer (OSZ_t) order flows expressed in Millions of Euros. The sample spans the time interval from 3 January 2016, 17:55 (5.55 pm.), to 30 December 2016, 21.55 (9.55 pm.). It includes relevant events related to the final spasms of the EMU crisis, which interact with the vagaries of the British Brexit referendum campaign. A painstaking synchronization of the time series has reduced the sample length to 36,207 observations per continuous time series. Weekends, holidays and late-evening / night periods are excluded.

Fig. 2 exhibits the EUR/USD exchange rate in levels (p_t) and in 5 minutes first differences expressed (Δp_t), together with the 5 minutes difference between the total values of the offer and bid order flows, which is our measure of net demand: $z_t = OSZ_t - BSZ_t$. The exchange rate p_t is clearly non - stationary, as corroborated by the unit root analysis of Table 1, and is affected by the outcome of the Brexit referendum (June 24) and by the election of Donald Trump as president of the US (November 11), which are marked with a thick vertical line in the graph. The appreciation of the dollar in the wake of the two events clearly stands out, although we see that the price and demand movements which are close in time to these events are by no means the largest in magnitude of the dataset. Indeed, the largest price swings followed the announcement by the ECB of a further expansion of Quantitative Easing on March 10² and the announcement of weak US jobs data on June 3³, which lead to a stark appreciation of the Euro. The largest demand swing instead was triggered by the Euro falling below the critical 1.04 \$ benchmark for the first time in 14 years on December 15⁴.

It should be noted that large demand swings do not necessarily correspond to large price swings and viceversa. The explanation is twofold: on the one hand, most public information is incorporated directly into price and, on the other, the EUR/USD market is liquid enough to absorb large demand imbalances without large price effects. This goes without denying that the two variables influence each other: indeed the Pearson correlation coefficient between Δp_t and z_t is 0.305. The sources of this correlation are twofold. On the one hand, we have the behavioral interaction of market makers and speculators as depicted by the model of Sec. 2. On the other we have a common factor affecting both price and demand, namely the exogenous flow

²<https://www.ft.com/content/02ec97ea-e6d9-11e5-bc31-138df2ae9ee6>

³<https://www.ft.com/content/eb77d7d6-2937-11e6-8ba3-cdd781d02d89>

⁴<https://www.ft.com/content/25fba186-fc9f-3bf2-a020-82efe29f1f7b>

of information which acts simultaneously on both market makers and speculators. The first source is captured by the coefficients $A_{0,t}, A_{1,t}, \dots, A_{q,t}$ of the model (19), the second one is captured by the coefficients of a parametric representation of the process which drives H_t . The purpose of the estimation described the next section is to distinguish and quantify each of the two sources. At the same time we must be aware of a potential factor of confusion, which might affect the impact of the exogenous flow of information on the endogenous variables. This is represented by the fact that we are dealing with a wholesale market whose participants do not necessarily trade at their own initiative. Indeed, according to the practice of “hot potato” trading mentioned in the previous section, dealers replicate on the wholesale market the trades of their customers in order to offload their inventory (King *et al.*, 2013). Thus the trades in our dataset do not necessarily represent the goals and beliefs of those who submit them. Instead they might represent goals and beliefs of different classes of traders, among which liquidity / noise traders might prevail⁵, triggering very different reactions to the external flow of information depending on the circumstances.

The statistical properties of the exchange rate first differences Δp_t , expressed in USD cents, and of net demand $z_t = OSZ_t - BSZ_t$, expressed in Millions of Euros, are analyzed in Table 2. They are stationary and display the main stylized facts of high frequency financial time series, with highly persistent serial correlation and severe volatility clustering. Skewed and leptokurtic, they have non normal distributions, as suggested by the very large values of the corresponding Jarque Bera statistics. In particular, the significance of ARCH effects justifies our parametric approach to model the conditional volatility of our data.

4 Estimation Strategy

A SVAR model such as Eq. (19) cannot be estimated with the methods which are currently available. Primiceri (2005) has proposed a Bayesian approach to estimate a time-varying SVAR, where both the coefficients of the SVAR and the covariance matrix of the shocks are allowed to vary over time. But this comes at the cost of imposing the usual *a priori* identification restrictions on $A_{0,t}$, which are unfortunately unwelcome in our case. Indeed feedback

⁵We recall the following definition of liquidity or noise traders provided by Dow and Gorton (2008): “Noise traders are economic agents who trade in security markets for non-information-based reasons”. One of the possible rational motivations of liquidity trading is insurance / hedging but other explanations have been proposed (Vitale , 2000). See Ramiah *et al.* (2015) for a recent review of the noise-trader literature.

p_t	Δp_t	z_t
-2.099	-32.629**	-29.11**
(c, 33)	(nc, 33)	(nc, 33)

ADF(nc, (n)c): Augmented Dickey Fuller unit root test statistic, with (no) constant term and nth order autoregressive component. **: significant at the 1% level.

Table 1: Unit root analysis

	Δp_t	z_t
Mean	-0.0002	1.0655
Std. dev.	0.0466	37.909
Skewness	0.5633	-0,7674
Kurtosis	37.829	58.576
JB	1831900 [0.00]	4663300 [0.00]
AR1	40.675 [0.00]	445.52 [0.00]
AR5	55.317 [0.00]	624.28 [0.00]
AR20	92.985 [0.00]	851.99 [0.00]
ARCH1	212.67 [0.00]	112.35 [0.00]
ARCH5	937.24 [0.00]	171.19 [0.00]

Notes: probability values in square brackets; JB: Jarque-Bera normality test; AR*k*: Ljung-Box test statistic for *k*-th order serial correlation of the time series; ARCH*k*: Ljung-Box test statistic for *k*-th order serial correlation of the squared time series.

Table 2: Descriptive Statistics

trading (which implies $\gamma_{0,t} \neq 0$) is ubiquitous in foreign exchange market since it is associated with price-contingent trading strategies such as stop-loss and profit-taking orders (Osler, 2003, 2005; Daniélsson and Love, 2006). We cannot realistically suppose that $\gamma_{0,t} = 0$ nor we can suppose that $\omega = 0$, since the latter restriction would imply that market makers are not adverse to inventory risk, contrary to a widespread evidence (King *et al.*, 2013). Finally, we cannot assume that $\omega - \frac{1}{\gamma_{0,t}} = 0$, since this would contradict our basic intuition about the behavior of market maker as discussed at the end of Sec. 2.

At the same time, a bivariate approach appears unavoidable. Evans and Lyons (2002) proposed to estimate the impact of order flows on price increments with a simple OLS approach using the following equation:

$$\Delta p_t = \beta z_t + u_t \quad (21)$$

As underlined by Vitale (2007), this approach is plagued by a simultaneity bias if exchange rate movements cause order flows. In this case, the ordinary least squares (OLS) estimate of the coefficient $\hat{\beta}$ is biased, and the results reported by Evans and Lyons (2002) are misleading. The usual identification strategy to overcome this limitation is to impose restrictions on A_0 , in particular forbidding z_t to depend from Δp_t . But from our arguments above we see that we cannot accept the restriction $\gamma_0 = 0$ in our model.

The most adopted alternative, when a priori restrictions cannot be made, is to find good instruments. But this is impossible to achieve in our case given that market efficiency rules out by principle that good predictors of one-period-ahead returns are available. This is why we choose a different approach, which is based on the observation that, if the variances of the structural shocks change over time, the coefficients of interest can be identified from the reduced-form covariances at different points in time (Rigobon, 2003). Under this approach, in order to obtain identification we need to impose some parametric restrictions on the structure of the covariance matrix, which might take the form of variance regimes (Rigobon, 2003), restricted or unrestricted multivariate ARCH or GARCH processes (Rigobon, 2002; Sentana and Fiorentini, 2001), constant or dynamic conditional correlation processes (Weber, 2010), stochastic volatility (Lewis, 2018)⁶. In particular, Herwartz *et al.* (2019) show by means of a simulation study that a GARCH based Gaussian ML approach provides the best results when the underlying d.g.p. is conditionally heteroskedastic even if the distribution of the structural errors is misspecified⁷. This is the approach we follow in our study.

⁶A detailed review is provided by Kilian and Lütkepohl (2017).

⁷In particular they employ skewed and leptokurtic distributions which reflect the char-

For the sake of simplicity, we restrict the coefficients to be constant⁸, with the idea that the effect of heterogeneity (which is the driver of the time evolution of those coefficients according to our model) can be tested by checking the stability of the coefficients of this restricted estimation. Thus the resulting model is

$$A_0 y_t = A_1 y_{t-1} + \dots + A_q y_{t-q} + \epsilon_t \quad (22)$$

$$A_0 = \begin{bmatrix} 1 & -\left(\omega - \frac{1}{\gamma_0}\right) \\ -\gamma_0 & 1 \end{bmatrix} \quad (23)$$

$$\epsilon_t = H_t^{\frac{1}{2}} v_t \quad (24)$$

$$v_t \sim \mathcal{N}(0, I) \quad (25)$$

$$H_t = \text{diag}(h_{1,t}, h_{2,t}) \quad (26)$$

$$h_t = \omega + \text{diag}(\alpha) \epsilon_{t-1} \odot \epsilon_{t-1} + \text{diag}(\beta) h_{t-1} \quad (27)$$

where ω , α and β are vectors of parameters and \odot stands for the Hadamard product. The original time varying coefficient become constant if we assume that there is no heterogeneity of speculators (i.e. $S = 1$) and that the participation to the market is constant ($N_t = \bar{N}$), because in this case we obtain

$$\gamma_k \equiv D\bar{N}g_{h,k} \quad (28)$$

We can test the hypothesis of no heterogeneity by looking at the reduced form residuals of eqs. (22) as explained in the next section. Since, as we have already pointed out at the end of Sec. 2, the original model contains both time varying and constant coefficients, one of the equations in (22) might pass the stability test, while the other might not. In particular, following the arguments at the end of Sec. 2, we expect that the price equation pass the stability test since its coefficients are linked to microstructural effect that do not change over time, while the demand equation doesn't since its coefficients are impacted by the time varying participation of the heterogeneous speculators to the market.

In a second estimation we relax the assumption of orthogonality of H_t and specify the covariance matrix with a Dynamic Conditional Correlation (DCC) model (Engle, 2002):

$$H_t = D_t R_t D_t \quad (29)$$

acteristics of financial time series.

⁸ Although there are in principle no impediments to assume that the reduced form coefficients vary over time, the time invariance of A_0 is required by all the existing identification schemes based on heteroskedasticity (Lewis, 2018). See also the observations above on Primiceri (2005).

$$D_t = \text{diag}(\sqrt{h_{0,t}}, \sqrt{h_{1,t}}) \quad (30)$$

$$R_t = \text{diag}(Q_t)^{-\frac{1}{2}} Q_t \text{diag}(Q_t)^{-\frac{1}{2}} \quad (31)$$

$$Q_t = (1 - a - b)\bar{Q} + a (D_{t-1}^{-1}\epsilon_{t-1}) (D_{t-1}^{-1}\epsilon_{t-1})' + bQ_{t-1} \quad (32)$$

where a, b are scalars and $h_t = (h_{0,t}, h_{1,t})$ is still given by eq. (27). This extension overcomes the identification restrictions of the approaches based on conditional heteroskedasticity, which require that the structural errors are uncorrelated (Rigobon, 2002, 2003; Ehrmann *et al.*, 2011). As explained below, this is accomplished thanks to the two step estimation procedure of the DCC model, which at the first step treats the structural errors as if they were independent.

The estimation proceeds through the following steps. Firstly, the reduced form parameters are estimated with a VAR model:

$$y_t = A_1' y_{t-1} + \dots + A_q' y_{t-q} + u_t \quad (33)$$

where

$$u_t = A_0^{-1} \epsilon_t \quad (34)$$

In the second step, according to the assumptions of the model, the following log-likelihood is maximized:

$$L(\theta) = -\frac{1}{2} \sum_{t=1}^T (n \log 2\pi + \log |H_t| + u_t' A_0' H_t^{-1} A_0 u_t) \quad (35)$$

where the parameter vector θ contains ω and γ_0 plus the parameters of the GARCH equation (27) which takes the following form:

$$h_t = \omega + \text{diag}(\alpha) (A_0 u_{t-1}) \odot (A_0 u_{t-1}) + \text{diag}(\beta) h_{t-1} \quad (36)$$

Following the terminology of Lütkepohl (2005), the model (33)-(36) is an A-model. This means that the diagonal elements of A_0 are normalized to one while the unconditional variances of the structural shocks are unrestricted. Alternatively, Lanne and Saikkonen (2007) and Lütkepohl and Milunovich (2016) adopt a B-model where the structural errors are uncorrelated with unconditional identity covariance matrix and the elements of A_0 are unrestricted but they must be chosen in such a way that $A_0 A_0' = \Sigma_u$, where Σ_u is the unconditional covariance matrix of the reduced form errors. Since the unconditional volatilities of the structural errors are normalized, the conditional volatilities are parametrized with a GO-GARCH model (Van Der Weide, 2002).

The main difference between the two models is that in a B-model the condition $A_0 A_0' = \Sigma_u$ implies equivalently that the estimated unconditional

variances of the structural errors are equal to the sample unconditional variances of the reduced form errors. In our model instead the unconditional variances of the structural errors are obtained from the GARCH parameters without imposing any additional conditions apart from those of the GARCH model itself. For this reason we consider our model to be less restrictive than the alternative one. Furthermore, given the assumption that the structural errors are uncorrelated, the maximization of the likelihood of our model can be performed consistently equation by equation without imposing the orthogonality conditions required by the B-model (Lanne and Saikkonen, 2007).

The details of the estimation procedure together with a discussion of the identification conditions of the structural parameters of the model are provided in the Appendix. Here we limit ourselves to provide the following general considerations. The identification conditions of Rigobon (2003) easily extend to the GARCH setting of this paper, since adding extra covariance regimes which satisfy, for each pair of regimes, the condition (A.12), simply makes the model overidentified. In particular, the non proportionality of the different volatility regimes is a. s. assured by the existence of ARCH effects across time periods. The same reasoning applies to the full rank condition that Sentana and Fiorentini (2001) require from the matrix containing the full path of conditional variances or to the conditions proposed by Milunovich and Yang (2013) which state that in a GARCH setting no structural shock be degenerate (i.e. $\omega \neq 0$ in eq. (27)) and that at most one structural shock i be homoskedastic (i.e. $\alpha_i = 0$ for at most one shock in (27)).

In principle it is not straightforward to check some of the conditions mentioned above as they are expressed in terms of population quantities, while we observe only their estimators⁹. To overcome this problem Lanne and Saikkonen (2007) provide three different test statistics which are based on the autocovariance of the structural errors of a B-model (see also Lütkepohl and Milunovich (2016) for an application). Lewis (2018) provides instead an identification criterion which is based on the application of the Cragg and Donald (1996) statistics and the corresponding rank test to the unconditional autocovariance matrix of the reduced form errors. This approach is more general, since it is consistent with a variety of different models (regime based, markov switching volatility, GARCH, stochastic volatility), but is rather demanding since it requires to estimate the standard errors of the estimator of the autocovariance matrix. Here we choose a less rigorous approach and

⁹ Lanne and Saikkonen (2007) observe that under the null hypothesis the model is unidentified, which implies nonstandard properties of conventional tests such as the likelihood ratio and motivates the quest for alternative tests. Curiously, this problem does not prevent the LR test from being used in applications, see e.g. Lütkepohl and Netšunajev (2014).

regard the existence of ARCH effects, which is attested in the data (see the previous section), as a sufficient proof that the identification conditions mentioned above are fulfilled. This claim is vindicated by the results of our estimation which show an improvement of the likelihood attained with our SVAR-GARCH model as compared with a simple VAR-GARCH estimation (see the Appendix).

The final step of our analysis is DCC estimation. It is known that the loglikelihood of this model can be decomposed in two parts which can be maximized sequentially:

$$L_V(\theta_1) = -\frac{1}{2} \sum_{t=1}^T (n \log 2\pi + 2 \log |D_t| + u_t' A_0' D_t^{-1} D_t^{-1} A_0 u_t) \quad (37)$$

$$L_R(\theta_2) = \frac{1}{2} \sum_{t=1}^T [(D_t^{-1} A_0 u_t)' (D_t^{-1} A_0 u_t) + \log |R_t| + (D_t^{-1} A_0 u_t)' R_t^{-1} (D_t^{-1} A_0 u_t)] \quad (38)$$

The parameter vector θ_1 contains the same parameters of the bivariate GARCH model and indeed a careful look at (37) shows that it is identical to (35), given the definition of D_t in eq. (30). The parameter vector θ_2 contains the unknown parameters of eq. (32), a , b and \bar{Q} . By adding the maximization of L_R as a third step after the two steps described above, we obtain an estimate of the time varying correlation of the structural shocks. All the considerations regarding identification apply to this extension exactly in the same way as discussed above.

5 Results

The analysis is performed on data which are aggregated on a time window of 5 minutes. The first step is to estimate the reduced form VAR. The number of lags is selected as the maximum across a number of information criteria and set equal to 24. We obtain the confidence intervals for the IRFs with a bootstrap procedure performed according to the recursive-design wild bootstrap scheme used by Lütkepohl and Schlaak (2019), which preserves the conditional heteroskedastic innovations of AR models.

From the cumulative IRFs (Fig. 3) we see that demand is much more autocorrelated than price as expected (see Sec. 2), although in both cases the VAR coefficients explain only a small fraction of the overall variance of the dependent variables¹⁰. The lagged cross-effects between price and

¹⁰The values of the adjusted R^2 are respectively 0.004 and 0.019.

demand are either small or non significant, as we could expect from market efficiency considerations. In particular, considering that price variations are expressed in USD cents and the unit of demand is one million Euros, we see that a positive demand shock of 1 million Euros at t causes a (negative) price variation worth as little as -0.0001 USD in the subsequent 2 hours, equivalent to the 0.19% of the sample volatility of price. Instead, a positive EUR/USD price shock worth 1 cent causes an extra demand of up to 57 millions Euros in the subsequent 2 hours, equivalent to the 149% of the sample volatility of demand.

A negligible price impact of past demand shocks is what we expect in an efficient and extremely liquid market, like the one we consider. Instead, a more pronounced impact of past price shocks on demand is what we expect in a market which is populated by speculators who try to react as fast as possible to price shocks but may do so with some delay and in addition are forced by the microstructure of the market to split their trades over different periods. At the same time, it is interesting to observe that the effect of demand on price is significant while the converse is not true. This points to the possibility that the former, smaller effect is linked to microstructural effects, which are much less volatile, while the latter, larger effect is made more volatile by the shifting expectations of heterogeneous speculators over time. This is confirmed by the Granger causality test that shows that the null of no causality can be rejected for z_t but not for Δp_t .

In order to assess the effect of heterogeneity we perform a stability analysis of the VAR estimation using a set of suitably chosen empirical fluctuation processes of the residuals. Within the framework of generalized fluctuation tests, we compute an empirical process that captures the fluctuation in residuals, and for which the limiting processes are known. In this way boundaries can be computed, which are crossed with probability α under the null hypothesis that the coefficients of the VAR model are constant over time. This means that, if the empirical process crosses the boundaries, the null hypothesis should be rejected (at significance level α). In particular, the processes we employ for the analysis are the cumulative sum (CUSUM) or the moving sum (MOSUM) of either the ordinary residuals (OLS) or the recursive residuals (Rec).

The results of this analysis are unequivocal for Δp_t since the null hypothesis is never rejected according to the measures mentioned above (data not displayed). This is in line with the model of Sec. 2 which suggests that price is only indirectly affected by the heterogeneity of speculators through the shifts of demand itself, provided that $\gamma_{0,t}$ is large enough in absolute terms. The results for z_t instead are contradictory as depicted in Fig. 4, which shows that the null hypothesis of constant coefficients for the VAR model is rejected

for recursive residuals, while the opposite holds for the ordinary residuals. Thus the claim that demand is directly affected by the shifting expectations of heterogeneous speculators and their changing participation to the market, which should shift the magnitude of the VAR coefficients over time, is not unequivocally vindicated by the data. This result lends some support to our estimation approach which is based on fixed VAR coefficients.

The timing of the structural breaks should correspond to either the beginning of a shift (MOSUM) or a peak (CUSUM) in the empirical process. In the graph we signal with a vertical line the following critical days (see Sec. 3): March 10 (QE expansion); June 3 (US Jobs announcement); June 24 (Brexit referendum); November 9 (Trump election); December 15 (Euro falls below \$ 1.04)¹¹. We see that the recursive errors converge to identify a potential structural break ahead of the landmark date of the Brexit referendum, mostly coinciding with the announcement of weak job data in the US on June 3.

The residuals of VAR estimation strongly reject the null hypothesis of homoskedasticity for a number of tests. In particular, both series display ARCH effects as required by the identification conditions discussed in Sec. 4. Then we are justified in proceeding to the second step of the estimation as described in the same section and in the Appendix. In order to proceed with the estimation, we suppose that the variances of the structural errors follow a GARCH(1,1) process. The results are summarized in the following equations for the volatility of the structural residuals of the price ($h_{0,t}$) and demand ($h_{1,t}$) equations respectively:

$$h_{0,t} = \frac{0.00011}{(0.00002)} + \frac{0.16961}{(0.02008)} \left[u_{1,t-1} - \left(\frac{0.00467}{(0.00001)} - \frac{1}{233.273} \right) u_{2,t-1} \right]^2 + \frac{0.79030}{(0.02030)} h_{1,t-1} \quad (39)$$

$$h_{1,t} = \frac{75.3427}{(4.50189)} + \frac{0.21843}{(0.01308)} \left(u_{2,t-1} - \frac{233.273}{(0.00000)} u_{1,t-1} \right)^2 + \frac{0.76365}{(0.00003)} h_{2,t-1} \quad (40)$$

All the coefficients are highly significant, as we expected from the mentioned analysis of the VAR residuals. The standardized residuals and squared residuals obtained from the GARCH estimation show no signs of serial correlation.

The economic implication of our estimates of ω and γ_0 is interesting, since it shows that the exchange rate is stabilized by the existence of momentum traders. Indeed, from Eq. (3) we see that the simultaneous price impact of

¹¹This date is depicted only in the CUSUM panels.

demand is $\hat{\omega} - \frac{1}{\hat{\gamma}_0} = 0.00038$. If the overall demand reaction of speculators to a simultaneous price variation had a negative sign, the simultaneous price impact of the same amount of net orders would be at least as great as $\hat{\omega}$, i.e. 10 times larger. The economic rationale of this result is straightforward: thanks to the fact that net orders are instantaneously increased by an increase of the price quoted by market makers, the profits of market makers are maximized with a smaller price increase than if momentum trading was absent.

The IRFs obtained from our estimates are depicted in Fig. (5). We see that the cross effects between price and demand are both positive and significant. The most appropriate comparison is with the cumulative impact of ordinary SVAR estimations. After two hours, the cumulative impact on price of a demand shock worth one unconditional standard deviation (i.e. 65 million Euros based on the GARCH estimates) is equal on average to an increase of 2.06 cents in the EURO/USD price according to the SVAR-GARCH model, which is almost double the 1.11 cents increase predicted by the ordinary SVAR model. Conversely, the cumulative impact on demand of a price shock worth one unconditional standard deviation (i.e. 5,23 cents based on the GARCH estimates) is worth on average 22.7 millions Euros, which is higher than the 19.6 millions predicted by the standard SVAR model.

The last step is the estimation of the conditional correlation of the structural shocks by means of Eq. (38). The results are summarized in the following equations:

$$Q_t = \left(1 - \begin{matrix} 0.01767 \\ (0.00271) \end{matrix} - \begin{matrix} 0.84294 \\ (0.02501) \end{matrix} \right) \begin{bmatrix} 0.99770 & -0,26746 \\ -0,26746 & 1.05377 \end{bmatrix} + \\ + \begin{matrix} 0.01767 \\ (0.00271) \end{matrix} (D_{t-1}^{-1} A_0 u_{t-1}) (D_{t-1}^{-1} A_0 u_{t-1})' + \begin{matrix} 0.84234 \\ (0.02432) \end{matrix} Q_{t-1} \quad (41)$$

$$A_0 = \begin{bmatrix} 1 & -0.00038 \\ -233.273 & 1 \end{bmatrix} \quad (42)$$

$$D_t = \begin{bmatrix} \frac{1}{\sqrt{h_{0,t}}} & 0 \\ 0 & \frac{1}{\sqrt{h_{1,t}}} \end{bmatrix} \quad (43)$$

where $h_{0,t}$ and $h_{1,t}$ are given by Eqs. (39) and (40). The time evolution of the conditional variances $h_{0,t}$ and $h_{1,t}$ and the conditional correlation ρ_t are depicted in Fig. 6. The vertical lines are the same of Fig. 4. We see that price volatility is more clearly related to exogenous events (March 10, June 3, June 24, Nov 9) as we expected, while demand volatility is related to endogenous events like the crossing of a critical support price (December

15). More in general, volatility spikes of the two variables show a tendency to move in the same direction (with a Pearson correlation of 0.46) but do not overlap perfectly. At the same time we see that the unconditional correlation of the structural shocks is negative (-0.26), and the conditional correlation is mostly negative with only some rare exceptions. These results are robust to an alternative estimation performed assuming that the errors follow a Student distribution. Indeed the exogenous events do not appear to have large effects on ρ_t , with the big exception of June 3, when ρ_t becomes strongly positive. Moreover, the movements of ρ_t result to be negatively correlated to those of $h_{0,t}$ and $h_{1,t}$.

A good candidate for explanation of the negative correlation of the structural errors is that the model of Sec. 2 does not take into account the full heterogeneity of demand because it overlooks the role of liquidity / noise traders. The latter in fact are generally assumed to trade with a loss, a circumstance which might help to explain why in our results a positive (negative) price shock might occur along with a negative (positive) demand shock. In particular, as explained above (see Sec. 3), we must take into consideration that many of the transactions we observe on the EBS platform occur because dealers wish to clear the trades of their customers, many of which are likely to be noise traders. These transactions are likely to appear as exogenous shocks in our estimations as they don't result from the speculative behavior we assume in our model.

In order account for the correlation of the structural errors, we make use of the well known generalized impulse response functions proposed by Koop *et al.* (1996) which are defined as follows:

$$\Theta_h = \Phi_h A_0^{-1} \Omega \Sigma \quad (44)$$

where Φ_h is the standard forecast error impulse response (FEIR) function, Ω is the unconditional covariance matrix of the structural errors, $\Sigma \equiv \text{diag}\left(\frac{1}{\sigma_0}, \frac{1}{\sigma_1}\right)$ and σ_0, σ_1 are the unconditional standard deviations of the structural errors. The results are depicted in Fig. 7 where we report only the cross effects between Δp_t and z_t . We see that the cumulative impact of a demand shock on price is still positive although the magnitude of the impact is greatly diminished, since the negative correlation of the two shocks moderates the simultaneous positive reaction of price variations to demand shocks. At the same time the negative correlation of the structural shocks turns the cumulative impact of a price shock on demand to negative, overturning the result of the previous estimation. These results reinforce the interpretation advanced above, namely that the negative correlation of the structural errors stems from an additional source of heterogeneity of demand. Indeed, we see

that the effect of the negative correlation is to make the market more stable since the price adjusts less to demand shocks and the demand curve takes a negative inclination thus favouring the balancing of the market. This stabilizing role is exactly the one assigned to liquidity / noise traders from the market microstructure literature (Evans and Lyons, 2002; King *et al.*, 2013). We will elaborate more on this point in the conclusions.

The cumulative impact depicted in Fig. 7 depends on the unconditional covariance of the structural errors. If we replace the latter with the conditional covariance matrix we obtain a time dependent IRF function $\Theta_{h,t}$. The results of Fig. 8 show that the cross cumulative impact of price and demand shifts dramatically its magnitude and direction over time, as a consequence of the fact that the conditional correlation of the structural errors changes its direction from negative to positive under specific circumstances and in particular when exogenous shocks occur. In particular, we see that the impact of demand shocks on price variations is often negative. As we explain in the next section, this result might be explained as a reaction of market makers to liquidity trading.

6 Conclusions

Using high-frequency transaction data on the EUR/USD market in 2016, we estimate a SVAR-GARCH-DCC model which shows that there is a positive simultaneous effect of the exchange rate variation on demand and vice-versa. After adjusting for the negative correlation of the structural errors, the cumulative effect of a price shock on demand turns out to be negative. In particular, our estimates imply that a positive price shock worth one unconditional standard deviation (which is equal to 5.23 cents based on the GARCH estimates) triggers on average, within the next 2 hours, net additional sell orders worth 4.6 million Euros, which is 12% of the standard deviation of z_t . The cumulative effect of a demand shock on price is instead positive. In particular, a demand shock worth one unconditional standard deviation (which is equal to 65 million Euros based on the GARCH estimates) triggers on average, within the next 2 hours, a positive price variation of 0.59 cents, which represents 12.6% of the standard deviation of Δp_t .

The reduced form VAR coefficients explain only a tiny fraction of the volatility of endogenous variables. This is by no means a surprise because we are dealing with a highly liquid market, where we expect that lagged effects are either negligible or non significant. In particular, a negligible price impact of past demand shocks is what we find. Instead, we find a more pronounced but non significant impact of past price shocks on demand,

which is consistent with heterogeneous speculators who try to react as fast as possible to price shocks but may do so with some delay and in addition are forced by the microstructure of the market to split their trades over time.

The stability analysis performed on the VAR coefficients shows no signs of structural breaks in the price equation. This result is in line with the model of Sec. 2 which suggests that price is only indirectly affected by the heterogeneity of speculators through the shifts of demand itself. Instead the results for the demand equation are different depending on the measure employed. Thus the claim, typical of HAMS, that demand is directly affected by the shifting expectations of heterogeneous speculators and their changing participation to the market, which should shift the magnitude of the VAR coefficients over time, is not unequivocally vindicated by the data. On the other hand, this result lends some support to our estimation approach which is based on fixed VAR coefficients.

Finally, our results suggest that one important source of heterogeneity in demand might be missing from the model described in Sec. 2. It is not easy to justify that the structural price and demand shocks are negatively correlated if we assume that speculators are the only source of demand. Indeed, the common factor influencing at the same time both price and demand shocks is the exogenous flow of information, and any piece of information that is likely to raise / lower price is also likely to raise /lower the demand of speculators. A reasonable explanation is that the model of Sec. 2 overlooks the role of noise traders. From this point of view, it is interesting to observe that allowing for the negative correlation of structural errors makes the market more stable since the price adjusts less to demand shocks and the demand curve takes a negative inclination. This stabilizing role is exactly the one assigned to noise traders from the market microstructure literature (Evans and Lyons, 2002; King *et al.*, 2013), following the original intuition of Black (1986).

Indeed we must take into consideration the fact that many of the transactions we observe on the EBS platform occur because dealers wish to clear the trades of their customers according to the well known “hot potato” scheme. These transactions are likely to appear as exogenous demand shocks in our estimations to the extent that they don’t result from the speculative behavior we assume in our model. The alternative hypothesis that demand shocks might come from better informed traders seems to be rejected by the fact that in the latter case the reaction of a rational market maker would be to raise and not to lower prices as we observe. Instead, lowering (raising) the price is rational for the market maker, if the buyer (seller) is considered to be uninformed, in order to prevent liquidity based speculation to catch on (Vitale , 2000; Jeanne and Rose, 2002). This supposed behavior of the mar-

ket maker would provide an explanation for the result mentioned at the end of Sec. 5, that the impact of a positive demand shock on price variations, although positive on average, turns out to be negative in many periods. We might speculate that in those periods noise trading prevails on the market.

We thank the participants to the International Conference “Minsky at 100”, held in Milan on December 16th-17th 2019, for their useful comments. All the usual disclaimers apply.

A Identification conditions and estimation procedure

It is known that, if the identification conditions based on conditional heteroskedasticity hold, the solution is unique up to a reordering, change of sign and renormalization of the columns of A_0 (Rigobon, 2003; Ehrmann *et al.*, 2011; Lewis, 2018). This means that we need additional information to identify the structural coefficients of the model.

In order to see why this is necessary, let’s start from a generic bivariate simultaneous system:

$$x_t = a y_t + \epsilon_{0,t} \tag{A.1}$$

$$y_t = b x_t + \epsilon_{1,t} \tag{A.2}$$

where x_t and y_t are observed variables and $\epsilon_{0,t}$ and $\epsilon_{1,t}$ unobserved errors. After exchanging the order of the equations, the system can be rewritten equivalently as follows:

$$x_t = \frac{1}{b} y_t + \epsilon_{0,t} \tag{A.3}$$

$$y_t = \frac{1}{a} x_t + \epsilon_{1,t} \tag{A.4}$$

Indeed, since we don’t observe the errors, we cannot distinguish between $\epsilon_{0,t}$ and $\epsilon_{1,t}$, between $\epsilon_{0,t}$ and $\epsilon'_{0,t} = \frac{\epsilon_{0,t}}{a}$ or between $\epsilon_{1,t}$ and $\epsilon'_{1,t} = \frac{\epsilon_{1,t}}{b}$. From the two equivalent formulations we obtain the same reduced form:

$$x_t = \theta_0 y_t + \epsilon_{0,t} \tag{A.5}$$

$$y_t = \theta_1 x_t + \epsilon_{1,t} \tag{A.6}$$

Following Rigobon (2003), we assume that there are two different regimes $s = 0, 1$ for the variance of $\epsilon_{0,t}$ and $\epsilon_{1,t}$. In this case the system (A.1)-(A.2)

satisfies exactly the order condition under the assumption that $\epsilon_{0,t}$ and $\epsilon_{1,t}$ are uncorrelated under each regime. In particular, a and b satisfy the following couple of equations:

$$a = \frac{w_{0,1,s} - w_{1,1,s}b}{w_{0,0,s} - w_{0,1,s}b} \quad s = 0, 1 \quad (\text{A.7})$$

where, under each regime, $w_{0,0,s}$, $w_{1,1,s}$, $w_{0,1,s}$ are the variances and covariance of x_t and y_t , which can be estimated from the data. In order to obtain (A.7), these estimates are equated to their theoretical counterparts which are on their part obtained from the following reduced form solutions:

$$x_t = \frac{\epsilon_{0,t} + \epsilon_{1,t}a}{1 - ab} \quad (\text{A.8})$$

$$y_t = \frac{\epsilon_{0,t}b + \epsilon_{1,t}}{1 - ab} \quad (\text{A.9})$$

From the ordering (A.3)-(A.4) we obtain a different reduced form solution:

$$x_t = \frac{(\epsilon_{0,t}b + \epsilon_{1,t})a}{1 - ab} \quad (\text{A.10})$$

$$y_t = \frac{(\epsilon_{0,t} + \epsilon_{1,t}a)b}{1 - ab} \quad (\text{A.11})$$

Nevertheless, subjecting (A.10)-(A.11) to the same algebraic manipulations of (A.8)-(A.9), it's possible to show that a and b still satisfy the equations (A.7).

Equating the r.h.s. of the system (A.7) we obtain a quadratic equation which yields two distinct real and finite solutions for b when the following holds (Rigobon, 2003):

$$w_{0,0,0}w_{0,1,1} - w_{0,0,1}w_{0,1,0} \neq 0 \quad (\text{A.12})$$

If we solve the equations (A.7) for b we obtain instead the following system:

$$b = \frac{w_{0,0,s}a - w_{0,1,s}}{w_{0,1,s}a - w_{1,1,s}} \quad s = 0, 1 \quad (\text{A.13})$$

Equating the r.h.s of these equations it's easy to check that, if b^* is a solution for (A.7), then $a^* = \frac{1}{b^*}$ is a solution for (A.13). This shows that the two solutions obtained from either (A.7) or (A.13) correspond to the values of the reduced form parameters θ_0 and θ_1 of eqs. (A.5)-(A.6), so that only the latter are identified while a and b are not.

Most of the existing literature addresses this problem by imposing inequality restrictions on the coefficients of the model (Ehrmann *et al.*, 2011;

Herwartz and Lütkepohl, 2014; Lütkepohl and Netšunajev, 2014; Lanne and Luoto, 2020). In our case, we can show that the constraint imposed by the model (22)-(23) on the reduced form coefficients allows the identification of one of the two structural coefficients. In order to illustrate this point, let's write in vector form the two ways in which the reduced form errors of the model (22)-(23) can be written:

$$\begin{cases} u_t = A_0 \epsilon_t \\ A_0 = \begin{bmatrix} 1 & \frac{1}{\gamma_0} - \omega \\ -\frac{1}{\gamma_0} & 1 \end{bmatrix} \end{cases} \quad (\text{A.14})$$

$$\begin{cases} u_t^* = A_0^* \epsilon_t \\ A_0^* = \begin{bmatrix} 1 & -\frac{1}{\gamma_0} \\ \frac{\gamma_0}{1-\gamma_0\omega} & 1 \end{bmatrix} \end{cases} \quad (\text{A.15})$$

Now let's suppose that θ_0 and θ_1 are the reduced form coefficients of either A_0 or A_0^* . This yields the following two systems:

$$\begin{cases} \theta_0 = \omega - \frac{1}{\gamma_0} \\ \theta_1 = \frac{1}{\gamma_0} \end{cases} \quad (\text{A.16})$$

$$\begin{cases} \theta_0 = \frac{1}{\gamma_0} \\ \theta_1 = \frac{\gamma_0}{\gamma_0\omega - 1} \end{cases} \quad (\text{A.17})$$

Let's solve the two systems for the structural parameters. In the first case, we obtain:

$$\begin{cases} \omega = \theta_0 + \frac{1}{\theta_1} \\ \gamma_0 = \theta_1 \end{cases} \quad (\text{A.18})$$

In the second case instead we obtain:

$$\begin{cases} \omega = \theta_0 + \frac{1}{\theta_1} \\ \gamma_0^* = \frac{1}{\theta_0} \end{cases} \quad (\text{A.19})$$

Thus we see that ω is identified independently from the ordering of the system (22)-(23).

In order to obtain the identification conditions similarly to Rigobon (2003), we write explicitly the reduced errors (A.14):

$$u_{0,t} = \frac{\gamma_0 \epsilon_{0,t} + (\gamma_0 \omega - 1) \epsilon_{1,t}}{\gamma_0 (2 - \gamma_0 \omega)} \quad (\text{A.20})$$

$$u_1 = \frac{\epsilon_{0,t} \gamma_0 + \epsilon_{1,t}}{2 - \gamma_0 \omega} \quad (\text{A.21})$$

Considering that $E[\epsilon_{0,t} \epsilon_{1,t}] = 0$, the reduced form covariance matrices are

$$W_s = \frac{1}{(\gamma_0 \omega - 2)^2} \begin{bmatrix} \sigma_{0,s}^2 + \frac{(\gamma_0 \omega - 1)^2 \sigma_{1,s}^2}{\gamma_0^2} & \gamma_0 \sigma_{0,s}^2 + \frac{(\gamma_0 \omega - 1) \sigma_{1,s}^2}{\gamma_0} \\ \cdot & \sigma_{0,s}^2 \gamma_0^2 + \sigma_{1,s}^2 \end{bmatrix} \quad s = 0, 1 \quad (\text{A.22})$$

Since the model satisfies the order condition for identification as above, the estimates of the parameters can be obtained by solving the following six moment conditions:

$$\hat{W}_s \equiv \begin{bmatrix} w_{0,0,s} & w_{0,1,s} \\ \cdot & w_{1,1,s} \end{bmatrix} = W_s \quad s = 0, 1 \quad (\text{A.23})$$

These conditions yield the following couple of equations:

$$\frac{w_{0,0,s} \gamma_0^2 - w_{1,1,s}}{(\gamma_0 \omega - 1)^2 - 1} = \frac{(w_{0,0,s} \gamma_0 - w_{0,1,s}) \gamma_0}{(\gamma_0 \omega - 1) (\gamma_0 \omega - 2)} \quad s = 0, 1 \quad (\text{A.24})$$

If we solve the eqs. (A.24) for γ_0 we obtain, for each regime, two solutions for γ_0 :

$$\gamma_0 = \frac{1}{2} \frac{w_{1,1,s} \omega \pm \sqrt{4w_{0,0,s} w_{1,1,s} - 4w_{0,1,s} w_{1,1,s} \omega + w_{1,1,s}^2 \omega^2}}{w_{0,1,s} \omega - w_{0,0,s}} \quad s = 0, 1 \quad (\text{A.25})$$

If we equate across the two regimes the r.h.s. of each of the two solutions we obtain the following unique solution:

$$\omega = \frac{w_{0,0,0} w_{1,1,1} - w_{0,0,1} w_{1,1,0}}{w_{0,1,0} w_{1,1,1} - w_{0,1,1} w_{1,1,0}} \quad (\text{A.26})$$

If we solve instead (A.24) for ω we obtain the following couple of equations:

$$\omega = \frac{w_{0,0,s} \gamma_0^2 - w_{1,1,s}}{(w_{0,1,s} \gamma_0 - w_{1,1,s}) \gamma_0} \quad s = 0, 1 \quad (\text{A.27})$$

By equating the two r.h.s. of (A.27) we obtain a quadratic equation in γ_0 which, under appropriate conditions on \hat{W}_s , yields two independent solutions for γ_0 ¹². It's clear that the unique solution (A.26) for ω corresponds to the one obtained from the reduced form coefficients above:

$$\omega = \theta_0 + \frac{1}{\theta_1} \quad (\text{A.28})$$

and that the two solutions for γ_0 (A.27) correspond to the two alternative values obtained in the same way:

$$\begin{cases} \gamma_0 &= \theta_1 \\ \gamma_0^* &= \frac{1}{\theta_0} \end{cases} \quad (\text{A.29})$$

So, except in the unlikely case that $\theta_1 = \frac{1}{\theta_0}$, γ_0 is not identified.

In order to distinguish between γ_0 and γ_0^* the sequence of the estimation of the equations matters. In fact, we see that if we wish to estimate γ_0 according to the ordering (A.14) of the system (22)-(23), we have to maximize firstly the likelihood of the second equation, i.e. the demand equation. We must follow instead the reverse order if we wish to obtain the other estimate γ_0^* .

Since we know that γ_0 and γ_0^* yield by construction the same value of the likelihood, we choose to avoid the potential numerical precision problems of estimating ω and γ_0 directly according to one or another of the two orderings (A.14) or (A.15) and begin estimating θ_0 and θ_1 instead. In detail, we take the following steps (in parentheses we report the likelihood obtained at the corresponding step):

1. Estimate the VAR (33) where Δp_t is ranked as first and z_t is ranked as second dependent variable
2. Estimate equation-by-equation the bivariate GARCH model (36) on the reduced form errors of the VAR model, i.e. set $\theta_0 = \theta_1 = 0$ ($L(\hat{\theta}) = -114,441.40$)
3. Use the GARCH parameters obtained under step 2 as starting values and maximize the likelihood of the structural residuals of each equation to obtain the estimates $\hat{\theta}_0, \hat{\theta}_1$ and the updated GARCH parameters ($L(\hat{\theta}) = -110,013.97$)

¹²It's possible to show that (A.26) and (A.27) are obtained also for the alternative reduced errors (A.15).

4. Use the parameter values obtained under step 3 as starting values for the maximization of the bivariate loglikelihood (35) ($L(\hat{\theta}) = -109,983.50$)
5. Compute the estimates $\hat{\omega}$ and $\hat{\gamma}_0$ according to the ordering (A.14): these are the ones reported in the main text, see Eqs. (39) -(40) ($L(\hat{\theta}) = -109,983.50$)¹³
6. Finally, compute the estimates $\hat{\omega}$ and $\hat{\gamma}_0^*$, according to the ordering (A.15) ($L(\hat{\theta}) = -109,983.50$)¹⁴

The results obtained under step 6 are as follows:

$$h_{0,t} = \underset{(0.00002)}{0.00011} + \underset{(0.02008)}{0.16961} \left(u_{0,t-1} - \frac{1}{\underset{(0.00000)}{2629.79}} u_{1,t-1} \right)^2 + \underset{(0.02030)}{0.79030} h_{1,t-1} \quad (\text{A.30})$$

$$h_{0,t} = \underset{(4.50189)}{75.3427} + \underset{(0.01308)}{0.21843} \left(u_{1,t-1} - \frac{1}{\underset{(0.00007)}{0.00467} - \frac{1}{\underset{(0.00000)}{2629.79}}} u_{0,t-1} \right)^2 + \underset{(0.00003)}{0.76523} h_{2,t-1} \quad (\text{A.31})$$

From the comparison of eqs. (A.30)-(A.31) with (39)-(40) we see that the value of $\hat{\omega}$ is identical, as we expected from the previous arguments, while the value of $\hat{\gamma}_0^*$ is much larger than the value of $\hat{\gamma}_0$ reported in the main text.

As explained above, the current literature tries to overcome this source of indeterminacy relying mainly on restrictions imposed on A_0 . For instance, if we deal with a standard demand-supply system, we might impose sign restrictions on the structural coefficients which translate into sign restrictions of the reduced form coefficients, making the former identifiable. In our model, the only theoretical restriction of this sort we might think of is the positivity of the adjustment coefficient in the price equation, i.e.

$$\omega - \frac{1}{\gamma_0} > 0 \quad (\text{A.32})$$

which is strongly supported by our knowledge of how market makers adjust prices. This constraint could have been informative if one of the

¹³ In particular we employ a variety of procedures which yield identical estimates: firstly we compute jointly $\hat{\omega}$ and $\hat{\gamma}_0$ conditioned on the value of the GARCH parameters obtained under Step 4; secondly, we perform a joint maximization of all the parameters; thirdly, we compute jointly $\hat{\omega}$ and $\hat{\gamma}_0$ conditioned on the value of the GARCH parameters obtained under Step 2.

¹⁴We employ the same procedures of the previous step and the estimates are identical across them as well.

two reduced form coefficients was negative and the other was positive. But unfortunately this is not the case with our estimation.

Alternatively, researchers have proposed to identify equations by supposing that the structural shock contributing most to the innovations of a given variable is the shock of interest for that variable (Lewis, 2018). This approach runs the risk of being circular, because it is very likely to end up identifying the residuals of the equation where, say, price is the dependent variable as the structural shock of the price equation in (16), thus circumventing the problem.

The common denominator of the different approaches in the literature is that they do not provide more than plausible arguments to identify the “right” ordering of equations¹⁵. Indeed, if the issue is one of forecasting the response of dependent variables to structural shocks, this identification is unimportant since the IRFs depend only on the reduced form coefficients θ_0 and θ_1 . It might become important in those contexts where the response to a policy shock is of interest, and thus we must know which is the structural error corresponding to a policy shock. Although this is not our case, nevertheless it is somewhat disturbing that we end up with two different estimates of γ_0 which are orders of magnitudes distant one from the other. Thus we venture to advance two plausible arguments in favour of the estimate we report in the main text.

The first argument relies on the magnitude of the coefficients. Taking into consideration that the unconditional standard deviations are respectively $\sigma(\Delta p_t) = 0.04661$ and $\sigma(z_t) = 37.90$ and considering the two alternative estimates of γ_0 , we obtain the following two alternative values for the instantaneous response of z_t to a price shock equal to $\sigma(\Delta p_t)$:

$$\begin{aligned} z_t &= \gamma_0 \times \sigma(\Delta p_t) = 10,84 \\ z'_t &= \gamma'_0 \times \sigma(\Delta p_t) = 121,76 \end{aligned}$$

We see that the first value is worth 28% of $\sigma(z_t)$, while the second one is worth more than three times $\sigma(z_t)$. We consider the second one to be the less plausible value since it entails that the average response of z_t to an average price shock is much larger than the average volatility of z_t itself.

The second, possibly more rigorous, argument relies on the correlation between the structural errors obtained from our model and those obtained from a restricted counterpart which uniquely identifies the ordering of the equations. To begin with, we observe that the only possible zero restriction on A_0 is $\theta_0 = \omega - \frac{1}{\gamma_0} = 0$, while no zero restrictions can be made on A_0^* .

¹⁵For a discussion see Lewis (2018)

Thus the standard SVAR model estimated adopting the restriction $\theta_0 = 0$ uniquely identifies the structural errors of the equation where Δp_t occurs as dependent variable, as those of the price equation, i.e. the first equation in (16). In fact the equivalent reordering of the equations is not possible in this restricted model. Let's denote the structural residuals of the price equation in the standard SVAR estimation just described as $\epsilon_{\Delta p_t}$. In the ordering (A.14) the price equation is the first one, while in the ordering (A.15) it is the second one. Let's denote the structural residuals of the first equation in (A.14) as $\epsilon'_{\Delta p_t}$ and those of the second equation in (A.15) as $\epsilon''_{\Delta p_t}$. Correspondingly, let's denote the structural residuals of the second equation in (A.14) as ϵ'_{z_t} and those of the first equation in (A.15) as ϵ''_{z_t} . Since the structural residuals are the same for both orderings, of course we have that $\epsilon''_{\Delta p_t} = \epsilon'_{z_t}$. Then we can identify the "correct" equation ordering in the unrestricted model as the one which yields the highest correlation of the assumed structural residuals of price equation (i.e. either $\epsilon'_{\Delta p_t}$ or $\epsilon''_{\Delta p_t} = \epsilon'_{z_t}$) with the structural residuals which are for sure those of the price equation in the restricted model ($\epsilon_{\Delta p_t}$). The values of the Pearson coefficients are respectively $\rho(\epsilon_{\Delta p_t}, \epsilon'_{\Delta p_t}) = 0.95$ and $\rho(\epsilon_{\Delta p_t}, \epsilon''_{\Delta p_t}) = 0.02$ suggesting that the more reliable estimate of γ_0 is the one reported in the main text. Indeed the estimate of γ_0 which is obtained from the restricted SVAR estimation is very close to the one in Eq. (40).

References

- Black, F. (1986), Noise, *The Journal of Finance* 41: 529–43.
- Brock, W. A., Hommes, C. A. (1998), Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control*, Volume 22, Issues 8–9, 1235-1274.
- Cont, R. (2001), Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, 1, 223-236.
- Cragg J. G., Donald S. G. (1996), On the Asymptotic Properties of LDU-Based Tests of the Rank of a Matrix, *Journal of the American Statistical Association*, 91:435, 1301-1309
- Daniélsson, J., Love, R. (2006), Feedback trading. *Int. J. Fin. Econ.*, 11: 35-53.
- Dow J., Gorton G. (2008), Noise Traders. In: Palgrave Macmillan (eds) *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, London

- Ehrmann, M., Fratzscher, M. and Rigobon, R. (2011), Stocks, bonds, money markets and exchange rates: measuring international financial transmission. *J. Appl. Econ.*, 26: 948-974.
- Engle, R. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics*, 20(3), 339-350.
- Evans, M. D. D., Lyons R. K. (2002), Order Flow and Exchange Rate Dynamics, *Journal of Political Economy*, vol. 110, no. 1, 2002, pp. 170–180.
- Herwartz, H., Lütkepohl, H. (2014). ‘Structural vector autoregressions with Markov switching: combining conventional with statistical identification of shocks’, *Journal of Econometrics*, Vol. 183, pp. 104–116.
- Herwartz, H., Lange, A. and Maxand, S. (2019), Statistical Identification in Svars - Monte Carlo Experiments and a Comparative Assessment of the Role of Economic Uncertainties for the US Business Cycle. CEGE Discussion Paper 375- July.
- Jeanne, O., Rose, A. (2002). Noise Trading and Exchange Rate Regimes. *The Quarterly Journal of Economics*, 117(2), 537-569.
- Kilian, L., Lütkepohl, H. (2017). Identification by Heteroskedasticity or Non-Gaussianity. In *Structural Vector Autoregressive Analysis (Themes in Modern Econometrics)*, pp. 491-531). Cambridge: Cambridge University Press.
- King, M. R., Osler, C. L., Rime, D. (2013), The market microstructure approach to foreign exchange: Looking back and looking forward, *Journal of International Money and Finance*, Volume 38, pp. 95-119.
- Koop, G., Pesaran, M. H., Potter, S. M. (1996). ”Impulse response analysis in nonlinear multivariate models,” *Journal of Econometrics*, vol. 74(1), pages 119-147.
- Lanne, M., Saikkonen, P. (2007) A Multivariate Generalized Orthogonal Factor GARCH Model, *Journal of Business & Economic Statistics*, 25:1, 61-75.
- Lanne, M., Lütkepohl, H., Maciejowski, K. (2010), Structural vector autoregressions with Markov switching, *Journal of Economic Dynamics and Control* Volume 34, Issue 2, February, pp. 121-131.

- Lanne, M., Luoto, J. (2020), Identification of Economic Shocks by Inequality Constraints in Bayesian Structural Vector Autoregression. *Oxf Bull Econ Stat*.
- Lewis, D. J. 2018, Identifying shocks via time-varying volatility, Staff Reports 871, Federal Reserve Bank of New York, revised 01 May 2019.
- Lütkepohl, H. (2005), *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin.
- Lütkepohl, H., Netšunajev, A. (2014), Disentangling Demand and Supply Shocks in the Crude Oil Market: how to Check Sign Restrictions in Structural VARS. *J. Appl. Econ.*, 29: 479-496.
- Lütkepohl, H., Milunovich, G. (2016), Testing for identification in SVAR-GARCH models, *Journal of Economic Dynamics & Control* 73 (2016) 241–258.
- Lütkepohl, H., Schlaak, T. (2019). "Bootstrapping impulse responses of structural vector autoregressive models identified through GARCH" *Journal of Economic Dynamics and Control*, vol. 101(C), pp. 41-61.
- Madhavan, A., Smidt, S. (1993), An Analysis of Changes in Specialist Inventories and Quotations, *The Journal of Finance*, 48(5), pp. 1595-1628.
- Milunovich, G., Yang, M. (2013). On identifying structural VAR models via ARCH effects. *Journal of time series econometrics*, 5(2), 117-13
- Osler, C.L. (2003), Currency Orders and Exchange Rate Dynamics: An Explanation for the Predictive Success of Technical Analysis. *The Journal of Finance*, 58: 1791-1819.
- Osler, C. L. (2005), Stop-loss orders and price cascades in currency markets, *Journal of International Money and Finance*, Volume 24, Issue 2, pp. 219-241.
- Primiceri, G. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *The Review of Economic Studies*, 72(3), 821-852.
- Ramiah, V., Xu, X., Moosa, I. A. (2015). "Neoclassical finance, behavioral finance and noise traders: A review and assessment of the literature," *International Review of Financial Analysis*, vol. 41(C), pages 89-100.
- Rigobon, R. (2002), The Curse of non-investment Grade Countries, *Journal of Development Economics*, 69, pp. 423-449.

- Rigobon R. (2003), Identification through Heteroskedasticity, *The Review of Economics and Statistics*, 85(4), pp. 777-792.
- Sentana, E., Fiorentini, G. (2001), Identification, estimation and testing of conditionally heteroskedastic factor models, *Journal of Econometrics*, Volume 102, Issue 2, pp. 143-164.
- Vitale, P. (2000), Speculative noise trading and manipulation in the foreign exchange market *Journal of International Money and Finance* 19 (5), 689-712.
- Vitale, P. (2007), A Guided Tour of the Market Microstructure Approach to Exchange Rate Determination. *Journal of Economic Surveys*, 21: 903-934.
- Weber, E. (2010), Structural Conditional Correlation, *Journal of Financial Econometrics*, Vol. 8, No. 3 , 392-407.
- Van Der Weide, R. (2002), GO-GARCH: A multivariate generalized orthogonal GARCH model (2002) *Journal of Applied Econometrics*, 17 (5), pp. 549-564.
- Zeileis, A., Leisch, F., Hornik, K., Kleiber, C. (2002). strucchange: An R Package for Testing for Structural Change in Linear Regression Models. *Journal of Statistical Software*, 7(2), 1 - 38.
- Zhou, B. (1996). High-Frequency Data and Volatility in Foreign-Exchange Rates. *Journal of Business & Economic Statistics*, 14(1), 45-52.

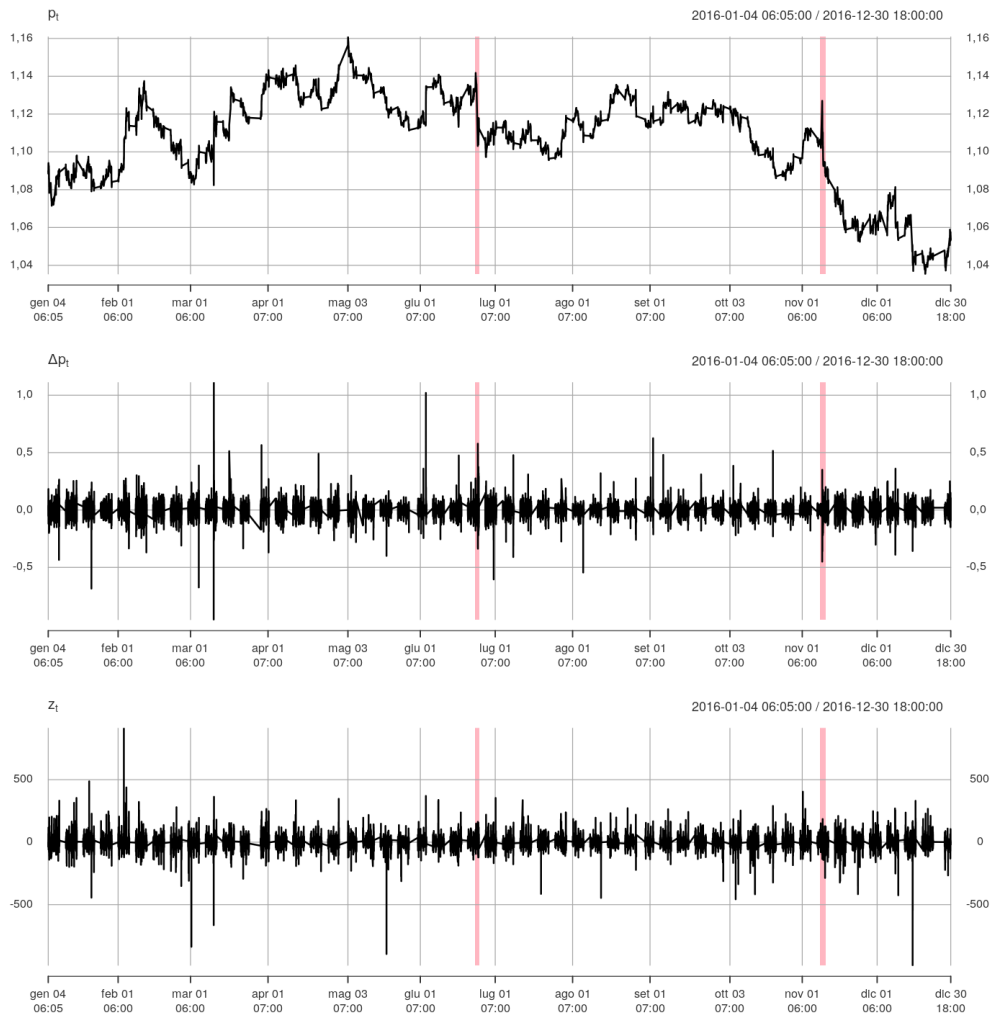


Figure 2: Time series plots




Figure 3: VAR cumulative impulse response functions

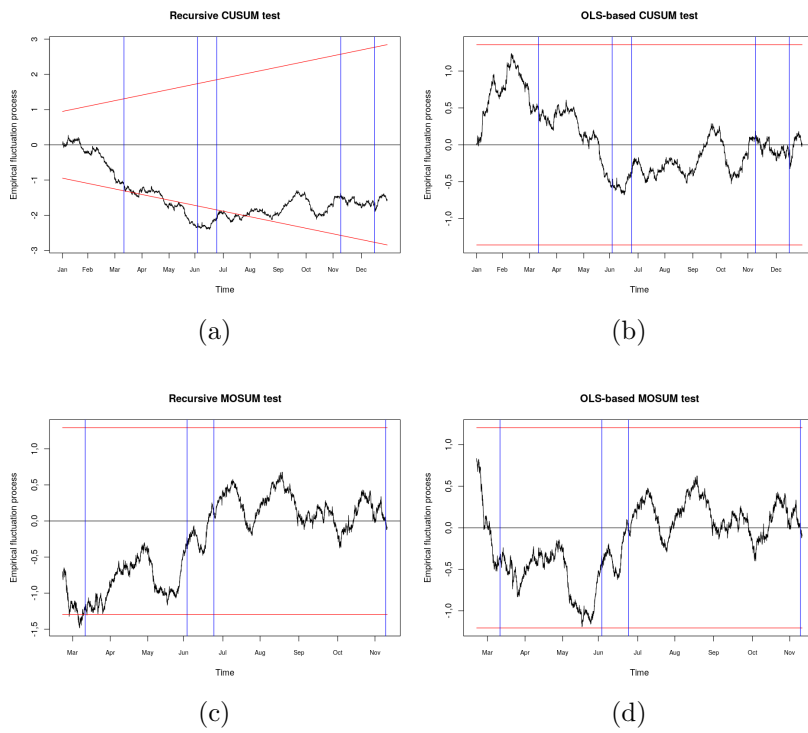


Figure 4: Empirical fluctuation processes of the VAR residuals for z_t . The boundaries (red horizontal line) are drawn for $\alpha = 0.05$. The vertical blue lines correspond to a set of “critical” days as described in the text.




Figure 5: SVAR-GARCH cumulative impulse response functions

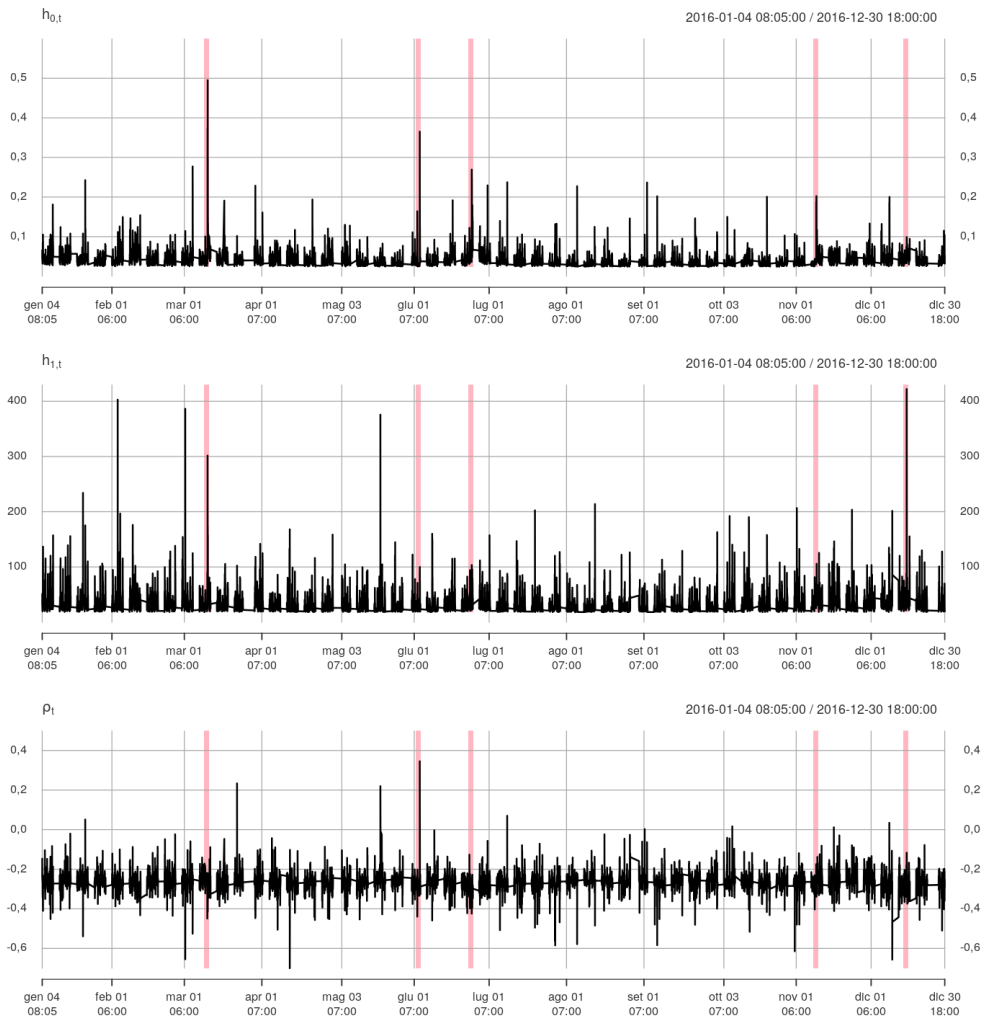


Figure 6: GARCH-DCC filters. The vertical red lines correspond to a set of “critical” days as described in the text.




Figure 7: SVAR-DCC-GARCH cumulative impulse response functions (only cross effects displayed)




Figure 8: Time dependent IRF function $\Theta_{h,t}$ with $h = 24$ (only cross effects displayed). The vertical red lines correspond to a set of “critical” days as described in the text.